

JNMM

Journal of Nuclear Materials Management

- Introduction to the JNMM Special Issue on Open Source and Geospatial Information Analysis 4
Jacques Baute
- A Note From The Editors of The Special Issue 8
Zoe N. Gastelum, Joshua Rutkowski, Yana Feldman
- A Visual Atlas on Strategic Trade 10
Cristina Versino, Simone Cagno, Peter Heine, Julie Carrera
- Automated Processing of Open Source Information for Nonproliferation Purposes 21
Ian J. Stewart, Alexander Lee, Ahmed ElGebaly
- Inferring the Operational Status of Nuclear Facilities with Convolutional Neural Networks to Support International Safeguards Verification 37
Zoe N. Gastelum, Timothy M. Shead
- Site Monitoring with Sentinel-1 Dual Polarization SAR Imagery Using Google Earth Engine 48
Joshua Rutkowski, Morton J. Canty, Allan A. Nielsen
- Enhancing the Geospatial Exploitation System within the IAEA Department of Safeguards 60
Antero Keskinen, Jacques Baute, Mark Carey, Jonetta Ng, Fatjon Ujkani
- Toward a Multimodal-Deep Learning Retrieval System for Monitoring Nuclear Proliferation Activities 68
Yana Feldman, Margaret Arno, Carmen Carrano, Brenda Ng, Barry Chen
- A Definitive Pipeline to Display Multivariate Geotechnical Data in a 3D Environment 81
James B.A. Palmer, J. Stevanović
- Geo-based Data Integration (GDI) in the IAEA Department of Safeguards 97
M. Barletta, A. Yeoh, A. Marques-Patricio, N. Sumarokov, S. Robb, M. Fowler, M. Carey
- The IAEA's Physical Model: Fine-Tuning Nuclear Fuel Cycle Understanding for Robust State-Level Safeguards 107
Brian D. Boyer, James Disser, Sebastien Richet, Daniel Gagne, Stephanie Poirier, Jeremy Whitlock, Agatha Walczak-Typke, Claude Norman

SPECIAL
ISSUE
on Open Source
and Geospatial
Information Analysis

Technical Editor
Markku Koskelo
Assistant Technical Editor
Brian Boyer

Technical Editor Emeritus
Dennis Mangan

Associate Editors
Sarah Frazar, Education and Training
Alicia Swift,
Facilities Operations
Irmgard Niemeyer and Carrie
Mathews, International Safeguards
Louise Worrall and Rian Bahrn,
Materials Control and Accountability
Leslie Fishbone, Nonproliferation
and Arms Control
Felicia Durán, Nuclear Security and
Physical Protection
Glenn Abramczyk, Packaging,
Transportation and Disposition

Book Review Editor
Mark L. Maiello

Taking the Long View Editor
Jack Jekowski

INMM Executive Committee
Cary Crawford, President
Susan Pepper, Vice President
Chris Pickett, Secretary
Robert Curl, Treasurer
Corey Hinderstein, Immediate Past President

Members At Large
Michael Baker
Irmgard Niemeyer
Therese Renis
Teri Leffers

Advertising Contact
Amy Chezem
1120 Route 73, Suite 200
Mt. Laurel, NJ 08054 USA
Phone: +1-856-380-6813
Fax: +1-856-439-0525
Email: achezem@inmm.org

JNMM (ISSN 0893-6188) is published four times a year by the Institute of Nuclear Materials Management Inc. The Institute of Nuclear Materials Management (INMM) is an international professional society dedicated to development and promulgation of practices for the safe, secure and effective stewardship of nuclear materials through the advancement of scientific knowledge, technical skills, policy dialogue, and enhancement of professional capabilities.

DIGITAL SUBSCRIPTION RATES:
Annual (United States, Canada, and Mexico) \$200 for individuals. Institutional subscriptions are \$600 per year. Single copy of the proceedings of the Annual Meeting (United States and other countries) \$200. Send subscription requests to JNMM, 1120 Route 73, Suite 200, Mt. Laurel, NJ 08054 USA. Make checks payable to INMM.

DISTRIBUTION and delivery inquiries should be directed to JNMM, 1120 Route 73, Suite 200, Mt. Laurel, NJ 08054 USA, or contact INMM Headquarters at +1-856-380-6813; fax, +1-856-439-0525; or email, inmm@inmm.org.

Opinions expressed in this publication by the authors are their own and do not necessarily reflect the opinions of the editors, Institute of Nuclear Materials Management, or the organizations with which the authors are affiliated, nor should publication of author viewpoints or identification of materials or products be construed as endorsement by this publication or by the Institute.

© 2018 Institute of Nuclear Materials Management

Topical Papers

Introduction to the JNMM Issue on Open Source and Geospatial Information Analysis (OSGI) 4
Jacques Baute

A Note From The Editors of The Special Issue 8
Zoe N. Gastelum, Joshua Rutkowski, Yana Feldman

A Visual Atlas on Strategic Trade 10
Cristina Versino, Simone Cagno, Peter Heine, Julie Carrera

Automated Processing of Open Source Information for Nonproliferation Purposes 21
Ian J. Stewart, Alexander Lee, Ahmed ElGebaly

Inferring the Operational Status of Nuclear Facilities with Convolutional Neural Networks to Support International Safeguards Verification 37
Zoe N. Gastelum, Timothy M. Shead

Site Monitoring with Sentinel-1 Dual Polarization SAR Imagery Using Google Earth Engine 48
Joshua Rutkowski, Morton J. Canty, Allan A. Nielsen

Enhancing the Geospatial Exploitation System within the IAEA Department of Safeguards 60
Antero Keskinen, Jacques Baute, Mark Carey, Jonetta Ng, Fatjon Ujkani

Toward a Multimodal-Deep Learning Retrieval System for Monitoring Nuclear Proliferation Activities 68
Yana Feldman, Margaret Arno, Carmen Carrano, Brenda Ng, Barry Chen

A Definitive Pipeline to Display Multivariate Geotechnical Data in a 3D Environment 81
James B.A. Palmer, J. Stevanovic

Geo-based Data Integration (GDI) in the IAEA Department of Safeguards 97
M. Barletta, A. Yeoh, A. Marques-Patricio, N. Sumarokov, S. Robb, M. Fowler, M. Carey

The IAEA's Physical Model: Fine-Tuning Nuclear Fuel Cycle Understanding for Robust State-Level Safeguards 107
Brian D. Boyer, James Disser, Sebastien Richet, Daniel Gagne, Stephanie Poirier, Jeremy Whitlock, Agatha Walczak-Typke, Claude Norman

Institute News

President's Message 2

Technical Editor's Note 3

Departments

Author Submission Guidelines 117

Calendar x117

Mission Statement

The Institute of Nuclear Materials Management is dedicated to the safe, secure and effective stewardship of nuclear materials and related technologies through the advancement of scientific knowledge, technical skills, policy dialogue, professional capabilities, and best practices.



President's Message

Cary Crawford
INMM President



Another INMM Annual Meeting has come and gone and we are already gearing up for the 2019 Annual Meeting in Palm Desert, California (USA)! For those who were able to attend in Baltimore this past July, it was a pleasure seeing you and having the opportunity to network and hear about the exciting areas of research you are engaged in. A summary of this year's annual meeting will be highlighted in the next JNMM later this year, so stay tuned for that issue.

As we transition into the new Fiscal Year for the INMM, I would especially like to thank outgoing President Corey Hinderstein for her leadership of the INMM these past four years both as Vice President and President. While she will continue serving on the Executive Committee for the next two years as the Immediate Past President, it's important to note Corey's initiative in 2016 in leading us in a strategic planning effort that has provided insight, focus, and direction for the future of the INMM. As I was serving as a Member-at-Large at the time this was initiated, it was

refreshing to learn that Corey and I were strongly aligned in the desire to not only claim, "The Institute of Nuclear Materials Management will be the leading international professional society for the stewardship of nuclear materials and related technologies to enhance global security" (our vision statement), but to truly accomplish that vision.

While this vision statement typically leads you to notice the "leading society" aspect, it's also important to stress the "international" aspect. Over the past sixty years, the INMM has expanded our reach from what was originally a small core group of US experts in nuclear Safeguards and Security to include more than 40 chapters in 14 countries. This has led to rich international collaborations, an increasingly international leadership team, and several strong partnerships to include WINS, ESARDA, WNTI, and others. Take a look at the INMM Global Presence map on our website (<https://www.inmm.org/Join-INMM/Why-Join>) and you can see how broad our reach truly is. While there is still

room to grow, I believe we are on the right track and would encourage continued ideas on how to further meet the needs of unreached communities in our field as well as how to enhance our impact in your area.

Whether it's holding chapter workshops in your region of the world, expanding our communications and social media presence, bringing educational opportunities to our meetings, or our current plans to try to bring the Annual Meeting outside the United States on a periodic basis, we would love to hear how we can serve your needs better. We welcome your feedback, and even more, your participation on our strategic planning efforts, leadership, and administration of the institute. Please feel free to reach out to any of our leadership with your ideas.

Until then, please enjoy this latest special issue of the JNMM on Open Source and Geospatial Information Analysis (OSGI).



The Importance of Special Issues

Markku Koskelo
JNMM Technical Editor

The JNMM editorial team is pleased to be able to publish this special issue that was put together by the INMM International Safeguards Technical Division's Open Source and Geospatial Information Working Group on the analysis of open source and geospatial data. The issue contains an introduction to the topic by Jacques Baute, Director, Division of Information Management, Department of Safeguards, at the International Atomic Energy Agency in Vienna. I will not try to summarize the content of the various contributed technical papers since it is very competently covered by Zoe Gastelum, Joshua Rutkowski, and Yana Feldman in their

more detailed introduction to this special issue.

Besides writing this one article for this issue, I want to recognize Zoe and Joshua and Yana who did the hard work of coordinating everyone's efforts. This included arranging for each of the manuscripts to be peer reviewed before submitting the collection of the papers to the JNMM editorial staff. Putting together a special issue is a large undertaking and I am grateful for their effort to make this happen. Special issues are a wonderful reference on a specific topic because they highlight the various aspects of the topic, including some of the latest research on the problems that are yet to be

solved. I hope you will enjoy this one.

I am looking forward to seeing the next special issue in the near future. It is on a different topic and already in the works based on a workshop that was held at PNNL last spring. I am also happy to report that we have several manuscripts that have gone through the peer review process and are ready to be published in the next issue. The next issue will also include material from the INMM Annual Meeting in Baltimore this summer.

Should you have any comments or questions, feel free to contact me.

JNMM Technical Editor, Markku Koskelo can be reached at mkoskelo@aquilagroup.com

DEPARTMENT OF ENERGY NATIONAL NUCLEAR SECURITY ADMINISTRATION

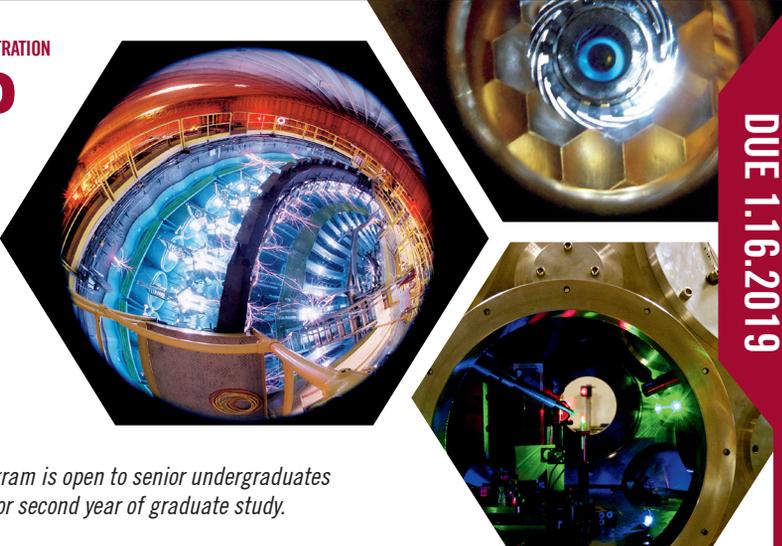
STEWARDSHIP SCIENCE GRADUATE FELLOWSHIP



The DOE NNSA SSGF program is open to senior undergraduates or students in their first or second year of graduate study.

The Department of Energy National Nuclear Security Administration Stewardship Science Graduate Fellowship (DOE NNSA SSGF) program provides outstanding benefits and opportunities to students pursuing a Ph.D. in areas of interest to stewardship science, such as **properties of materials under extreme conditions and hydrodynamics, nuclear science, or high energy density physics**. The fellowship includes a 12-week research experience at Lawrence Livermore National Laboratory, Los Alamos National Laboratory or Sandia National Laboratories.

This equal opportunity program is open to all qualified persons without regard to race, gender, religion, age, physical disability or national origin.



**APPLICATIONS
DUE 1.16.2019**

www.krellinst.org/ssgf

BENEFITS

- + \$36,000 yearly stipend
- + Payment of full tuition & required fees
- + \$1,000 yearly academic allowance
- + Yearly program review
- + 12-week research practicum
- + Renewable up to four years







Introduction to the JNMM Issue on Open Source and Geospatial Information Analysis

Jacques Baute

*Director, Division of Information Management, Department of Safeguards,
International Atomic Energy Agency*



Introduction

Over a quarter of a century ago, in 1991, the discovery of Iraq's clandestine program made the international community realize that an approach to nuclear verification that focused only on verifying declared nuclear materials had a major limitation. For about a decade, an undeclared series of activities, ranging from research and development to nuclear material production, many of which were supported by the construction of undeclared facilities and installations, were conducted in parallel to the declared activities subject to verification under Iraq's 1972 Comprehensive Safeguards Agreement under INFCIRC/253.

The International Atomic Energy Agency (IAEA) embarked on a comprehensive review of its verification approaches with its "Programme 93 + 2," aimed at strengthening the effectiveness and improving the efficiency of IAEA Safeguards. An essential change identified to address the existing weaknesses was to abandon a verification approach addressing solely the correctness of state declarations and to include all measures necessary for the IAEA to assess the completeness of state declarations — that is, the absence of undeclared nuclear material and activities.

By May 1997, the Board of Governors had approved the "Model Protocol Additional to the Agreement(s) between State(s) and the International Atomic Energy Agency" (referred to as the Additional Protocol [AP]), providing the IAEA with broader access to information through augmented state declarations and better physical access for inspectors in the field.

Over the past two decades, the collection and analysis of safeguards-relevant Open Source and Geospatial Information (OSGI) have become undisputable contributors to the process of drawing safeguards conclusions, particularly with regard to the assessment of the completeness of states' declarations. At a time when more information is available more quickly than ever before, this special issue of the *JNMM* is a timely contribution.

Lessons Learned from Historical Cases

The discovery of Iraq's 1980s undeclared program was a wakeup

call for the international community as far as the prevention of nuclear proliferation is concerned. The mandate that the IAEA was provided, as defined by United Nations Security Council (UNSC) Resolution 687 (1991) and other relevant resolutions, allowed the development of approaches and methodologies, many of which were unexplored before. These measures included the provision of declarations covering not only nuclear material and related facilities, but also all specially designed and dual-use items listed in the Ongoing Monitoring and Verification Plan Annex 3 approved by UNSC Resolution 715 (1991).

It quickly became apparent that had the IAEA included safeguards-relevant open source information collection in the 1980s, a number of indicators could have been collected in relation to the existence of undeclared activities in Iraq. The Director General's Standing Advisory Group on Safeguards Implementation considered as early as 1992 the issue of "the use of information from the public press." Consequently, the IAEA established soon thereafter an open source team aimed at collecting safeguards-relevant information that would serve as a contribution to the simultaneously established state evaluation process.

An interesting workshop titled "Open Sources Indicators of Iraq's Nuclear Weapons Program Prior to the Outbreak of War on 17 January 1991" took place in September 2000 at the IAEA headquarters. That workshop allowed comparison of what was by then known about Iraq's past program (later comprehensively documented in report S/1997/779 of 8 October 1997 by the Director General to the UNSC) and what could have been found in open sources. This workshop highlighted the power but also the limitations of open sources with regard to, for instance, their credibility.

The capabilities presented by remote sensing or overhead imagery were also quickly understood thanks to verification-relevant images from U-2 high altitude planes made available to the United Nations in the context of the UNSCR 687 mandate. Such high-resolution images were essential for the preparation of field activities before reaching a site for the first time. In the 1990s, commercially available satellite images did not have sufficient



resolution to allow a proper assessment of relevant sites' infrastructure. By the end of 1999, and with the availability of the first commercial imagery with better than 1 m resolution, the IAEA Iraq Action Team started to use such new images. This was quickly followed by the Department of Safeguards for other states.

Given the ability to implement wide-area environmental monitoring in Iraq with the survey of water bodies and the associated collection of water, sediments, and biota samples, and the pilot establishment of a network of atmospheric samples, the use of a geographic information system became a valuable tool to correlate possible findings with known industrial infrastructures.

The Libya revelations in December 2003 and the discovery of a "nuclear black market" highlighted the need for the safeguards community to reassess its information collection approach to include the effects of globalization. It became obvious that assessments of the declarations' completeness had to account for the possible shortcuts that an undeclared proliferation program could contemplate through external contributions, including from nonstate actors. It was actually quite easy to jump start trade analysis competence in the Department of Safeguards through the transfer of the Export-Import team of the Iraq Action Team (earlier renamed the Iraq Nuclear Verification Office), which had been established to fulfill the IAEA's responsibility of reviewing all of Iraq's Oil for Food program contracts.

Over the years, these major changes were followed by incremental improvements generated by the use of OSGI to contribute to all state evaluations.

Today's Essential Contribution of OSGI

As explained to member states in the context of the communication associated with the state-level concept and captured in the IAEA document GOV/2014/41, it is recognized that in drawing safeguards conclusions, the IAEA relies on three pillars of information:

- Information provided by the state, including the declarations associated with their obligations, with or without an AP
- Information from IAEA Safeguards activities resulting from inspections, design information verification, complementary access when appropriate, and the subsequent evaluations (material balance evaluation, environmental sampling evaluation, etc.)
- Other safeguards-relevant information

Open sources, including satellite imagery, represent the overwhelming majority of the other safeguards-relevant information,

substantially complementing the two other sources of information. Listed below are a few illustrations of how such information contributes to safeguards verification, with the view to highlight the importance of constantly improving its acquisition and use.

For states with significant nuclear fuel cycles, the assessment of the correctness of declarations or the preparation of field activities benefits from the collection and analysis of government and operator information available in the public domain (e.g., annual reports from nuclear facilities or regulatory authorities). Scientific and technical literature review offers significant support to the technical assessment of nuclear fuel cycle research and development activities. Acquisition path analysis (a component of state evaluation essential in establishing the technical objectives of verification activities for a state) is, for instance, substantially supported by the examination of trade publications, individual company websites, and world trade-related data. The analysis of satellite images provides an effective means to assess the nature, extent, and technical connections of nuclear sites. Satellite imagery also provides unique "remote access" to sites in states where field activities are restricted.

For states with small quantity protocols, in which information coming from the states and from verification activities conducted by the IAEA is limited, open source information, including satellite imagery, combined with information provided by other states as part of their obligations, becomes the main tool to assess the states' compliance with their safeguards agreements.

Past Evolution and the Way Forward

Steady improvement took place with regard to strengthening the effectiveness and efficiency of safeguards verification through the use of OSGI between the early initiatives to address the unexpected challenges of Iraq and Libya and the current environment that supports safeguards conclusions. Remaining ahead of the curve in terms of information management is a systematic process with five specific aspects:

- Identifying new safeguards-relevant data sources
- Enhancing methodologies to more efficiently extract safeguards-relevant information from the data collected
- Bringing additional competences — that is, subject matter experts to areas significant to safeguards verification, including advanced nuclear fuel cycle technologists — into the OSGI analytical process
- Improving processes in the areas of collection, evaluation, information dissemination, and utilization of collaborative analysis that is required to draw sound



safeguards conclusions

- Experimenting with and successfully implementing new information technology tools to enhance analytical capabilities

Given that the IAEA's financial resources are primarily focused on the implementation of its legal obligations, much of the IAEA's progress achieved in OSGI collection and analysis would not have occurred without the contributions of Member States Support Programmes, which build on work conducted in a variety of research centers, laboratories, and universities, and through the provision of Cost Free Experts to assist in the IAEA's internal reflections.

The rapid development of the Internet has provided access to an amount of information sources that presents both an opportunity due to its variety in covering many domains of safeguards interest and also challenges due to the volume of information, the variety of formats, and the language diversity. Much of the safeguards-relevant OSGI can be found electronically with both past and present documentation being published online. The extent of historical information released online is an asset from both a research and development and an overhead imagery standpoint. Over a decade ago, automatic collection based on safeguards-relevant keywords was implemented, with the intent to provide analysts with an improved environment for daily news monitoring and for the preparation of their contributions to state evaluation. The recent effort to better exploit multimedia (photos, audio, videos) is expected, when mature, to similarly improve the effectiveness and efficiency of open source information analysis for state evaluation.

Integrating and properly structuring safeguards-relevant data became a necessity to support the identification of inconsistencies, which are essential triggers for follow-up actions at the heart of the state evaluation process. In this context, the introduction of the Physical Model in 1998 provided a common framework for analysts to structure collection and analysis related to the nuclear fuel cycle. Link analysis tools, introduced in the department after Libya's revelations in 2003, and other data visualization tools also assist analysis. While these tools help analysts assess complex relationships, significant effort can be required to preprocess unstructured data.

A number of initiatives over a decade have sought to incorporate automatic entity extraction from unstructured data. This work continues to be explored both by the IAEA and research entities associated with Member State Support Programmes. The prospect of contributions from machine learning remains

promising. However, given the sophistication of the contextual understanding that is necessary to fully assess the safeguards relevance of a given piece of information, it is anticipated that experienced analysts will remain in the loop for the foreseeable future. The specialized terminology related to the nuclear fuel cycle and proliferation issues, the contextual assessment of OSGI, and the diverse source languages will likely require additional investments to help the algorithms reliably achieve satisfactory results.

With regard to addressing globalization, new approaches seem promising to better master the understanding related to transfers of nuclear materials and scientific collaborations. For example, one area of research is on the use of Distributed Ledger Technology (DLT). The assessment of this technology is in its early stages and is primarily focused on developing use cases in which DLT can be gainfully employed. Implementation in the area of safeguards verification will certainly need support by motivated external contributors.

In the geospatial area, georeferencing of all safeguards-related entities would be quite useful, starting with state declarations. Unfortunately, the legally agreed formats do not include such attributes. A range of IAEA staff, from imagery analysts to inspectors, will benefit from improved access to such information. However, there currently is no efficient way to automatically generate georeferencing information for the wide range of generally unstructured information available to the IAEA.

Until recently, the main obstacle to the amount of satellite images being processed by a team of IAEA dedicated analysts was financial (the actual cost of images) and the limited number of commercial images associated with a given site (the repeat rate). Now there are also potential challenges associated with the new approaches provided by high revisit rate constellations of sensors. If these constellations are assessed to be of added substantive value for safeguards issues, analysts may require automatic preparation and extraction of the information of interest from the large volume of data (e.g., through change detection able to reliably separate man-generated changes from those resulting from satellite parameters, sun position, and weather conditions). New sensor sources and associated methodologies utilized by the IAEA have recently included Synthetic Aperture Radar, which provides certain additional information that assists in the state evaluation process. More progress is expected there, as well as from the potential for commercial availability of higher resolution thermal infrared and hyperspectral images.



Conclusion

Addressing the IAEA's responsibility with regard to its contribution to the prevention of nuclear weapons proliferation relies heavily on its ability to collect and evaluate all safeguards-relevant information. Within this safeguards-relevant information, the support that OSGI has provided through the years — in particular with regard to effectively assessing the completeness of state declarations, but also in efficiently evaluating their correctness — has been instrumental to the strengthening of safeguards.

The constant expansion of information sources and the rapid advance of associated methodologies and tools (e.g., data analytics) for exploiting the information present an opportunity as well as a challenge for the safeguards community that can be addressed only through collaborative means. In other words, our community needs to ensure that no relevant information is missed, while at the same time providing a thorough assessment of the possible added value of any new technology and the estimation of the

associated effort to make it successful. It is crucial to utilize the highly constrained resources that are available to the IAEA in the most optimal manner.

After over a quarter-century dedicated to nuclear verification, with a strong focus on the information angle, I wholeheartedly hope that this *JNMM* special issue will motivate many, including young professionals, to explore the ways nuclear nonproliferation can be strengthened through the further development of the contribution of OSGI.

Acknowledgments

The assistance to the progress made by the IAEA in the area of OSGI over a quarter of a century by so many Member State Support Programmes, foundations, academics, and individuals has been so essential that it would have been unfair to list a few and impossible to list all. We address our highest gratitude to all of those who have contributed.



A Note From the Editors of the Special Issue

Zoe Gastelum

Sandia National Laboratories, Albuquerque, NM, USA

Joshua Rutkowski

Jülich Research Center, Jülich, Germany

Yana Feldman

Lawrence Livermore National Laboratory, Livermore, CA, USA



Zoe N. Gastelum



Joshua Rutkowski



Yana Feldman

The information age. Big data. All-source analysis. No matter its name or context, the field of data analytics is going through a period of tremendous technical advancement and increased visibility. The ongoing transformation directly impacts the methods for data analytics (also known as information analysis) in support of nuclear nonproliferation. The unique cross-section of skills required for nonproliferation data analysis spans international relations, law, nuclear engineering, physics, geography, and now data science. The nonproliferation information analysis community is necessarily a hybrid of researchers who can navigate both technical and policy dimensions of this multidisciplinary field. And while there have been various platforms for policy-focused discussions on the topic of information analysis/data analytics for nonproliferation, few such platforms have existed to share novel technologies, data sources, approaches, and methods applied thoughtfully to the nonproliferation problem space outside of a handful of technical working groups. The Institute of Nuclear Materials Management's Open Source and Geospatial Information (OSGI) Working Group is one such arena, and has organized this special issue of the *Journal of Nuclear Materials Management* to provide a forum for technical exchange on the topic. Therefore, we hope that this special issue of the *Journal* can provide a constructive technical forum for experts in this field to share their most recent research while also offering readers a broad overview of the current developments in data analytics for nonproliferation, present applications, and ongoing challenges.

In the research presented here, we see evidence of multidisciplinary teams working together collaboratively to achieve unique contributions within the international nuclear safeguards and nonproliferation domain. Effective communication and close collaboration among team members builds on the perennial question, How can the data analytics research community best support the analysis? Technology should aid in achieving an

optimal utilization of the analyst's time and effort along the information analysis cycle. We are pleased to include manuscripts in this special issue that describe research into novel tools, methodologies, and data sources aimed at enhancing nonproliferation analysis.

In this special issue, we attempt to collate research from across the information analytic cycle. While the manuscripts presented in this issue could have been organized in a number of ways, we elected to arrange them by the specific data type that they address (text, imagery, and multimodal) in an effort to provide proximity to authors working with similar data.

The special issue begins with textual data analytics. In Versino et al., the authors present novel data visualization techniques that support the human analyst in recognizing patterns, detecting anomalies, and understanding broad data sets at a glance in strategic trade data. The authors' country- and commodity-based visualizations offer rich environments for customs authorities and trade analysts to better understand trade patterns and trends and to explore the ripple effect of changes in trade policies related to strategic commodities based on open source trade data. Stewart et al. take a broader approach, surveying a variety of textual data analytical methodologies and capabilities and providing nonproliferation-relevant examples in their assessment of each. The authors seek to explore the boundaries between automated text data analytics and the expert human analyst across seven unique analytical workflows.

The special issue then moves to imagery analysis. It opens with Gastelum and Shead, in which the authors fine-tune existing machine learning models to recognize images with nuclear cooling towers and steam plumes to support the verification of a facility's operational status. The authors determine that fine-tuning models for specific international safeguards' objects of interest is necessary due to the lack of incentive from the information



and communication technology industry to develop specialized models for a niche market, but that such development is both possible and useful. The imagery portion of the special issue then takes on a higher elevation, with Rutkowski et al.'s demonstration of Synthetic Aperture Radar (SAR) for site monitoring by using case studies relevant for nuclear nonproliferation. The authors assess the Sentinel-1 SAR imagery via the Google Earth Engine platform to gauge the practicality of using such a system to conduct change detection assessments at sites worldwide. Then, Keskinen et al. provide an update on enhanced capabilities of the International Atomic Energy Agency's (IAEA) internal Geospatial Exploitation System, which is intended to serve as a collaborative platform for safeguards analysts and inspectors to store, manage, analyze, disseminate, and view satellite imagery throughout the Department of Safeguards. The authors describe recent architectural, security, and interface enhancements to the system, supporting more effective workflows.

The special issue then moves into multimodal data. The multimodal section begins with Feldman et al., in which the authors describe the development of deep neural networks to search and triage textual, image, and video data related to the nuclear fuel cycle. The authors use unsupervised and supervised machine learning approaches to calibrate models developed on large-scale open source datasets to locate multimodal data in a semantic feature space proximal to nuclear fuel cycle-related seed data of interest. Palmer and Stevanović approach multimodal data from a different perspective, integrating multimodal data into single 3D visualizations for better analyst comprehension. The authors use data from two recent multilateral nonproliferation inspection exercises to demonstrate the capability to integrate geolocated data in a 3D geospatial platform for large-area situational awareness and assessments. Barletta et al. also emphasize the importance of geospatially linked data integration with their manuscript, describing the IAEA's internal Geo-based Data Integration platform. The authors illustrate how the system integrates multiple sources and types of safeguards data within the Integrated Safeguards Environment and how it is being used to support state evaluation by the IAEA. The special issue closes with Boyer et al.'s description of recent developments for the IAEA's internal Physical Model resource, a compendium of knowledge on the nuclear fuel cycle. The authors demonstrate how linking documents to their appropriate fuel cycle steps and associated geolocation within the IAEA's Collaborative Analysis Platform plays a key role in safeguards information analysis.

From the manuscripts in this special issue, we see examples

across analytical approaches and across data types and sources. Although these are not definitive processes to be mandated upon the nonproliferation analysis community, they serve as exemplars of research being done in this moment. As such, we hope that readers will not just read those articles that are directly pertinent to their work but will instead explore all the articles in this special issue. For many of these, the approaches can be used across multiple data modalities, information cycle stages, and various applications.

We have attempted to include timely contributions from the field as it currently stands. This is not a comprehensive literature review but a topical snapshot that exposes some of what is happening at present. We hope this gives the reader insight into recent developments in the field, and we aspire that this communication can be used both as a technical reference and a starting point for future collaboration. We are grateful to the many authors and peer reviewers who have shared their time and expertise to bring this special issue to life. We have enjoyed curating this collection and hope you enjoy reading it.



A Visual Atlas on Strategic Trade

*Cristina Versino and Simone Cagno, European Commission, Joint Research Centre, Nuclear Security Unit, Italy
Peter Heine, Argonne National Laboratory, Decision and Infrastructure Sciences, Argonne, IL, USA
Julie Carrera, Argonne National Laboratory, Strategic Security Sciences, Argonne, IL, USA*

Abstract

Strategic trade concerns the global trade flows of items with strategic military value. Items include goods with exclusive military use, but also dual-use items that can have both military and civil uses. Many countries require licenses for trade-in of strategic items as a measure of nonproliferation. For example, items listed for trade control by multilateral export control arrangements, such as the Nuclear Suppliers Group (NSG), typically are considered strategic goods.

Analysts and officials involved in strategic trade control often lack a quantitative global understanding of relevant trade patterns that would help them focus control and enforcement efforts, select companies for outreach and audit, target transactions for verification and inspection, and assess the potential economic impact of trade control policies. Many are unaware of available open source data on international trade that can inform these issues. Many do not know how to make use of such data or lack the tools necessary to do so.

A book of charts dubbed a Strategic Trade Atlas was developed to illustrate global trade flows of strategic goods and raise awareness of the availability and utility of open source trade data. The Atlas was prepared for the World Customs Organization (WCO) and its member states to support their planning and execution of strategic trade control enforcement (STCE) operations. It offers macroscopic graphical representations of global trade flows classified under Harmonized System (HS) codes associated by the WCO with commodities described in its STCE Implementation Guide. While these HS codes are understood to also encompass many nonstrategic items, they nevertheless provide a good overview of potentially strategic trade flows. These include HS codes covering nuclear and nuclear-related commodities listed by the NSG, such as uranium, heavy water, graphite, carbon fiber, nickel powder, beryllium, aluminum tubes, vacuum pumps, valves, machine tools, cameras, and many others. The Atlas uses trade data from the Base pour l'Analyse du Commerce International (BACI), a statistical elaboration of public United Nations

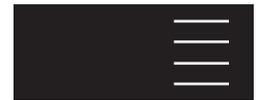
(UN) Comtrade data. BACI data reconciles trade flows, taking into account asymmetries in import and export trade valuation.

The data visualization approach facilitates rapid sensemaking, pattern recognition, and anomaly detection. Country-based profiles summarize vast amounts of data into readily understandable graphical representations of each country's primary strategic imports and exports and their main trading partners. Commodity-based profiles provide a global view of major exporters, importers, and trade patterns for each commodity. These insights can inform the risk management approach employed by customs to allocate enforcement and compliance efforts.

Introduction

Open source trade data offers a wealth of information that can inform and enhance strategic trade control (STC) implementation.¹ Customs authorities, for example, benefit from understanding typical trade flows in goods with potential military or weapons of mass destruction uses, because such knowledge can help focus resource-constrained efforts in STC outreach, audit, and targeting. However, officials tasked with implementing or influencing STC regulations often are unaware of the availability of open source trade data or lack the analytic and visualization tools to glean meaningful, actionable data from very large data sets. This represents a missed opportunity to enhance international security, especially with respect to the nuclear threat; many transactions involving nuclear and nuclear-related materials and equipment of proliferation concern are visible in trade data because many of these goods are well encoded in the Harmonized System (HS) nomenclature used to describe goods in international trade.

The Strategic Trade Atlas ("the Atlas") was developed in response to the challenges and opportunities afforded by open source trade data. The authors produced an initial version of the Atlas to support the World Customs Organization's (WCO's) Operation Cosmo, the first global operation to target strategic trade, which was conducted in 2014. That initial Atlas provided summary visualizations of each country's exports and imports



classified under HS codes associated with strategic goods. The Atlas was described in a paper presented at the Annual Meeting of the Institute of Nuclear Materials Management in 2015² and was later published in 2016.³ This article describes a fundamental revision of the first Atlas, prepared to support the second WCO Cosmo Operation in 2018. This new Atlas includes several improvements over the initial version. Most noticeably, the new Atlas provides commodity-based visualizations⁴ in addition to country-based visualizations.⁵ It further provides richer and more information-dense visualizations to allow for more observations, including trend analysis and indications of Revealed Comparative Advantages (RCAs),⁶ which highlight commodities especially significant to each country's export and import baskets. Finally, the underlying trade data for the revised Atlas includes more recent data (2012–2016) and uses mirrored and reconciled data from the United Nations Comtrade Database (UN Comtrade⁷) from the Base pour l'Analyse du Commerce International⁸ (BACI). This new Strategic Trade Atlas therefore provides views of international trade in potentially strategic goods with an unprecedented scope and level of detail.

This paper describes the development, organization, and content of the Strategic Trade Atlas. It explains the trade data sources and HS codes used, as well as the visualization approach to the data. A discussion follows on the utility of the new visualizations, concluding with remarks about potential applications of the Atlas for enhancing STC implementation.

Description of the Work

Trade Data Sources

Open source data on international trade originate from goods declarations made by importers and exporters to customs authorities. Customs data are collected, processed, and aggregated by national statistical offices. Data then are shared, by countries' decision, in databases on international trade, such as UN Comtrade⁷, the largest free-access data resource on world trade.

A UN Comtrade record includes the following main dimensions:

- The country reporting the data
- The country partner in trade
- The commodity traded, as described in the Harmonized System⁹ nomenclature
- The trade flow, as import or export, re-import or re-export
- The time period in which the trade took place, specified at the yearly or monthly interval

Trade is aggregated along the above dimensions and measured by value (e.g., in U.S. dollars) and quantity (e.g., in kilograms).

Data underlying the Strategic Trade Atlas are derived from the BACI database developed by the Centre d'Etudes Prospectives et d'Informations Internationales (CEPII). BACI applies a data harmonization procedure¹⁰ to the raw UN Comtrade data to increase both the number of trade records available per country and their robustness to data reporting errors. Specifically, the BACI harmonization procedure makes use of data mirroring and reconciliation. Data mirroring applies when a trade flow is reported only by the exporter or the importer. In this case, a mirror record is created to compensate for the unreported data by one of the two traders. By contrast, when both the exporter and the importer report trade data on a flow, the trade measures are reconciled to derive a single figure for the flow. First, asymmetries between import and export due to cost, insurance, and freight (CIF) associated with imports are estimated and deducted from import values to obtain Free on Board (FOB) values. Second, the reconciled figure is computed as a weighted value of the FOB measures, taking into account the data reporting reliability of the exporter and the importer. CEPII assesses a country's data reporting reliability on the basis of reporting *distances* between the country and its partners in trade. A reporter is rated as reliable when the data it contributed in UN Comtrade are generally close to those reported by its partners in trade. Vice versa, when a country's reported data are distant to those declared by the partners, its reliability is rated as low. Hence, the harmonized data in BACI provide trade figures reported by the world about each country's trade, not just the trade data reported by each individual country.

Commodities in Focus of the Strategic Trade Atlas

World statistics of trade in all goods, including nuclear and nuclear-related items, are described in terms of the HS, the commodity classification system designed and maintained by the WCO. The HS is the reference taxonomy of commodities adopted by States adhering to the Harmonized System Convention, but also by trade associations and statistical offices in the majority of countries. It is based on approximately 5,000 commodity groups organized within 22 sections in a hierarchy made up of chapters, headings, and subheadings. Each level in the hierarchy is identified by a HS code and a description, as illustrated in Table 1.



Table 1. Examples of HS codes and their descriptions

Level	HS	Description
Section	XVI	Machinery and mechanical appliances; electrical equipment; parts thereof; sound recorders and reproducers, television image and sound recorders and reproducers, and parts and accessories of such articles
Chapter	84	Nuclear reactors, boilers, machinery, and mechanical appliances; parts thereof
Heading	8401	Nuclear reactors; fuel elements (cartridges), nonirradiated, for nuclear reactors; machinery and apparatus for isotopic separation
Subheading	8401.10	Nuclear reactors

As part of its Strategic Trade Control Enforcement Programme¹¹ (STCE), the WCO identified HS headings and subheadings associated with strategic commodities. The Strategic Trade Atlas is built from a subset of items from that list of commodities. The HS codes associated with nuclear and nuclear-related items are shown in Table 2. HS code abbreviations used in Table 2 are expanded in Table 3. The textual descriptions of the HS codes listed in Table 2 are included in the commodity-based views of the Atlas. Comparing the HS self-explanatory text¹² included in commodity-based views of the Atlas with the associated strategic commodity presented in the implementation guide of the STCE enables the reader to assess the closeness or distance between the strategic commodities listed for control (with all the relevant technical parameters that make them strategic), and the less detailed HS descriptors most used by exporters and importers to declare traded items to customs authorities.

Table 2. List of HS codes associated with nuclear-relevant commodities

HS	Associated Strategic Commodity	HS	Associated Strategic Commodity
2612	Uranium and thorium ore and concentrates	8414.10	Vacuum pumps
2844.10	Natural uranium	8456.30	Electrical discharge machines
2844.20	Enriched uranium, plutonium	8457	Machining centers
2844.30	Depleted uranium, thorium	8458.11	Turning machines
2844.40	Radioactive elements	8459.21+	Milling machines
2844.50	Spent fuel	8460.11+	Grinding machines
2845.10	Heavy water	8463	Spinning lathes, flow-forming machines
3801	Graphite	8477	Filament winding machines
5501+	Synthetic fiber (Aramid)	8479.50	Robots
6815	Carbon fiber	8479.89	Isostatic presses
6903	Ceramic crucibles	8481	Valves
7019	Glass fiber	8504.40	Power supplies
7020	Radiation-shielding windows	8514.10+	Furnaces
7224+	Specialty steels	8525.80	Cameras (high speed, radiation hardened, or night vision)
7504	Nickel powder	8532	Capacitors
7604.21+	Aluminum tubes	8535	Switches
8101	Tungsten	8543.20	Pulse generators
8103.90	Tantalum crucibles	9013.20	Lasers
8104.11	Magnesium	9022.19	X-ray systems
8109	Zirconium	9026.20	Pressure transducers
8112.12+	Beryllium	9027.80	Mass spectrometers and toxic gas monitors
8401	Nuclear reactors, machinery for isotope separation, nuclear fuel	9031.10	Balancing machines
8413	Pumps	9031.80	Vibration test systems



Table 3. Full list of HS codes included in shortened HS codes (+) appearing in Table 2

Group	Member HS codes	Group	Member HS codes
5501+	5501, 5502	8459.21+	8459.21, 8459.31, 8459.51, 8459.61
7224+	7224, 7225, 7226, 7228, 7304.51, 7304.59	8460.11+	8460.11, 8460.21
7604.21+	7604.21, 7608.20	8514.10+	8514.10, 8514.20
8112.12+	8112.12, 8112.13, 8112.19		

Besides the nuclear-related items shown in Table 2, the STCE list and Atlas also include items relevant for the production of chemical weapons, biological weapons, delivery systems, and the military equipment (e.g., munitions) and delivery systems themselves. The HS codes used in the Strategic Trade Atlas for these and the nuclear-related items above are based on the STCE Implementation Guide. For the selected HS codes, BACI data was downloaded for 221 countries in the reference period 2012–2016, which was the most recent data period available for download. There is wide variation in the quantity of data for import and export across these countries, with some countries having tens of thousands of records, while others have fewer than 100. Because trade indicators, such as the RCA, cannot be estimated in a reliable way for countries with a small number of trade records,¹³ the Atlas retains in scope, with 188 countries having at least 50 export records in the selected HS codes over the 5-year period considered. The data was organized in a database and enriched by other metadata, including the HS hierarchical structure and International Organization for Standardization (ISO) country codes¹⁴ to name countries in compact space in the charts and reference them automatically on geographical maps.

Total trade data by reporting countries in all commodities was also downloaded to provide a comparison to the countries' potentially strategic trade. Total trade data is also used to compute each country's RCA in the trade of potentially strategic commodities. The RCA is computed as the ratio of the share of a commodity in a country's export basket to the share of that commodity in world trade. The RCA thus measures whether a country's trade in a commodity is higher or lower than the share of trade in that commodity in world trade. A high RCA in export for a commodity by a country is associated with high competitiveness by the country in the production and export of that commodity.¹⁵ High RCA in import for a commodity by a country may reveal relatively intense use by or dependence on the import of that commodity.

Strategic Trade Profiles by Data Visualization

The trade data visualizations included in the Atlas offer high-level summaries of the main trade flows classified under HS codes related

to strategic commodities, the primary origins of such imports, and the primary export destinations for these types of goods. The profiles provide insights into the nature of global commodity trade flows involving the HS codes listed in Table 2 and others listed in the STCE Implementation Guide.

Profiles have been designed to provide both country- and commodity-based views of strategic trade. While country-based⁵ views take individual countries as the focus of data presentation, commodity-based⁴ views aggregate global trade data for specific HS groups. Country- and commodity-based views are compact representations of potentially strategic trade, in which priority is given to top traders for a given HS group, or top HS groups for a given country.

The views include value and quantity of trade presented by country or by HS, time series data over the reference period of 5 years, flow data between main origins and main destinations, and an indication of the RCA in relation to countries or commodities. Specifically, use is made of the Normalized Revealed Comparative Advantage¹⁶ (NRCA) index, an RCA measure ranging from -1 to +1, with positive NRCA values associated with comparative advantage in trade.

The views presented in the next sections follow a few graphical conventions. A blue color is associated with traded values, whereas a red color is associated with traded quantities. Positive NRCA values are represented by grades of blue (the darker the shade, the higher the comparative advantage), and by grades of red for quantities. Negative NRCA values are represented by grades of gray, indicating that there is no comparative advantage. Finally, numbers appearing in charts are rounded percentages; for example, 10.7 percent is indicated as 11, and 0.4 percent is indicated as 0.

Country-based views

The Strategic Trade Atlas⁵ contains trade profiles for 188 countries. In the understanding that the HS codes represented in the Atlas also encompass nonstrategic items, the main take-away messages of country views are the following: (1) which countries are significant¹⁷ and competitive¹⁸ traders worldwide of HS codes associated with strategic commodities (HSSC); (2) which HSSC are traded more in-



tensively by countries, and with which partners; and (3) whether trends can be seen in exports and imports of HSSC over time.

A country-based view consists of two facing pages, one for export and one for import, plus a header, presenting high-level trade statistics in HSSC for the country. Figure 1 shows the header (top part of the figure) and the export page for an example country. The import page (not shown in Figure 1) contains the same fields as the export page, only for import data.

The header of a country view features the following information:

- [A] The country name, the ISO country code, and its position on the world map
- [B] The years of data included in the profile
- [C] The world ranking for export and import in HSSC
- [D] The trade balance by value in HSSC expressed as percentage
- [E] The value of HSSC expressed as a percentage of the country's total trade
- [F] The export-import trend by value in HSSC over the five-year period

The remaining part of the country-based view contains information about the HSSC traded by flow [G], export, or import. Specifically, the view includes the following: [H] The number of BACI data records available for country and flow. [I] The top 25 HS

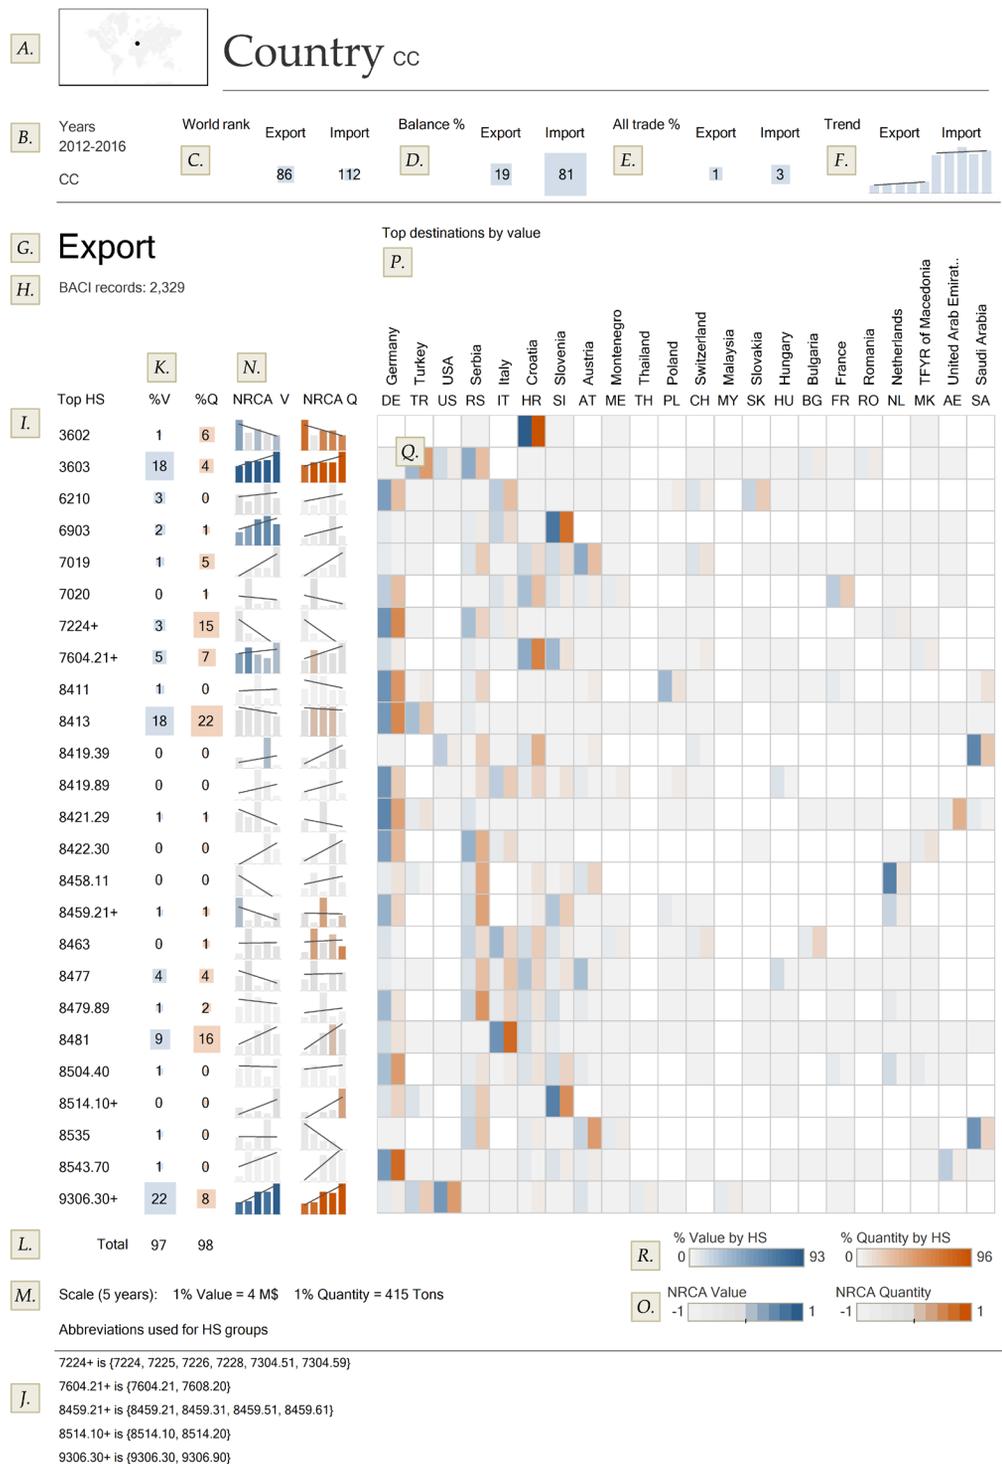


Figure 1. Country-based view: header [A–F] and export data [G–R]

traded by the country presented in HS order. [J] When abbreviations are used for groups of HS, these are expanded at the bottom of the page. [K] For each of the 25 HS groups, the value and quantity traded by the country are presented as a percentage of the



Harmonized System

8414.10

Years 2012-2016

BACI records: 18,574

All trade %: Value 0.02% , Quantity 0.00%

Trend by Value & Quantity



Top exporters

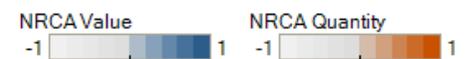
		%V	%Q	NRCA V	NRCA Q
Germany	DE	25	23		
Japan	JP	11	7		
USA	US	8	6		
Rep. of Korea	KR	7	6		
Italy	IT	7	9		
France	FR	7	7		
China	CN	6	11		
Czech Rep.	CZ	5	4		
Mexico	MX	4	4		
Switzerland	CH	4	4		
United Kingdom	GB	4	4		
Canada	CA	1	1		
Taiwan	TW	1	2		
Netherlands	NL	1	1		
Romania	RO	1	1		
Belgium-Luxembourg	BE	1	1		
Singapore	SG	1	1		
Malaysia	MY	1	1		
Spain	ES	0	1		
Hong Kong	HK	0	0		
Sweden	SE	0	0		
Austria	AT	0	1		
Brazil	BR	0	0		
Poland	PL	0	0		
India	IN	0	0		
Total		95	97		

Top Importers

		%V	%Q	NRCA V	NRCA Q
USA	US	18	15		
China	CN	11	9		
Germany	DE	7	8		
Rep. of Korea	KR	5	4		
Taiwan	TW	4	3		
United Kingdom	GB	4	4		
France	FR	4	4		
Japan	JP	3	3		
Singapore	SG	3	2		
Canada	CA	3	2		
Italy	IT	2	4		
Netherlands	NL	2	3		
Mexico	MX	2	2		
Spain	ES	2	3		
Russian Federation	RU	2	2		
Austria	AT	2	2		
Thailand	TH	1	1		
Switzerland	CH	1	1		
India	IN	1	1		
Czech Rep.	CZ	1	1		
Belgium-Luxembourg	BE	1	1		
Poland	PL	1	1		
Turkey	TR	1	1		
Saudi Arabia	SA	1	2		
Hungary	HU	1	2		
Total		86	81		

Scale (5 years): 1% Value = 188 M\$ 1% Quantity = 6,139 Tons

HS self-explanatory text



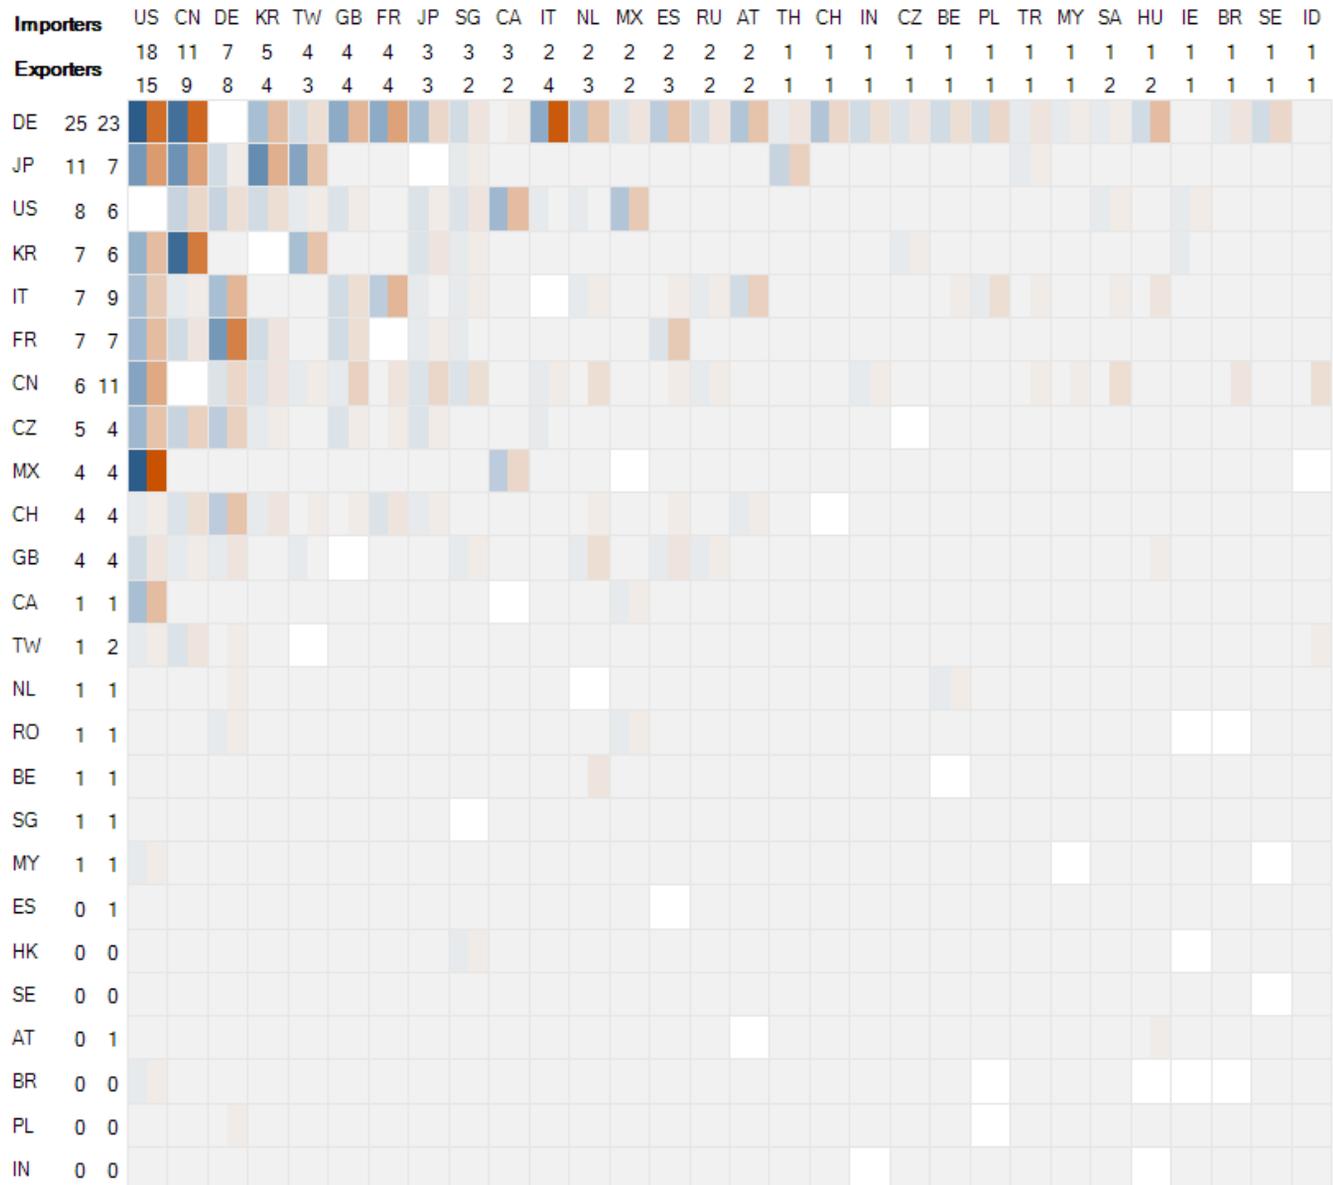
8414.10 Vacuum pumps.

Figure 2. Commodity-based view: header [A, E-H], footer [B], and top exporters and importers [I-N]



Vacuum pumps

Top trade flows



Scale (5 years): 1% Value = 165 M\$ 1% Quantity = 5,090 Tons

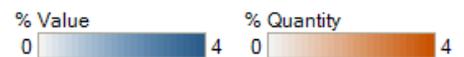


Figure 3. Commodity-based view: header [C, D] and main trade flows [O–T]



country's trade in the HSSC basket, [L] with their total coverage. [M] The relationship between the charts and the actual trade figures is given by the scale indication, defining the equivalent of 1 percent of the value expressed in millions of US dollars, and 1 percent of quantity expressed in tons. [N] Bar charts, scaled by HS group, show the time evolution and trend of trade value and quantity. [O] The color of the bars indicates the NRCA measured for the HS group in a given year, as per the color legend. [P] The top 22 trading partners by value are presented in a matrix that shows their respective flow per HS group, with destination countries (for export) and origins (for import) sorted left to right by decreasing trade value. [Q] Each cell in the matrix is split into two halves, indicating the trade value and quantity for that HS and the partner country. [R] Color grades in cells are scaled per each HS group, with a color legend indicating the relationship with the trade value and quantity percentages presented below.

Commodity-based views

The Strategic Trade Atlas⁴ contains trade views for 70 HS groups associated with strategic commodities, including the 46 HS groups related to nuclear-relevant commodities listed in Table 2. In the understanding that the HS codes represented in the Atlas also encompass nonstrategic items, the main take-away messages from commodity-based views are: (1) the size of world trade in HSSC codes; (2) its breakdown by countries, in export and import flows; (3) the RCA of trading countries; (4) trade trends over time; and (5) trade flows between main exporters and main importers.

A commodity-based view consists of two facing pages, shown in Figures 2 and 3. They include a header and a footer spanning over the two pages, with trade data presenting top exporters and importers on the left page (Figure 2) and main trade flows on the right page (Figure 3).

The header and footer feature the following information:

- [A] The list of HS codes included in the selected strategic HS group
- [B] A definition of the HS codes according to the HS self-explanatory text¹²
- [C] The name of the strategic commodity associated with the selected HS group (e.g., see Table 2)
- [D] The reference to the page of the STCE Implementation Guide¹¹ where the associated strategic commodity or HS group is described in some detail
- [E] The years of data included in the profile
- [F] The number of BACI trade records available for the selected HS group

- [G] The world's value and quantity of trade in the selected HS group expressed as a percentage of the world's trade in all commodities
- [H] The world's value and quantity of trade in the selected HS group over five years.¹⁹

Other parts of the profile feature the following information (Figure 2):

- [I] The top 25 exporters and importers by value and quantity for the selected HS group, with exporters and importers listed by decreasing value of trade.
- [J] The value and quantity exported (imported) by the corresponding country to (from) the world; figures express the percentage contribution of the country's export (import) to the world trade in the selected HS group.
- [K] Totals indicate the data coverage presented on the page for the world trade in the selected HS group.
- [L] The scale defines the equivalent of 1 percent of value expressed in millions of US dollars and of 1 percent of quantity expressed in tons.
- [M] Bar charts, scaled by the corresponding country, show the time evolution and trend of trade value and quantity.
- [N] The color of the bars indicates the NRCA measured for the country in a given year, as per the color legend.
- [O] The trade flows matrix (Figure 3) encodes export–import flows between top exporters and top importers listed under point I.
- [P] Exporters are listed by decreasing value of trade (top to bottom); their percentage value and quantity of export to world are reported here.
- [Q] Importers are listed by decreasing value of trade (left to right); their percentage value and quantity of import from world import are reported here.
- [R] A cell in the matrix is split into two halves, encoding by color the value (blue grades) and quantity (red grades) traded between corresponding exporters and importers, with color grades scaled to the entire matrix.
- [S] Color encodings for the trade distribution shown in the trade flow matrix (see points O, P, and Q); the color scales represent percentages of trade (value and quantity) over the entire matrix.
- [T] The scale applies to trade percentages in the trade flow matrix, defining the equivalent of 1 percent of value expressed in millions of US dollars and of 1 percent of quantity expressed in tons. The scale allows translation of trade percentages in the matrix into approximate values



and quantities traded between countries over the reference period.

Results

Customs officers working in a particular country can use the country- and commodity-based views in the Strategic Trade Atlas to develop a better understanding of trade flows that might provide cover for unlicensed trade in strategic commodities (SC). Knowing the fraction of total trade covering potentially strategic exports allows one to estimate the size of the associated export search space, and hence the maximum effort required for controlling them. The Atlas country-based views show that the majority of countries have limited trade in HSSC.²⁰ In a few highly industrialized countries, this fraction of trade is considerable, reaching up to 17 percent of their total export by value. However this large upper bound covers exports by a country in all HSSC toward all destinations. By using the different levels of trade detail accessible in the Atlas, the trade space for scrutiny can be narrowed down considerably — for example, by removing impractically broad HS codes or by excluding trade with low-risk countries.

When planning STCE operations, an improved knowledge of strategic trade patterns permits a better planning of controls and better prioritizing of HS codes to target. The “dense visualization” approach developed for this version of the Atlas provides the analyst with complementary “context and detail” views about the trade of HSSC. The choice of graphical encodings for the underlying data allows for the displaying of a rich set of data in a compact space to enhance understanding and the spotting of anomalies. Examples of this are the 100 bar charts organized in a *small multiples*²¹ arrangement in the commodity-based views (Figure 2, point M), and the ordered columns of heat maps used in the country-based views (Figure 1, point R). For depicting trade flows, the use of a matrix layout was preferred because of its clearly structured bidimensional order, by rows and columns, which enables the precise and systematic visiting of cells of interest. Despite the geographical nature of trade, no geographical maps were used in the Atlas other than the one locating the country in focus in country-based profiles. In previous work,³ it was found that trade data are hard to represent on maps in a legible way, and the benefit of seeing a geographical pattern, when it exists, is outweighed by the amount of space on a page required to show a map.

A key feature of the new Atlas is the availability of NRCA indicators. A high NRCA in export for a commodity by a country is associated with high competitiveness by the country in production and export of that commodity. For import data, it may reveal

intense use by or dependence on the import of that commodity. The embedding of the NRCA indicator in the time series provides a different angle than that offered by a simple comparison of trade volumes. For instance, this might suggest, when planning STC activities, that one should prioritize the monitoring of commodities featuring positive NRCA, in the presence of overall similar trade volumes.

The global perspective offered by the Atlas also allows the identification of anomalies. Examples of such anomalies in commodity-based views are peculiar value-to-quantity ratios for certain countries, possibly pointing to erroneous reporting or to products of different quality and value classified under the same HS code. The total trade per year and its trend over the 5-year period can also give indications on the reliability of high figures for a specific country or commodity: when trade is high only for a single year, this might represent incorrect reporting. In these cases, this Atlas provides a basis for focusing further investigations. For instance, UN Comtrade monthly data could be queried on selected data dimensions to obtain further details, or Customs authorities might look up the detailed export declarations in their possess.

A number of issues and points for further improvement have become evident in the course of this work. The main known issue is related to the HS codes, which, except for a few commodities, encompass many goods that are not SCs. Trade data are reported according to the HS taxonomy, which makes this issue unavoidable in the absence of additional information. However, other data available on customs declarations, such as weight, price, and units of shipment could be used to better identify potentially strategic goods in trade data. Furthermore, until the WCO defines in the HS codes more specific to strategic goods, it would be beneficial to differentiate between HS codes that are very selective for strategic goods and those that are extremely broad.²²

The use of reconciled BACI data in this Atlas aims at improving the quality and the world coverage of the data available. As detailed in the section on trade data sources, the Atlas makes use of mirrored data when reported in UN Comtrade as a single flow, either by the exporter or by the importer. In this case, when incorrect reporting occurs (e.g., by wrong HS code, value, or quantity), the data is no longer weighted by the BACI reconciliation procedure, and will appear in the Atlas as such. The visualization approach already helps to identify potential cases of wrong reporting. In future versions, this could be enhanced by adding statistical plots, such as a box plot of the distribution of value quantity ratios, to support the identification of outliers.

Conclusions

While the trade data used for this project are publicly available, most officials charged with the responsibility to create or implement STC policies lack the data processing capacities and capabilities to make use of them, except where narrowly defined queries can be formulated. The Strategic Trade Atlas distills a vast amount of trade data into a concise but rich visual format.

Customs authorities responsible for enforcing STC regulations have used these snapshots to guide their risk management strategies, which allocate resources to risky trade while facilitating low-risk trade. It is common for customs authorities to have a very low level of familiarity with strategic goods, or not to know which strategic goods may be most relevant to their countries. Customs officers known to one of the authors have used the Strategic Trade Atlas to quickly identify relevant sectors of their national economy for further research, including identification of relevant companies, outreach to those companies to understand the nature of their products, and targeting campaigns to verify (through preclearance cargo examination and postclearance audit) that potentially strategic exports comply with national STC regulations. In addition, the Atlas helps Customs authorities identify common, high-volume trade flows that can be facilitated once steps have been taken to verify or ensure they are low risk. Doing so contributes to a risk-informed effective allocation of resources.

In addition to its value for customs authorities, other analysts wishing to understand trade patterns and trends involving strategic goods can readily gain insight and spot anomalies for further investigation. These insights can help inform not only enforcement operations, but also outreach to and audit of exporters and technology holders. Visualizations such as those in the Atlas could also afford the opportunity to assess potential impacts of changes in trade policies. Beyond these potential applications, the Atlas illustrates the value of open source trade databases and their potential to support trade policy implementation and development.

Acknowledgments

Cristina Versino and Simone Cagno (European Commission, Joint Research Centre) are supported by the project Strategic Trade Analysis for Non Proliferation (STANP, 756) funded by the European Commission's Euratom Horizon 2020 Research and Training Programme. Peter Heine and Julie Carrera (Argonne National Laboratory) are supported by the U.S. Department of Energy (DOE) National Nuclear Security Administration under interagency

agreement, through the U.S. DOE contract DE-AC02-06CH11357.

Keywords

Nonproliferation, strategic trade, customs controls, data visualization

Author Biographies

Cristina Versino is senior scientist in the Nuclear Security Unit of the European Commission (EC) Joint Research Centre. The focus of her work is data analysis in support of EU policies. Her work also involves design and development of new tools and approaches for use by nuclear safeguards inspectors and non-proliferation analysts. Previously she worked as project officer in the Future and Emerging Technologies Unit in EC DG CONNECT, funding cutting-edge research in information technologies. Earlier she was FNRS researcher in machine learning at IDSIA, Switzerland. Cristina holds a MS in information science from the State University of Milan.

Simone Cagno received his PhD in Chemistry from the University of Antwerp (Belgium) in 2012. Since then, he worked as a postdoc at the Centre for Environmental Radioactivity at NMBU in Ås (Norway), with a focus on advanced analytical techniques for plutonium, uranium and thorium detection. In 2014 he joined the Radiochemistry group at the Belgian Nuclear Research Centre SCK-CEN, where he worked on new chemical processes for the front- and the back-end of the fuel cycle and for the production of medical radioisotopes. In 2017 Simone joined the Nuclear Security Unit of the European Commission Joint Research Centre in Ispra (Italy). There, in the framework of the activities regarding non-proliferation of nuclear weapons, his contribution falls mainly in the area of trade analysis and open source monitoring.

Peter Heine is a Director of the Center for Strategic Security in Argonne National Laboratory's Global Security Sciences Division. Pete received his BS in Applied and Engineering Physics from Cornell University and his MBA from the University of Chicago. Pete has worked on strengthening measures to prevent the proliferation of weapons of mass destruction (WMD) for over 25 years. His personal focus is on effective control of trade in materials and equipment related to the development or production of WMD. He has worked to advance trade data analysis strategies for identifying high-risk shipments and developed and conducted export control training programs for the outreach programs of the U.S. Department of Energy and the World Customs Organization. He has also worked with the International Atomic Energy Agen-



cy to promote the use of export and import information in international safeguards. Pete spent most of his career at Argonne National Laboratory, most recently as Director of the Center for Strategic Security. He started his professional career serving as an officer in the U.S. Navy's nuclear propulsion program. Pete also worked for a management consulting firm specializing in the strategic application of technology.

Julie Carrera is Group Leader for Chemical/Biological Proliferation Analysis and Detection in the Strategic Security Sciences Division at Argonne National Laboratory. In this role, she oversees a multidisciplinary portfolio of technical analysis and modeling aimed at mitigating chemical and biological weapons threats. Julie's research interests include applying supply chain solutions to WMD nonproliferation challenges, and she leads a cross-organizational effort on Advanced Supply Chain Analytics. She earned her B.S. in Chemistry from Carroll College and her M.S. and Ph.D. in Chemistry from the University of Chicago.

References

1. Versino, C., Tsukanova, M., & Cojazzi, G.G.M. (2010). *Catalogue of web data services on global trade*. ISBN 978-92-79-18919-7.
2. Heine, P. & Versino, C. (2015). A strategic trade atlas for non-proliferation and export control. *56th Annual Meeting of INMM*. Indian Wells, CA.
3. Versino, C. & Heine, P. (2016). *Strategic trade atlas: Country-based views*. ISBN 978-92-79-55003-4.
4. Versino, C., Heine, P., & Carrera, J. (2018). *Strategic trade atlas: Commodity-based views*. ISBN 978-92-79-84031-9. Retrieved from <http://publications.jrc.ec.europa.eu/repository/handle/JRC111471>
5. Versino, C., Heine, P., & Carrera, J. (2018). *Strategic trade atlas: Country-based views*. ISBN 978-92-79-84030-2. Retrieved from <http://publications.jrc.ec.europa.eu/repository/handle/JRC111470>
6. Balassa, B. (1965). Trade liberalisation and "revealed" comparative advantage. *The Manchester School*, 33, 99–123.
7. United Nations Statistics Division. United Nations Commodity Trade Statistics Database (UN Comtrade). Retrieved from <http://comtrade.un.org>
8. Centre d'Etudes Prospectives et d'Informations Internationales. Base pour l'Analyse du Commerce International (BACI). Retrieved from www.cepii.fr/CEPII/en/bdd_modele/presentation.asp?id=1
9. World Customs Organization. Harmonized System. Retrieved from www.wcoomd.org/en/topics/nomenclature/overview/what-is-the-harmonized-system.aspx
10. Gaulier, G. & Zignago, S. (2010). BACI: International trade database at the product-level. The 1994–2007 Version. CEPII Working Paper N. 2010-23.
11. World Customs Organization. Strategic Trade Control Enforcement (STCE). Implementation Guide. Retrieved from www.wcoomd.org/en/media/newsroom/2014/june/~//link.aspx?_id=A3AD81E1BFC742F2806DCAC55C54B743&_z=z
12. Eurostat. (2012). RAMON-Reference and Management of Nomenclatures. Harmonized Commodity Description and Coding System. Retrieved from http://ec.europa.eu/eurostat/ramon/nomenclatures/index.cfm?TargetUrl=LST_CLS_DLD&StrNom=HS_2012&StrLanguageCode=EN&StrLayoutCode=HIERARCHIC
13. For a general explanation of why this is the case, illustrated by examples unrelated to trade studies, see Wainer, H. (2007). The most dangerous equation. *American Scientist*, 95, 249–256.
14. International Organization for Standardization. Country Codes-ISO 3166. Retrieved from www.iso.org/iso/home/standards/country_codes.htm
15. Haussmann, R., Hidalgo, C., Bustos, S., Coscia, M., Chung, S., Jimenez, J., ... Yildirim, M. (2013). *The atlas of economic complexity: Mapping paths to prosperity*. ISBN 978-0-262-52542-8. Cambridge: MIT Press.
16. World Trade Organization. (2012). *A Practical Guide to Trade Policy Analysis*. ISBN 978-92-870-38128. Geneva: WTO Publications.
17. In terms of values and quantities traded in their import–export basket.
18. As measured by the Revealed Comparative Advantage.
19. Note that A–B and E–N are on Figure 2, and C and D are on Figure 3.
20. See point E in the header of country profiles.
21. Tufte, E. R. (1990). *Envisioning information*. ISBN 978-1930824140. Cheshire, CT: Graphics Press.
22. Chatelus, R., & Heine, P. (2016). Rating correlations between customs codes and export control lists: Assessing the needs and challenges. *Strategic Trade Review*, 2(3), 2016.



Automated Processing of Open Source Information for Nonproliferation Purposes

Ian J. Stewart and Alexander Lee
King's College London, London, UK

Ahmed ElGebaly
King's College London, Egypt

Abstract

This paper examines automated approaches to assist human analysts in systematic utilization of open source information in nonproliferation efforts. The paper begins by highlighting the value of big data analysis to nonproliferation efforts. It then introduces a number of approaches that can be used for data collecting and storage, enriching, and visualization. These are structured as workflows that arguably are a useful construct when considering how large analytical organizations can incorporate automation processes in their analytical processes. Finally, the paper examines a number of architectural approaches to integrate the capabilities examined in this paper for practical application. The paper argues that such platforms and capabilities can augment but not replace human analysts in nonproliferation efforts. This paper draws on a test-bed platform constructed at King's College London.

Introduction

For many years, the main thrust of international nonproliferation efforts has been to verify peaceful nuclear energy bona fides.¹ Defined negatively, these verification efforts are driven by the identification of undeclared activity. Verifying the peaceful nature of state nuclear programs consists of assessing the correctness of a state's declarations about its nuclear program, including looking for indicators of undeclared activity. In practice, states submit their declarations to the International Atomic Energy Agency (IAEA) for the purposes of verification, and the IAEA holds these in strict confidence. The IAEA's task in this context is to comb through and combine a vast amount of data to compare it with the state's declaration and spot anomalies that must be further investigated. Actors outside the IAEA, including Project Alpha at King's College London (KCL) and state intelligence organizations, are equally focused on identifying undeclared nuclear activity. However, such

organizations do not have access to the state's declaration other than what has been described publicly by the state and are thus forced to focus on the identification of clandestine nuclear activity through the analysis of vast amounts of open source information.

Somewhat helpfully, nuclear programs are large multidisciplinary scientific and engineering undertakings. Even though nuclear programs are hard intelligence targets, a secret program will produce detectable discrepancies between what is declared and what is the actual case. In practice, it has proven difficult for states to hide all indicators of a nuclear program. A nuclear weapons program must inevitably spring out of a civilian one. It requires a staff of scientists and technicians from a range of disciplines, many of whom are from recognized tertiary education establishments, and some of whom have a publicly observable academic research footprint. Scientific and technical literature may reveal patterns of illicit activity, and a "follow-the-people" approach might lead to revealing disclosures on social media and curriculum vitae-type media. Furthermore, virtually all nuclear programs rely on a global system of trade in highly specified, and often export-controlled, equipment. The procurement of necessary technologies may produce detectable signatures enabling trade analysis. The manual searching for and analysis of such "nuggets" at such depth and scale is time-consuming and resource intensive.

Since 1991, the IAEA has been the principal internationally respected actor with the responsibility to conduct such examinations. The IAEA has a broad mandate to identify undeclared nuclear activities in any state with a Comprehensive Safeguards Agreement. This mandate has seen the IAEA utilize an ever-growing number of tools and sources in its safeguards mission. The allowed tools are expanded further by the IAEA's Additional Protocol (AP), which affords the Agency access to more trade-related data as well as additional tools for on-the-ground inspection. As of early 2018, the IAEA has AP agreements in force with 129 states,



with an additional 12 not yet in force and a separate agreement with Euratom.² This provides the IAEA with wide access to data that needs to be verified both in the field and using open sources to ensure that the state's declarations are correct and complete. The necessity for wide-remit examination of all available information is a nontrivial burden on analysts and available resources.

Recent advances in search technology, semantic technologies, machine learning and automation offer the promise of reducing the information processing burden, improving not only the intake of a high quantity of data but improving the depth and quality of information analysis. Most importantly, automation will improve the timeliness of information collection, potentially with positive impacts on intervention and other enforcement measures.

However, it is important to note that the upward trend of technological improvements occurs against a backdrop of an evolving information environment and an increasingly acute public debate around electoral demographics and the use of highly granular private–public data. More information relevant to the safeguards mandate is becoming publicly available. This includes ever-growing amounts of trade and business information, satellite imagery, and more novel sources and tools, such as aircraft and maritime tracking capabilities and other data that can add to the broader picture. These pose practical design problems for the collection of discrimination dilemma (that is, whether to collect in a focused manner and potentially miss an important finding or collect everything and spend time processing for relevance), and there are deep ethical and legal concerns spanning valid data subjects, data retention and admissibility/end use. Furthermore, although secret intelligence derived from national technical means and other sources may always retain some salience, it is possible that the proportion of its relevance compared to information obtained through public access will diminish, particularly in the post-Snowden environment, in which the public are more attentive to risks associated with privacy and information security.³ It is unclear how these dynamics will affect international organizations such as the IAEA and the national agencies with a routine responsibility for proliferation detection.

The purpose of this paper is to scope out the key steps and technologies involved in information monitoring relevant to nonproliferation efforts in verification. For clarity, this paper does not argue that automation can replace human analysis. It argues that automation can fill a key role in gathering and processing data that composes the routine stream of safeguards-relevant information. Instead of replacement,

automated systems need to be carefully integrated into existing schemes of work and new human–machine “Centaur”⁴ workflows to maximize both the use of machines for that which is routine or mundane and the time and effort available for human creativity and analysis.

This paper seeks to demonstrate the following:

- Automation can enable a higher volume of data to be collected and processed for subsequent analysis.
- Automation can enable the discovery of linkages between relevant data-points that would not otherwise be achieved through human cognition.
- New tools can greatly aid the analysis with regard to both of these actions.

These findings are the product of an iterative development process by Project Alpha, KCL, that culminated with a test-bed platform, the Alpha Proliferation Open Source Tool (Alpha-POST). Project Alpha was founded in 2011 with UK government funding and originally focused on the utilization of well-established manual approaches to data collection and analysis. Over time, Project Alpha has grown and evolved to focus on what could be described as big data analysis for nonproliferation purposes and has conducted detailed open source analysis of various nuclear programs. This analysis has, in some cases, resulted in the identification of hitherto unknown nuclear activity using a quasi–big data approach.⁵

This paper aims to capture the lessons learned from this work at KCL so that they can be applied more broadly in the nonproliferation community. This paper proceeds as follows: First, we introduce the types of data sources that are useful in nonproliferation analysis and workflows that can integrate automated data collation into a human-centric analytical process. These workflows might be incomplete, but from the author's perspective, they provide a solid foundation on which to examine the utility of different automated processes. Second, we introduce several cross-cutting data processing capabilities, including natural language processing (NLP), machine translation, geoparsing and visualization, and provide our assessment of the maturity of these capabilities. Third, we examine several architectural approaches to building a comprehensive automation process.

The paper concludes that some automated approaches are mature enough to be included in the analytical workflow. However, with each workflow, there continue to be caveats that will require human analysts to understand the limitations of each processing approach to reach a well-rounded assessment of any given topic.



Table 1. Types of data available for nonproliferation efforts

Type	Examples	Structured/ unstructured	Typically retrieved by ⁶	Quantity
News articles	Reuters, BBC, GDELT, ⁷ company websites	Semi-structured	RSS, BigQuery, API	Thousands per day
Academic papers	Scopus, Web of Science, university websites	Semi-structured	Direct download, API	Thousands per month
Trade data	Tender data, customs data	Structured	API, direct download	Millions per day
Entity data	Ownership structures (Panama papers, company registers)	Structured	API, direct download	Millions updated regularly
Patents	Google patents, WIPO, national databases	Structured	API, RSS, direct download	Thousands per month
Other	Vessel databases	Structured	API, direct download	Hundreds per month
Media	Photos (Google photos), videos (YouTube)	Not structured	API, direct download	Millions per day
Social media	LinkedIn, Facebook, Twitter	Structured	API	Millions per day

Types of Data Available for Nonproliferation Efforts

There are a wide variety of sources of information that can contribute to nonproliferation efforts. Although we do not attempt to provide a comprehensive list, many of the sources used at KCL can be categorized into one of the categories shown in Table 1.

Enabling the Analyst

Despite the considerable interest in the machine analytic domain, human analysts are unlikely to be replaced entirely. Competent analysts, adequately trained and prepared, remain the cornerstone of many organizations' resources. Excellent examples of where human analysis has laid bare clandestine activity through a combination of careful scrutiny and recollection — and an element of serendipity — has surfaced in the public domain.⁸ Yet, human cognition is limited, and despite a near-century's worth of valuable improvements in analytic tradecraft, certain conditions restrict our ability to process and consistently analyze complex problems over extended periods of time. Chief among these issues is the human limits of memory and power of recollection. As the quantity and types of incoming information increase, traditional organizations reliant on human analytic throughput will struggle with the new information environment. Classification and filtering algorithms will play a key role in delivering information takes to an

analyst's desk for analysis. Furthermore, the ability to connect the dots, especially across siloed information in incompatible formats, relies heavily on human memory and the ability to see patterns. Pattern recognition is founded on the ability to recognize exact or similar matches — for example, recognizing the recurrence of an individual's name — but the human mind is computationally limited. Repeating this named entity recognition and linking process consistently will be best performed by machines.

Mitigating the limitations of human cognition can benefit from the use of automation to manage the routine, boring-but-necessary steps preceding analysis — that is, collecting and processing information — coupled with the involvement of humans to guide and correct. As an aspiration, automation will in time maximize human creativity and minimize involvement in the mundane, leading to better analytic outcomes overall. Analytic organizations are likely to engage with new algorithmic approaches to well-established schemes of work, while retaining a large involvement of humans in the loop. The appropriate combination of algorithm-human fusion will be the subject of much analyst management effort in the coming years.⁹

To this end, a research project at KCL has identified several workflows with which automated and semiautomated information collection approaches can be used to augment human analysts. These workflows are not necessarily new and have been used in practice in various forms. However, the purposeful creation of a



platform to perform these workflows in tandem and consecutively is claimed by the authors to be novel. These workflows reflect the previous manual workflows developed by KCL to incorporate big data analysis into its efforts to identify activity related to proliferation. Although the work undertaken by governments and international organizations like the IAEA differs from those at KCL to some extent, we argue that the following workflows can be used by any organization in the nonproliferation sphere to discover hitherto unknown activity:

- Workflow 1: Advanced search
- Workflow 2: Natural language processing for link analysis
- Workflow 3: News and media monitoring
- Workflow 4: Multilingual academic publications for link analysis
- Workflow 5: Mapping entities geographically
- Workflow 6: Machine classification
- Workflow 7: Multimedia object detection

Advanced Search

The large volumes of data available to the analyst require effective search strategies and tools for information retrieval. A simple search for a topic like “Iran nuclear” (interpreted by the search engine as “Iran” and “nuclear”) could yield hundreds of thousands of results and cannot, on its own, be considered an effective search strategy in a big data environment. Instead, advanced search tools can aid the analyst in both retrieving known information and finding new connections. Some of the tools that are explored below are in wide use in the commercial sector for specific website search boxes but are not integrated into major search engines (such as Google). As such, it is useful to dwell on the types of tools that exist, because from the authors’ experience, current information technology capabilities in governments and international organizations generally focus on straightforward Boolean operators rather than the more advanced features that can be made available.¹⁰

Faceted search: Faceted search, as found on sites like Amazon.com, integrates content classification and count information on the search page. It becomes possible to dynamically browse the search results to find new and unexpected linkages. For example, browsing the search results might reveal articles about the Democratic People’s Republic of Korea (DPRK) in an Iran search results page, which, when clicked, might further reveal a piece of information highlighting a nuclear connection. Faceting shows the number of articles under each categorization at each stage. It should be noted that these advanced search

processes can be integrated with many of the other workflows described below, so that, for example, it is possible to produce a link analysis chart from the facets generated in the search results. This can be a useful tool for showing where a link exists based on search terms.

Spatial proximity search: Location information, where available, including latitude and longitude, can be indexed for search. This allows for proximity searching, including via facets. For example, it is possible to search for any results originating within 5 miles of Isfahan, Iran. Most mainstream Internet search engines do not directly allow for special proximity search, although mapping services, such as Google Maps, provide this capability for points of interest such as coffee shops, gas stations, and other businesses. OpenStreetMap’s Nominatim search engine also provides a proximity search capability to search its mapping database.¹¹ For social media, the platform Echosec geofences social media posts, including untagged posts, using geoparsing and geocoding, which allows for spatial queries by area of interest, pin drop, or multifaceted keyword search.¹² In particular, if geoparsing scripts are used to generate location data, it is likely that the location will be somewhat inaccurate. Furthermore, some geoparsing services are inherently limited or flawed. For example, Google and Bing, and likely all other major non-Chinese mapping providers, tend to be inaccurate in China and Russia.

N-gram searches: In these searches, the query string (for example, “Iran nuclear”) is interpreted based on the proximity of words to one another. So, a user might search for records in which “Iran” appears within three words of “nuclear.”

Synonym searches: These are searches in which a thesaurus lookup is automatically used alongside the original search term to find results based on alternative terms — for example, a search for “frequency changer” will also search for “frequency converter.”

Fuzzy search: These are searches in which the search system compensates for optical character recognition (OCR) errors. It is language dependent and could retrieve large amounts of results if misused.

Phonic search: In these searches, the search system looks for words that seem similar phonetically. This is useful when searching for non-Latin names, such as Ahmed, Ahmad, Achmed, Ahmet, Ahmat, Achmat, and so on.

Word stemming: Word stemming is also useful so that the search index returns not only exact matches, but also variations in words. For example, if the search term entered was

“stemming,” the search term that would be used would be “stem,” and the results returned would include “stem,” “stemming,” and “stemmed.” For Romance languages rooted in regular expressions in Latin, words are constructed in a fashion that allows for stemming; however, this does not necessarily bear out especially well in other languages such as Arabic, where the “vowels” of a word vary and impart different meanings depending on the culture and context. To hedge against this, phonic searches understand the common variations of a given word — Ahmed, Ahmad, Ahmet — but essentially are contextually blind.

Natural Language Processing for Link Analysis

Link analysis has long been a tool used in the nonproliferation domain. A number of efficient tools exist to create link charts from structured data, including IBM's I2 Analyst's Notebook (which generates .anb files¹³), Maltego, and, in a more generic way, Gephi. However, these tools require very structured data to produce link charts. Indeed, the first step in importing data to any of these tools is to clean and organize the data. Link analysis has many uses in broader data analysis. However, two main challenges must be overcome first. In our experience, none of the above-mentioned software packages are capable of addressing these challenges on their own.

First, unstructured data must be made structured. This could be done manually; staff could read through all source text and highlight and then extract, for example, names, organizations, places and times to be fed into the link analysis software. However, given the large amounts of data now available, this manual approach is perhaps suited only for the most important tasks, because of the high human cost (salary cost and analyst hours) associated with manual approaches. Instead, NLP, which is described further below, provides a relatively useful substitute. NLP can extract organizations, places, dates, times, email addresses, and other information from text and record this information in a structured way for further processing.

Second, the data must be dynamically queryable. The problem with the link analysis software mentioned above is that data is usually ingested manually (in a comma-separated values format). If this approach is used, it will not be possible to expand the data set from a dynamic source database without reimporting a fresh batch of data. A dynamic link to a database might require some add-ons to be implemented with additional license costs and hardware setup. This might not always be possible.

Other than the platform developed at KCL, we are unaware of any tool that completes this workflow. Instead, we have each

undertaken manual or semimanual annotation for link analysis. This approach is inherently limited in scale.

News and Media Monitoring

News and media monitoring is not a new topic. Most news Internet sites have operated a Really Simple Syndication (RSS) feed for decades, and many applications are available to read such RSS feeds. Several capabilities also exist that allow for the generation of RSS feeds from sites that do not natively produce an RSS feed. A type of news media monitor is built by essentially combining various RSS feeds into an interface, which can also be multilingual. At KCL, for example, we used the Google-supported GDELT project as an RSS source. GDELT purports to cover the majority of the world's media sources in 65 languages. Querying this for terms like “DPRK nuclear,” “Iran nuclear,” “Iran missile,” “uranium,” and “plutonium” produced on average 5,000 new results per day in many different languages. If such a large volume of content was to be effectively monitored, it would have to be (1) reliably collected and stored, (2) possibly machine translated (perhaps to English, depending on the analyst's language skills), (3) passed through NLP to extract key entities, and (4) indexed for search. A system of rules-based matching would also be required to move items into analyst's workbaskets, perhaps based on country or topic.

Through the use of such capabilities, it is possible for a team of analysts to monitor the majority of media articles worldwide in any language.

Multilingual Academic Publications for Link Analysis

Another workflow identified at KCL utilizes link chart analysis of academic publications in which the source language is not English. This is different from Workflow 2 mentioned above in that it relates to languages other than English, for which most tools are tuned and ready out of the box. Furthermore, academic publications are considered somehow semistructured, because a publication record contains tagged information such as the title, authors, organizations, and sometimes tools, materials, and equipment used in the research.

Although analysis of academic articles for nonproliferation purposes is a well-established approach, many relevant publications will not be available in English. Considerable data processing and preparation is required prior to the ingestion of the data into link analysis. Furthermore, despite the fact that academic publications typically include the same kinds of information,



and that there are some bibliographic formats, we found substantial variation in the formats used by publishers. Automated approaches reduce this burden by cleaning and processing semi-structured data and can level the playing field by also allowing for the inclusion of non-native content. In particular, the collection and processing of large quantities of technical non-English material can utilize a variety of mechanisms, such as bulk translation, gazetting,¹⁴ and semantic annotations, to index and make such large data sets queryable and understandable. These mechanisms rely on human-generated knowledge that is transferred to the machine, such as through ontologies, which are examined later.

Mapping Entities Geographically

As our geographic information systems have become more sophisticated, the amount of data that is relatable to geospatial maps and globes has increased. Much of this data was designed to include geolocational metadata (typically, latitude and longitude), and relative to non-locational data, geolocated data is increasing. However, this surge in georeferencing is relatively new and is disproportionately found in certain forms of non-traditional digital media, such as those found in Facebook and Twitter.¹⁵ A sizeable amount of relevant data existing in semantic/textual forms but relating to real physical spaces remains locked within the intricacies of human language. Utilizing NLP (described in Workflow 2), such as named entity disambiguation and named entity linking, locations are extracted from blocks of text and related to geographical coordinates. Overall, this is called *geoparsing*.¹⁶

Geoparsed data, originally in a semantic form, that describes events occurring in time and space, can then be further passed on for more complex queries (such as the proximity searches described in Workflow 1) and can enable new forms of analysis. Activity Based Intelligence, for example, insists on the use of data that is “georeferenced to discover” relevant information.¹⁷ For the purposes of network analysis, georeferenced entities can form the knowledge base for spatial queries and, in combination with adequate overhead imagery, can enable an analyst to discover hitherto unknown sites of interest.

Machine Classification

The machine classification workflow is an important one that has not yet been effectively deployed at KCL. It involves the automatic classification of articles — or any other object — based on how previous articles were classified. There are a wide number

of algorithms and approaches that can be taken for classification, ranging from naive Bayesian classifiers, which predict the likelihood of a classification based on how objects with similar subproperties were previously classified, to Subject Vector Machines, which bound clusters in a multidimensional array.

The key challenge in classification relates to the training of the algorithm. Training is generally undertaken in two ways. First, a known corpus can be fed into the classifier, perhaps after manual annotation. The classifier will then determine what elements are statistically the most relevant in affecting the classification, before using these in future prediction. The second, and complementary, approach is to provide feedback based on the predictions as to whether the classification is correct or incorrect. The challenges associated with classification relate primarily to the need for a substantial training corpus, particularly given that the corpus has to be language specific. If successfully implemented, classification can provide a confidence value that any article or object is related to any specific activity or issue. It could thus be used to map news articles and academic publications against a model of the nuclear fuel cycle, which is an end use under development at the IAEA. However, this approach could also be used more broadly to categorize every news article for proliferation relevance, which could greatly reduce the burden on a media-monitoring analyst.

Multimedia Object Detection

Another workflow of significant importance is detection and classification of objects of interest in non-text media. Such media includes images, sound and video, and is collectively called “multimedia.” Optical Characterization and Text to Speech are relatively mature fields, so they are not explored here. However, images and video are worthy of examination.

Images

Advances in neural networks have improved machines’ ability to recognize objects in images, and they offer promising opportunities for new analytical capabilities. The Google Vision API example shown in Figure 1 (left) demonstrates such a capability. Here, an image is analyzed via the API and automatically tagged with labels based on Google Vision’s machine learning capability. The tags, shown in Figure 1 (right), correctly identify the fact that the pictured device relates to pressure measurement. This device is also called a pressure transducer or capacitance manometer, and is manufactured by MKS Instruments.



Figure 1. Google Vision API

This capability could be readily applied to the nonproliferation sphere. It could be used to screen images posted to Internet trading platforms, such as Alibaba or eBay, to identify possible controlled items, for example. If it performs well on user-taken photos, it might also be an effective way for customs officials to determine whether an item that they have in front of them is likely to be controlled. It is unclear whether this will be practical, however, as when the author tried this with a photo of a stapler, the service identified it most likely as a computer mouse. Accuracy might thus depend on the service being able to match the image to another that was used in its training that was sourced from an existing image on the Internet — which, when it comes to controlled dual-use items, may be difficult to consistently find. As such, it is clear that such capabilities will need to be trained on a domain-specific image set.

Video

Video is a category of media that has been difficult to analyze up to this point without manual effort. Although it has been possible to extract the sound from video and run it through text to speech libraries, the capability to automatically identify and categorize objects or text in video has been lacking. However, this appears to be changing. Microsoft now offers APIs for extracting text from videos. Screenshots of items in videos could be treated as images and could be passed through services similar to Google Vision, as noted above.

A set of capabilities — video splitters — has been devised that splits videos into image segments for analysis. The analysis can take the form of OCR, object identification, or some other service. Such tools have utility in the nonproliferation domain, but perhaps only when applied to videos of higher potential relevance.

Facial recognition

Facial recognition is not a new topic. However, again, Microsoft

has recently launched API services for facial recognition that purport to be able to identify faces, recognize them, and even judge the emotions of the person. This could have real application to the counter proliferation domain, too. For example, there are hundreds of images of North Koreans taken with Kim Jong Un. But analysis of these media images is laborious — they must be geo-located and referenced, the images must be carefully situated in their context, and the individuals' identifying features must be recognized and resolved across multiple pictures.¹⁸

This list of workflows might not be complete, but it reflects the workflows identified as being relevant for Project Alpha at KCL in its work to identify previously unknown proliferation activities and facilities using a big-data approach. These workflows would complement those used by an international organization like the IAEA, even though the IAEA's mandate differs slightly in that it uses all-source information to verify a state's declaration. The primary adaptation required to introduce such workflows to an international organization like the IAEA would be to develop an additional workflow around alternate hypothesis generation. That is, they would be used to determine whether insights identified through these workflows provide new insight to prove or disprove a state's declaration. From this brief review of the workflows, it is apparent that a number of specialist processing capabilities is required. The next section of this paper examines these in turn.

Processing Data

For the above-mentioned workflows to be effectively implemented, a number of cross-cutting processing capabilities are required. The purpose of these processing capabilities is to take some form of data and enrich it in some way. The key technologies include NLP, machine translation, visualization, and geoparsing. The next sections discuss these capabilities, as explored at KCL.

Two processing capabilities that have not yet been fully exploited at KCL are also examined: machine learning and the use of ontologies. It should be noted that we have not devised sound statistical tests for each of these capabilities; as a result, our analysis of these capabilities is based more on our impression from using the tools rather than on a quantitative assessment of their performance. KCL will be developing more formalized performance metrics alongside its introduction of machine learning approaches to data processing. In fact, performance assessment is an inherent part of the machine learning process.



Natural Language Processing

An active effort was undertaken to compare the information extraction capabilities of different information extraction tools (see Table 2). The textbox below shows the text chosen for this test.¹⁹

An Iranian Aerospace Industries Official told the London-based newspaper Al-Sharq Al-Awsat that the Military Industry High Commission has reordered priorities in many major missile projects because of domestic and foreign factors. Some countries with strong ties to Iran alerted Tehran that its Zelzal ballistic missile programme, which includes Shehab-3, has caused regional and international concerns. Iran had reported successful tests of Shehab-3. The United States warned Iran from developing its missiles and required stopping the development of Shehab-4 with a range of 2400 km.

Ali Shamkhani, Iranian Defense Minister, decided to postpone the Shehab-4 project for an indefinite period of time and focus on shorter range and more accurate missiles such as Oghab, Nazeat, Fajr, Iran-130, Shehab-2 and Shehab-3 with the support of China and North Korea.

According to an Arab diplomat, postponing the Shehab-4 project is also a result of Iran's will "to lower tension, particularly in the Gulf." He added, "the Iranians are confident and feel that if they wait they will control the entire Persian Gulf without a fight."

The Iranian Defense Ministry decided to increase the budget of the Aerospace Industry Institution to focus on less costly programmes and to meet needs of the army, the air force and the special forces of anti-armor, air-to-surface and air-to-air missiles. Importing such missiles represents a financial burden and a logistic hindrance to the armed forces.

Thanks to eminent experts such as the late Ali Mahmoudi, the Iranian Aerospace Industries managed to produce prototypes of different missiles by upgrading old models of Raad, Toufan-1 Tondar and Towsan. The last version of the Towsan missile is a copy of the Russian "Choncors 113." The Raad-1 and Toufan-1 missiles are already in service in the Iranian regular army and the Islamic Revolution Guard Corps.

The Saeqeh-1 missile underwent changes to improve its guidance system. According to the Iranian Defense Ministry, the missile is very accurate. The Ministry also said that American and Russian technologies were used in its development. The missile looks like the US-made Dragon but its components are copies of Russian products.

Source: Rodan, S. (2001, August 2). Iran changes direction in missile development. Middle East Newline.

From this examination, it was apparent that the Rosette API provided the best out-of-the-box capability for the processing of this specific text. Each of the above services implements different NLP libraries with different training sets on their side. The above differences in results serve to highlight the technical barriers to full implementation of NLP algorithms. Without additional training and feedback from the user side, the Alchemy API, for example,

identified Shehab-4 and Towsan as cities, when they are actually items (weapons). Towsan is a *metonym*, a word that can refer to a place (a location in Iran with an oilfield) or an item (an antitank weapon), depending on the context.

Machine Translation

The large quantity of data available to nonproliferation analysts will not always be available in the languages read by the analyst. Indeed, it could be argued that it would be a mistake to overly rely on only those sources understood by the analyst when attempting to identify undeclared activity, because typically more information is available in the native language of the place being examined (if that language is not English) than is available in English. At the same time, it is understood that at least in the foreseeable future, machine translation will not be as good as a native speaker or a professional translator, particularly in relation to highly technical subject matter and terminology relevant to the nuclear domain.

Presently, two services are in most common use: Google Translate, through its Translate API, and the Bing Translator, provided by Microsoft. Of the two, Google Translate proved to be more accurate compared to the Bing Translator, but both appear poor in comparison to human translation by native speakers.

Another problem with machine translation is inconsistency in results. A review of translations repeated sometime after the original translation found that the translated text was not the same. The reason for this appears to be that the translation model is statistically based and resolved from user feedback from the crowd (crowdsourced). As such a feedback loop-based model improves, the choice of translation will differ, meaning that a machine translation captured at a particular time may differ. Although improved performance is to be welcomed, predictability and repeatability are also important to an evidence-based organization. Changing translations have an immediate impact on non-native speakers, creating hidden problems for analysis and potentially invalidating evidence items that were created in this manner.

Furthermore, for sensitive investigations, secure and localized access to information retrieval will be a strong priority. As a result, the use of cloud-based public services provided by Google and Bing for storage may not be appropriate. Perhaps a system comprising both cloud access and specialized access to training materials specific to the required technological domain (the nuclear domain, in our case) could be the answer to such a problem. An example of such a service is provided through SDL's Language Cloud (previously called Language Weaver). This



Table 2. Results comparison of different information extraction tools used on a sample text

Result Type	OpenCalais API	Gate API	Rosette API	Alchemy API
City	Fajr, London, Tehran	(no result)	(no result)	Shehab-4, Towsan
Company	Iranian Aerospace Industries	(no result)	(no result)	Iranian Aerospace Industries
Location	(no result)	China, North Korea, Iran, Middle East, Tehran, United States, the Gulf	North Korea, Iran, London, Middle East, Russian, American, Tehran, Gulf, China, Persian Gulf	Persian Gulf
Country	North Korea, China, Iran, Islamic Republic of Iran, United States	(no result)	(no result)	Iran, Tehran, United States, China, North Korea
Organization	Iranian Defense Ministry, Iranian Air Force, Iranian Army, Islamic Revolution Guard Corps, Military Industry High Commission, Aerospace Industry Association	Iranian Aerospace Industries, Military Industrial High Commission, Iranian Defense Ministry, Islamic Revolution Guard Corps	Iranian Defense Ministry, Islamic Revolution Guard Corps, Military Industry High Commission, Aerospace Industry Institution, Iranian Aerospace Industries Official, Al-Sharq Al-Awsat, Oghab Iranian Aerospace Industries, Ministry	Iranian Defense Ministry, Military Industry High Commission, Aerospace Industry Institution, Islamic Revolution Guard Corps
Person	Ali Shamkhani, Steve Rodan	Al-Sharq Al-Awsat, Ali Shamkhani, Ali Mahmoudi, Steve Rodan	Ali Shamkhani, Steve Rodan, Ali Mahmoud	Steve Rodan, Al-Sharq, Ali Mahmoudi, Ali Shamkhani
Position	Defense Minister, Official, Diplomat	(no result)	Defense Minister	Defense Minister
Product	Saeqeh	(no result)	Zelzal, Nazeat, Fajr, Toufan-1 Ton-dar, Raad-1, Toufan-1, Saeqeh-1, Dragon	(no result)
Nationality	(no result)	(no result)	Iranians, Iranian, Arab	(no result)
Dates	(no result)	2 August 2001	2 August 2001	(no result)
Distances	(no result)	(no result)	2,400 km	2,400 km, 113

service could also be localized, but with a commercial license that is based on language pairs.

Another approach to translation is the use of pretranslated keywords structured through an ontology. Substantial and multilingual dictionaries now exist in the nuclear space. However, there are evident problems with this approach in relation to coverage particularly for specialist terms. While there are a number of glossaries and thesauri already in the nuclear domain,²⁰ in our experience, these tend to be too high level and poorly structured for use in data processing. Furthermore, while such thesauri often are available in the major world languages, they often are not available in all languages or localized based on linguistic differences at the national or regional level.

Visualization Tools

There are any number of visualization tools designed to enable the analyst to understand and interpret data. Commercial plat-

forms such as Tableau, Microsoft Power BI, and Qlik are increasingly seen in the nonproliferation sphere as complements to the graph capabilities of Microsoft Excel to convey complex statistical data.²¹ There are also a large number of custom visualization libraries that play a more specialist role and that generally can be integrated into webpages. It is useful in the context of this paper to examine link analysis in some depth.²²

One particularly interesting area of the visualization tool development relates to link analysis. Link analysis has become increasingly important as data volumes have grown, and it has proven to be a powerful tool to understand the relationship between different data sets and between different entities. For example, KCL maintains a database of suspected DPRK-linked entities in China, including the telephone numbers, email addresses, and locations of these entities, and link analysis is used to show where two entities share the same details (thus suggesting that one is a front company for the other; see Figure 2).

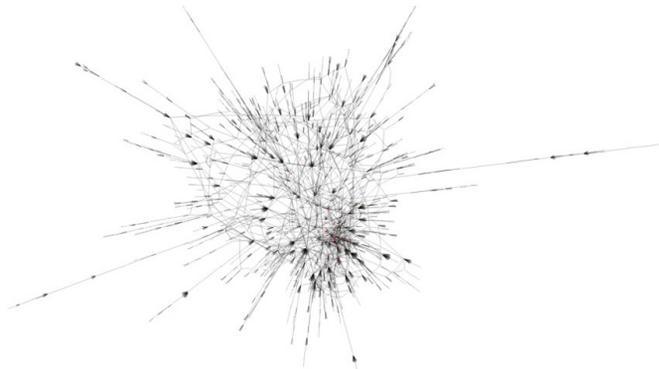


Figure 2. Link chart illustrating links between proliferation actors in an actual proliferation case

As we mentioned earlier, the traditional tools for link data analysis are I2 Analyst's Notebook, Maltego CE (an open source service), and Gephi. Palantir is known to have a mature capability in this area, too, although it was not available to the research team. There are limitations with the first three platforms because their default deployments are stand-alone software applications rather than web-based solutions that can integrate with structured databases.²³

There are also open source visualization libraries that allow for the dynamic creation of link charts based on any source of structured data. These libraries simply pass an array of data through a visualization algorithm to produce a chart, and many thousands of types of charts can be produced. Selection of the libraries appears to be as much down to the preferences of the user as to the merits of each library. Two libraries stand out: Data Driven Documents (D3) and vis.js.²⁴ KCL opted to use vis.js for its network diagrams in part because it proved to be better at handling large numbers of nodes.

Such libraries have benefits over tools such as Analyst's Notebook as described above. One particular benefit is that they can be rendered directly from webpages. The libraries do little more than render arrays of data; however, this highlights the importance of collecting data into structured stores. Recent developments in this space are promising for nonproliferation activity. The advent of graph databases such as Neo4j has brought new visualization libraries from which dynamic queries can be targeted simply by double clicking on the node on the graph. These capabilities provide analysts with dynamic new ways to explore data to discover new insights.

Geoparsing

A number of services can be used to extract (parse) location data

from text. These methods fall under the term "geoparsing." One method of data geoparsing involves location lookup by address. It is this approach that is used by services provided by Google, Bing, and Yandex. The second approach is more of a gazetted list approach and involves extracting place names directly from the text and correlating them to predefined geocoordinates. Perhaps the most widely used platform in this area is GeoNames, which describes itself as follows: "The GeoNames geographical database covers all countries and contains over 11 million place names that are available for download free of charge."²⁵ The authors have not derived a firm assessment on the accuracy of the GeoNames gazetter, which remains an opening for additional research.²⁶

As with NLP, there are trade-offs between the service-based and gazetted list approaches. The gazetted list approach can be implemented for free locally, meaning that data does not have to be sent to an external service provider. Other services are likely to be more accurate but have an associated cost and require transmission of data to service providers.

Over and above these considerations, it was apparent from KCL's experimentation with the services that no one service has the most accurate results on a global basis. Indeed, the Google map data for some countries, including China, appears to be inaccurate. As a result, it is likely that different services will perform more accurately for different countries, and further work is required to understand the relative merit of each.

Machine Learning

Machine learning offers substantial potential in assisting human analysts in extracting meaning from large quantities of data. Machine learning can have a direct role in classification and prediction. Several services mentioned herein rely to a greater or lesser extent on machine learning services. These out-of-the-box solutions, provided externally by commercial organizations, do not involve the user in the generation of the service.

Generally, it is a concern that the training of these external services is blind to the user — that is, by the time the model is being used by the analyst, it will already have been trained by the providers' specialists. However, providers of machine learning services have recognized this and offer bespoke training and development with the end user as an active participant but this is costly and can be difficult to manage. Without such training, it is generally not possible for the user to train or further train the model. This will have ramifications for the deployment of the service as part of a broader organization-specific system.



Of the approaches described above, the following include a machine learning element:

- Machine Categorization
- Machine Translation
- Natural Language Processing
- Object Detection in Multimedia

Beyond the services described already, machine learning can play a more direct role in content classification. There are a number of types of machine learning approaches that are relevant, although this list is incomplete.

Recommendation engines: Recommendation engines are common on Internet sales platforms. The algorithms attempt to guess at what items might interest the consumer based on their previous browsing history, and often also combining trends from other consumers. A recommendation in nonproliferation application space might offer up articles to analysts based on the confidence score that the article contains relevant information. The specific factors relevant for the algorithms, and the specific algorithms to be used are beyond the scope of this paper. In general terms, there is clearly a risk of selection bias in this approach.

Summarization engines: Another interesting class of tools summarizes text. These tools can be helpful in reducing the content to be read by the analyst. However, there is clearly a risk that such tools will remove the pertinent details unless sufficiently trained.

Sentiment analysis: A final class of interest tools relates to sentiment analysis. These tools can determine whether an article, sentence, or paragraph is positive or negative about a particular subject. It is not yet clear how these tools can be used for nonproliferation efforts, although it is possible that these tools could be used to understand, for example, popular sentiment for the state acquiring nuclear weapons.

Although the exact use cases for machine learning are still poorly defined in the nonproliferation domain, it is clear that some of these tools will have utility in nonproliferation in the future. It will be necessary to provide sufficient training sets so that these processes can work accurately in the nonproliferation domain. The capabilities described in other sections of this paper can provide useful content for training such algorithms in the future.

Storing Data/Architecture

The workflows described above require the integration of large amounts of data of many different types. Given the strong interconnection between the different types of data and the desire for automation, a single unified architecture is probably necessary.

There are perhaps many approaches to building such an architecture. However, some lessons have been learned from KCL's efforts to build such a platform.

Databases

The most common databases used worldwide are relational databases, which comprise sets of tables that are interlinked through some common data. In light of the huge amount of data acquired daily, such databases quickly reach their capacity limits. To overcome this problem, information management specialists have developed alternative database models to relational databases that are worthy of examination. The principle alternatives relevant for this paper are triplestores and graph databases, Lucene-based search indexes, and ontologies.

Triplestores

Triplestores were designed to enable a concept called the Semantic Web, in which information could be federated but integrated at the point of use. There are many advantages to Semantic Web technologies, including the ability to reconcile entity IDs because every entity has a unique reference identifier (URI). All data is stored in a triplet format, wherein each element is also a URI. So, in the example given in Figure 3, a URI exists for Shakespeare, Wrote, and King Lear.²⁷

Subject	Predicate	Object
Shakespeare	Wrote	King Lear

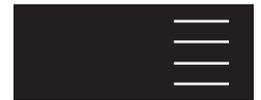
Figure 3. Triple structure

The concept of the triplestore and the Semantic Web has direct benefits to nonproliferation efforts. It would be easy to query all available data sources for all information on nuclear facilities in Iran, for example, with data from any location automatically being integrated into the correct schema. KCL experimented with the integration of a triplestore in its test-bed platform. However, we found that most of the types of data identified previously are not currently available in triple format, so this approach has not yet been pursued beyond proof of concept. It should be noted that the main data set that was useful — DBpedia — was available anyway as an API, allowing results to be stored in a relational database. Although this did not provide a solution to large-scale integration of triple-format information, it did effectively provide a way to utilize DBpedia entries as URIs in the test bed, as demonstrated in Table 3.



Table 3. Named entity recognition and DBpedia entries²⁸

Matched Term	DBpedia Entry	Assessment
Asharq_Al-Awsat	<i>Asharq Al-Awsat</i> (Arabic: <i>اَلْاَسْوَاطُ الْقَرْشِيَّة</i> , meaning “the Middle East”) is an Arabic international newspaper headquartered in London. A pioneer of the offshore model in the Arabic press, the paper is often noted for its distinctive green-tinted pages.	Correct
Tehran	Tehran (Persian: تِهْران – Tehrān, pronounced [teh r n]) is the capital of Iran and the Tehran Province. With a population of about 9 million in the city and 16 million in the wider metropolitan area, Tehran is the largest city and urban area of Iran, the second largest city in Western Asia and the third largest in the Middle East. It is ranked 29th in the world by the population of its metropolitan area. Large-scale demolition and rebuilding began in the 1920s, and Tehran has been a destination for the mass migrations from all over Iran since the 20th century.	Correct
Ali_Shamkhani	Ali Shamkhani (Persian and Khuzestani Arabic: <i>آلی شامخانی</i> , born September 29, 1955) is an Iranian Rear Admiral Upper Half. He is the secretary of the Supreme National Security Council of Iran.	Correct
Ministry of Defence and Armed Forces Logistics (Iran)	The Ministry of Defence and Armed Forces Logistics is the defense ministry of the Islamic Republic of Iran. It was reestablished in 1989. The ministry is responsible for planning, logistics, and funding of the Armed Forces of the Islamic Republic of Iran, whereas the General Staff, a separate institution, has control over the forces. The ministry is one of the three sovereign ministerial bodies of Iran due to nature of its work at home and abroad.	Correct
Oghab	The Iranian Oghab (Eagle) missile is an unguided 230 mm (9 in) artillery rocket with a range of 34–45 km (20–28 mi). It is spin-stabilized in flight but has a proven circular error probable in excess of 500 m (1,500 ft), making it a highly inaccurate weapon. It carries a 70 kg (150 lb) high explosive fragmentation warhead, though it may also be able to carry chemical warheads. According to U.S. sources, a modified version has been developed that could be carried and fired from Iranian Air Force F-14 Tomcat and F-4 Phantom II aircraft. The launcher is an elevatable triple-rail launcher assembly fitted to a Mercedes-Benz LA 911B 4x4 truck chassis.	Correct
IAIO_Toufan	The IAIO Toufan or Toophan (Persian: توفان, “typhoon”) is a series of combat helicopters produced by the Iran Aviation Industries Organization. Based on the U.S.-built AH-1J SeaCobra, the Toofan is produced in two variants: the Toufan I, unveiled in May 2010, was followed by the improved Toufan II, unveiled in January 2013.	Wrong
Tondar	Tondar (known as the Kingdom Assembly of Iran) is an organization of Iranian exiles based in Los Angeles.	Wrong
9M113 Konkurs	The 9M113 Konkurs (Russian: 9М113 «Конкурс»; cognate of French: <i>Concours</i> ; English: “Contest”) is a semi-automatic command to line of site (SALCOS) wire-guided antitank missile of the Soviet Union. “9M113” is the air Missile and Artillery Directorate of the Ministry of Defense of the Russian Federation (GRAU) designation of the missile. Its North Atlantic Treaty Organization (NATO) reporting name is AT-5 Spandrel.	Correct. Also note that the spelling differs.
Raad-1	Raad-1 (“Thunder-1”) is an Iranian self-propelled howitzer.	Probably wrong. Article refers to Raad as a missile.
Army of the Guardians of the Islamic Revolution	The Army of the Guardians of the Islamic Revolution (Persian: <i>سپاه پاسداران انقلاب اسلامی</i> <i>Sepāh-e Pāsdārān-e En-qelāb-e Eslāmi</i> , or <i>Sepāh</i> for short) — often called Revolutionary Guards or Islamic Revolutionary Guard Corps by the English-speaking or U.S. media — is a branch of Iran’s Armed Forces, founded after the Iranian revolution on May 5, 1979. Whereas the regular military (<i>artesh</i>) defends Iran’s borders and maintains internal order, according to the Iranian constitution, the Revolutionary Guard (<i>pasdaran</i>) is intended to protect the country’s Islamic system. The Revolutionary Guards state that their role in protecting the Islamic system is preventing foreign interference as well as coups by the military or “deviant movements.”	Correct
HESA Saeqeh	The HESA Saeqeh (Persian: <i>هتق عاص</i> , “thunderbolt”) — alternatively spelled Sa’eqeh, Saegheh, or Saeqeh-80 — is an Iranian-built, single-seat jet fighter, derived from the American Northrop F-5. A joint product of the Islamic Republic of Iran Air Force and the Iranian Ministry of Defence, it is the second-generation of the Iranian Azarakhsh fighter. Saeqeh aircraft were tested successfully in Iran on September 20, 2007.	Wrong — it’s a missile, not a fighter.
Rostec	Rostec (Russian: <i>Ростех-Ростekh</i>), formerly Rostekhnologii (Russian: <i>Ростехнологии</i>), is a Russian state corporation (a nonprofit entity type) established in late 2007 to promote development, production, and export of high-tech industrial products for civil and defense sectors. It brings together 663 entities which form 13 holding companies. Eight holding companies operate in the defense-industry complex, and five of them are involved in civil sectors. Rostec’s organizations are located in 60 constituents of the Russian Federation and supply goods to over 70 countries worldwide.	Correct
Steve_Rodan	Born on April 19, 1954, in Glasgow to Robert W. Rodan and Betty Rodan, he is a qualified chemist. He resides in Laxey, the most populated village in Garff. Politically, he is liberal, having chaired both the Heriot-Watt University Liberal Club and Scottish Young Liberals. He was also a member of the National Executive of the Scottish Liberal Party and stood unsuccessfully for the House of Commons in 1979 for the constituency of Moray and Nairn, being defeated by the Conservative Alexander Pollock.	Wrong
BBC Newsline	<i>BBC Newsline</i> is the BBC’s regional television news service for Northern Ireland. The program is broadcast on BBC One Northern Ireland from BBC Northern Ireland’s headquarters in Broadcasting House, Ormeau Avenue, Belfast.	Wrong



Graph Databases

Graph databases have received a great deal of attention in recent years as a more scalable alternative to relational databases. Facebook operates a graph database for user data. The Panama Papers (offshore leaks) are also available in a graph database. Graphs differ from triples in that every piece of data has two elements: the node and the edge. The node defines the piece of data, whereas the edge defines how this node links to other data. There are several graph database technologies. Perhaps the most exciting is Neo4j, the graph database that powers the Panama Papers,²⁹ because it is easy to deploy and is offered as an open source application.

Graph databases are attractive because of their scalability and speed at querying large data sets. However, there are drawbacks with graph databases. At present, there is no direct way to integrate graph data with a relational database that hosts the current data sets. This results in the need for a relatively complex integration using the graph database API to return results that can then, in turn, be saved to a relational database. Graph databases also do not integrate directly with the major search services, such as Solr or Elasticsearch (because Neo4j, for example, is a search index in its own right). This creates challenges in integrating data stored in relational databases with data stored in graph databases, particularly if a search should reconcile the results.

Lucene-Based Search Architectures

A third architectural approach of interest relates to Lucene-based search engines. Such engines can be run locally or on servers and can be made to index large quantities of structured or unstructured data stored locally or remotely. These tools also integrate with powerful web crawling tools, such as Apache Nutch and ManifoldCF, which can index either local data repositories or portions of the Internet. These tools can also be integrated with many of the capabilities examined in this paper, including NLP, geoparsing capabilities and machine learning capabilities, such as clustering.

The two leading open source search engines at present are Elasticsearch and Solr search, both of which can cope with large volumes of data with scalability. Both support the advanced search tools described earlier, including faceted search, Boolean operators, and stemming. Solr appears to be more widely used and proved to be more reliable when used in the KCL test bed. Elasticsearch, by contrast, has more integration with advanced tools for data analysis, including Kibana, which is produced by the

same company that maintains Elasticsearch.³⁰ As a result, Elasticsearch is often referred to as a tool for big data analysis in a way that Solr is not.

These extra integrations of Elasticsearch appear to be relevant for nonproliferation analysis. The Kibi platform from Siren solutions, for example, which integrates with Elasticsearch, provides a business intelligence analysis and visualizing platform based on data in the Elasticsearch index using the link visualization library, Keylines.³¹ These capabilities have not yet been fully explored at KCL. However, KCL has found that these tools can be integrated with common web platforms such as Drupal relatively easily, and KCL has found both to be comparable for indexing vast quantities of data.

Ontology- and Knowledge-Graph-Based Approaches

A final architectural option available to support big data analysis relates to ontologies, which are structured lists of domain-specific data. Ontologies can be defined as a formal, explicit specification of a shared conceptualization.³² In the computing science community, the phrase conveys a specific meaning of a structured vocabulary that shows how entities are related to one another. An extension of the ontology is the knowledge graph, which not only maps out how categories map to one another, but which also contains all the relevant entities for a domain. Knowledge graphs can be very large in size. For example, DBpedia, a knowledge graph based on Wikipedia, has nearly 400 million entries describing objects (that is, triples).³³ Ontologies and knowledge graphs are useful tools for text categorization and potentially for machine categorization, because the structure of the ontology can be used to provide context. However, it is a substantial undertaking to develop a functional ontology for any specific domain. Such an ontology must also be maintained and updated. As such, although a nonproliferation ontology offers promise for improving data categorization and analysis in the nonproliferation domain, the merits and costs are not yet fully understood. Nonetheless, development of such a knowledge graph for the nonproliferation domain would provide a foundation for implementing the capabilities examined in this paper.

While this section has outlined four different approaches to systems architecture for big data analysis in the nonproliferation domain, it should be noted that each approach has different pros and cons and that integration of these different approaches is far from straightforward from either a conceptual or practical perspective. As such, the systems architect must decide which



approaches to pursue for different purposes.

Conclusions

Open source information has long been an important but under-tapped resource for counter proliferation efforts, even though it is heavily relied on by organizations like the IAEA. To fully realize the potential of open source information, a great deal of domain-specific development work is required in order to leverage the advanced tools being created by the leading data analysis companies, many of whom will have little — if any — knowledge of the problem of proliferation. A key problem is that of integration of different tools and platforms. However, as this paper has demonstrated, the increasing availability of web services through which the latest capabilities can be offered as a service provides a simple and relatively cost-effective way to conduct analysis of large data sets. Such services appear likely to grow in number and capability in the future. Furthermore, local search platforms such as Solr and Elasticsearch provide platforms to index and process large volumes of data.

The growing capabilities of automated systems to harvest open source information leads to important questions about how organizations and communities of analysts can cooperate to maximize the utility of such systems. This paper has made the case for the use of automation to assist human analysts and has identified a number of specific workflows and tools for this purpose. It is clear that these tools cannot replace the human analyst. Despite this, what is clear is that automation capabilities have the potential to greatly enhance information collection toward nonproliferation objectives in the years ahead.

Keywords

Open source collection, information fusion, natural language processing, link analysis, geoparsing, machine translation

Author Biographies

Ian J. Stewart is a senior researcher in the Department of War Studies at King's College London (KCL), where he heads Project Alpha. He is also the training advisor for the consortium implementing the European Union's Partner-2-Partner program on dual-use goods. He is a specialist on issues related to export controls, sanctions, and nonproliferation. An engineer by training, Ian has many years of experience in government and working with companies and international organizations in relation to nonproliferation matters. He was also a research fellow at Harvard University's Belfer Center for Science and International Affairs.

Ahmed ElGebaly is a nonproliferation expert with more than 18 years of expertise with the International Atomic Energy Agency (IAEA). He was part of the team of experts contributing to the inspection team in Iraq, and he was part of the Safeguards task force that tracked and monitored the activities of the A.Q. Khan proliferation network. He is a specialist in monitoring the movement of commodities of potential use in weapons of mass destruction programs. At the IAEA, he also developed and implemented its safeguards concept for performing industrial infrastructure studies to identify industries supporting nuclear activities in different states. He managed several cooperation projects between the IAEA and member states as well as international organizations such as the European Union. Ahmed is a chemist by education and holds a MSc in knowledge management. He has extensive experience in data collection and management. He is also experienced in development and implementation of web-based applications for different purposes in the nuclear domain.

Alexander Lee is a research assistant at Project Alpha, KCL, where he is working on illicit trade and proliferation networks. He has relevant background in open source monitoring and investigations in Syria and Yemen. He holds a BA in politics from the School of Oriental and African Studies and an MA from the Department of War Studies at KCL. His research focuses on scientific and technical mapping, event detection and verification, and data fusion models for mixed-type data, especially for the generation of Common Operational Pictures.

References

1. On the matter of indicators of nonpeaceful use and minimum and elevated indicators of peaceful use, see Seward, A., Matthews, C., & Kessler, C. (2008). Evaluating nonproliferation bona fides. In J. E. Doyle (Ed.), *Nuclear safeguards, security, and nonproliferation* (263–281). Burlington: Butterworth-Heinemann.
2. Additional protocol. (2018). Retrieved from www.iaea.org/topics/additional-protocol
3. On the “80/20 rule” of secret to open intelligence, see Gibson, S. D. (2014). Exploring the role and value of open source intelligence. In C. Hobbs et al. (Eds.), *Open source intelligence in the twenty-first century: New approaches and opportunities* (9–23). London: Palgrave Macmillan.
4. Coined by Gary Kasparov; literature on said concept includes McAfee A. & Brynjolfsson E. (2017). *Machine, plat-*



form, crowd: *Harnessing our digital future*. New York: W.W. Norton & Company.

5. Project Alpha undertook a quasi-big data analysis approach in using trade data to identify front companies thought to be acting on behalf of Pakistan's nuclear program. This work culminated in the discovery of a suspected centrifuge facility under construction. See Project Alpha. (2018). Pakistan's strategic nuclear and missile industries: A baseline study for non-proliferation efforts ('public' version). Retrieved from <http://projectalpha.eu/wp-content/uploads/sites/21/2018/06/31may2018-public.pdf>
6. Retrieval means include the following: RSS: RDF Site Summary, a form of web feed that outputs machine readable updates; BigQuery: Google's proprietary serverless data query and storage service; API: Application Programmable Interfaces, a set of defined software interactions that lets users connect and access software; and Direct Download: a simple interaction whereby a user retrieves data from a host (client-to-server) in a predefined format.
7. GDELT describes itself as "the largest, most comprehensive, and highest resolution open database of human society ever created. Creating a platform that monitors the world's news media from nearly every corner of every country in print, broadcast, and web formats, in over 100 languages, every moment of every day and that stretches back to January 1, 1979 through present day, with daily updates, required an unprecedented array of technical and methodological innovations, partnerships and whole new mindsets to bring this all together and make it a reality." GDELT. (2018). Retrieved from www.gdeltproject.org/about.html
8. For a recent example, see Panda, A. (2018, July 13). Exclusive: Revealing Kagson, North Korea's first covert uranium enrichment site, *The Diplomat*. Retrieved from <http://thediplomat.com/2018/07/exclusive-revealing-kangson-north-koreas-first-covert-uranium-enrichment-site/>
9. Eldridge, C., Hobbs, C., & Moran, M. (2018) Fusing algorithms and analysts: open-source intelligence in the age of "Big Data." *Intelligence and National Security*, 33(3), 391-406.
10. Boolean operators in search are now common and might no longer be considered "advanced." However, their importance can help in refining search results from "Iran nuclear" to "Iran AND nuclear MINUS Bushehr MINUS negotiations." Boolean searches may be complemented by date range searches, such as those available on the Google advanced search page.
11. Openstreetmap Nominatim. (2018). Retrieved from <http://nominatim.openstreetmap.org>
12. Echosec. (2018). Retrieved from www.echosec.net
13. The .anb format is associated with IBM Analyst's Notebook application.
14. Gazetting is the preparation of text for further processing using other tools and techniques and storing such a preparation for future use.
15. Abernathy, D. (2017). *Using geodata and geolocation in the social sciences*. London, England: SAGE.
16. Gritta, M., Pilehvr, M.T., Limsopatham, N., & Collier, N.H. (2017). What's missing in geographical parsing? *Language resources and evaluation*. Retrieved from <http://doi.org/10.1007/s10579-017-9385-8>
17. Biltgen P. & Ryan, S. (2016). *Activity based intelligence: Principles and applications*. Norwood, MA: Artech House.
18. The UN Panel of Experts and entities, such as NK Leadership Watch, track and watch for the presence of certain individuals of interest in public images. See, for example, *Report of the UN Panel of Experts* established pursuant to Resolution 1874, S/2017/150, point 52.
19. This text is an extract from a CNS publication for NTI entitled "Iran Missile Chronology." It is available online at www.nti.org/media/pdfs/iran_missile.pdf?_=1316474223
20. See, for example, <http://untermportal.un.org/UNTERM/portal/welcome>
21. See, for example, "The CNS North Korea Missile Test Database," Centre for Non-proliferation Studies, available online at www.nti.org/analysis/articles/cns-north-korea-missile-test-database
22. Libraries are precompiled code that perform a specific function. In this case, libraries produce a chart based on an input datastream.
23. There might be exceptions to this. Analyst's Notebook is supposed to link to the database, for example, to allow dynamic querying, but this integration was never successfully executed at KCL.
24. For an example of D3 charts, see <http://github.com/d3/d3/wiki/Gallery>. For an example of vis.js charts, see <http://visjs.org>.
25. Geonames. (2018, June 26). Retrieved from www.geonames.org
26. Ahlers, D. (2013). Assessment of the accuracy of GeoNames



- gazetteer data. *Proceedings of the 7th Workshop on Geographic Information Retrieval* (pp. 74–81). Orlando: ACM.
27. Allemang, D. & Hendler, J. (2011). *Semantic web for the working ontologist: Effective modelling in RDFs and Owl*. Burlington: Elsevier.
 28. DBpedia Spotlight also matched the following terms, but they were considered too generic to warrant inclusion in this table: High Commissioner, Missile, Iran, International, China, North Korea, Arab League, Diplomacy, Persian Gulf, Special Forces, Surface to Air Missile, Russian_Armed_Forces, Ali, Middle East.
 29. Hunger, M., & Lyon, W. (2016). "Analyzing the Panama papers with Neo4j: Data models, queries & more." Retrieved from <http://neo4j.com/blog/analyzing-panama-papers-neo4j>
 30. Kibana. (2018). Retrieved from www.elastic.co/products/kibana
 31. See Keylines. (2018). Retrieved from <http://cambridge-intelligence.com/keylines/>
 32. Borst, as read in Guarino, N., Oberle, D., & Staab, S. (2009). What is an ontology? In S. Staab and R. Studer (Eds.), *Handbook on ontologies* (1–17). Berlin Heidelberg: Springer-Verlag.
 33. DBpedia data sets. (2018). Retrieved from <http://wiki.dbpedia.org/Datasets>



Inferring the Operational Status of Nuclear Facilities with Convolutional Neural Networks to Support International Safeguards Verification

Zoe N. Gastelum, Timothy M. Shead

Sandia National Laboratories, Albuquerque, New Mexico, USA

Abstract

International nuclear safeguards analysts use images in myriad ways to support verification analysis tasks, from analyzing the design and construction of a facility to understanding the scope and capacity of work performed therein. Potentially relevant ground-based imagery from open sources has increased significantly in the past 10 years as individual users with smart phones have become “citizen sensors,” posting geolocated content to social media platforms in near real-time. While this is an exciting new source of data for analysts, it is impractical to review unaided. The authors use machine learning to make image search and prioritization more efficient for safeguards analysts in three potential workflows. In this paper, the authors demonstrate the successful use of cooling towers and steam plumes as a signature that can indicate a facility’s operational status and describe a convolutional neural network modeling approach that yields over 90 percent accuracy for identification of cooling towers and steam plumes from open source ground-based images.

Introduction

The International Atomic Energy Agency (IAEA) verifies that states are acting in accordance with their international commitments to maintain nuclear materials and facilities for solely peaceful purposes. They accomplish this through a variety of field activities, including nondestructive measurements, sampling of nuclear materials and particulates, facility record audits, and providing continuity of knowledge over containers and areas with containment (tags, seals, and tamper-indicating devices) and surveillance measures. Activities performed at IAEA headquarters include analysis of states’ declarations of their nuclear activities, measurement and sampling results, notes and observations from in-field activities, and other information available to the IAEA, such as news media, state, and nuclear industry promotional materials, open source ground-based and satellite imagery, trade information, and scientific and technical journal publications.

While analysis of visual information such as photographs

has been recognized as an important aspect of the IAEA’s open source information analysis activities,¹ the vast stream of open source image data available on the Internet is too broad for safeguards analysts to review unaided. Machine learning algorithms, particularly those developed for image classification and labeling, hold promise for supporting safeguards analysts in their search and prioritization of images for further analysis.

We describe our experiments using convolutional neural networks (CNNs) to classify open source photographs of hyperboloid-shaped cooling towers at nuclear facilities to determine the plants’ operational status, and we discuss the validity of cooling towers and steam plumes as a signature for the operational status of a nuclear power plant.

Open Source Information Analysis for International Safeguards

In 1992, following the dismantlement of South Africa’s nuclear program, challenges verifying nuclear materials in the Democratic People’s Republic of Korea, and the discovery of a nuclear weapons program in Iraq, the IAEA recognized a need to verify the “completeness” of state safeguards declarations in addition to their “correctness.” This meant verifying the absence of undeclared or clandestine nuclear activities within a country. As part of the IAEA’s 93+2 Program, multiple strengthened safeguards measures and obligations were proposed — some of which were implemented immediately for all states (“Part 1” measures), while others required additional legal authority for the IAEA and were implemented as part of the Model Additional Protocol (“Part 2” measures). A Part 1 measure that is now a routine IAEA safeguards verification activity is “improved analysis and evaluation of all relevant information available to the Agency,”² including open sources of information relevant to a state’s safeguards agreement.

A 1995 IAEA report on the development and progress of the 93+2 program, GOV/2784, describes the “comprehensive approach towards the acquisition, processing and evaluation of information available to the Agency about a state’s nuclear



activities” as coming from three general sources:

Information provided by a state as part of its comprehensive or strengthened/expanded safeguards declarations

Information derived from in-field safeguards activities, including “strengthening measures” such as environmental sampling

Any other information available to the Agency, “including the public media, scientific publications and existing Secretariat databases ... as well as other information made available by Member States.”³

As a result, the IAEA Department of Safeguards performs expert, specialized collection and analysis, which “requires knowledge of information sources, search techniques, collection and filtering strategies, and finally collation, storage and retrieval methodologies.”⁴

Multimedia Analysis Challenges

Open source safeguards information collection and analysis has traditionally focused on text, but a growing share of open source data is multimedia information. Multimedia includes “photographs, video and audio recording, artist illustrations, computer-generated imagery (CGI), site and facility diagrams, process flowsheets” and other visual information, and can convey important information that would not otherwise be known based on text-only content.⁵ It appears frequently in open sources, including news media, government and industry website and print information, scientific and technical publications, patent databases, and social media platforms. Photographs of a nuclear facility may tell a different or more nuanced story than the accompanying text and may “provide key contextual information regarding their use in a location relevant to safeguards” that might not appear in text-based information sources.⁶

Although it has significant potential to add value to safeguards information analysis, multimedia information is more difficult to collect, store, and process than text-based information. For this reason, the IAEA Department of Safeguards’ most recent Research & Development Plan prioritizes strengthening and integrating all-source safeguards collection and analysis. The IAEA calls out the need to enhance its tools, further integrate all-source information and safeguards information analysis, and to make use of “new sources of openly available information, including from multimedia, and address the associated information management needs.”⁷

Compounding the need for analytical tools to support multimedia analysis is the sheer volume of potentially relevant multimedia information available from open sources. According to

Internet Live Stats, as of 2018 there were 836 photos uploaded to Instagram per second.⁸ The goal of our research is to develop machine learning techniques that can be used by an analyst to search for and prioritize the fraction of those images that are safeguards-relevant as part of a multimedia data analytic workflow.

Use Case: Operational Status of a Nuclear Facility

One example of a safeguards information analyst’s duties is to assess a nuclear facility’s operational status based on evidence from open sources. This can be accomplished, for instance, by cross-referencing press releases, statements, and other information from utilities, nuclear fuel cycle service providers, and nuclear regulators. To supplement this text-based assessment, and as a proof of concept for open source multimedia analysis, we hypothesized that the combination of a hyperboloid-shaped cooling tower and a steam plume within an image could provide a useful signature, signaling that a facility was operating. We imagined that an analyst could use machine learning models that were trained to recognize safeguards-relevant signatures in their workflow in multiple ways:

- To retrieve images of cooling towers (or steam plumes) from internal or external data streams for completely manual review
- To prioritize images for analyst review based on the confidence of automated assessments, allowing the analyst to focus on difficult (low-confidence) images
- To automatically compare photographic evidence of the operational status of a facility to a state’s declaration, “flagging” images for analyst follow-up when the two disagree

To validate these ideas, we created machine learning models that implement the proof-of-concept *cooling tower + steam plume* signature. Of course, we recognize that in practice the signatures and models deployed for use in a safeguards environment will likely involve much more nuanced processes, equipment, and analysis.

Data

For our experiments, we used the Flickr API⁹ to search for images using the query strings “cooling tower,” “nuclear cooling tower,” “nuclear generating station,” “nuclear power,” and “nuclear power plant,” which were chosen for their relevance to cooling towers and nuclear energy. We retrieved the top 1,000 images from each query, for a total of 5,000 images. After removing duplicate imag-

es that matched more than one query, we were left with a total of 4,460 images. A majority of the images contain structures and content related to nuclear power plants. However, there are a handful of exceptions, including photographs of unrelated subjects (e.g., a close-up of an insect), photographs that may be relevant to a facility but do not contain the facility itself (e.g., a sign indicating the facility name and direction), or photographs that may be near a nuclear facility but do not contain the facility itself (e.g., a landscape not including the facility), along with images accidentally mislabeled by the user.

We manually labeled each image based on whether it contained a hyperboloid cooling tower and whether it contained a steam plume. One subtlety that arose during labeling was that a significant subset of our cooling tower images (11.8 percent) were photographs taken inside decommissioned cooling towers. Because we were concerned only with images of cooling towers at operational facilities with the potential to produce steam plumes, our *cooling tower* label was reserved for images of cooling tower exteriors. In total, 1,803 (40.4 percent) of our images contained cooling tower exteriors, and 1,193 (26.7 percent) contained steam plumes.

Sample images from our dataset can be seen in Figures 1–3.



Figure 1. Sample hyperboloid cooling tower from the Drax coal-fired power station in the United Kingdom. This image is a reminder that not all cooling towers are nuclear cooling towers, and it reinforces the importance of location information for operational status assessment. Image credit: Jonathan Brennan, via Flickr. Image number 3635530311, 6/12/2009.



Figure 2. Interior of a cooling tower. Although interior images were not used for this work, they could be included in future research as indicators that the facility is not operating. Image from the cancelled Satsop nuclear plant in Washington, USA. Image credit: camknows, via Flickr. Image number 14117610685, 5/01/2014.

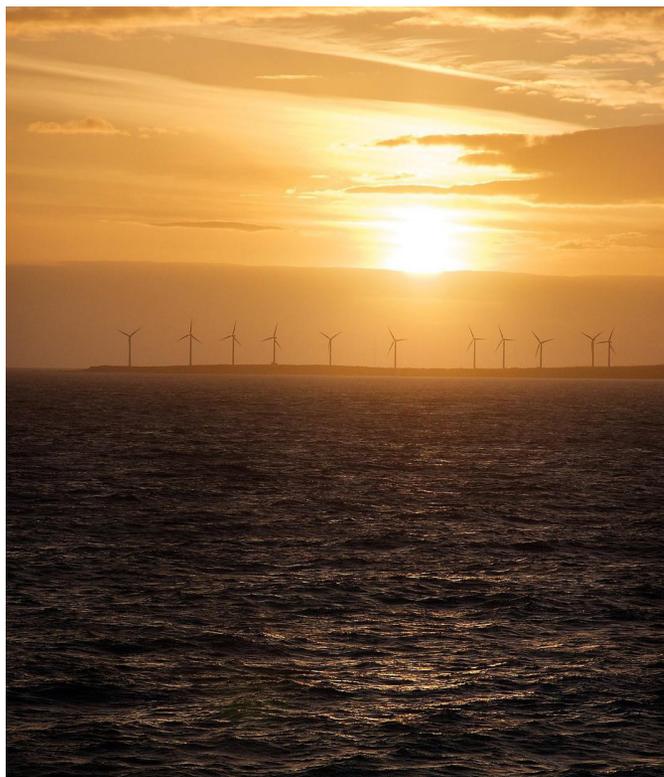


Figure 3. An image returned by our “cooling tower” search, but which is not a cooling tower. Image credit: Michael Foley via Flickr. Image number 13444809543, 3/24/2014.



Signature Validation

With the images labeled, our first task was to validate our assumption that the *cooling tower + steam plume* signature could be effective at determining the operational status of a nuclear facility. To do this, we identified the subset of images that contained cooling towers (based on our manual labeling) and were geolocated within the United States, based on the explicit location included in the Flickr metadata accompanying each image. This set contained 272 images. We did not validate locations, nor extrapolate them from other metadata features such as image tags or titles. We then used a 40-mile search radius and overhead images from Google Earth¹⁰ to identify the corresponding nuclear power plants. Some of the cooling towers in our images were associated with coal-powered plants, and others did not have a recognizable facility within the search radius, which reduced our validation data to 236 images taken near 22 distinct U.S. nuclear power sites. We defined a “site” as a nuclear power facility with one or more units with the same name. For example, we treated Units 1 and 2 of the Byron Nuclear Generating Station as a single site.

‘A substantial percentage (27.1 percent) of the images came

from sites where the nuclear power plants were shut down, decommissioned, undergoing decommissioning, or cancelled, but still contained cooling towers. For example, Figures 2 and 4 were taken at the cancelled Satsop Nuclear Power Plant site in the state of Washington in the United States. The prevalence of images of shutdown sites in our dataset was expected, due to access restrictions at nuclear power plants that can make photography of operational facilities challenging. Open source images of shutdown or decommissioned reactors could be used as an indicator of continued shutdown status of those sites.

Because all 236 images contained cooling towers, any image that also had a steam plume could be considered “plant operating” according to our signature, and any image without a steam plume was “plant not operating.” For ground truth information to evaluate these predictions, we used the U.S. Energy Information Agency’s (EIA) daily nuclear outage statistics,¹¹ counting a facility producing 50 percent or more of its operational capacity as “plant operating.” Comparing this ground truth to the *cooling tower + steam plume* predictions, we determined the predictive power of the signature (see Table 1).

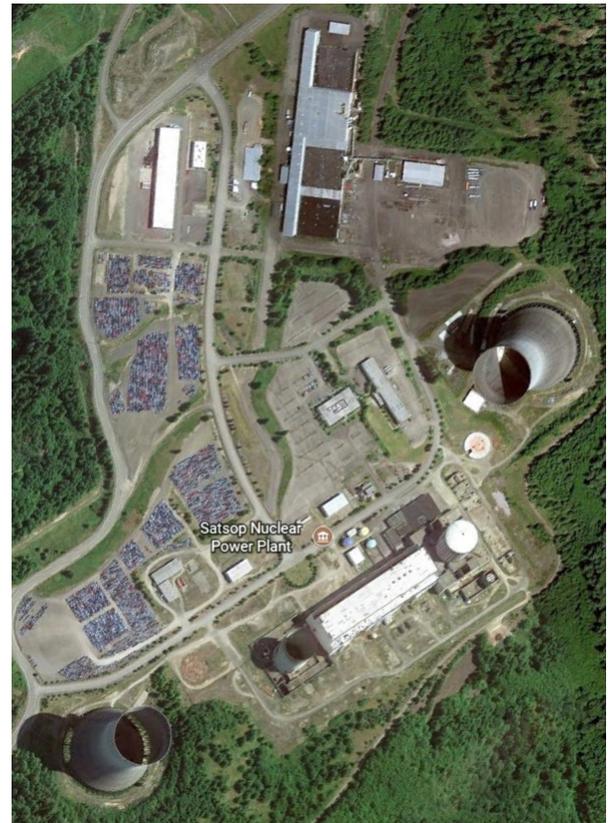


Figure 4. The cancelled Satsop Nuclear Power Station. **Left:** An intact cooling tower at the Satsop site. Image credit: Sharkhats, via Flickr. Image number 6115042441, 9/4/2011. **Right:** Overhead image of the Satsop site. The tightly parked vehicles indicate that the site is being used as a junkyard. Image credit: Google Earth, 2018.

Table 1. Cooling tower + steam plume predictive power

	Precision	Recall
Plant operating	0.96	0.95
Plant not operating	0.89	0.91

Here, we use precision and recall to assess our performance, following common practice within the machine learning community. Precision is the number of true positives divided by the sum of true positives and false positives, and is a measure of the ability to avoid false positives. Recall is the number of true positives divided by the sum of true positives and false negatives (i.e., the relevant elements), and is a measure of the ability to avoid false negatives.

We see from Table 1 that when a U.S. nuclear power plant was operating, the signature correctly predicted it was operating in 95 percent of the images (recall). Conversely, when the signature predicted that the plant was operating, it was correct 96 percent of the time (precision). The “Plant not operating” results, where the absence of a steam plume indicates that a plant was not operating, were slightly less accurate, likely due to atmospheric conditions, time of day, or other factors that prevented a visible steam plume at an otherwise operating facility. Other factors that could impact the performance of the *cooling tower + steam plume* signature include incorrect date or location metadata and misattribution of a cooling tower from a nearby site. For an example of the latter, see Figure 5, in which the Beaver Valley Nuclear Power Station (bottom left) has two hyperboloid-shaped cooling towers and is located adjacent to the Bruce Mansfield coal-powered plant (top right) with three hyperboloid-shaped cooling towers.



Figure 5. Adjacent facilities with hyperboloid-shaped cooling towers. **Left:** The Beaver Valley Nuclear Power Station. **Right:** The Bruce Mansfield Coal-Powered Plant. Image credit: Google Earth, 2018.

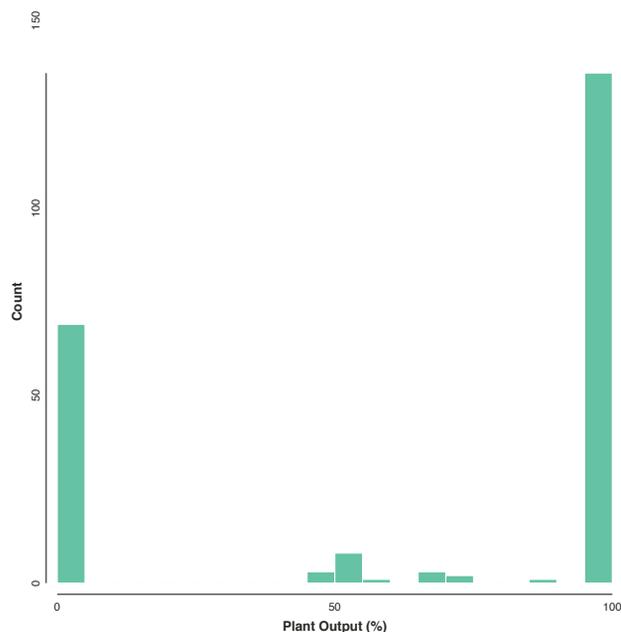


Figure 6. EIA plant output from the 236 images in our validation set, which includes a time period ranging from 2001 to 2016.



Changing the threshold at which we considered a plant to be operating had little effect on these results, because most nuclear power plants in the United States are operated as base load energy near full capacity (98–100 percent) or not at all (0 percent). Only a handful of plants in the U.S. EIA data were operating near 50 percent, presumably when one of two units at a site was operating at full capacity (see Figure 6). Although the focus of this work was on identifying operating sites, we note that verification of nonoperation could be equally interesting from a safeguards standpoint, both for operational and decommissioned reactor sites.

From these results, we concluded that the *cooling tower + steam plume* signature could provide a useful signal for identifying the operational status of a facility. Note that the accuracy of the *cooling tower + steam plume* signature puts an upper bound on the utility of the machine learning experiments that follow — that is, a machine learning model that correctly identifies cooling towers and steam plumes 100 percent of the time will only be correct 96 percent of the time when predicting that a site is operating.

Method

Convolutional Neural Networks for Image Classification

Broadly, *machine learning* is the art of teaching a machine to make predictions based on past observation. In *supervised machine learning*, a model is trained using a set of input observations, each of which has one or more ground truth labels. During the training process, a learning algorithm identifies correlations between the observations and their labels. Later, the trained model can be used to predict the labels for new, unlabeled observations, a task we refer to as *classification*.

In our case, we wanted to train a machine learning model that would correlate images of cooling towers (or steam plumes) with their labels, so that it could later identify cooling towers (or steam plumes) in unlabeled, previously unseen images. For this image classification task, the current state of the art uses *artificial neural networks*, which are based on a naive model of the interactions between neurons in a living brain. For each observation sent to the network, a vector of input features passes through layers containing activation units, producing a vector of output features. Typically, these networks are trained using gradient descent and backpropagation to learn a set of parameters that minimize a loss function. The loss function maps the output of the network to an error value, and the network is trained by iteratively

adjusting its parameters to minimize that error. Gradient descent and backpropagation efficiently determine how the parameters should be adjusted.

Observations typically are submitted to the network in small batches during training to amortize the cost of the backpropagation process, until the network has been updated with every available observation, completing a single epoch. The process then repeats until a stopping criterion determines that model training is complete. Often, experimenters train more than one model for a given problem, varying hyperparameters that parameterize the architecture of the model and learning algorithm, to identify those that produce the best models.

Neural networks have been used to great advantage in many domains, including speech recognition,¹² computer-based translation,¹³ new content generation,¹⁴ and automated image captioning.¹⁵ Some of the best-known applications of neural networks are for image-classification tasks, in which the output features of a network encode a fixed set of labels and the network predicts which label(s) apply to an input image. Popular examples include recognition of handwritten digits¹⁶ and photographs¹⁷ as part of image classification contests.¹⁸

Much of the current work on image classification relies on CNNs,¹⁹ specialized variants of artificial neural nets that are further inspired by the organization of the mammalian visual cortex. CNNs replace general-purpose activation layers with special-purpose convolutional and pooling layers that respond to spatial structure in images, organized to extract low-level features from an image (edges, colors), which are combined in subsequent layers to create higher-level features of increasing abstraction and complexity. These high-level features are then fed to fully connected activation layers that implement the logic to classify an image based on the combination of features that it contains.

There are many high-quality image classification platforms, both open source and commercial, and these platforms are continuously improving. However, the labels provided by these services when we began our research were too ambiguous for our use case. For example, as of February 2017, an ad hoc (neither exhaustive nor statistically sound) experiment conducted using the Google Cloud Vision API²⁰ demonstrated that a picture of a nuclear cooling tower might be assigned any of a variety of irrelevant labels, including “power station,” “outdoor structure,” “skyscraper,” “architecture,” “sky,” “atmosphere,” or “tower.” The labels varied considerably depending on the content of the photograph and whether the cooling tower was the focal point of the image. We consider it unlikely that these platforms would ever

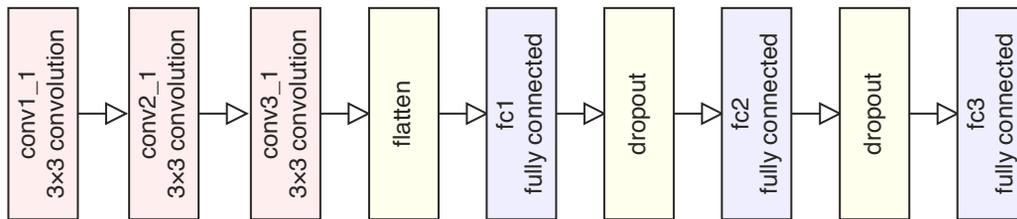


Figure 7. Proof-of-concept network with three convolutional (red) and three fully connected (blue) layers

have an incentive to provide the highly specialized labels needed for safeguards analysis tasks. Therefore, we decided to create our own CNN models that could specifically identify hyperboloid-shaped cooling towers and discern whether the towers had steam plumes.

Model Training

We used the Keras²¹ toolkit and its TensorFlow²² backend to train CNN image classification models using our nuclear cooling tower dataset. The first step was to train a model that could classify an image as “cooling tower” or “not cooling tower.” As a proof of concept, we experimented with a built-from-scratch network made up of three convolutional layers and three fully connected layers with dropout to avoid overfitting (Figure 7), using the hyperopt²³ library to conduct a hyperparameter search in which we trained more than 10,000 models.

A major concern during training is *overfitting*, a condition in which the model’s predictions correspond too closely to the training data, so that it fails to predict future observations reliably. Conceptually, an overfitted model has memorized the correct labels for its training inputs but cannot generalize that information to predict new labels for previously unseen inputs. To prevent overfitting, we used random sampling without replacement to partition our image data into separate “training,” “validation,” and “test” sets. Data from the “training” set was used to train the model, updating the network parameters through backpropagation. In addition, we augmented the training images by generating permutations using random rotation, mirroring, shearing, and scaling of each image. These permutations help avoid overfitting by exposing the network to a larger set of images with greater variance, forcing it to ignore irrelevant features. For example, observing the same cooling tower in many slightly different orientations

during training can help force a network to identify a more robust set of features that are indicative of cooling towers generally, while ignoring irrelevant features introduced by the viewpoint of the person taking a photograph.

At the end of each training epoch, we computed a loss value using the “validation” data and the same loss function used for training, and we used this information to determine when to stop training the model. Figure 8 illustrates this process: the green data series shows the loss function applied to the training data, which is driven steadily downward by the learning process as the model’s parameters are updated. The orange data series shows the loss function applied to the validation data. The point at which the validation loss begins to trend upwards (right of the dashed line) marks the point at which the information learned from the training data no longer generalizes to the validation data, and training should stop to avoid overfitting.

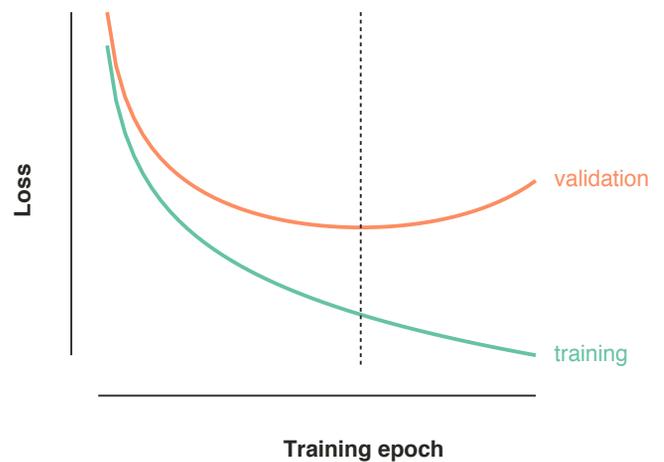


Figure 8. Training loss (green) versus validation loss (orange). Overfitting is likely when the validation loss rises (right of the dashed line).



Finally, we evaluated the model's performance using the "test" data, which ensured that we would have a conservative, realistic estimate of the model's performance on unseen future inputs. Because random chance in the partitioning process will affect the model that is produced, we repeated this entire process, training multiple models with different partitions, averaging together their performance scores to produce our final, cross-validated results.

Unfortunately, the accuracy of the proof-of-concept models ranged from 0.6 (only slightly better than random) to 0.8 (far from the state of the art) with mode = 0.74. (The standard to which we compare our results is the performance of the VGG-16 CNN on the ImageNet dataset with an accuracy, depending on evaluation method, of 90–92 percent.²⁴) Furthermore, no identifiable trends in the hyperparameter search that could explain this variance. This suggested that, despite the simplicity of the network, we did not have enough data to train it from scratch, so that random variations in the data partitions were driving the variance in accuracy. Stated differently, although our dataset containing 4,460 images might seem large in human terms, it did not provide enough variety for the learning algorithms to identify robust correlations between the input images and output labels.

To overcome this, we turned to transfer learning,^{25,26} where an existing model trained on an unrelated dataset is retrained (or "fine-tuned") to perform an alternate task. For this purpose, we compared several popular image classification networks, each of which was pretrained using the ImageNet dataset,¹⁴ a collection of 10 million images labeled using 1,000 widely varying classes, from animals and food to inanimate objects such as basketballs, canoes, and espresso makers. Notably, ImageNet does not include architectural photographs, but stone walls, totem poles, tile roofs, and other parts of buildings do appear.

We removed the final layers of each network that performed

the thousand-way ImageNet classification task and substituted our own fully connected layers to perform our cooling tower (or steam plume) classification. During training, we held most of the low-level, existing convolutional layers constant, while allowing some of the high-level convolutional layers and our classification layers to be trained. This allowed us to use the low-level features already known to the network while focusing the training of the remaining layers on our specific problem. For our hyperparameter search, we explored the following parameters:

- Choice of input image colorspace, including red, green, blue (RGB); hue, saturation, and value; luma and chrominance; and grayscale
- Image augmentation parameters (amount of random rotation, amount of random skew, etc.)
- The number of input images for each batch
- The choice of starting network, including VGG-16,²⁷ ResNet50,²⁸ and InceptionV3²⁹
- The number of existing layers to train in addition to our new classification layers
- The number of network layers used for classification and the number of activation units in each layer
- Regularization, including L1, L2, and dropout parameters (regularization imposes Occam's razor on a model; it helps prevent overfitting by producing simpler, sparser networks that generalize better to new data)
- Choice of optimization algorithm and learning rates

Once the search was complete, we selected the set of hyperparameters that produced the highest-performing models, which included using RGB images, modest random augmentation parameters, the VGG-16 network, a batch size of 32, two output activation layers to perform classification, and 50 percent dropout for regularization (see Figure 9).

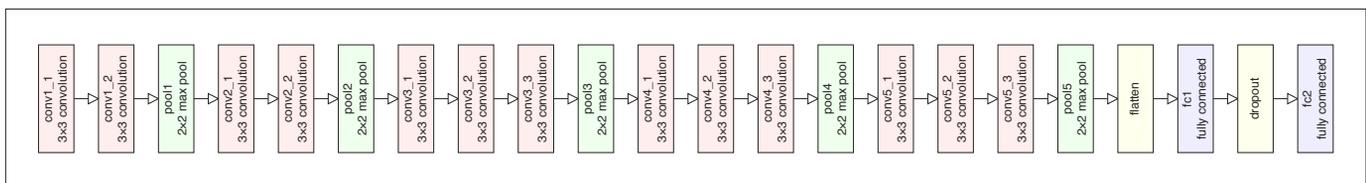


Figure 9. Final VGG-16-based architecture with 13 convolutional (red), five pooling (green), and two fully connected (blue) layers



Results

Using this architecture, we explored several approaches to implementing the *cooling tower + steam plume* signature, including (1) a *multilabel* approach in which we trained a single CNN model to produce two binary outputs — *cooling tower* and *steam plume* — then combined the two outputs with logical-AND to produce the signature, and (2) a *multiclass* approach in which we trained a single CNN model to produce four categorical outputs: *cooling tower*, *steam plume*, *cooling tower + steam plume*, and *null*, using the *cooling tower + steam plume* output directly as our signature.

First, we compared the predictions made by these models to the validated ground truth for our subset of 236 U.S. images, with the results as shown in Table 2.

Table 2. Validated model performance on U.S. nuclear power plants

	Multilabel model		Multiclass model	
	Precision	Recall	Precision	Recall
Plant operating	0.95	0.83	0.94	0.87
Plant not operating	0.66	0.87	0.71	0.86

Recall from Table 1 that the highest possible “Plant operating” precision and recall for the *cooling tower + steam plume* signature are 0.96 and 0.95. The results in Table 2 combine that error with errors in the predictions made by the model — for example, incorrectly labeling whether an image contains a cooling tower. Despite this, the “Plant operating” precision and recall for both models compares favorably with the best case. There is no clear winner between the two models; an analyst might prefer one or the other depending on his or her tolerance for false positives (lower precision) or false negatives (lower recall).

Next, we used both models to make predictions for the full 4460 image dataset, with the results shown in Table 3.

Table 3. Overall model accuracy on our full data set

	Multilabel model	Multiclass model
Cross validation accuracy	0.906	0.855

Because we had validation data only for U.S. nuclear facilities, the results in Table 3 are not validated based on operational ground truth only — that is, they reflect only whether a model correctly identified the presence or absence of cooling towers and steam plumes in an image, based on our manually labeled ground truth. Thus, we are able to report only their accuracy (percentage of correctly labeled images), for which the multilabel model has a clear advantage.

Conclusions and Future Work

From these results, we conclude that fine-tuned CNNs can be a useful technique when working with relatively small datasets containing the specialized subjects likely to be of interest to safeguards analysts. Models trained using these techniques could be used for search, filtering, and preliminary assessment of both internal IAEA safeguards data repositories and new, open data streams.

The Flickr API made it relatively easy to retrieve images, but that ease of use came at a cost: because Flickr appeals more to serious photographers than casual users, the majority of the images are artistic and have been digitally processed. Furthermore, few of the images were geotagged, which limited the number of images that we could use for validation; inferring the location of an image from its surrounding context (for example, the metadata tags, labels, or the text of a containing news report) is an obvious opportunity for future work. More broadly, this highlights the importance of understanding the user bias and community norms that inevitably shape the availability and credibility of information on any open platform the IAEA might use for collection.

The models we created for this work produced confidence values between zero and one for each output, which we converted into labels with a simple threshold. In the future, we would like to explore how adjusting that threshold could affect outcomes, particularly because different analysts might have different preferences for false positives or false negatives during image retrieval. Alternatively, analysts might prefer to accept high-confidence predictions with minimal review so that they can focus their attention on images with lower-confidence outputs.

A common complaint with artificial neural nets in particular and machine learning algorithms in general is the lack of rigorous uncertainty quantification: although our models produced output confidence values, no concrete interpretation can be assigned to those levels (that is, they are not probabilities). Some recent work has suggested methods for extracting meaningful uncertainty information from neural networks using dropout,³⁰ and we hope to explore these techniques in the near future.

Our *cooling tower + steam plume* signature, while simplistic, provides an example that the IAEA could build on to develop a wide variety of signatures inferring activity at nuclear facilities. Image signatures representing “patterns of life” could be developed for vehicular traffic, parking lot levels, personnel attire, the local environment, and more. These signatures could



become part of a broader, integrated approach to prioritizing and assessing the activities at a nuclear fuel cycle facility under IAEA safeguards.

Acknowledgments

This work was funded by the National Nuclear Security Administration's Office of Defense Nuclear Nonproliferation Research & Development (NA-22), in collaboration the Nevada Nuclear Security Site and the University of California, Berkeley.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2018-9087 J.

Keywords

Convolutional neural networks, machine learning, international safeguards, patterns of life

Author Biographies

Zoe N. Gastelum is a Principal Member of the Technical Staff at Sandia National Laboratories, in the International Safeguards and Engagements Department, where she researches the application of advanced analytical approaches and novel data sources for international nuclear safeguards and nonproliferation information analysis. Prior to Sandia, Zoe was a staff member at Pacific Northwest National Laboratory and the International Atomic Energy Agency.

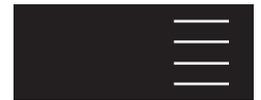
Timothy M. Shead is a Senior Member of the Technical Staff at Sandia National Laboratories. Tim's background is in computer graphics, scientific visualization, and – more recently – machine learning and data analysis.

End Notes

1. Barletta, M., Fowler, M., & Khaled, J. (2016). Integrating multimedia information into IAEA safeguards. *Proceedings of the American Nuclear Society Advances in Nuclear Nonproliferation Technology and Policy Conference: Bridging the Gaps in Nuclear Nonproliferation*. Santa Fe, NM.
2. International Atomic Energy Agency. (1995). Strengthening the effectiveness and improving the efficiency of the safeguards system: A report by the director general. GOV/2784.

para 14.(a)(v).

3. GOV/2784, para 56(i-iii).
4. Ferguson, M. & Norman, C. (2010). All-source information acquisition and analysis in the IAEA department of safeguards. *IAEA Safeguards Symposium 2010: Preparing for Future Verification Challenges*. Vienna, Austria.
5. Barletta, M., Fowler, M., & Khaled, J. (2016). Integrating multimedia information into IAEA safeguards. *Proceedings of the American Nuclear Society Advances in Nuclear Nonproliferation Technology and Policy Conference: Bridging the Gaps in Nuclear Nonproliferation*. Santa Fe, NM.
6. Barletta, M., Fowler, M., & Khaled, J. (2016). Integrating multimedia information into IAEA safeguards. *Proceedings of the American Nuclear Society Advances in Nuclear Nonproliferation Technology and Policy Conference: Bridging the Gaps in Nuclear Nonproliferation*. Santa Fe, NM.
7. International Atomic Energy Agency. (2018). Research and development plan: Enhancing capabilities for nuclear verification. STR-385. 12.
8. Internet Live Stats. Retrieved from www.internetlivestats.com/one-second
9. Flickr API. Retrieved from www.flickr.com/services/api
10. Google Earth. Retrieved from www.google.com/earth
11. U.S. Energy Information Administration. Status of U.S. nuclear outages. (2018). Retrieved from www.eia.gov/nuclear/outages
12. Dahl, G. E., Yu, D., Deng, L., & Acero, A. (2012). Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. *IEEE Transactions on Audio, 20*, 1, 30–42.
13. Lewis-Kraus, G. (2016). The great AI awakening: How Google used artificial intelligence to transform Google Translate, one of its more popular services—and how machine learning is poised to reinvent computing itself. *New York Times Magazine*.
14. Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative adversarial nets. *NIPS*.
15. Vinyals, O., Toshev, A., Bengio, S., & Erhan, D. (2015). Show and tell: A neural image caption generator. *IEEE Conference on Computer Vision and Pattern Recognition*. 3156–3164.
16. LeCun, Y., Bottou, L., & Bengio, Y. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*.
17. Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L.



- (2009). ImageNet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255.
18. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S. ... Fei-Fei, L. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3), 211–252.
 19. LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010). Convolutional networks and applications in vision. *ISCAS*, 253–256.
 20. Google. (2017). Google Vision API. Retrieved from <http://cloud.google.com/vision/docs>
 21. Chollet, F. (2015). keras. GitHub.
 22. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng, X. (2016). TensorFlow — large-scale machine learning on heterogeneous distributed systems. *Computing Research Repository*.
 23. Bergstra, J., Yamins, D., & Cox, D. D. (2013). Hyperopt: A Python library for optimizing the hyperparameters of machine learning algorithms. *Proceedings of the 12th Python in Science Conference*.
 24. Simonyan K. & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv, vol. cs.CV.
 25. Oquab, M., Bottou, L., Laptev, I. & and Sivic, J. (2014). Learning and transferring mid-level image representations using convolutional neural networks. *CVPR*, 1717–1724.
 26. Zheng, L., Zhao, Y., Wang, S., Wang, J., & Tian, Q. (2016). Good practice in CNN feature transfer. arXiv.org.
 27. Simonyan, K. & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv, vol. cs.CV.
 28. He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep residual learning for image recognition. arXiv.org.
 29. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2015). Rethinking the inception architecture for computer vision. arXiv.org.
 30. Gal, Y. (2016). *Uncertainty in deep learning*. Retrieved from [pdfs.semanticscholar.org](https://arxiv.org/pdf/1604.07982v1.pdf)



Site Monitoring with Sentinel-1 Dual Polarization SAR Imagery Using Google Earth Engine

Joshua Rutkowski

Jülich Research Center, Jülich, Germany

Morton J. Canty

Jülich Research Center, Jülich, Germany (retired)

Allan A. Nielsen

Technical University of Denmark, Lyngby, Denmark

Abstract

At no cost to the user, the Copernicus mission frequently releases Synthetic Aperture Radar (SAR) datasets collected by the Sentinel-1 sensors. These datasets are regularly pre-processed by the Google Earth Engine and made available to the scientific community for further processing and analysis. This paper describes the application of a recently developed sequential change detection algorithm for Sentinel-1 datasets based on an omnibus likelihood ratio test statistic within the Google Earth Engine platform. Change detection methods, such as the one described here, offer the nuclear non-proliferation community a new way to use remote sensing datasets for monitoring nuclear facilities worldwide.

Introduction

Synthetic Aperture Radar (SAR) is being widely used for nuclear nonproliferation activities and offers an imaging capability independent of cloud-cover and solar illumination conditions over areas of interest throughout the world. The Sentinel-1 SAR space mission, operated by the European Space Agency (ESA), offers free access to the sensors' data, unlike other SAR missions that charge fees for each data set. A convenient and near-real-time source of time-sequential SAR data is provided by the Google Earth Engine (GEE) platform,¹ which archives Sentinel-1 dual polarization SAR images as soon as they are made available by the ESA. The GEE also offers a platform to design and execute server-side algorithms on server-side hosted remote sensing data, thereby reducing local storage and computation requirements. Change detection analysis for nuclear nonproliferation and activities supports international verification regimes by offering a remote mechanism for treaty verification.²

This paper describes the application of a recently developed

sequential change detection algorithm for Sentinel-1 data sets. The algorithm is based on an omnibus likelihood ratio test statistic for the equality of several variance–covariance matrices following the complex Wishart distribution, with an associated p-value and a factorization of the test statistic.^{3,4} It enables geospatial and temporal isolation of changes in a time series of multilook SAR data in a simplified polarimetric matrix representation, with single polarization intensity data included as a special case. The ability of the sequential change detection method to detect and isolate regions of intense activity, together with easy access to Sentinel-1 SAR imagery on the GEE, suggests specific applications in the area of remote monitoring for nuclear nonproliferation treaty verification.

The change detection method is described in detail. Furthermore, an open source change detection software for interaction with the GEE has been developed and is explained in this paper. The client-side programs run in a local Docker container serving a simple Flask web application. Apart from the Docker engine and a browser, no software installation is required of the user for the change detection analysis. The capabilities of the method are demonstrated via an analysis of a long time series of Sentinel-1 dual polarization data acquired over uranium mines and a military site.

Theory

The term “multilook” in SAR imagery refers to the averaging of single observations to reduce the effect of speckle, a noise-like consequence of the coherent nature of the radar signal emitted from the sensor. In the case of polarimetric SAR, the observed signals are complex Gaussian-distributed signals,³ and their variance–covariance representations, when multiplied by the equiva-



lent number of looks (ENL) — an empirical estimate of the number of independent observations — are correspondingly complex Wishart-distributed signals. This distribution is the multivariate complex analogue of the well-known chi-square distribution for the variance of Gaussian-distributed scalar observations.

The complex Wishart distribution is completely determined by the parameter (the covariance matrix) and the ENL. The latter is the number of looks corrected for the lack of independence of the multilook observations. Given two measurements of polarized backscatter from the same region on the ground surface at different times from the same orbital path, one can set up a so-called hypothesis test to decide whether a change has occurred. The null hypothesis (H_0) is that — that is, the two observations were sampled from the same distribution, and no change has occurred — and the alternative hypothesis (H_1) is ; in other words, there was a change (the ENL is assumed to be the same in both images). Because the distributions are known, a likelihood ratio test can be formulated that allows one to decide to a desired degree of significance whether to reject the null hypothesis. Acceptance or rejection is based on the p-value — that is, the probability of getting a test statistic that is at least as large as the one observed, given the null hypothesis. The p-value may be derived from the (approximately known) distribution of the likelihood ratio test statistic. In the case of k observations, this procedure can be generalized to test a null hypothesis that all of the k pixels are characterized by the same Σ against the alternative that at least one of the, is different — that is, at least one change has taken place. Furthermore, this so-called omnibus test procedure can be factored into a sequence of tests involving hypotheses of the form:

$$\Sigma_1 = \Sigma_2 \text{ against } \Sigma_1 \neq \Sigma_2$$

$$\Sigma_1 = \Sigma_2 = \Sigma_3 \text{ against } \Sigma_1 = \Sigma_2 \neq \Sigma_3$$

and so forth. The tests are statistically independent under the null hypothesis. In the event that the null hypothesis is rejected at some point in the test sequence, the procedure restarts from that point, so that multiple changes within the time series can be identified.

Dual Polarization, Diagonal Only

To be more specific, we describe below the hypothesis tests for dual polarization (pol), diagonal-only matrices. The per-pixel observations are expressed as multilook polarimetric matrices which, in the case of vertical emission and vertical and horizontal reception, have the form

$$\bar{C} = \begin{pmatrix} \langle |S_{vv}|^2 \rangle & \langle S_{vv} S_{vh}^* \rangle \\ \langle S_{vh} S_{vv}^* \rangle & \langle |S_{vh}|^2 \rangle \end{pmatrix},$$

where S_{vv} is the complex scattering amplitude for vertically polarized emission and detection and S_{vh} is the vertical emission horizontal reception. The angular brackets denote multilooking. The Earth Engine database archives only the diagonal matrix elements, not the (complex) off-diagonal elements. The block diagonal elements are assumed to be independent.⁵ So, in fact, one works with observations of the form

$$\bar{C} = \begin{pmatrix} \langle |S_{vv}|^2 \rangle & 0 \\ 0 & \langle |S_{vh}|^2 \rangle \end{pmatrix}.$$

If we write

$$x = m\bar{C}$$

where m is the ENL, then the omnibus test for

$$H_0: \Sigma_1 = \Sigma_2 = \dots = \Sigma_k = \Sigma,$$

against

$$H_1: \Sigma_i \neq \Sigma_j \text{ for at least one pair } i, j$$

has the critical region

$$Q = k^{2km} \frac{|x_1|^m |x_2|^m \dots |x_k|^m}{|x_1 + x_2 + \dots + x_k|^{km}} \leq t, \text{ where } |\cdot| \text{ denotes the determinant}$$

The left side of the inequality is the ratio of the maximum likelihood of observing the data under the null hypothesis to that of observing the data irrespective of change. If the ratio is less than some threshold t , the null hypothesis is rejected. Taking logarithms:

$$\ln Q = m(2k \ln k + \sum_{i=1}^k \ln |x_i| - k \ln |\sum_{i=1}^k x_i|) \leq t'$$

We have the following approximation for the statistical distribution of the test statistic Q :

$$\text{prob}(-2 \ln Q \leq z) \approx P_{\chi^2, 2(k-1)}(z)$$

where $P_{\chi^2, f}(z)$ is the chi-square distribution with degrees of freedom. In practice, we choose a significance level α — for



example, $\alpha = 0.01$ — and reject the null or no-change hypothesis when the p value satisfies

$$p = 1 - P_{x^2;2}(-2 \ln R_j^1) < \alpha.$$

There is some empirical justification that the assumption of the independence of the diagonal elements does not seriously affect the validity of the hypothesis tests with GEE Sentinel-1 data. Previous research investigates histograms of the values,⁶ where is the likelihood ratio test statistic, in no-change regions and finds good agreement with the expected chi-square distributions. See previous work for a comparison of the distribution test statistic for no-change observations with the theoretical (chi-square) values for the quad pol case and for the dual pol diagonal case using Sentinel-1 data available in GEE.⁴

Sequential Omnibus Test

As mentioned earlier, the omnibus test can be factored into a sequence of tests involving hypothesis of the form $\Sigma_1 = \Sigma_2$ against $\Sigma_1 \neq \Sigma_2$, $\Sigma_1 = \Sigma_2 = \Sigma_3$, against $\Sigma_1 = \Sigma_2 \neq \Sigma_3$, and so forth. For example, to test

$$H_{0j}: \Sigma_1 = \Sigma_2 = \dots = \Sigma_{j-1} = \Sigma_j$$

against

$$H_{1j}: \Sigma_1 = \Sigma_2 = \dots = \Sigma_{j-1} \neq \Sigma_j$$

the likelihood ratio test statistic is given by

$$\ln R_j^2 = m[2(j \ln j - (j-1) \ln(j-1)) + (j-1) \ln \left| \sum_{i=2}^{j-1} x_i \right| + \ln |x_j| - j \ln \left| \sum_{i=2}^j x_i \right|]$$

for $j = 2 \dots k$, and

$$\text{prob}(-2 \ln R_j^1 \leq z) \approx P_{x^2;2}(z).$$

The null hypothesis is correspondingly rejected when

$$p = 1 - P_{x^2;2}(-2 \ln R_j^1) < \alpha.$$

Moreover, the constitute a factorization of such that, or

$$\ln Q = \sum_{j=2}^k \ln R_j^1$$

and the tests are statistically independent under the null hypothesis.

Now suppose that we conclude. We can then continue to look for additional changes by restarting the tests at,

$$\ln R_j^2 = m[2(j \ln j - (j-1) \ln(j-1)) + (j-1) \ln \left| \sum_{i=2}^{j-1} x_i \right| + \ln |x_j| - j \ln \left| \sum_{i=2}^j x_i \right|]$$

for $j = 3 \dots k$. This procedure is continued until the end of the image sequence.

Computation Strategy and Change Detection Products

Denoting the test statistics R_j^ℓ , $\ell = 1 \dots k-1$, $j = \ell+1 \dots k$, for a series of, say $k = 5$ images, we have the following tests to consider:

ℓ/j	2	3	4	5
1	R_2^1	R_3^1	R_4^1	R_5^1
2		R_3^2	R_4^2	R_5^2
3			R_4^3	R_5^3
4				R_5^4

For example, if the test R_2^1 does not reject $H_{0,}$, but test R_3^1 does, then the subsequent test R_4^1 and R_5^1 in the first row are not valid. We restart the sequence in row 3 at and continue to check for additional changes. From a computational standpoint, because the test sequences must be applied to every pixel, it is advantageous to precalculate all of the p-values (in this example, there are 10 of them), irrespective of their relevance. Then, in a second pass over the p-value arrays, the valid changes are identified. This is the strategy adopted in the implementation of the sequential omnibus algorithm on the GEE Python API: The p-value arrays are precomputed and then scanned at the desired significance level. Four 1-byte arrays having the spatial dimensions of the images are updated recursively to generate georeferenced change maps. These are:

- c_map: the interval in which the most recent significant change occurred (single-band)
- s_map: the interval in which the first significant change occurred (single-band)
- f_map: the frequency of significant changes (single-band)
- b_map: the interval in which each significant change occurred ((k-1)-band).

The f_map is particularly advantageous for identifying small regions of concentrated activity over the time sequence, whereas the b_map can be used to determine change profiles of specific regions of interest and to generate change animations.



Figure 1. The NATO Airbase near Geilenkirchen, Germany. (Copyright 2018, Google)

Case Studies

For site monitoring applications using the GEE Sentinel-1 data archive, one has to take both the spatial as well as the temporal resolution into account. The Sentinel-1 Ground Range Detected (GRD) scenes that are processed by the sequential change detection algorithm are acquired in Interferometric Wide (IW) swath mode. The multilook averaging of five azimuth looks and one range look results in a spatial resolution of about 20 meters, although the pixel dimensions are 10 x 10 meters. Thus, apart from large-scale construction/excavation activity, the most detailed changes which we might expect to detect would be associated with movements of large vehicles, aircraft, or ships. Images over any location are acquired by the two spacecraft, Sentinel-1a and

Sentinel-1b, operating together at revisit times of about six days. Keeping this in mind, we consider the following examples of monitoring a military airbase and two uranium mines.

Geilenkirchen

The North Atlantic Treaty Organization's (NATO) airbase near Geilenkirchen, Germany (Figure 1), is the Main Operating Base of the NATO Boeing E-3 Sentry Component, one of two operational elements responsible for managing the NATO Airborne Early Warning and Control System (AWACS) fleet of aircraft.⁷

The E-3 component operates 16 AWACS aircraft based on the Boeing 707 airframe and manned by about 30 multinational crews from various NATO countries. In addition to regular mission

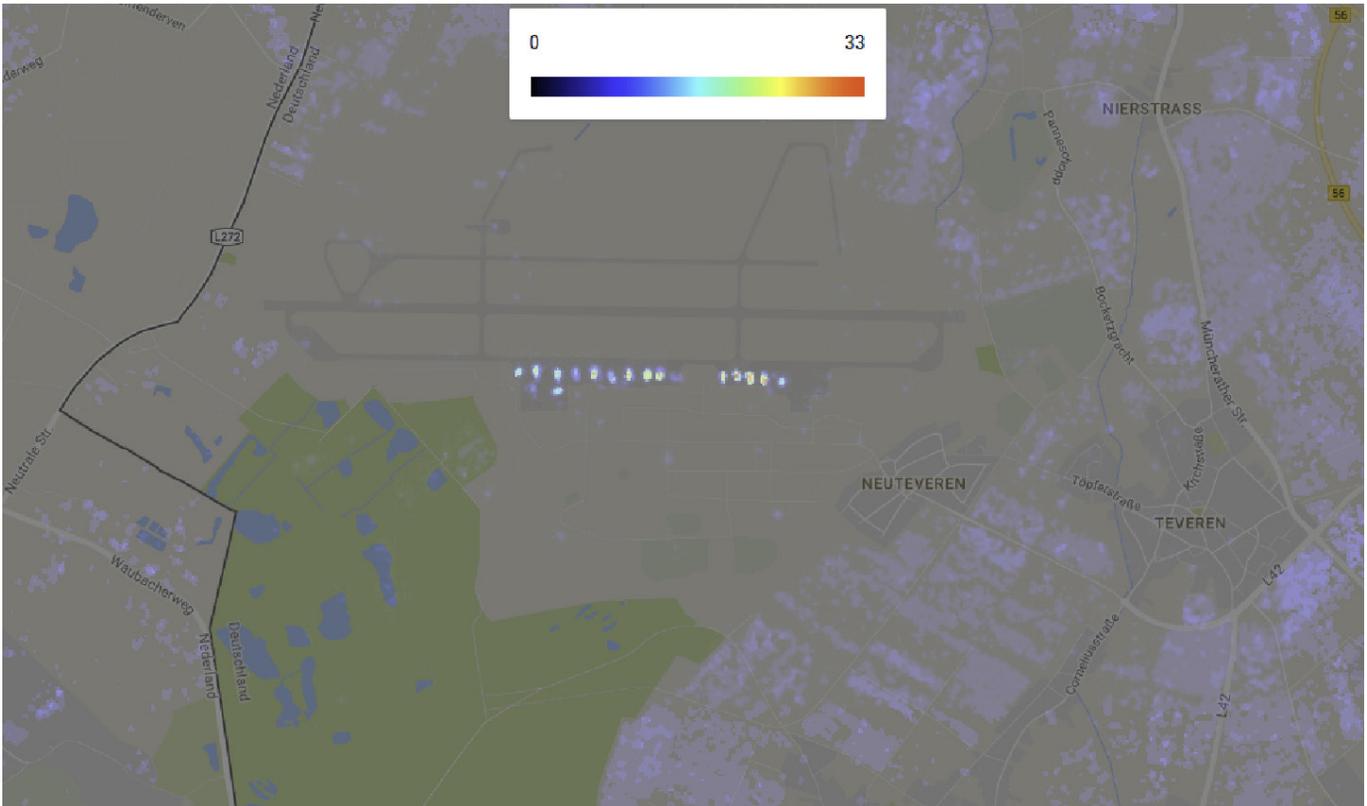


Figure 2. Changes observed per observation period for the NATO Airbase, 34 Sentinel-1 scenes from April to October 2017. Screenshot from Google Earth Engine Code Editor. (Copyright Google.)

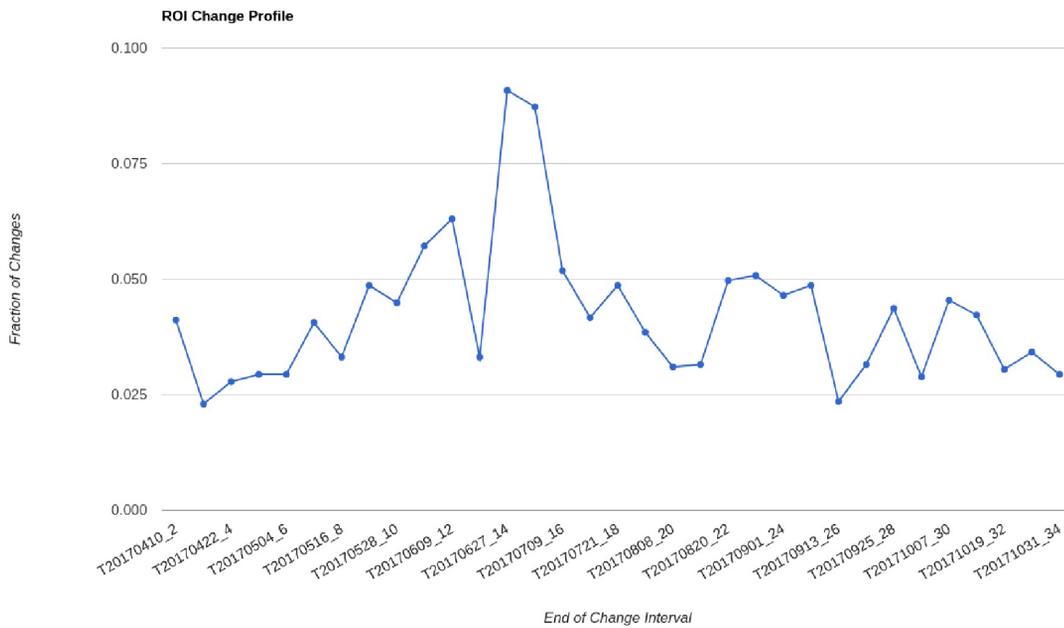


Figure 3. Activity profile for the aircraft parking positions. Note: the x-axis labels only show every second date. All available Sentinel 1a and 1b of the same relative orbit in the GEE archive are included.



Figure 4. The McArthur River uranium mine in northern Saskatchewan. (Copyright 2018, Google.)

starts and landings, a significant number of training flights are conducted from the airbase. This activity is accompanied by considerable noise pollution affecting the densely populated areas nearby in the Netherlands and Germany and has led to public protests. Future disarmament treaties would likely cover weapon delivery systems, including strategic bombers. The ability to regularly track airplane movement at airports would assist a disarmament treaty inspectorate through verification without site visitation. A change frequency map obtained with a series of 34 Sentinel-1 acquisitions between April and October 2017 is shown in Figure 2.

The parking positions of the aircraft are immediately apparent since the changes are strongly localized. By defining a polygon region of interest around the parking positions (not shown in any figures) to exclude changes from the surrounding agricultural regions, we can track the fraction of changes (number of changed pixels in the window divided by the total number of pixels) over time (see Figure 3). A peak in activity in the month of July 2017 is clearly evident.

McArthur River and Cluff Lake Uranium Mines

Satellite sensors can play an important role in the monitoring of remote uranium mining and milling sites, especially to verify un-

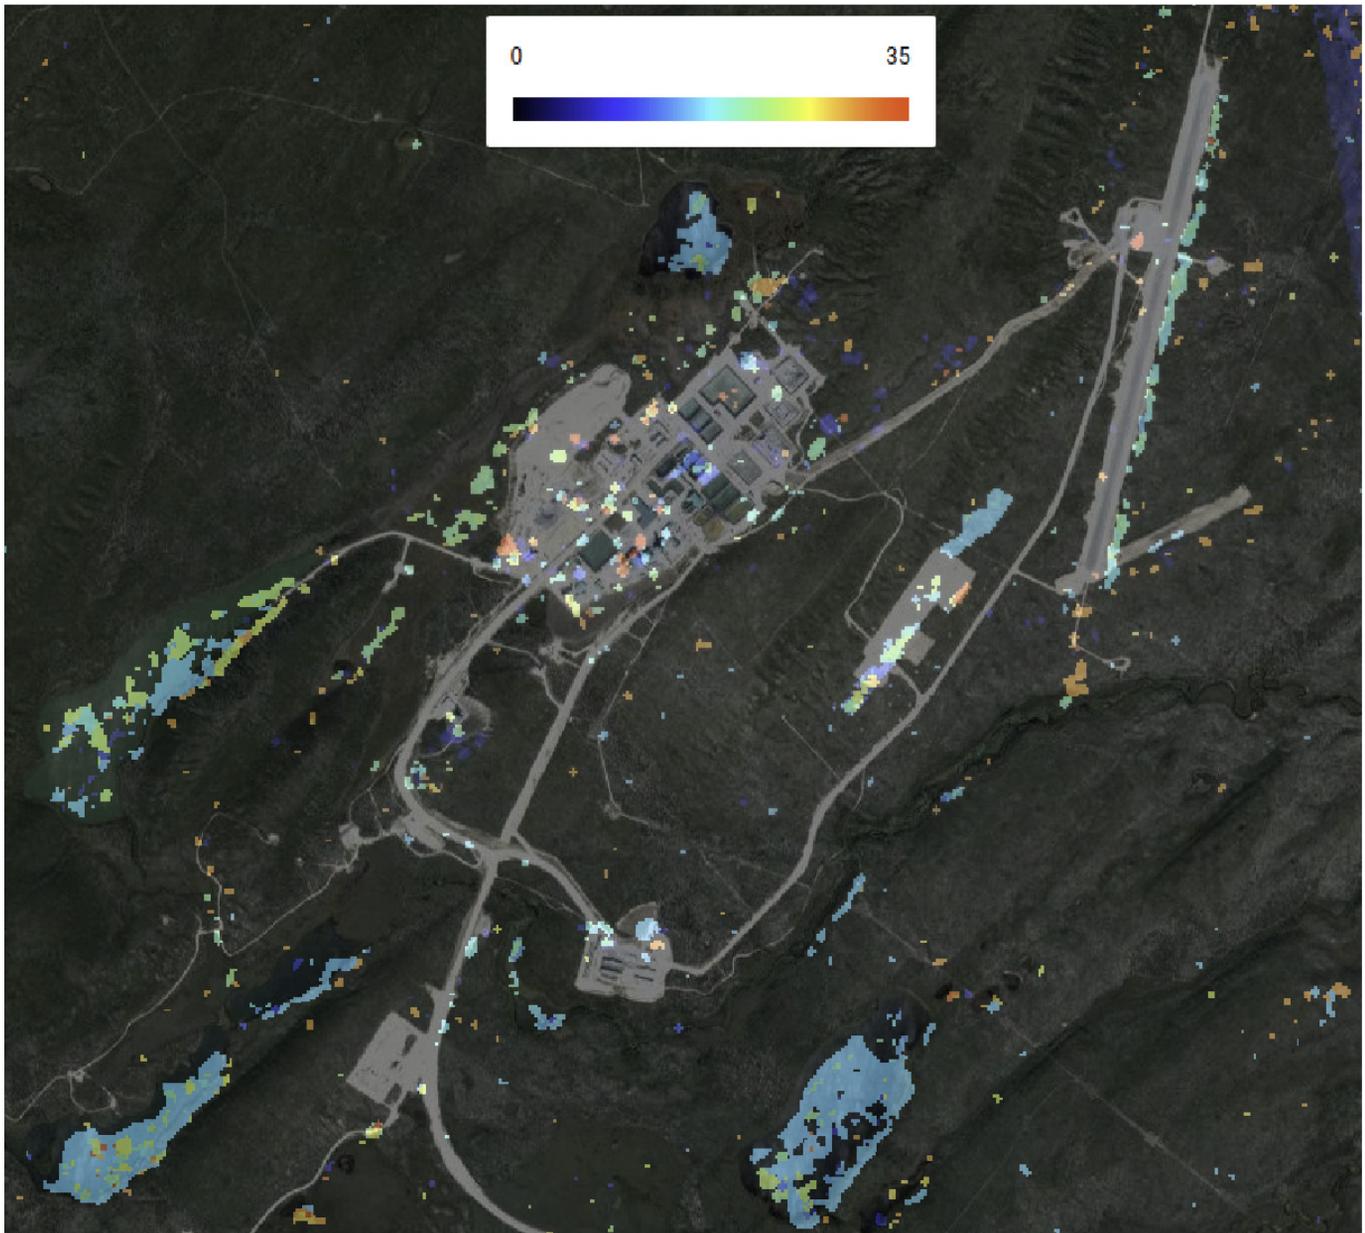


Figure 5. Most recent change map (dark blue early, yellow-red late) for the McArthur River mine using a series of 36 Sentinel-1 scenes from August 2016 to October 2017. (Screenshot from GEE Code Editor; copyright Google.)

declared activity where the amount of ore mined is in excess of that declared by the site operator. We first consider an example of a declared, operational uranium site, the McArthur River Uranium Mine in Canada (see Figure 4). The McArthur River deposit was discovered in 1988 and is located in Saskatchewan, Canada, approximately 70 kilometers northeast of the Key Lake mill.⁸ The 2016 Uranium Report (commonly referred to as the “Red Book”) published by the Nuclear Energy Agency and the International

Atomic Energy Agency on uranium resources, production and demand notes that “production at the McArthur River mine, the world’s largest high-grade uranium mine, was 7,684 tU and 7,312 tU in 2013 and 2014, respectively.”⁹

Figure 5 shows a change map color-coding the time of the most recent change detected in a time series of 36 Sentinel-1 images between August 2016 and October 2017. It is evident that changes are occurring throughout the time interval. The change

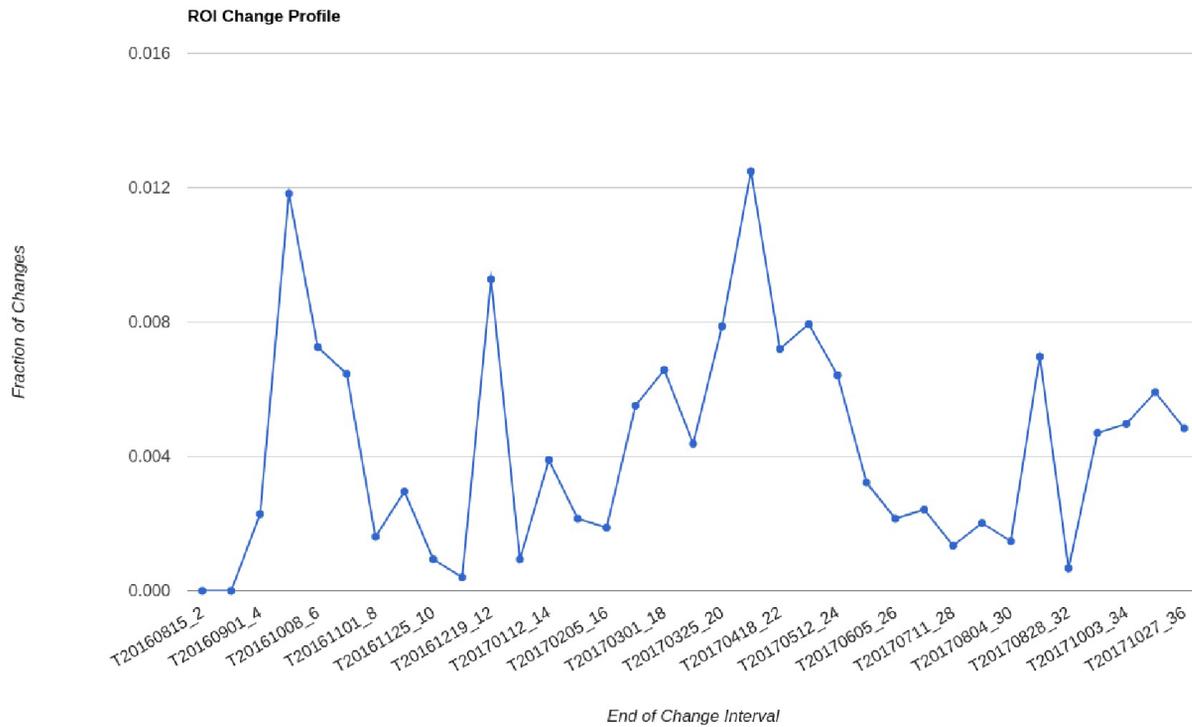


Figure 6. Change profile for a series of 36 Sentinel-1 scenes, from August 2016 through October 2017 over the McArthur River uranium mine.

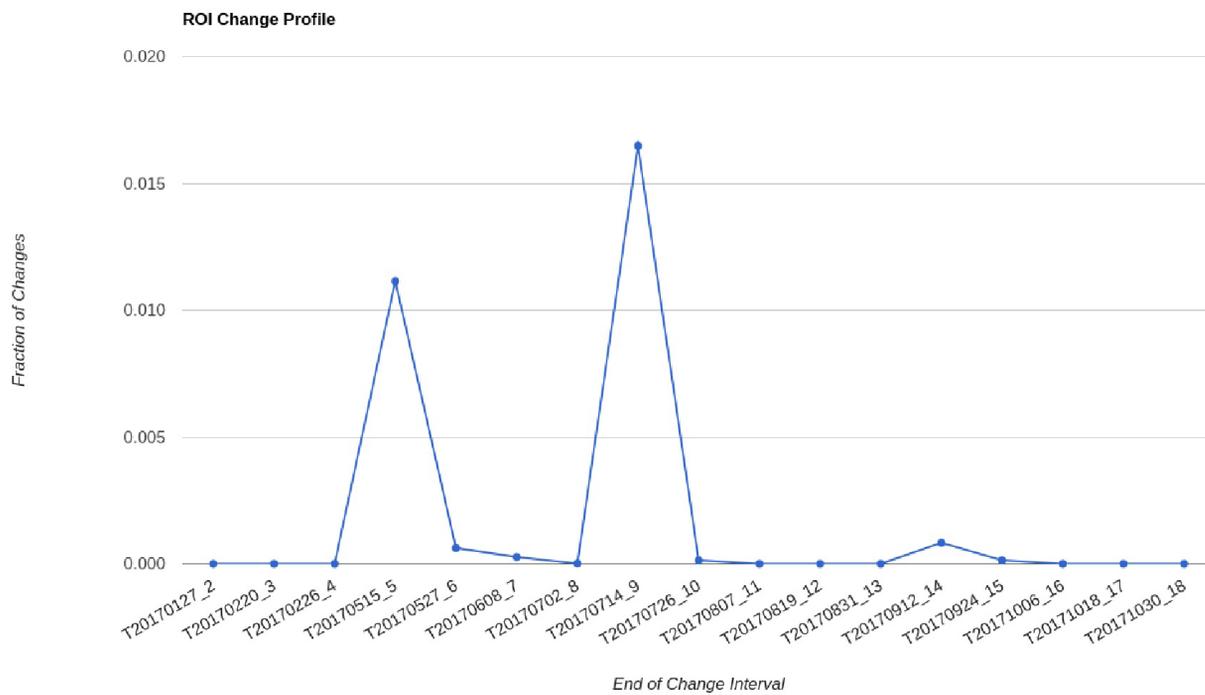


Figure 7. Change profile for a series of 18 Sentinel-1 scenes from January through October 2017 over the Cluff Lake uranium mine.



Figure 8. The decommissioned Cluff Lake uranium mine. (Copyright Google.)

profile for the main site can be seen in Figure 6. This is in contrast to a similar study of the decommissioned Cluff Lake mining site, also located in northern Saskatchewan (see Figure 7 and Figure 8). This mine ceased uranium production at the end of 2002 when the ore reserves were depleted. Apart from the two spikes in May and July 2017, no activity is evident in the change profile of Figure 8. The spikes can be attributed to environmental changes, most probably ground moisture following rainfall. Figure 9 shows the changes that occurred in May (orange/pink) and July (green), superimposed onto the Cluff Lake area. Whether or not these changes are related to undeclared activities would require further analysis that ought to include state-declared information.¹⁰ The change detection demonstrated here illustrates the utility of

SAR data to detect change but cannot uniquely determine the reason the change occurred, only that a change was observed.

Algorithm Distribution

The change detection algorithm presented in this paper offers analysts a quick and effective way to exploit Sentinel-1 data sets. To this end, the algorithm has been made available through two different platforms, a Docker container with a web browser or the Google Earth Engine (GEE) Code Editor. Each of these platforms illustrates emerging mechanisms for analysts to conduct remote sensing image analysis tasks with little local computation power and without relying on local remote sensing data stores. With the recent deluge of remote sensing data available to analysts from

many new and existing sensors, visual interpretation and single-scene analysis routines can no longer keep pace. The software used to execute the omnibus change detection algorithm is free, open source, and available on GitHub at <https://github.com/mortcanty/earthengine>.

Docker Container with Web Browser

Docker is a software containerization platform intended for large-scale web application deployment which may also be used to deploy an isolated operating system-level container within an existing operating environment so that specific software applications may run reliably on any user platform without the need for a full virtual machine to be installed on the users operating system. Using Docker, an analyst can, for example, run one or more Linux based applications on a Windows 10 operating system. The change detection algorithm described in this paper is available in such a Docker container. This implementation uses the simple Flask web application framework for performing change detec-

tion tasks with GEE imagery on the associated Google servers. This use of an ordinary web browser as a user interface demonstrates an option in which analysts are offered remote sensing analysis tools without using traditional remote sensing software packages. A website, whether hosted behind an enterprise firewall or open to the worldwide web, offers analysts a set of tools in a controlled setting. Such an interface can be programmed to limit the tool's input parameters, thus offering inexperienced analysts the option to compute change detection without user induced errors. For example, the change detection algorithm should only consider one orbital path for each change detection routine since the look geometry needs to be the same. Novice users should be prevented from simple mistakes like this but still offered the underlying power of the change detection algorithm. In addition, website designers can offer a more structured approach to integrating a full suite of remote sensing tools for analysts to use. Tool integration into a broader analyst working environment is a topic mentioned in another article in this issue.^{11,12} An example of

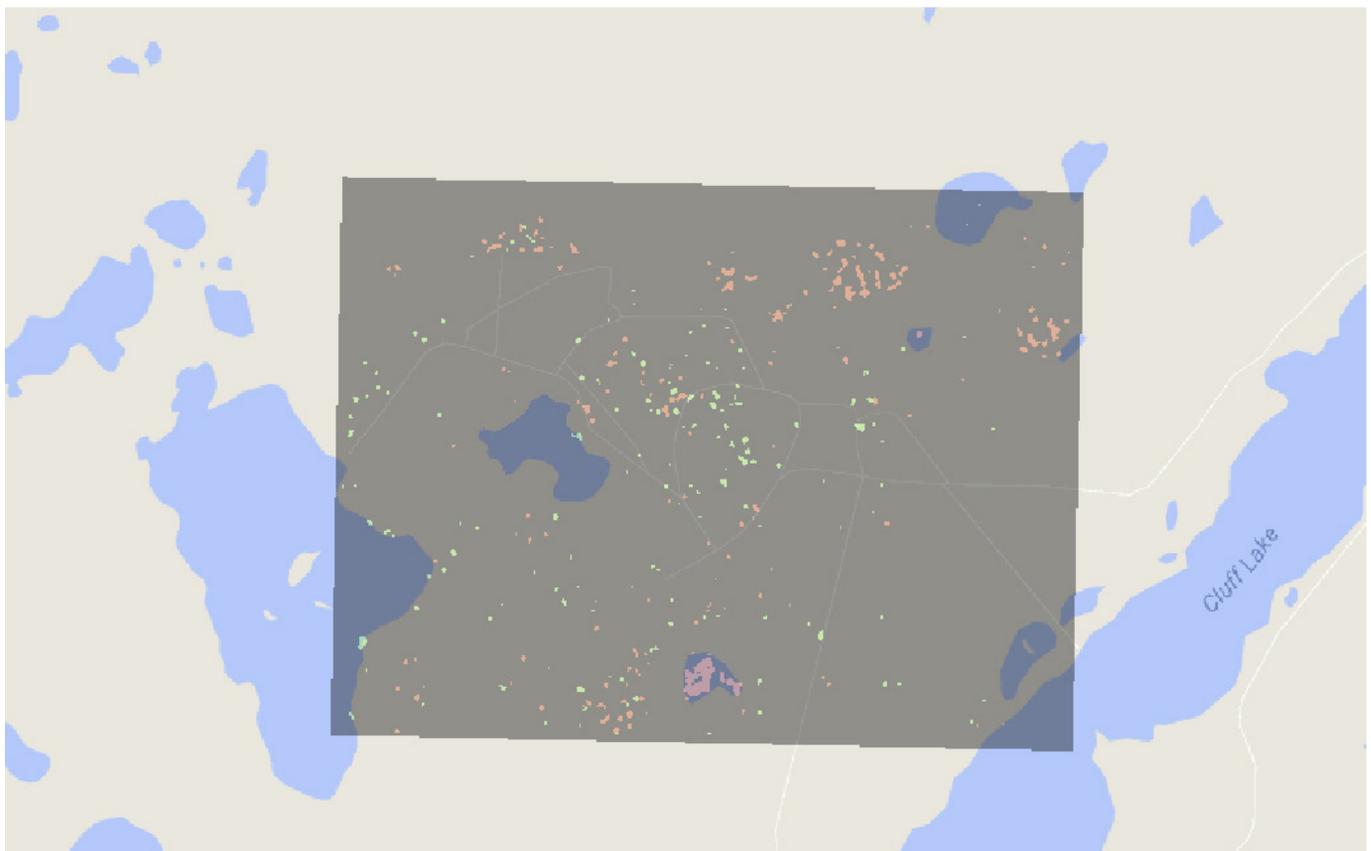


Figure 9. The region of the Cluff Lake uranium mining site, with significant changes detected in May and July 2017 shown in orange/pink and green, respectively. (Screenshot from GEE Code Editor; copyright Google.)



this style of deployment can be seen at <http://ms-image-analysis.appspot.com/static/index.htm>.

Google Earth Engine Code Editor

Analysts can also use a JavaScript implementation of the change detection algorithm within the GEE Code Editor. This offers more flexibility than using the python implementation for the analyst to modify or extend the change detection scripts and to control processing through the interactive JavaScript code editor. Google describes this platform as “a web-based IDE [Integrated Development Environment] for the Earth Engine JavaScript API. Code Editor features are designed to make developing complex geospatial workflows fast and easy” and has the following elements: “JavaScript code editor, Map display for visualizing geospatial data sets, API reference, documentation, Git-based script manager, Console output, Task manager to handle long-running queries, Interactive map query, Search of the data archive or saved scripts, Geometry drawing tools.”¹³

Conclusions

The Sentinel-1 data repository is freely accessible to remote sensing analysts. Satellite-based SAR sensors offer a remote sensing data source free from cloud-cover and sun illumination concerns and with a regular revisit time. Such data sources have proven valuable for change detection algorithms given the data sources’ reliable geographic coverage and accuracy. The omnibus algorithm running on GEE servers provides analysts with the opportunity to carry out fast spatial/temporal change detection on an easy-to-use platform. The monitoring of activities related to the nuclear fuel cycle and nuclear arms were used to illustrate some of the advantages of this approach and how remote sensing data can be robustly exploited. The Cluff Lake uranium mine example demonstrates how initial change detection often requires follow-up investigations to determine the cause of the change and any implication on treaty verification. Access to the Sentinel-1 data sets in the Google Earth Platform relieves analysts from data storage burdens, while executing the computations in the cloud reduces local hardware demands. These conditions enhance the capability of a single analyst to conduct change detection over many sites worldwide with a minimal amount of local resources — an option that has been uncommon in the remote sensing field in the past.

Keywords

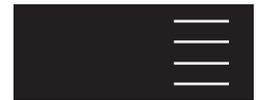
Remote sensing, synthetic aperture radar, statistical testing,

change detection, cloud computing

Author Biographies

Joshua Rutkowski is a research staff member in the Safeguards and Security Department at Jülich Research Center in Germany. He received his M.A. in geography from West Virginia University, USA, in 2005. His research focus is on using geographic information for nuclear nonproliferation and disarmament. Previously, Joshua was a cartographer at the U.S. Department of State’s Intelligence and Research Bureau. Between 2008 and 2015, he worked as a Safeguards information analyst in the International Atomic Energy Agency’s Department of Safeguards, with a focus on processing satellite imagery and developing geospatial information systems for use within the inspectorate. Joshua sits on the editorial committee for the European Safeguards Research and Development Association (ESARDA).

*Morton John Canty was a senior research scientist in the Institute for Bio- and Geosciences at the Jülich Research Center in Germany (now retired). He received his PhD in nuclear physics from the University of Manitoba, Canada, in 1969. In 1979, he began working at Jülich, where his principal interests have been the development of statistical and game-theoretical models for the verification of international treaties and the use of remote sensing data for monitoring global treaty compliance. He has served on numerous advisory bodies to the German federal government and to the International Atomic Energy Agency in Vienna and was a coordinator within the European Network of Excellence on Global Monitoring for Security and Stability, funded by the European Commission. Morton is the author of three monographs in the German language, in nonlinear dynamics (*Chaos und Systeme*, Vieweg, 1995), neural networks for classification of remote sensing data (*Fernerkundung mit neuronalen Netzen*, Expert, 1999) and algorithmic game theory (*Konfliktlösungen mit Mathematica*, Springer, 2000). The 2000 text was revised and published in English as *Resolving Conflicts with Mathematica* (Academic Press, 2003). Morton coauthored a monograph on mathematical methods for treaty verification (*Compliance Quantified*, Cambridge University Press, 1996) and authored the textbook *Image Analysis, Classification and Change Detection in Remote Sensing* (Taylor and Francis, 2006). He has published many papers on the subjects of nuclear physics, nuclear safeguards, applied game theory, and remote sensing. He has lectured on nonlinear dynamic growth models and remote sensing digital image analysis to graduate and undergraduate students at universities in Bonn, Berlin, Muenster, Freiberg/Saxony and Rome.*



Allan Aasbjerg Nielsen is an associate professor in the Department of Applied Mathematics and Computer Science at the Technical University of Denmark (DTU). He received an MSc from the Department of Electrophysics at the DTU in 1978 and a PhD from IMM: Informatics and Mathematical Modelling at DTU. From 1977 to 1978, he was with the Danish Defense Research Establishment. From 1978 to 1985, he worked on energy conservation in housing with the Thermal Insulation Laboratory at DTU. He was with the section for image analysis from 1985 to 2001 and with the section for geoinformatics from 2001 to 2006, both at IMM. From 2007 to 2013, he was with the Danish National Space Center's section for geodesy. Since 1985, he has worked on many national and international projects on the development, implementation, and application of statistical methods and machine learning, and on remote sensing in mineral exploration, mapping, geology, agriculture, environmental monitoring, oceanography, geodesy, and security funded by industry, the European Union, the Danish International Development Agency, and the Danish National Research Councils.

References

1. Gorelick, N., et al. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202, 18–27.
2. Lafitte, M. & Robin, J. (2015). Monitoring nuclear facilities using satellite imagery and associated remote sensing techniques. *ESARDA Bulletin*, 52, 53–59.
3. Conradsen, K., Nielsen, A.A., & Skriver, H. (2016). Determining the points of change in time series of polarimetric SAR data. *IEEE Transactions on Geoscience and Remote Sensing*, 54(5), 3007–3024.
4. Nielsen, A.A., Conradsen, K., Skriver, H., & Canty, M.J. (2017). Visualization of and software for omnibus test-based change detected in a time series of polarimetric SAR data. *Canadian Journal of Remote Sensing*, 43(6), 1–11.
5. Conradsen, K., et al. (2003). A test statistic in the complex Wishart distribution and its application to change detection in polarimetric SAR data. *IEEE Transactions on Geoscience and Remote Sensing*, 41(1), 4–19.
6. Nielsen, A.A., et al. (2017). Change detection in multi-temporal dual polarization Sentinel-1 data. *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)* (pp. 3901–3904). Fort Worth, TX: IEEE.
7. NATO. (2018). HQ airborne early warning & control force: Geilenkirchen. Retrieved from <http://awacs.nato.int/default.aspx>
8. Marlatt, J., et al. (1992). The discovery of the McArthur River uranium deposit, Saskatchewan, Canada. *Technical committee meeting on new developments in uranium exploration, resources, production and demand*. Vienna: IAEA.
9. Organisation for Economic Co-operation and Development. (2016). Nuclear Energy Agency and the International Atomic Energy Agency, Uranium 2016: Resources, production and demand.
10. Under a safeguards agreement with the IAEA, the state declares information about its nuclear material, facilities, and activities.
11. Keskinen et al. Enhancing the geospatial exploitation system within the IAEA Department of Safeguards.
12. Rutkowski, J., Warner, T.A., & Niemeyer, I. (2017). A semi-automated change detection model for nuclear fuel cycle related activities using large remote sensing time series datasets and geographic information systems. *Institute of Nuclear Materials Management 58th Annual Meeting*. Indian Wells, CA: INMM.
13. Google. (2018). Earth engine code editor. Retrieved from <http://developers.google.com/earth-engine/playground>



Enhancing the Geospatial Exploitation System within the IAEA Department of Safeguards

*Antero Keskinen, Jacques Baute, Mark Carey, Jonetta Ng, Fatjon Ujkani
International Atomic Energy Agency, Vienna, Austria*

Abstract

The Geospatial Exploitation System (GES) is an enterprise-wide, collaborative platform that uses Geographical Information System (GIS) technology to exploit commercial satellite imagery (CSI) and geospatial data within the Department of Safeguards of the International Atomic Energy Agency. In 2011, the GES was first deployed into the Department's secure Integrated Safeguards Environment. It provided users across the Department with up-to-date access to CSI, geospatial data and analytical products. The application demonstrated the successful adoption of commercial off-the-shelf software integration with customized tools to ensure secure and efficient storage, management, analysis, and dissemination of safeguards-relevant data to authorized users.

In 2017, the Department deployed a new version, GES Web 2.0, in full coordination with the Modernization of Safeguards Information Technology (MOSAIC) project and integration with its products, including the recently introduced Authorization Management system. The significant upgrade incorporated new user requirements, resolved existing technical constraints by redesigning the back end of the application, and introduced further enhancements to improve interoperability with other MOSAIC applications, including the new Geo-based Data Integration platform. GES Web 2.0 has streamlined the security framework, significantly improved performance, enhanced usability, and introduced advanced analytical capabilities, primarily for imagery and geospatial analysts. The improved web-based interface can exploit new CSI sensor types and associated metadata and interact with expert analytical software.

In addition, the GES streamlines analysts' workflows from imagery and geospatial data acquisition and creation through processing, editing, and analysis to dissemination. Standard operating procedures and quality control mechanisms are applied to workflows to ensure the integrity, correctness, and completeness of information shared to users across the Department.

Future plans for improving and expanding the functionalities of the GES include the incorporation of evolving satellite sensor

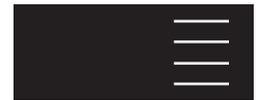
systems, development of a task management system, ingestion of state-declared digital site maps, and integration with the Additional Protocol System. These capabilities will enhance interoperability between different applications within the Department, making the GES an integral part of the Safeguards analytical environment.

This paper outlines the evolution in the development and deployment of GES Web 2.0. It presents major technical improvements and innovative solutions as well as streamlined processes in the exploitation of CSI and geospatial data. The paper also highlights the future plans for further improving the use of CSI within the Department.

Introduction

During the past 20 years, satellite imagery has become a significant component of safeguards-relevant information supporting safeguards implementation. A vision for establishing an integrated information technology (IT) environment aimed at receiving, disseminating, and evaluating all safeguards-relevant information in a manner commensurate with an effective and efficient implementation of safeguards was established in the mid-2000s. That led, inter alia, to the definition of a way forward to integrate all georeferenced data using Geographic Information System (GIS) technology. As a result, the Geospatial Exploitation System (GES) was designed as the first tool aimed at supporting both the team of imagery and geospatial analysts and the community of inspectors within the IAEA's Department of Safeguards.

Since its deployment in late 2011, GES has been used to successfully exploit and disseminate commercial satellite imagery (CSI), geospatial data, and a range of imagery-derived analytical products. The GES is an enterprise-wide, collaborative platform that uses commercial off-the-shelf (COTS) GIS technology to provide a secure gateway into the Department's extensive CSI repository and geospatial data holdings. With the increasing demands for CSI and geospatially attributed information to assist the IAEA's nuclear nonproliferation and verification mission, the



GES has become an important enabler within the Department. However, with the demands on the GES to incorporate and exploit increasing volumes of CSI and geospatial data, ingest data from new commercial earth observation satellites, adapt to evolving user requirements, and improve interoperability with an array of customized Modernization of Safeguards Information Technology (MOSAIC) applications, the GES required modernization, and GES Web 2.0 was launched. This paper will provide a brief overview of the concept, development, and implementation of the GES, followed by a description of the key enhancements and new functionalities for the core users of the platform — the IAEA imagery and geospatial analysts. The paper further outlines major improvements in the security framework, software, database architectures, web-based interface, and the introduction of enhanced analytical tools and workflows. With significant advances in CSI and geospatial technologies anticipated in the future, the paper discusses the ongoing efforts to further optimize and expand the capabilities of the GES to meet the evolving safeguards analytical requirements of tomorrow.

GES Context

A comprehensive account of the vision, development and deployment of the original GES is discussed by Rutkowski et al. The concept of the GES was born out of a departmental need to develop and improve workflows for exploiting CSI and authoritative geospatial data in support of nuclear Safeguards. Funded by a combination of voluntary contributions from member states and the IAEA regular budget, the Department of Safeguards initiated a project to design and develop a customized GIS-based platform for the management, analysis, and dissemination of thousands of acquired commercial satellite images, hundreds of imagery-derived authoritative geospatially attributed site plans, and an archive of satellite imagery analysis reports and products. New tools, standard operating procedures, and quality control (QC) workflows were implemented in GES to ensure that the integrity of the data and analytical holdings was maintained over time.

Today, the GES provides reliable, up-to-date access to a valuable repository of CSI, authoritative geospatially attributed site plans, and imagery-derived analytical products for hundreds of nuclear fuel cycle (NFC) sites around the globe to authorized users in the Department of Safeguards. Commercial satellite imagery (CSI) has become an increasingly valuable and unique source of safeguards-relevant information to the Department for identifying and remotely monitoring NFC-related sites, preparing

in-field verification activities, and contributing to the drawing of sound safeguards conclusions.

Since the initial launch of the GES, commercial satellite imaging and geospatial technologies have continued to evolve rapidly and today offer the IAEA a more advanced capability in support of its nuclear nonproliferation and verification mission. At the same time, the IT environment in which the Department operates has also become more complex, with ever-increasing demands to improve the performance and interoperability of IT software and systems to allow more efficient, timely, and meaningful integration and analysis of Safeguards-relevant information while respecting “need-to-know” principles. The Department recognizes that to maintain and improve its processes in the specialist disciplines of imagery and geospatial analysis, it needs to be proactive and strategically commit research and development (R&D) resources, with the essential support of the IAEA Member State Support Programmes (MSSPs), to further shape and sustain its capabilities. Upgrading the GES to meet the safeguards priorities and challenges of the future has been officially recognized in the Department’s Strategic R&D Plan under the Satellite Imagery Analysis project (SGIM-002) of the Development and Implementation Support Program for Nuclear Verification 2018–2019.

The GES operates on the Department’s Integrated Safeguards Environment (ISE) — a secure computing environment that hosts a range of customized applications, data, and services. ISE was established to facilitate a Department-wide, strengthened all-source collaboration environment, while securely hosting, processing, and delivering Safeguards data, information, and applications in a timely manner. The GES was the first analytical capability to be deployed in the ISE and has been an effective mechanism for promoting collaboration within the Department.

In 2015, the constant upgrade of the Department’s IT system was accelerated. The MOSAIC project’s objective was to develop customized software applications to take Department of Safeguards IT into the future. Following 3 years of development, MOSAIC was completed in May 2018. This modernized IT system supports most implementation processes, allowing better planning, conducting, reporting, and quality assessment of Safeguards activities. As the demands on the Department have continued to escalate over time, the requirement for the GES platform to integrate seamlessly with the modern IT environment in the Department has been critical.



Evolving Satellite Imaging and Geospatial Technologies

Since the initial launch of the GES in 2011, there have been significant developments in the CSI industry in the type, volume, and quality of CSI available to the IAEA. The commercial marketplace has seen a larger number of competing vendors offering, at lower cost, new multisensor earth observation satellites, expanding satellite constellations, improved imaging capabilities, and greater collection opportunities around the globe. Systematically acquiring, processing, storing, disseminating, and exploiting this ever-increasing volume and diversity of CSI through the GES has set new, more demanding and unforeseen burdens on the GES and the Department's IT infrastructure. Ensuring that users have all the available CSI and geospatial resources at their fingertips when they need it is crucial in fulfilling the IAEA's mission.

In recent years, one of the most pressing issues for the GES was the compatibility and inclusion of new commercial satellite sensor technologies that had become available. Incorporating the technical parameters of the CSI from new sensors was often a slow, tedious process in the GES that could entail a lengthy procedure and the subsequent release of updated GES versions only after extensive IT system testing. Furthermore, higher spatial resolutions captured by the new satellite imagery sensors usually translate into larger data volumes that require more storage space, increased processing time, and the need for different technical workflows to exploit in specialist imagery and GIS software. Over time, the performance of the GES gradually deteriorated, and maintaining the platform became more problematic due to the technical constraints and limitations of the outdated software and increased data holdings.

Not only has the spatial resolution of conventional optical satellite sensors improved over time, but Synthetic Aperture Radar (SAR) satellite sensors are also increasingly available to the IAEA through commercial vendors. These sensors provide the Department with new satellite imagery capabilities that allow the IAEA to significantly improve its ability to effectively monitor NFC-related sites around the globe. While the number and spatial resolutions of these commercial SAR satellite sensors are lower than the more conventional optical satellite sensors, the Department recognizes that these new sensors offer unique capabilities previously not available. Seamlessly integrating the increasing volume of SAR imagery into the GES is imperative if the Department is to take full advantage of these emerging sources of new safeguards-relevant information.

In recent years, there has also been a rapid development of innovative high-frequency or high-revisit-rate nanosatellite constellations. These technologies will further evolve and expand in the future, with new commercial providers and improved sensor capabilities. The availability of these relatively inexpensive satellite constellations may transform the way remotely sensed earth observation data is collected and exploited for nuclear safeguards purposes. Although the technology is new and the spatial resolution of these nanosatellites is still significantly lower than the most advanced conventional optical satellite sensors, nanosatellites have been successfully deployed since 2013 and today allow unprecedented daily imaging of the Earth's surface. The adaptation of these emergent satellite imaging technologies for nuclear safeguards purposes is currently in the early stages of being assessed and will likely require new processing and analytical approaches within the Department to be effectively utilized. The development of semiautomatic change detection processes and application of advanced machine learning to efficiently exploit the large volumes of new CSI that will be available from these nanosatellite constellations will require innovative technological solutions to be applied through the GES.

However, it is the transformational advances in geospatial technologies since the GES was launched in 2011 that have highlighted to many users within the Department the opportunities that now exist to greatly improve the discovery, evaluation, and analysis of geospatially attributed safeguards-relevant information. With the increasing availability and lower cost of CSI, geospatial data sets and mapping information have created a growing awareness of how geospatially attributed data from diverse sources can now be seamlessly integrated, analysed, and visualized through a modern GIS architecture. The integration of multisensor CSI, foundational geospatial data sets, georeferenced maps, 3D landscapes, and authoritative geospatially attributed site plans with other safeguards-relevant information derived from diverse open sources, state declarations, and IAEA in-field verification activities will offer the Department the opportunity to greatly enhance its ability to ensure that Safeguards implementation continues to be effective in the future. The GES needs to evolve and adapt to the changing technological landscape in such a data-rich world — the amount of geospatially attributed information and CSI that is increasingly accessible and streamed through GES Web 2.0 will feed into new innovative and advanced analytical capabilities in the Department. This capability is already available through specialist COTS big data analytics software or is being developed in-house, such as with the recently deployed Geo-based Data Integration (GDI) platform.

Description of the GES Web 2.0

In 2016, the IAEA initiated the GES Web 2.0 project to modernize the incorporation, exploitation, and distribution of CSI, geospatial data, and imagery-derived analytical products within the Department. The upgrade applied best practises from the development of the GES to date, as well as incorporating new and enhanced requirements that reflected the current and future needs of the users. The main objective of the project was to build the GES Web 2.0 for the future: that is, to identify and satisfy emerging end-user requirements, ensure better compliance with mapping and imagery industry standards, improve system scalability, strengthen IT security, and streamline system interoperability with other applications in ISE. Achieving these objectives required a comprehensive redesign of the outdated system architecture to be a multitier

system that includes a modern, feature-rich web-based interface as well as robust, secure web services and database layers to allow integration with specialist analytical software and other applications in ISE.

Architectural Changes

The new GES Web 2.0 is based on industry-leading COTS geospatial technology, which provides the back-end technological capabilities to publish dynamic web services. Furthermore, a client-side JavaScript application programming interface (API) is used to visualize CSI and geospatial data in the web browser. The GES Web 2.0 interface is intuitive, user friendly, and requires no previous experience using more sophisticated imagery or GIS software. ESRI's ArcGIS for Server and Image Server technology

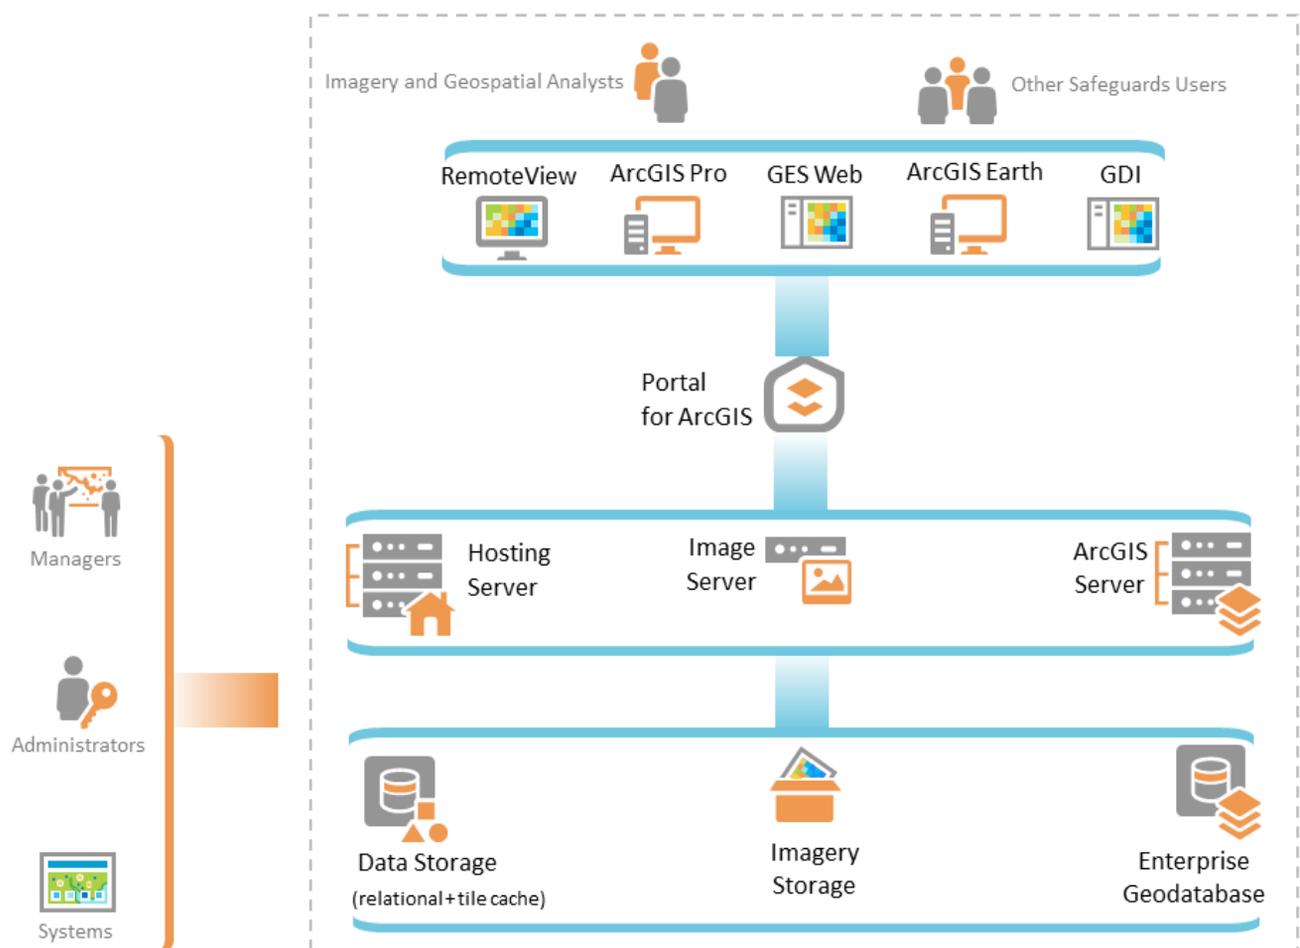


Figure 1. Architectural overview of the GES Web 2.0 system



gies are used to host and publish the two web services: an image service and a map service. The former is responsible for serving a growing catalogue of optical and SAR imagery in a fast and efficient way, whereas the latter enables serving reference map data and geospatially attributed site plans. The image and map services, along with other services available through the GES Web 2.0 platform, conform to the Open Geospatial Consortium map service industry standards. Conforming to modern industry standards has improved system maintenance and made possible a number of capabilities, including the ability to implement standard system security mechanisms, such as Integrated Windows Authentication (IWA); access map and image services from various COTS and open source applications; and upgrade essential core server software without breaking user software applications.

The GES system can support hundreds of concurrent users within the Department and be further expanded to support additional users simply through increasing server hardware resources. Upgrading the system architecture allows for the anticipated increased future usage of the GES Web 2.0 application and services with other software applications in ISE. Furthermore, GES Web 2.0 was developed in such a way that it is easy to implement future system upgrades. The state-of-the-art GIS platform ensures a high level of continuity and compatibility with the latest updates and versions of specialist analytical COTS software, such as RemoteView, ERDAS IMAGINE, and a suite of ArcGIS applications well into the future.

Figure 1 schematically shows the high-level architecture of the GES Web 2.0 system. Essential to the system architecture is the enterprise geodatabase, which serves as the data storage and indexing for all geospatial data. At the server tier for services, ESRI's ArcGIS Server and Image Server products are used to publish image and map services within the Department, whereas Portal for ArcGIS provides the visualisation, analysis, and collaboration platform for the users. All communication to system servers is conducted through standard web services that streamline the integration and interoperability with other IT platforms, including the GDI application. As part of the architecture changes, the data schema of the geodatabase and imagery storage for all geospatial and nonspatial data was redesigned to improve the performance and administration of the GES Web 2.0 system. These significant changes resolved outstanding technical constraints and performance issues of the previous GES by simplifying the overall data model and storage locations, removing duplicate data records, refining the attributed data structure of geospatial sites plans, and consequently optimising the data administration processes. The

new enhanced workflows better serve the evolving needs of the core users, the IAEA imagery and geospatial analysts, and have streamlined the handling of increasing volumes of data in the geodatabase and imagery storage.

Security Framework

Information security within the Department remains a high priority, and controlling access for authorized users to confidential information through the GES is critical. The system security for GES Web 2.0 applies industry standards and processes for user authentication and authorization, such as IWA. Importantly, the GES Web 2.0 platform is aligned with the Departments' new Authorization Management (AM) system, which effectively manages and coordinates access to information in ISE. Through different enterprise and task user roles, the GES Web 2.0 uses the AM system to dynamically query users' data access permissions and then filter outgoing requests accordingly. This is made possible through a technology available in the ArcGIS Server platform that can process all requests for data and filter the outgoing results to only the data that the requesting user or system is permitted to access, as defined by the AM system (Figure 2). Standard IT system monitoring and log auditing are also implemented against the GES Web 2.0 platform to provide another level of security awareness to system administrators. The AM system enables the GES Web 2.0 users to securely access all available CSI and geospatial data that they are authorized to use in their work.

Web-Based Interface

The GES Web 2.0 platform has a light, map-centric interface written entirely in the JavaScript programming language and accessible through a web browser (Figure 3). Unlike the Geobrowser user interface of the original GES, the new web-based interface does not require the separate installation of multiple software applications on user workstations, and thus, the system maintenance and upgrade processes have been streamlined. GES Web 2.0 can be accessed by users who have installed a web browser and who have been granted appropriate access privileges through the centralized AM system. The design layout of the interface integrates the best tools and practises from the previous GES as well as introduces new enhancements and functionalities. The central component of the new interface is the Map Viewer, which allows the user to easily navigate around the globe using an imagery base layer or a gallery of reference base maps. NFC locations of interest that a user is authorized to view, based on their AM permissions, are visually depicted as cartographic map

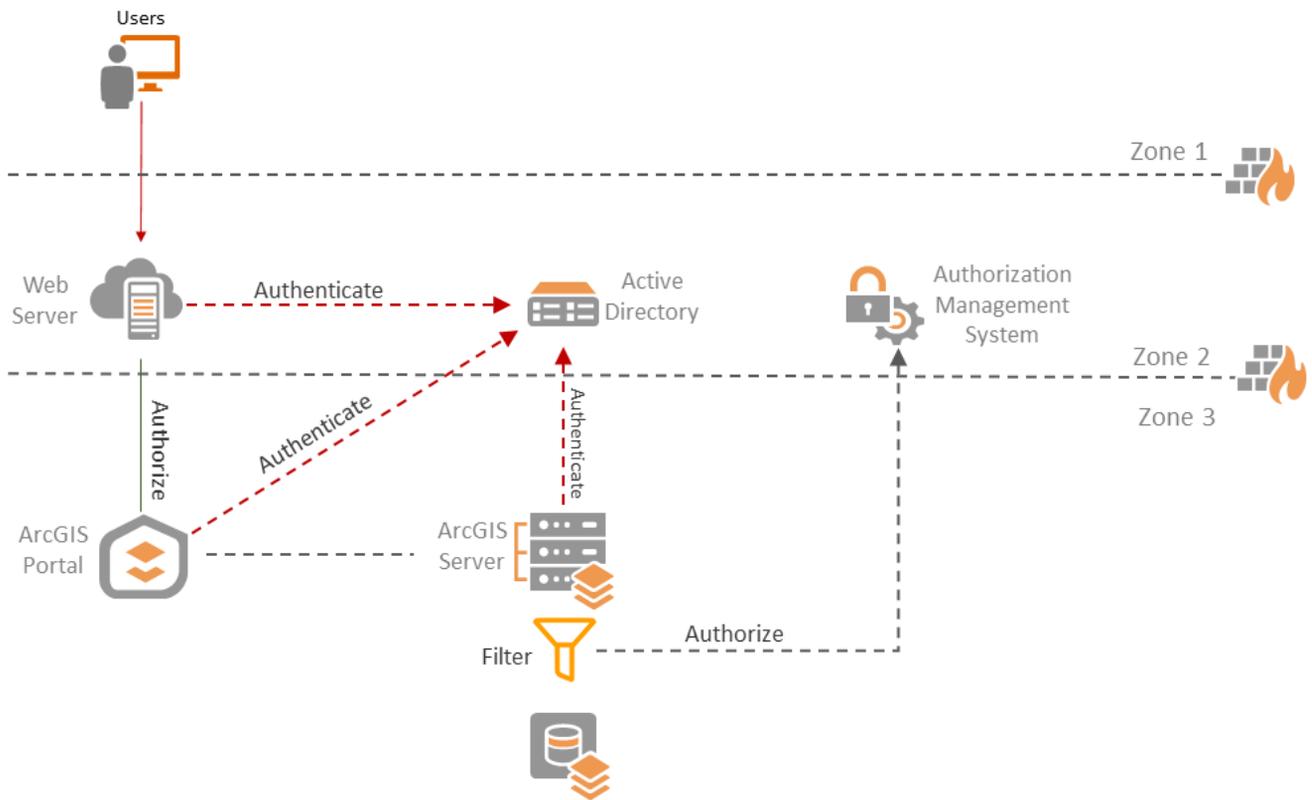


Figure 2. The principles of the GES Web 2.0 security framework

icons. A site selection panel lets users quickly access specific NFC-related sites available in the GES. Once a site is loaded, all CSI acquired by the Department for that location is available with the corresponding geospatially attributed site plan data and imagery-derived analytical products for viewing and further exploitation if required. Searching all accessible data holdings has also been greatly enhanced with the ability to now conduct more advanced queries for specific CSI, geospatial data, or derived analytical products.

A main driver of the GES upgrade has been to improve, streamline, simplify, and semiautomate various in-house imagery exploitation and geospatial analysis processes. These workflows typically consisted of standard manual processes developed over time, and while effective, they were often time-consuming, repetitive, and prone to human error, which could impact the integrity, correctness, and completeness of the data. Therefore, customized tools were introduced that significantly improved the QC processes for imagery exploitation and site plan creation

workflows and are now readily available through the web-based interface.

GES Web 2.0: Future Enhancements

The future enhancements of the GES 2.0 will be essential to the ongoing ability of the Department to optimize the use of CSI and geospatial analysis in its complex Safeguards mission. The GES Web 2.0 must remain flexible and adaptable to technological changes to strengthen its ability to deliver timely CSI, geospatially attributed data, and imagery-derived analytical products to its users within the Department. Preparing the GES Web 2.0 for the future changing requirements of the Department is both difficult and demanding; however, several key areas for future development in the short term are currently being implemented.

GES Task Management System. The TMS is being developed to streamline the internal administration of all workflows associated with the Department's satellite imagery requirements. The TMS will be fully integrated with the existing satellite imagery



Figure 3. The web-based user interface of the GES Web 2.0

analysis tasking request application in ISE, introduced as part of the Departmental MOSAIC Document Management and Workflow Support (DMWS) project. With the increasing demands for more insightful and timely imagery analysis of NFC-related sites around the globe, the TMS will allow for the streamlined management of CSI requests from imagery acquisition and processing, through to analysis and dissemination of analytical reports and products. The TMS will be an integral component of GES Web 2.0, which will enhance the prioritisation and productivity of the satellite imagery capability within the Department of Safeguards.

GES interoperability with the enhanced Additional Protocol System. The APS stores and disseminates state-declared information to authorized users. Previously, valuable state-declared information for NFC-related sites in the APS was integrated with geospatially attributed site plans available in GES through a time-consuming manual workflow. The improved interoperability between GES Web 2.0 and the APS will allow users to link state-declared information semiautomatically through standard data services to the digitized, attributed site plans in the GES.

GES interoperability with the Digital Declaration Site Maps (DDSM) initiative. States that have implemented the AP to their comprehensive safeguards agreement with the IAEA are currently obliged to provide the Department with site maps of NFC-related locations that include a general description of each building. These site maps, which are submitted in hard copy or in a variety of “static” electronic formats (e.g., PDF, JPEG), are stored in the APS. In an initiative to significantly improve the utility of this state-declared information, the DDSM project was launched with a few member states to test and evaluate the provision of site maps to the IAEA in a more usable digital, standardized geospatial format. The enhanced capabilities of GES Web 2.0 will allow the IAEA to automatically ingest the DDSM data into the GES for exploiting with other safeguards-relevant information.

Maintain and strengthen ongoing interoperability with the GDI platform. The GDI is a new collaborative platform for information integration, analysis, and activity planning involving geospatially attributed information. The Department’s extensive library of historical CSI, authoritative geospatially attributed site plans, and imagery-derived analytical products will be increasingly accessible from GES Web 2.0 through GDI. This will provide a critical foundation from which users can better share and integrate other types of safeguards-relevant, geospatially attributed data and analyses through the layered, interactive maps in GDI, and contribute to planning and evaluating the results of in-field verification activities.⁸

Conclusion

Since its launch in the ISE system in 2011, the GES has been an essential collaborative IT platform for the effective long-term storage, management, exploitation, and dissemination of CSI, authoritative geospatially attributed site plans, satellite imagery analysis reports, and other analytical products. The GES has been a successful mechanism in promoting the use of CSI and geospatial analysis within the Department. CSI has become an increasingly valuable source of independent Safeguards-relevant information for the IAEA. However, with the increasing demands for timely multisensor CSI and geospatial analysis to support the Department’s nonproliferation and verification mission, a significant upgrade to the GES was needed, and the new GES Web 2.0 application was deployed. With commercial satellite imaging and geospatial technologies rapidly evolving and user requirements ever changing, the GES needs to remain a responsive, efficient, secure, and scalable platform within the modern Safeguards IT environment into the foreseeable future.

The major system and database architectural changes have significantly improved the GES performance and ability to deliver up-to-date CSI, geospatial data, and analytical products efficiently, securely, and on demand to authorized users within the Department. Improved workflows in GES Web 2.0 to ingest, manage, and disseminate increasing volumes of multisensor CSI from a range of vendors enable faster retrieval of CSI and superior search and analytical tools to allow users to discover and exploit CSI, authoritative geospatial data, and analytical products in a timely and reliable manner.

Interoperability of the GES with other MOSAIC applications in ISE, such as the DMWS, AM, APS, and GDI, is a priority with further development planned in the near future, to ensure that the GES remains a critical collaborative platform of the Safeguards



IT architecture of tomorrow. Ensuring that the GES is able to reliably communicate and seamlessly integrate with a diverse range of Safeguards-relevant information sources and collaborative platforms within the Department will allow significantly better planning, conducting, reporting, and quality assessment of safeguards activities.

Keywords

Commercial satellite imagery, CSI, geospatial data, geospatial exploitation system, IAEA, nuclear safeguards

Author Biographies

Antero Keskinen is a Geospatial Analyst in the Department of Safeguards at the IAEA, Vienna, Austria.

Jacques Baute is the Director of the Division of Information Management in the Department of Safeguards at the IAEA, Vienna, Austria.

Mark Carey is an Imagery Analyst in the Department of Safeguards at the IAEA, Vienna, Austria.

Jonetta Ng is an Information Integration Specialist in the Department of Safeguards at the IAEA, Vienna, Austria.

Fatjon Ujkani is a Geospatial Technologist in the Department of Safeguards at the IAEA, Vienna, Austria.

References

1. Baute, J. (2007). Information management for nuclear verification: How to make it work in a sustainable manner. *Journal of Nuclear Materials Management*, Summer(4).
2. Rutkowski, J., Steinmaus, K., Dahl, F., Robb, S., & Kirkgoeze, R. (2016). The implementation and use of the geospatial exploitation system within the IAEA's department of safeguards. *Journal of Nuclear Materials Management*, XLIV(2), 39–44.
3. IAEA. 2018. Research and development plan: Enhancing capabilities for nuclear verification. SRT-385.
4. IAEA. 2018. Development and Implementation Support (D&IS) programme for nuclear verification 2018-2019. STR-386.
5. IAEA. 2016. MOSAIC — the Modernization of Safeguards Information Technology: Completing the picture. Retrieved from www.iaea.org/sites/default/files/17/01/mosaic.pdf
6. IAEA. 2018. <https://www.iaea.org/newscenter/pressreleases/iaea-completes-3-year-project-to-modernize-safeguards-it-system>, accessed May 18, 2018.
7. Rutkowski, J., Warner, T. A., & Niemeyer, I. (2017). A semi-automated change detection model for nuclear fuel cycle related activities using large remote sensing time series datasets and geographic information systems. *Proceedings of the 58th INMM Annual Meeting*. Indian Wells, CA.
8. Barletta, M., Yeoh, A., Marques-Patricio, A., Sumarokov, N., Robb, S., Fowler, M., & Carey, M. (2018). Geo-based Data Integration (GDI) in the IAEA department of safeguards. *Special issue of the Journal of Nuclear Materials Management* 46(3): Open Source and Geospatial Information Analysis, 2018.
9. IAEA. (1997). Model protocol additional to the agreement(s) between state(s) and the International Atomic Energy Agency for the application of safeguards. INFCIRC/540 (corr.).
10. IAEA. (2004). Guidelines and format for preparation and submission of declarations pursuant to articles 2 and 3 of the model protocol additional to safeguards. *Services Series* 11. Vienna.
11. Rutkowski, J., Keskinen, A., Steinmaus, K., Rialhe, A., Idinger, J., Nussbaum, S., & Balter, E. (2014). Digital declarations: The provision of site maps under INFCIRC/540 Article 2.a.(iii). IAEA-CN-220 Book of Abstracts, Presentations and Papers, p.115. *Symposium on International Safeguards: Linking Strategy, Implementation and People*. Vienna.



Toward a Multimodal-Deep Learning Retrieval System for Monitoring Nuclear Proliferation Activities

*Yana Feldman, Margaret Arno, Carmen Carrano, Brenda Ng, Barry Chen
Lawrence Livermore National Laboratory, Livermore, CA, USA*

Abstract

As the amount of available data increases, the human ability to locate, process, and analyze it is strained and eventually overwhelmed. To address this challenge for nonproliferation analysts, the authors have been creating a large-scale multimodal retrieval system to help analysts triage and search open source science, technology, and news data for indicators of nuclear proliferation capabilities and activities. The system relies on a set of deep neural networks (DNNs) trained to evaluate conceptual similarities across data modalities, such as text, image, and video. These DNNs can be used to search and prioritize data, according to a nuclear fuel cycle (NFC) process template, that are conceptually closest to the seed query items regardless of data modality. The DNNs have been trained to map conceptually related words, sentences, and images to nearby points in a multimodal feature space, enabling intra- and intermodal retrieval via nearest-neighbor calculations to seed query points. The authors evaluate the system's ability to retrieve NFC-related data that have been purposely hidden in collections of unrelated background data. They demonstrate quantitative and qualitative results for text-to-image, image-to-image, and image-to-video retrievals. This paper discusses data challenges confronting nonproliferation analysts, describes the DNNs designed to retrieve multimodal data that are proximal in the semantic space, and demonstrates their effectiveness by applying the DNNs to nonproliferation-specific multimodal data sets.

Introduction

Nonproliferation analysts are faced with growing challenges as the volume, rate, and variety of information outpaces and overwhelms the human ability to find and prioritize incoming data in a timely manner. Manual information processing is further complicated by (1) the potential for a low density of highly valuable data, (2) unlabeled or improperly labeled data, and (3) a growing number of sources that are difficult to systematically monitor, including increasing amounts of non-text data such as images and video.

To address these challenges, we have been creating large-scale multimodal retrieval systems that help analysts triage and search open source science, technology, and news data for indicators of nuclear proliferation capabilities and activities. Our system relies on a set of deep neural networks (DNNs) trained to search and prioritize data that are conceptually closest to the seed query items selected by analysts, regardless of data modality.

To help train parts of our DNNs, subject matter experts (SMEs) curated nuclear proliferation–relevant data sets in multiple modalities (e.g., text, image, and video), which were categorized with respect to a process template based on a generic model of the nuclear fuel cycle (NFC). The specialized subject matter made large-scale data set annotation efforts impractical; therefore, this research explored unsupervised and supervised pretraining on large-scale open source data sets like ImageNet and Wikipedia, followed by fine-tuning using the curated data sets. Here, fine-tuning means calibrating general models trained on open source data to be more proficient in NFC concepts. Our DNNs learn to map individual words, sentences, and images to points in feature spaces. Proximity in these feature spaces indicates conceptual similarity enabling unimodal and bimodal retrieval via nearest-neighbor calculations to query seeds (also referred to as seed queries or search queries).

Once trained, we used these DNNs to retrieve NFC-related data that we hid among large sets of unrelated background data and evaluated how well our approach was able to find the hidden data. We demonstrated the effectiveness of our system for retrieving images from image query seeds (image to image and stock image to real image), images from text query seeds (text to image), and video clips from image query seeds (image to video).

This paper discusses data challenges confronting nonproliferation analysts, describes the DNNs designed to retrieve multimodal data proximal in the feature space according to a process template (also referred to as a process model), and demonstrates their effectiveness by applying the algorithms to

nonproliferation-specific multimodal data sets. The paper concludes by discussing future research opportunities.

Developing a Large-Scale Multimodal Information Retrieval System

Imagine that you are given a large set of data that could fill multiple hard drives and you are tasked with finding everything related to your area of expertise by the end of the day. Such a task would be easy if every file was already labeled with a set of keywords describing its content. Even if the data were unlabeled, if you had some idea of the content of the data, and had representative examples of items in each category of data you were interested in, you could train your favorite machine learning model (i.e., a classifier) to automatically retrieve available data related to a given category of interest. However, when dealing with massive, unstructured data, not only are categories (also referred to as keyword tags) often nonexistent or too complicated or costly to create at scale, but they may even change over time.

Alternatively, if there was a way to map all the unlabeled data to a feature space where conceptually related multimodal data are proximal, then you could retrieve relevant data by looking for the nearest neighbors to your search query (see Figure 1). For example, if you wanted to find data related to uranium enrichment by gas centrifuge, you would examine this multimodal feature space and see where images, video, and text descriptions of gas centrifuges are mapped by the system. This research designed a system to do just that — to take a data instance (image, text, or video) and return other data instances that are proximal in the multimodal space. The system also makes advanced definition of categories of interest unnecessary, allowing us to search data that are unlabeled and/or from unknown categories.

In our work, we use DNNs to learn these multimodal feature spaces. Figure 1 depicts DNNs mapping text descriptions, an image, an audio waveform, and a video of a gas centrifuge to proximal locations in a multimodal feature space. The creation of this system of DNNs involves two main steps: (1) train DNNs to learn high-quality feature representations of unimodal data, then (2) learn mappings from these unimodal feature spaces to the multimodal feature space.

Since the specialized subject matter makes large-scale data set annotation efforts impractical, our two-step approach is designed to reduce the system's dependence on large amounts of NFC annotated data for training. In our DNN system, most of the trainable parameters are dedicated to learning unimodal representations, so we can train individual DNNs separately, in

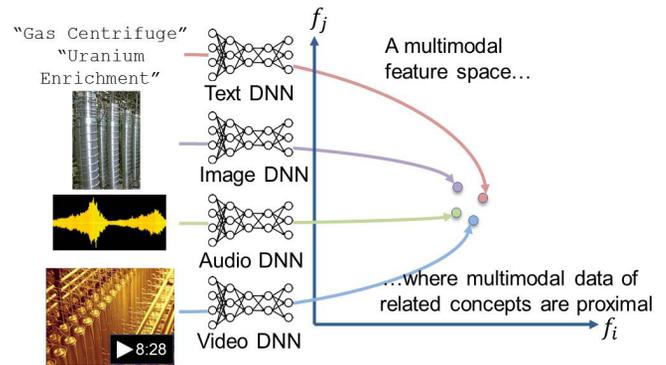


Figure 1. A system of deep neural networks (DNNs) that maps text, imagery, signal, and video into a multimodal feature space where conceptually related data are proximal. In this example, text tags, an image, an audio file, and a video related to a gas centrifuge enrichment process are mapped to nearby points in the multimodal feature space by a system of DNNs.

parallel, on massive amounts of non-subject-specific labeled and unlabeled data of the appropriate modality, and further fine-tune using small, curated data sets focused on proliferation-related data. Once the unimodal DNNs are trained to understand their own modalities, the system requires much fewer subject-specific data pairs (e.g., images and corresponding text descriptions) for training the “merger” DNN layers responsible for mapping from unimodal features to the joint multimodal feature space. In fact, we have found that good retrieval performance can be attained by including as few as a dozen examples per NFC category of interest when training the merger DNN layers.

Process Models

How the data is structured in the feature space, and what drives data retrieval, is guided by a process model that reflects expert knowledge on a specific subject or research question. In this research, the process model is a simplified version of the nuclear fuel cycle (NFC), as shown in Figure 2. The first-level process nodes represent the major NFC processes that could lead to the production of highly enriched uranium metal or plutonium metal (displayed in blue in Figure 2). Processes that could be accomplished using one or many methods, such as uranium enrichment, each with distinct equipment and visual signatures, were further subdivided into second-level nodes (displayed in green in Figure 2). In the case of reactors, third-level nodes were used to differentiate the reactor types within the second-level categories (displayed in yellow in Figure 2).¹

Using this process model, the neural networks are trained to map raw data into a semantic feature space. Each piece of

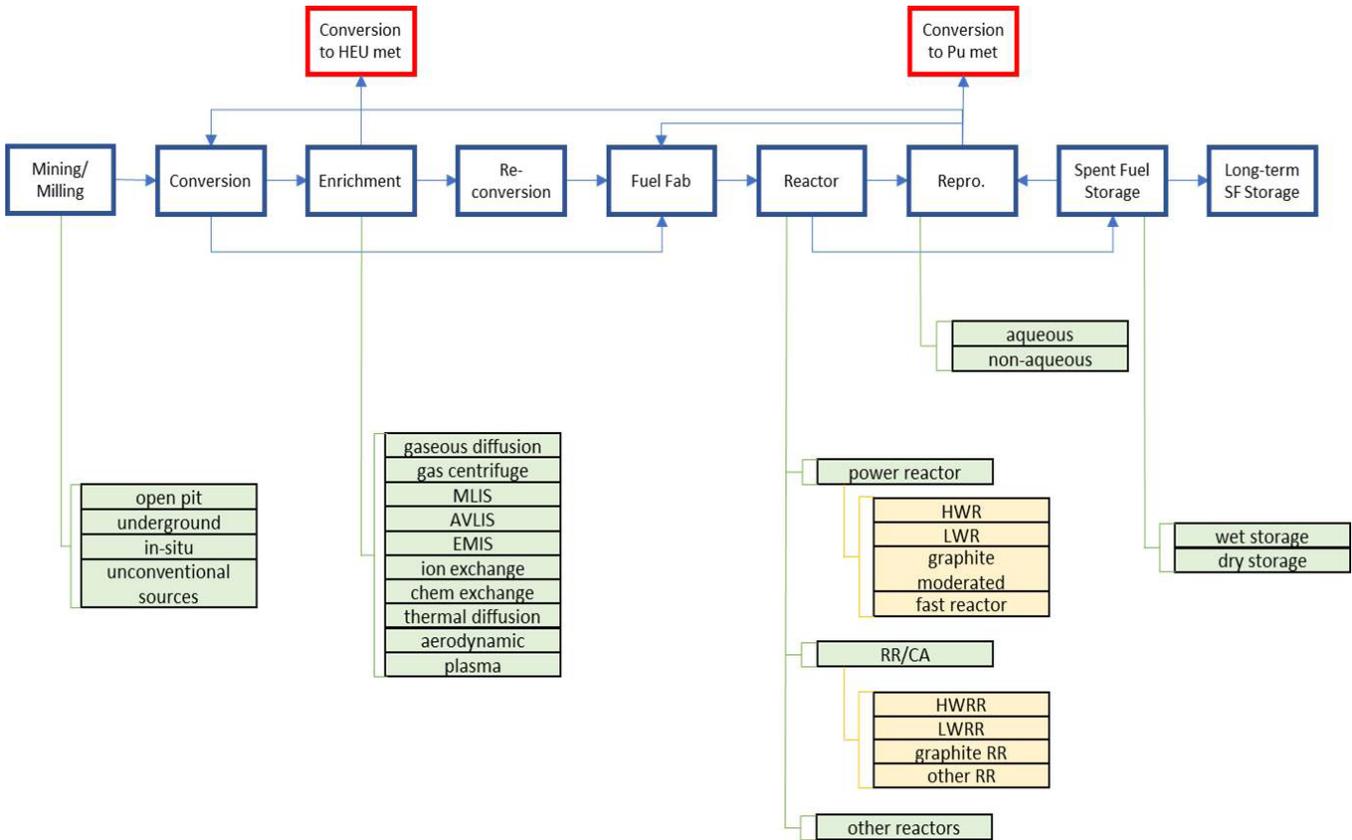


Figure 2. A process model that illustrates a simplified version of the nuclear fuel cycle

data in the curated proliferation-specific data sets is labeled using terms that correspond to one or several nodes of the process model. If a datum is representative of a subnode, it is labeled corresponding to the subnode and all related higher-level nodes. For example, an image of a gas centrifuge is labeled as “gas centrifuge” and “enrichment,” while a video of a fuel fabrication facility is labeled as “fuel fabrication.” This way the process model (1) guides the categorization of data within the feature space and places the query and search results in the context of a proliferation question, and (2) drives data retrieval in a way that directs the analyst’s attention to a particular process node where new data may indicate a change in status or development of a proliferation capability.

Process models, such as the one used in this research, are employed by analysts to describe sequences and relationships among activities, events, or entities relevant to proliferation. They may include representations of physical processes used to produce and transform nuclear material into weapon-usable forms, organization charts of key entities in a proliferation

procurement network, or graphs modeling a hypothetical strategy to conceal the construction of a clandestine enrichment facility. Which process model analysts use is driven by their specific research question — that is, “What evidence exists that a state has capabilities to reprocess spent fuel?”

Subject-Specific Multimodal Data Sets

Guided by this process model, we created nuclear proliferation-specific multimodal data sets containing NFC-related text, images, and videos.² In our approach, these data sets were split into two sets: a training set to fine-tune and improve the unimodal representations (also referred to as embeddings), and a test set in which its elements serve as query seeds for information retrieval. Each data set is detailed below.

Text Data Sets

We used two types of text data sets: a background data set and a series of curated proliferation-specific data sets. The background text training data set is a 20-year snapshot of *The New York*



Times, from 1987 to 2007, and a snapshot of Wikipedia, generated on April 3, 2015; all together these two data sets contain 149 million sentences. To achieve a more relevant text embedding, we needed to augment this general data set with subject matter-specific data. To that end, we created a focused collection of text documents (ranging from several sentences to multiple pages in length) drawn from technical documents on various fuel cycle processes,³ in which the extracted text was edited to be exclusively about a particular process or an element of the NFC. After we corrected for text errors due to object character recognition or regional spellings, as well as deleting any nonsensical text (e.g., misaligned rows from tables, orphan labels from figures), this subject-specific collection totaled 31,000 sentences. After excerpts and short articles were identified, they were labeled with one or more nodes that correspond to our process model.

Image Data Sets

Similar to the text data sets, we had two types of image data sets: a background data set and curated proliferation-specific data sets. For our background data set, we used ImageNet,⁴ an open source visual database consisting of more than 14 million hand-annotated images classified into 21,000 different categories. For our proliferation-specific visual data, our collection has focused on objects that are uniquely recognizable as being part of the NFC, such as gas centrifuges, nuclear fuel rods, and hot cells. The curated image data set so far consists of approximately 700 images of uranium gas centrifuges (individual centrifuges, centrifuge cascades, centrifuge schematics), flow forming machines, reactor fuel elements (Magnox, PWR, BWR, CANDU), reactor cooling towers, spent fuel pools, diagrams of reactor core fuel matrices, hot cells, and various types of UF₆ cylinders. As in the text data set, each image was labeled with one or multiple nodes corresponding to the process model.

In putting together the curated image data set, the “prominence” of the object of interest within the image, inversely related to the amount of background clutter, was taken into account. The collection in each category was organized according to the following scheme: (1) the cleanest images (those with only one instance of the item and with little to no background clutter), (2) images with bundles or multiples of the item with little to no background clutter, (3) multiples with clutter, (4) diagrams or schematic drawings of an item, (5) drawings or schematics in which the item of interest is only one of several elements, and (6) images where the object or objects of interest are obscured to some degree. This differentiation is useful in establishing the sensitivity of the

system’s performance to the level of clutter in the training images.

Video Data Set

The video data set assembled to date for this research consists of more than 8,000 annotated NFC-specific video scenes totaling about 21 hours. Each scene is given a detailed annotation stating what is happening in the scene and whether it is of NFC relevance. The annotations typically describe what is seen in the video (including objects, actions, and other visual elements, such as captioned text or posters and signs in the background), what is heard in the audio, and other SME observations. Moreover, each scene is assigned labels corresponding to the process model nodes. Roughly 90 percent of the scenes in the current data set were ultimately deemed relevant to an NFC process model, although this statistic may change as we continue curating and expanding this data set.

Videos were collected in a variety of languages, including Russian, Farsi, English, and French, making annotation difficult in some cases when an analyst was not familiar with the language. Overall, each video clip was watched roughly three times to understand what was being seen and to demarcate the start and end of a scene. Half of the videos were viewed by at least two analysts, and each annotation went through a quality control check, making video curation a time-intensive process that took more than 270 hours for the 21 hours of video in the current collection.

Training the Deep Neural Networks

As mentioned earlier, the creation of this system of DNNs that can project data of multiple modalities onto a feature space where conceptually related items are proximal involves first training DNNs to learn high-quality feature representations of unimodal data, then learning “merger” mappings from these unimodal feature spaces to the multimodal feature space.

Learning Unimodal Feature Representations

The goal of learning unimodal feature representations is to learn mappings from raw unimodal data to a feature space where similar unimodal data are proximal. These unimodal feature spaces are similar in spirit to principal components or basis vectors of data, but they can be more expressive — that is, have higher dimensionality and be nonlinear.⁵ To learn these features, DNNs that take raw unimodal data as input (e.g., pixel values for images or one-hot vectors for words) are trained. These unimodal DNNs are trained using modality-specific objectives, which may be super-



vised or unsupervised. For labeled data, the supervised objective may be to classify the raw unimodal data into some prespecified categories.⁶ For unlabeled data, the unsupervised objective may be to predict co-occurrences of context data instances given a target data instance,^{7,8} or even to reconstruct the unimodal data.⁹ The neural network learning algorithms seek to find the set of learnable parameters in the DNN that best achieves the objective of the training data set. After training is complete, one can feed a raw unimodal data instance to the DNN and take the outputs of the penultimate layer of the neural network as the unimodal feature vector for that input data instance. In essence, the training objective forces the learned model to map qualitatively similar input data instances to numerically proximal features.^{8,9} Thus, we can apply this feature space as a surrogate for the raw features for our information retrieval.

In the sections below, we show the retrieval results from our system that uses the following DNNs for unimodal feature representations. For images, we use the imagenet-21k-inception model (INC21k) from the Distributed (Deep) Machine Learning Community,¹⁰ an openly available convolutional neural network based on Google's Inception-BN network¹¹ trained on ImageNet,¹² which classifies images into 21,000 different categories. The model was then fine-tuned using a collection of NFC-related images, as described above, categorized into 12 categories. For text, we trained a custom FastText¹³ model on the text training set described above — that is, a combination of text from *The New York Times*, Wikipedia, and NFC-related documents. We take the output of the penultimate layer of the INC21k as our unimodal features for images, which gives us a 1,024-dimensional feature vector for each image. We use the standard FastText word embeddings as the unimodal text feature vectors, which results in a 200-dimensional feature vector for each word.

Learning the Projections from Unimodal Features to the Multimodal Feature Space

Once unimodal feature representations have been learned, the next step is to train another set of neural networks that map the unimodal features to the multimodal feature space where conceptually related multimodal data are proximal. To accomplish this, we require training data consisting of aligned (conceptually related) data from different modalities. In Figure 3, training data consisting of an image of a gas centrifuge and the associated keywords “gas centrifuge” and “centrifuge cascade” are fed into their respective unimodal neural networks for text and images. The resulting unimodal feature vectors, (for text) and (for images),

are then transformed into multimodal feature vectors and . The goal of this training process is to learn the neural network layers parameterized by and that perform the transformations (i.e., and) from the unimodal space to the multimodal space, such that the cost function is minimized. The cost function is designed to return high values in cases where multimodal feature vectors of conceptually related images and text are far from each other, and, conversely, low values when they are close. In other words, the cost function encourages the network to map text and image representations of conceptually related data close together. We use a cost function similar to that found in Zhang et al. (Fast0Tag),¹⁴ which was described more fully in Ni et al.¹⁵ Specifically, the cost function has four terms: the first term rewards the neural network if images are mapped closely to their text labels in the feature space; the second term penalizes otherwise; the third term ensures that similar text labels, such as “cooling pond” and “spent fuel pool,” would have similar features; and the last term ensures that dissimilar text labels would be “pushed” apart in the feature space. The neural network is trained using back-propagation,¹⁶ a standard algorithm for training neural networks.

This cost function encourages projecting related data along the same direction in the multimodal feature space. For training the multimodal transformations, we use the training partition of the ESP-Game data set,¹⁷ which consists of 53,525 images and corresponding tags. We processed the tags by filtering out words that occur fewer than 30 times to form a set of 1,200 unique keyword descriptors. In addition to the ESP-Game data, we use 12 labeled examples from each of the object categories in our NFC-related image data sets, where categories correspond to select nodes in our process model, for further training the multimodal projections. We found that adding even a few examples of NFC images for fine-tuning our model significantly improves our retrieval performance.

Once the projections from the unimodal feature space to the multimodal feature space have been learned, we can map images and words onto the multimodal feature space and retrieve like items simply by returning the nearest neighbors to a seed query. For example, if we wanted to find images of uranium enrichment gas centrifuges, we would project the seed query word “centrifuge” into the multimodal feature space and return the images that were closest. One of the major benefits of this approach is the flexibility it affords for new and previously unseen queries. Despite never having seen the co-occurrence of images of missiles and the tag “missile,” for example, we have seen how this approach was still able to retrieve missile images from the

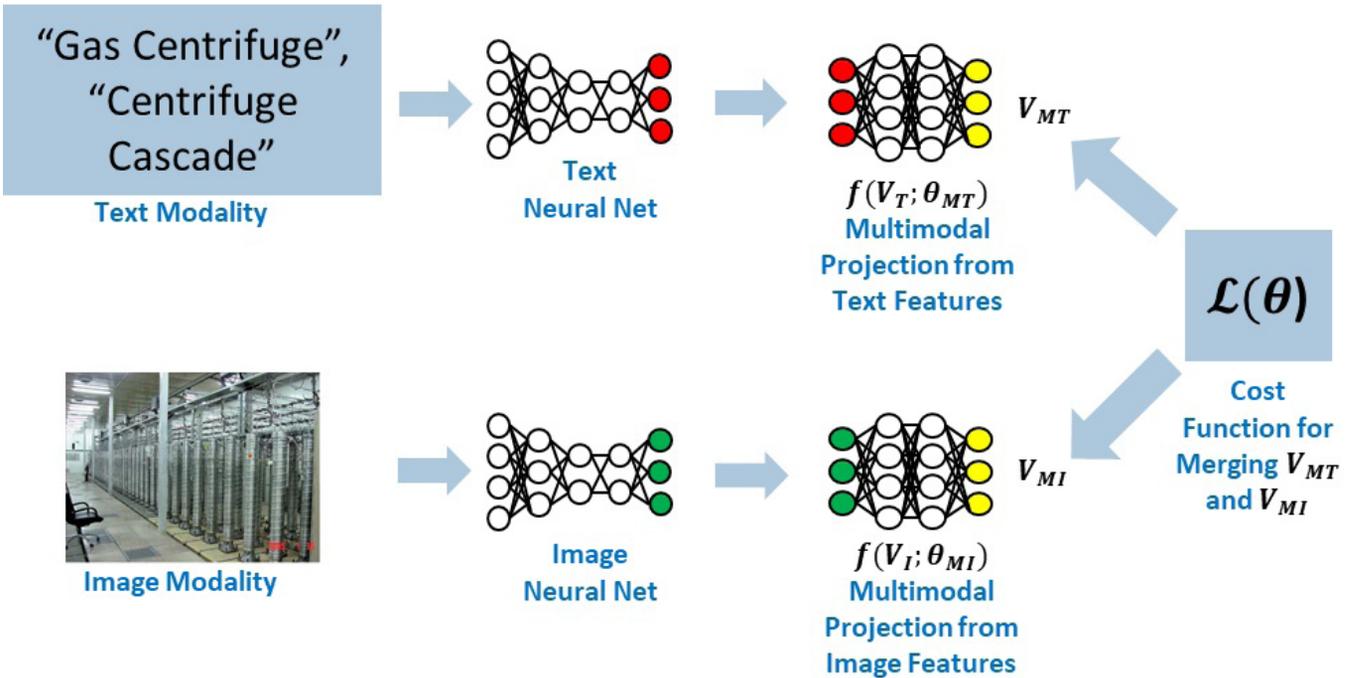


Figure 3. Conceptually related data are first projected into their respective unimodal feature representations. Multimodal transformations are learned by additional neural network layers that map image and text features onto the multimodal feature space. The cost function is responsible for encouraging the merger networks to map conceptually related images and text to proximal locations multimodal feature space.

query seed “missile” because the system learned to map images and words of concepts related to missiles to similar locations in multimodal feature space.

An additional benefit of this method is that it would permit rapid integration of new modalities of data. The overall framework can be viewed as a tightly coupled ensemble of supervised learning regression models whose targets (the individual unimodal representations) slowly gravitate toward semantically related concepts as these concepts are learned and refined across different modalities. This would allow for flexible construction and rapid incorporation of specialized modalities, enabling fast reconfiguration of the system to accommodate data from new sources and sensors for analysis and exploration.

Example Retrieval Results

This section describes system performance in the retrieval of unimodal (image-to-image, stock-image-to-real-image) and bimodal (text-to-image, image-to-video) data. Qualitative and quantitative results are presented.

Image-to-Image Retrieval

To evaluate how well our system performs, our general approach was to take our NFC-related images and embed them in a background set of more than 50,000 images from the ESP-Game data set. We then assess how well our system can retrieve the targeted NFC-related images out of this varied set of unrelated background images. This is akin to hiding needles in haystacks and assessing how well our system can find those needles. Quantitatively, we use standard retrieval metrics, such as average precision and mean average precision,¹⁸ that measure to what extent the relevant data (needles) are ranked higher than the background data (hay). Table 1 displays average precision results for retrieving images related to one-and five-seed query images (i.e., image-to-image retrieval) for categories of NFC-related items found in our NFC process graph. As the results in Table 1 demonstrate, providing the system with more seed images improves average precision scores because the multiple seed images encourage the system to find a more diverse set of results that are close in feature space to more than one specific query image.¹⁹



Table 1. Average precision results for image-to-image retrieval using one- or five-seed images.

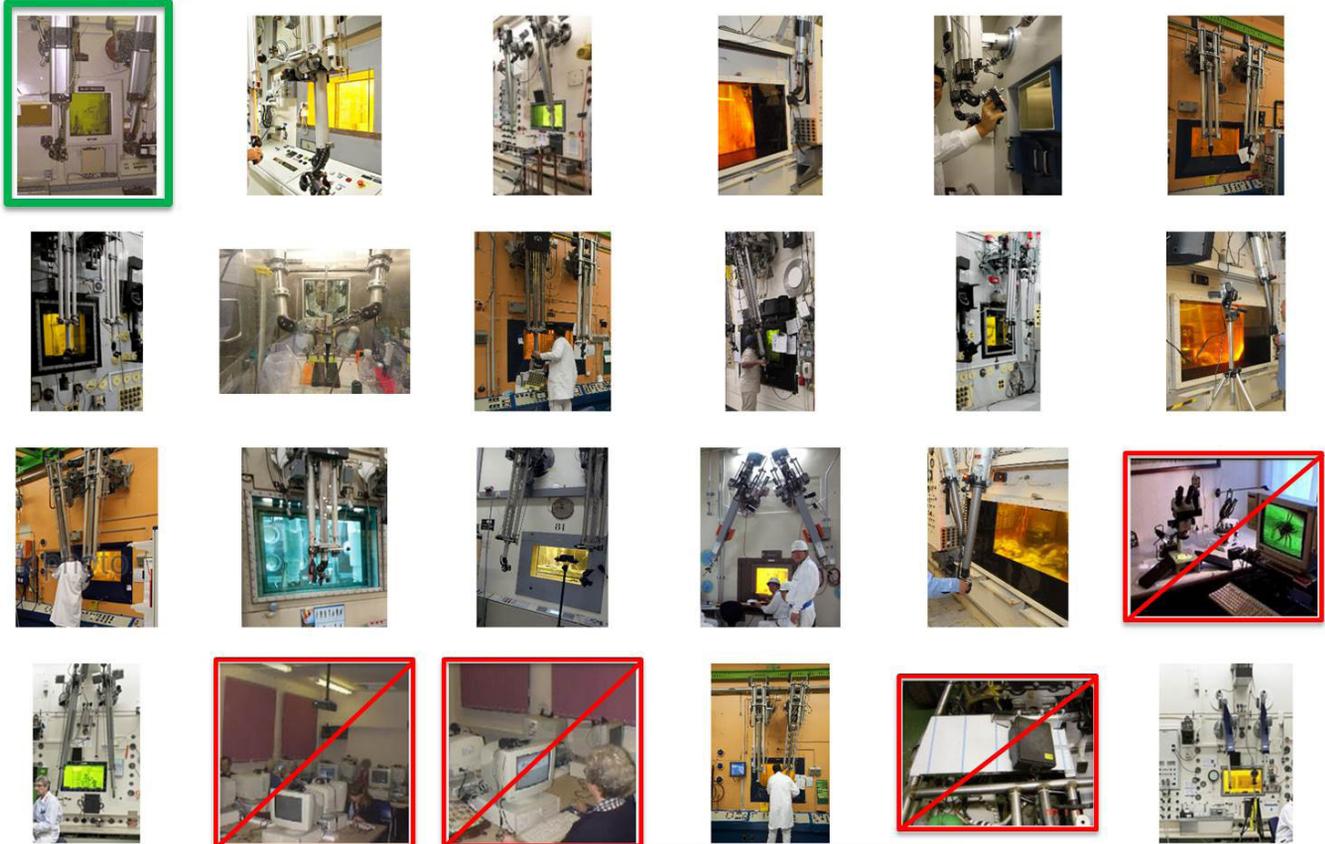
Retrieval Categories	Average Precision	
	Image-to-Image 1 seed image	Image-to-Image 5 seed images
centrifuge	0.13	0.29
centrifuge cascade	0.16	0.35
cooling pond	0.46	0.58
cooling tower	0.57	0.63
hot cell	0.37	0.51
pressurized water reactor	0.06	0.14
magnox fuel element	0.02	0.03
flow forming machine	0.34	0.46
uranium hexafluoride cylinder	0.32	0.5
centrifuge schematic	0.09	0.14
reactor core fuel matrix	0.27	0.47
candu fuel bundle	0.22	0.4
mean Average Precision (mAP)	0.25	0.38

Qualitatively, we can display the retrieval results in ranked order and see whether our needles are at the top of the hay stack. In Figure 4, we ask the system to find images that are closest in the feature space to the seed image of a single hot cell (highlighted in green). The system returned 19 out of the 54 total hot cell images hidden in the background set within the top 23 retrieval results.

Stock-Image-to-Real-Image Retrieval

Figure 5 demonstrates the ability to retrieve realistic images of an object, such as a piece of technology, by asking the system to find images that are closest to stock images of that object. Using two stock photos of flow-forming machines, our system was able to return three out of five realistic images of flow-forming machines (highlighted in yellow) that were hidden in a background set of more than 50,000 images in the top 23 results.

Seed Image



19 true hits out of 54 in top 23

Figure 4. Image-to-image retrieval: hot cell

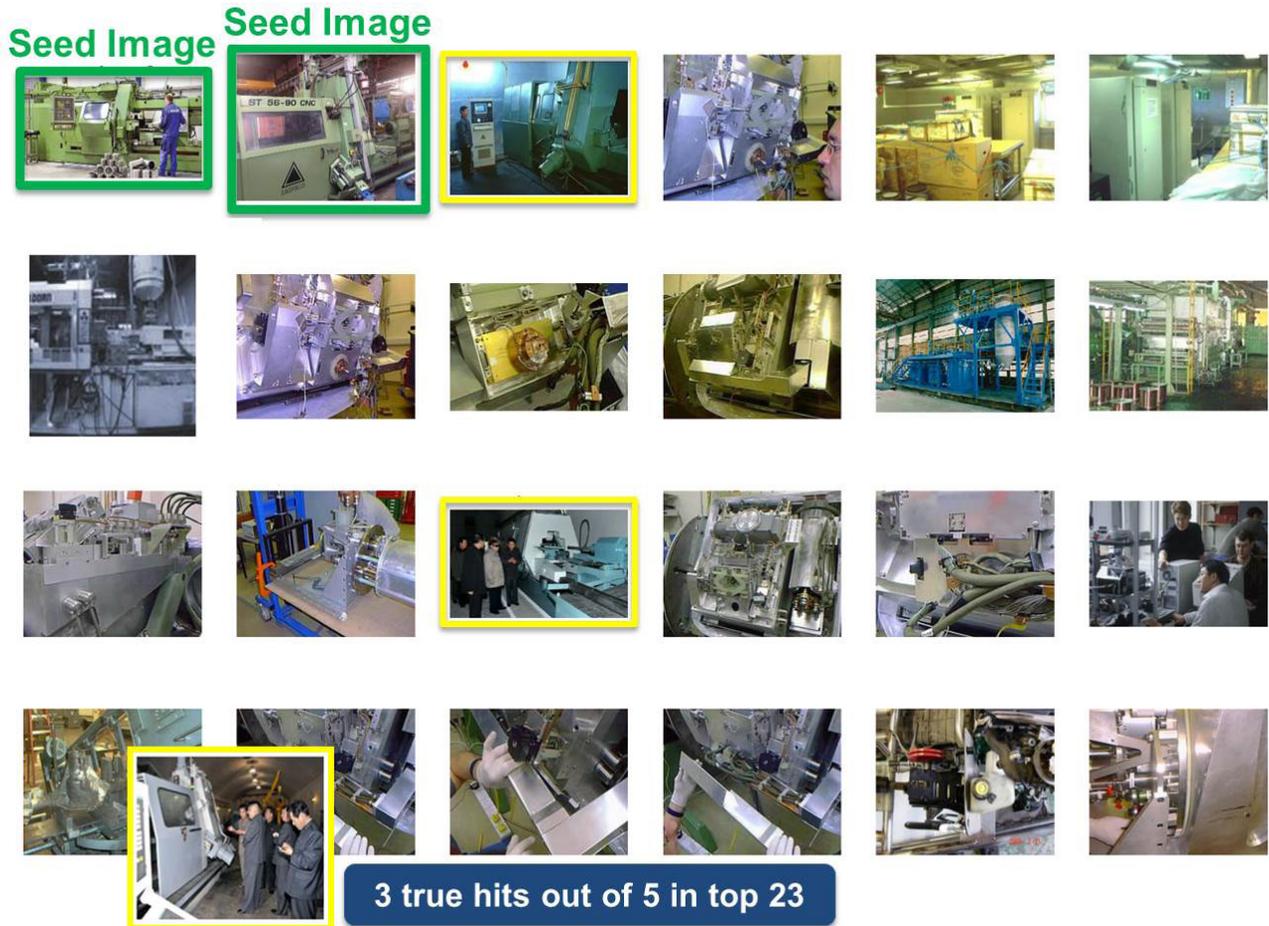


Figure 5. Image-to-image retrieval: flow-forming machine

Text-to-Image Retrieval

In Table 2, we display the average precision scores for retrieving images from text seed queries. In the “Zero Shot” case, our system is not trained on image and text pairs occurring in our NFC data sets. In the “Few Shot” case, we provide the system with a few examples of our NFC image/text pairs (on average, about 12 instances per category). Two points are worth noting: first, achieving nonzero average precision for the Zero Shot case is impressive, considering the system received no pretraining on the subject-specific topic; second, by adding even a few NFC image and text pairs for training, our system improves retrieval performance by approximately 40-fold. As mentioned earlier, in specialized subject matters such as nuclear nonproliferation analysis, we will have limited amounts of annotated data for training machine learning systems, making systems that can effectively leverage what little is available important.

Table 2. Average precision results comparing text-to-image retrieval performance of a system that was not exposed to NFC text and image pairs for training (Zero Shot) versus a system that had access to a few examples (about 12 per category) for training (Few Shot).

Retrieval Categories	Average Precision	
	Text-to-Image Zero Shot	Text-to-Image Few Shot
centrifuge	0.0011	0.55
centrifuge cascade	0.0006	0.49
cooling pond	0.0017	0.80
cooling tower	0.1526	0.91
hot cell	0.0012	0.72
pressurized water reactor	0.0011	0.37
magnox fuel element	0.0002	0.11
flow forming machine	0.0003	0.37
uranium hexafluoride cylinder	0.0062	0.89
centrifuge schematic	0.0020	0.42
reactor core fuel matrix	0.0006	1.00
candu fuel bundle	0.0016	0.66
mean Average Precision (mAP)	0.0141	0.61



Text seed:
'centrifuge cascade'

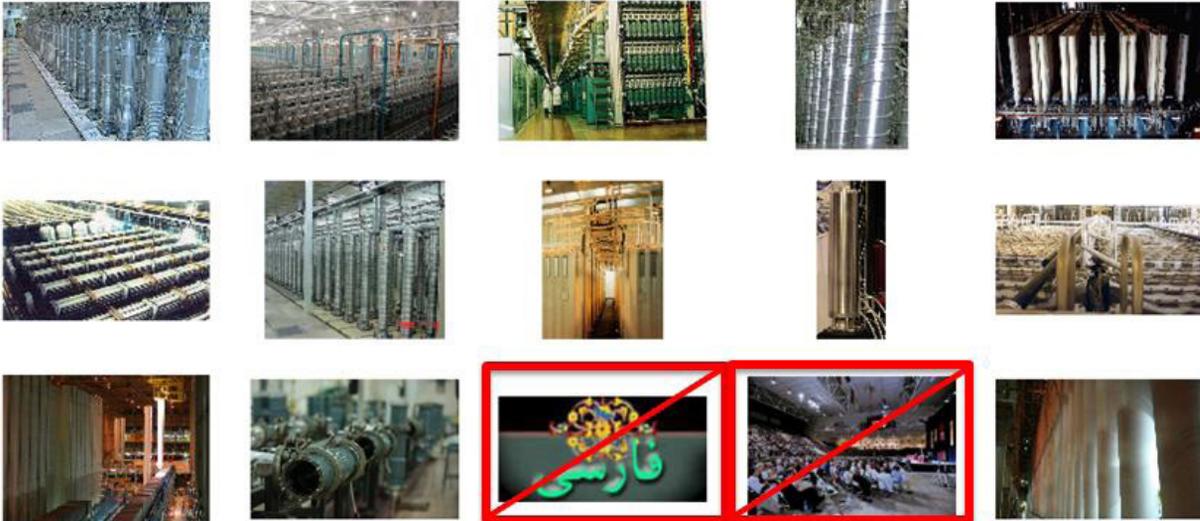


Figure 6. Text-to-image retrieval: "centrifuge cascade"

Text seed:
'centrifuge', 'schematic'

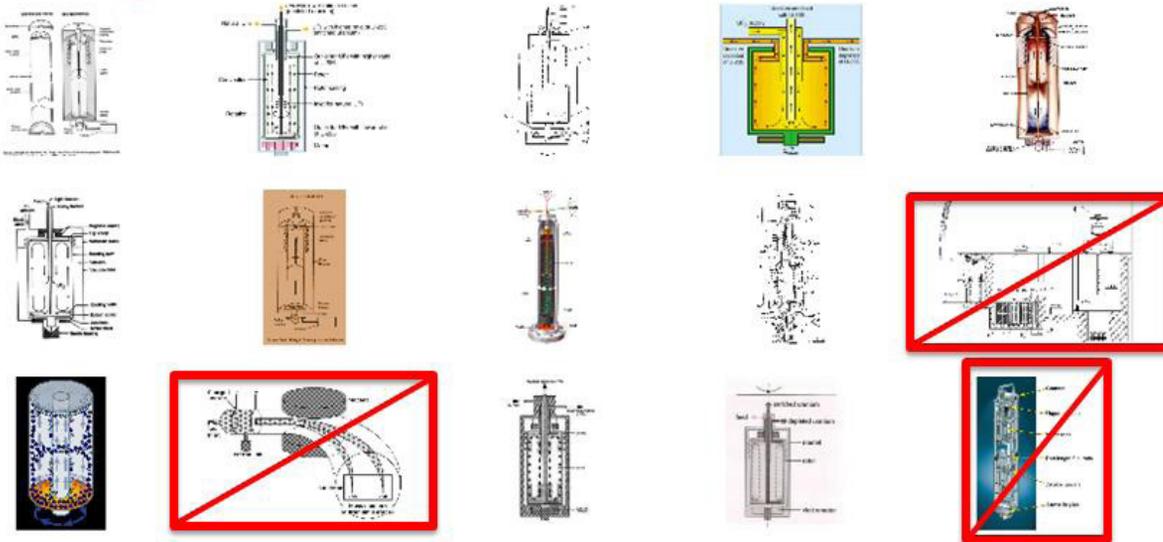


Figure 7. Text-to-image retrieval: "centrifuge," "schematic"



Figure 8. Image-to-video retrieval: cooling towers

Figure 6 demonstrates bimodal text-to-image retrieval where, given the phrase “centrifuge cascade,” we look for the nearest images in the multimodal feature space. In the top 15 retrieval results, only two images are unrelated to gas centrifuges, and one contains a single centrifuge not in a cascade. The qualitative results again demonstrate impressive retrieval performance, as substantiated by many true hits in the top 10 to 25 ranked results.

Our system also demonstrates the ability to find drawings and schematic images (Figure 7). In this text-to-image retrieval example, our system retrieves drawings and schematics of uranium enrichment gas centrifuges by using “centrifuge” and “schematic” as the seed text terms. In the top 15 results, only three images do not show a gas centrifuge schematic.

Image-to-Video Retrieval

Our last example (Figure 8) shows how we can use the multimodal feature space to look for videos that contain frames that are related to a chosen image seed. Here, we demonstrate finding scenes with cooling towers in a collection of videos by performing a frame-by-frame matching in the multimodal feature space with a seed image of a cooling tower. The top 3 videos all have frames that contain examples of cooling towers.

Discussion

While the multimodal retrieval system described in this paper is

still in development, the preliminary results demonstrate the efficacy of applying it to nuclear nonproliferation analysis, specifically in data identification, collection, and filtering. Analysts typically perform separate searches for text, images, and videos of interest in different data repositories, using different search strings. The multimodal feature space currently allows unimodal and bimodal data retrieval, and will eventually incorporate multimodal retrieval. By integrating and structuring different types of conceptually related information in one common space, our DNN retrieval system enables the user to identify and analyze untagged data across modalities in an integrated way, therefore streamlining data collection.

The lack of accurately and adequately labeled data presents a collection challenge for nonproliferation analysts. As shown by our efforts to put together NFC-specific image and video data sets, systematically scouring the open source domain for unlabeled visual data is difficult, underscoring the need for a multimodal retrieval system for locating data that has not been, or has been improperly, labeled. Being able to perform visual searches based on the semantic meaning of image content using targeted subject matter seed queries, as demonstrated by our preliminary results, would be a new capability for nonproliferation analysts. It would potentially provide access to troves of new open source visual information typically underexplored using text-based searching. While the practicalities of using this retrieval system on



all of the open web are beyond the current scope of this project, the viability and benefit of this approach with respect to mining information from an analyst's existing collections of partially labeled images is clear.

Our DNN system would also reduce the amount of time required by analysts to process videos. The current workflow relies heavily on manually reviewing videos to determine their relevance. For example, a video titled or described as "uranium trafficking" may include scenes not reflective of the title or its description — for example, about a local research reactor. The potentially relevant content would be discovered only if the video were manually reviewed (and subsequently properly categorized and tagged). The increasing rate at which videos are uploaded to the Internet, coupled with the inability to know if a video contains NFC-relevant data, makes it impossible for analysts to watch every potentially relevant video manually. Our image-to-video results suggest the possibility of querying videos based on visual elements both related and unrelated to the video's title or description. The multimodal retrieval capability described in this paper could help automate, and therefore accelerate, review of video data by identifying key scenes of interest for the analyst, obviating the need to manually evaluate every video frame.

Finally, using process templates to categorize data is an effective way to represent an expert's knowledge on a particular research topic. The ability to retrieve semantically similar data, without necessarily knowing in advance how it relates to the process model, will help data collectors and junior analysts identify data they may not have been aware was relevant to the research question.

Conclusion

This paper described the development of a large-scale system of DNNs designed to map and retrieve multimodal data proximal in a multimodal feature space, with respect to an analyst-derived process model. We demonstrated the effectiveness of our approach by applying the DNN algorithms to open source science, technology, and news-based multimodal data sets to identify possible indicators of nuclear proliferation capabilities and activities.

We plan to take our future research in several directions. First, we intend to continue growing our curated data sets, both in depth and in breadth of coverage of the NFC process model. The growth of these data sets would necessitate a more scalable approach to data management. As this work is a first step toward applying deep learning to multimodal nonproliferation data retrieval, we used the standard version of nearest-neighbor

search for our retrieval algorithm, which required that all items be stored in memory. We recognize that improvements can be made along the lines of scalability, by leveraging the optimizations for distributed search algorithms.^{20–22}

We also plan to develop and incorporate video-specific DNNs. While the system as described in this paper treated video as a set of static images, we believe that video-specific DNNs — those that model video as an evolving sequence of images and learn to follow objects, focus attention on dynamic cues, and segment objects and represent their state — will greatly improve video retrieval results. In particular, we expect improvements in cases where temporal information is important, such as videos where objects are moving behind obscurations or are on the periphery.

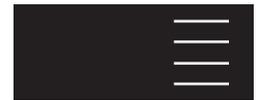
Another future area of research is to expand our range of modalities to include transactional data, such as statistical and individual transaction data on NFC-relevant commodity flows. This will require curation of an appropriate data set, development of process model(s) to represent transactional flows, and development of new vector representations for transactional data. We plan to integrate the transactional features into the multimodal feature space and to demonstrate cross-modal retrieval on our proliferation data sets.

We are continuing to evaluate how this approach to information retrieval compares to and improves on the traditional ways of locating proliferation-relevant open source information. We plan to begin integrating the technologies developed as part of this research into the generic workflows of nonproliferation analysts.

We believe our semisupervised machine learning approach — which combines unsupervised and supervised pretraining on large-scale open source data sets with fine-tuning of the DNNs on curated, subject-specific data sets — lays the groundwork for developing a large-scale multimodal retrieval capability for monitoring nuclear proliferation activities. Such a capability would significantly improve an analyst's chances of locating and retrieving hard-to-find nuggets of relevant information in a sea of unlabeled open source data.

Acknowledgments

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, appa-



ratus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

Keywords

IAEA safeguards, multimodal information retrieval, deep neural network, unsupervised learning, semantic feature space

Author Biographies

Yana Feldman is a Nonproliferation and International Safeguards Analyst at Lawrence Livermore National Laboratory (LLNL). Prior to coming to LLNL, Ms. Feldman worked as a Safeguards Information Analyst in the Department of Safeguards at the International Atomic Energy Agency, focusing on open source information analysis for safeguards state evaluation. She holds a Bachelor of Science in chemistry from UC Berkeley and a Master of Science in international relations from the London School of Economics.

Margaret Arno is a Nonproliferation Analyst at Lawrence Livermore National Laboratory. She holds a Bachelor's in mathematics from Occidental College and a Master's in nonproliferation and terrorism studies from the Middlebury Institute of International Studies at Monterey.

Carmen Carrano is an engineer working on Global Security projects in the Computational Engineering Division at Lawrence Livermore National Laboratory (LLNL). Her general research interests at present include multimedia research with a current emphasis on machine learning for video analytics and multi-modal learning applications. She is actively contributing to the image and video understanding efforts underway at LLNL (e.g. deep feature extraction, analysis, and search). Additionally, she leads

the long-range, high-resolution surveillance imaging through atmospheric turbulence with speckle imaging efforts at LLNL. Ms. Carrano received her B.S. and M.S. in electrical and computer engineering from the University of California at Santa Barbara.

Dr. Brenda Ng is the group lead for the Machine Learning Group in the Computational Engineering Division of Lawrence Livermore National Laboratory. She is an expert in machine learning with interests spanning deep learning, reinforcement learning, multi-agent systems and probabilistic graphical models. Her recent work includes uncertainty quantification and decision-making under uncertainty for carbon capture technologies, multi-modal deep learning for knowledge discovery, and active learning for design of laser experiments. She received her PhD in Computer Science from Harvard University.

Dr. Barry Chen is a Machine Learning researcher at the Lawrence Livermore National Laboratory (LLNL) with over 20 years of experience in developing and applying novel machine learning algorithms to a wide variety of applications including automatic speech recognition, image classification, and threat analysis and detection. He has deep expertise in Neural Networks, Random Forests, Graphical Models, Hidden Markov Models, and Support Vector Machines. Dr. Chen currently leads several research teams at LLNL focused on the development of new neural network learning algorithms that address recurring challenges in scientific applications including: the scarcity of labeled data, the multimodality of data types, and the scalability of large models on massive training datasets. He received a Ph.D. in Electrical Engineering and Computer Science from UC Berkeley.

End Notes

1. Reactors (and other fuel cycle steps) can be categorized in different ways. Here, for simplicity, we chose to organize reactors along moderator type.
2. In the absence of a reliable, subject-specific visual-based search — i.e., similar to what is being developed as part of this research, text-based querying was used to locate potential images and videos to be screened by subject matter experts for inclusion in the image and video data sets.
3. So far, this has been done for enrichment, reactors, reprocessing, and spent fuel management; collections for other processes of interest are also in progress.
4. Russakovsky, O., et al. (2014). ImageNet large scale visual



- recognition challenge. arXiv:1409.0575.
5. Bourlard, H. & Kamp, Y. (1988). Auto-association by multilayer perceptrons and singular value decomposition. *Biological Cybernetics*, 59(4–5), 291–294.
 6. Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*.
 7. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*.
 8. Mundhenk, T., Ho, D., & Chen, B. (2018). Improvements to context based self-supervised learning. *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*.
 9. Baldi, P. (2012). Autoencoders, unsupervised learning, and deep architectures. *Proceedings of the ICML Workshop on Unsupervised and Transfer Learning*, 27, 37–49.
 10. Campr, P. & Li, M. (2016). imagenet-21k-inception. *GitHub repository*. Retrieved from <http://github.com/dmlc/mxnet-model-gallery/blob/master/imagenet-21k-inception.md>
 11. Ioffe, S. & Szegedy, C. (2015). Batch normalization: accelerating deep network training by reducing internal covariate shift. *Proceedings of the International Conference on Machine Learning*.
 12. Russakovsky, O., et al. (2014). ImageNet large scale visual recognition challenge. arXiv:1409.0575.
 13. Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2016). Enriching word vectors with subword information. arXiv:1607.04606.
 14. Zhang, Y., Gong, B., & Shah, M. (2016). Fast zero-shot image tagging. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 15. Ni, K., Zaragoza, K., Gude, A., Tesfaye, Y., Carrano, C., Foster, C., & Chen, B. (2017). Sampled image tagging and retrieval methods in user generated content. *Proceedings of the British Machine Vision Conference (BMVC)*.
 16. Le Cun, Y. (1988). A theoretical framework for back-propagation. *Proceedings of the 1988 Connectionist Models Summer School*, 1988.
 17. von Ahn, L. & Dabbish, L. (2004). Labeling images with a computer game. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 319–326. Vienna: ACM.
 18. Manning, C., Raghavan, P., & Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge: Cambridge University Press.
 19. When multiple seeds are specified, the average distance to all the seeds is used for determining the nearest data points in the multimodal feature space.
 20. Patwary, M. A., et al. (2016). Panda: Extreme scale parallel k-nearest neighbor on distributed architectures. arXiv:1607.08220.
 21. Babenko, A. & Lempitsky, V. (2016). Efficient indexing of billion-scale datasets of deep descriptors. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 22. Kalantidis, Y. & Avrithis, Y. (2014). Locally optimized product quantization for approximate nearest neighbor search. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.



A Definitive Pipeline to Display Multivariate Geotechnical Data in a 3D Environment

James B.A. Palmer, Jennifer Stevanović

Atomic Weapons Establishment Plc, Aldermaston, Reading, United Kingdom

Abstract

On-site inspections are a vital component of verification regimes across a number of international treaties. Many agencies carry out on-site inspections or have the remit to do so, if required. Inspections produce large quantities of geospatial or geolocated data. Modern computer hardware and software makes it possible to recreate the location of such inspections and display the data collected in situ. The authors used the specialist visualization tool GeoVisionary to build such interactive 3D environments. Once built, the environments could be explored using standard computer displays and immersive virtual reality systems. This paper relates to the authors' work building environments for two real-world events: the Comprehensive Nuclear Test Ban Treaty Organisation's¹ large scale on-site inspection² exercise called the Integrated Field Exercise 2014^{3,4} in Jordan, and LETTERPRESS,⁵ a 2017 "quad" exercise by Norway, Sweden, the United States, and the UK for nuclear arms control verification research. This paper reports on the development of a process to guide the design of an optimal visual form for any geospatial (or geolocated) data type to be displayed in a 3D environment. The authors found few default display options for data to be displayed in 3D. The process was applied to test cases of data types produced during the events: seismic wave data, terrestrial laser scanning data, magnetometry, ground resistivity tomography, computer-aided design models, handheld photography, and aerial photography. The results are presented using screenshots taken from within the environments showing the visual forms in isolation and in the geographical context. The authors discuss the results, benefits, and challenges of applying the process. They conclude that the benefits justify the initial work required to design optimal forms, create them, and use them to build and make use of an interactive 3D environment. Furthermore, some types of information can be extracted from the data more rapidly by using a 3D environment than by using some conventional techniques — in particular, information that is revealed only by cross-comparing data from different techniques, such as magnetometry and aerial photography.

© British Crown Owned Copyright 2018/AWE

Introduction

The Comprehensive Nuclear Test Ban Treaty and On-site Inspections

On-site inspections are a vital component of verification regimes across a number of international treaties. Many agencies carry out on-site inspections or have the remit to do so, if required. Inspections produce large quantities of geospatial or geolocated data. This paper relates to the authors' work investigating the potential for using digital, interactive three-dimensional (3D) environments to aid on-site inspection activities. It reports on the development and use of a process to design an optimal visual form for geospatial data that is to be displayed in those 3D environments when no default form exists.

As part of the UK's support for the Comprehensive Nuclear Test Ban Treaty (CTBT)¹, research into developing and improving the CTBT Organisation's (CTBTO) on-site inspection (OSI)² capabilities has been carried out at the Atomic Weapons Establishment (AWE), UK, with funding from the UK Ministry of Defence. One area of support was to research how to use data from the imaging remote sensing techniques permitted by the CTBT. As that work progressed, and by witnessing the CTBTO's large-scale on-site inspection exercise, called the Integrated Field Exercise 2014 (IFE14),^{3,4} it became clear that collecting imaging data was of limited use if there was no effective way to present the information contained in that data in a form that the inspectors could use productively. The work then expanded into investigating the best options for displaying and exploiting the collected imagery data. The solution chosen was to display the information/data in an interactive 3D environment that recreated the actual inspection area and also displayed the imagery data gathered during an inspection. It followed that such a 3D environment could also be used to display many of the other geospatial data sets collected during the inspection. These included magnetometry, ground resistivity tomography, and seismic signals, among others.



The 3D visualization tool GeoVisionary⁶ 3.0 and 3.1 from Virtualis⁷ was chosen for the 3D visualization work. It was created to support the work of the British Geological Survey, which has many overlaps with on-site inspection activities, as both deal with large quantities of geospatial data and geotechnical data.

The wide range of options for how to display data provided by such 3D visualization tools incites careful consideration. For some data sets, there was a clear default or an existing convention for displaying them in 3D, but for others there was not. As part of this work, we developed some processes to guide the design of the visual form the data should take. Those processes and the results of applying them are presented in this paper.

Some data types relevant to this paper were not encountered with the IFE14 data set. To illustrate those data types, and for completeness, we present later work using data generated during LETTERPRESS,⁵ a quad exercise by Norway, Sweden, the United States, and the UK for nuclear arms control verification research, held in the UK in 2017.

The Challenge of Visualizing Geotechnical Data

During an inspection, lots of geotechnical data is collected, which is both spatial and geospatial in nature and can consist of surface measurements, vertical section measurements, volumetric measurements, and point measurements. Our experience is that it is easier to interpret the data when it is seen in its actual geographical context, along with related data from other techniques. Critically, some information is revealed only by cross-comparing different data sets when they are in their actual geographical context.

Displaying the data can be challenging for a number of reasons, including the following:

- Traditional ways of displaying the data may not translate well to a 3D environment
- Data may have a geospatial aspect but not be suitable for 3D display
- Options for how data can be displayed require a general knowledge of visualization tools and constraints of the specific visualization tool to be used, which other technical experts will not necessarily be aware of
- Data generally requires extra processing before it can be imported into the visualization tool, in addition to that done during conventional processing
- Technical issues exist relating to file format compatibility between programs

- The digital size of some data sets makes it difficult to display them without making compromises — for example, some of the image files used in this work were larger than 15 GB, which sets constraints on the minimum capability of the computer hardware used
- The limits of the display technology (e.g., dynamic range, use of only three color channels, and display resolution) and of the human eye–brain vision system (e.g., dynamic range; limited to seeing only three color channels: red, green, and blue).

Problem Statement

The objectives of this piece of work were to find optimal solutions to display complex geospatial data in an interactive 3D environment and, where necessary, to define a processing pipeline to enable the data to be displayed in a way that adds the most value to the inspectors.

Method of Work

This section describes the software, hardware, and data sources used. It describes the process created to guide the final design of the visual form each data type was to take in the 3D environment, and it states the data types that were test cases for the process.

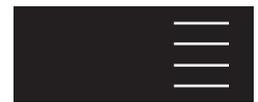
GeoVisionary

GeoVisionary was the commercial visualization tool chosen for this work after a review of the options available for displaying and interacting with large and diverse geotechnical data sets. It was to be used to reconstruct the IFE14 (after the exercise) as comprehensively as possible given the data available in the IFE14 data store and any other data products that could be generated from it. This led to the development of processes to decide how to display particular data types in a 3D environment, which is the topic of this paper.

Data created outside of the software is loaded into it and displayed in 3D. The display can be viewed on a standard monitor with or without depth perception or via a virtual reality system. The three principal reasons for using this software were that (1) it did not need to be connected to the Internet, (2) it could load very large terrain and imagery files yet retain and display their original resolution, and (3) it can use and display some standard geographic information system (GIS) data types, including shapefiles.

Using GeoVisionary

Once the GeoVisionary software had been sourced, a process



was developed to digest various data from the IFE14 exercise into the environment.

1. Learn how to use the GeoVisionary software and other tools used
2. Review, collate, and (where possible) organize the available data from the IFE14 data store
3. Outline the contents of possible environments
4. Create any extra data products as required
5. Decide what visual form each data type should take and what information that form should encode
6. For each data type, act on the results of Step 5 and produce data products that can be imported
7. Build the environments required:
 - a. The pre-inspection environment
 - b. The environment to assist in the general operation of the inspection
 - c. The environment to act as a repository for data collected
 - d. The environment to be used as briefing tool for multiple audiences
8. Explore each of the environments in Step 7 and assess how useful they are

These tasks identified the need for processes to help guide the creation of suitable visual forms for different technical data sets.

Hardware and Software Requirements

In addition to GeoVisionary, various tools were used for this work to prepare data to be imported into an environment, are listed in Table 1 and Table 2. Tools used to collect data or produce data products during IFE14 are not included.

Table 1. Software used for this work

Program	Purpose
GeoVisionary 3.0, Virtualis	3D data visualization
GeoVisionary 3.1 Alpha version, Virtualis	3D data visualization with virtual reality capability
SteamVR, Valve Corp.	Software required for HTC Vive virtual reality (VR) system
ArcGIS Desktop 10.3, ESRI	General GIS activities
ENVI+IDL, Harris Geospatial Solutions	Image analysis and data manipulation
Accutrans 3D, Micromouse Productions	3D model display, manipulation, and format changing
Photoscan Pro, Agisoft	Structure from motion processing to recreate terrain and objects from photographs
Windows Movie Maker, Microsoft	Editing video clips generated in GeoVisionary and producing animations

Data

Data sources for IFE14

The principal sources of data were the CTBTO website hosting the IFE14 data that AWE had been granted access to, data provided directly from the CTBTO, and data products generated by the authors at AWE.

The IFE14 data store holds approximately 300 GB of data collected or generated during the IFE14. Unfortunately for this work, a significant amount of the data had been stored in proprietary data formats that the authors did not have access to read. Consequently, not all of the data types could be worked with, but in some cases, conversion software was available for a data file.

The following data types were chosen to be displayed in the 3D environments (although not all are presented in this paper):

Table 2. Hardware used for this work

Item	Description	Purpose	Notes
Laptop computer	HP ZBook 17 G2, 2.7 GHz i7 processor, 32 GB RAM, Windows 7 64 bit, NVIDIA Quadro K5100M graphics card, 2.5 TB SSD storage	General computing requirements and to run the visualization software	GeoVisionary benefits from using a powerful graphics card
4k monitor	FlexScan EV3237, 3840 x 2160 pixels monitor, Eizo Ltd.	External display for laptop	The work benefited from the large pixel count of the 4k monitor
3D mouse	Various models of SpaceMouse from 3Dconnexion	Navigating environments in GeoVisionary	A 3D mouse allows a more natural interaction with the 3D environment
External hard drives	diskAshur Pro, 2TB hardware encrypted drives, USB3, iStorage Ltd.	Storing data	Large data sets could not always be stored on the laptop itself
Virtual reality system	Vive with audio strap accessory, HTC Corp.	Providing the user interface for immersive VR work	The system has head tracking and tracking for two handheld controllers



- Height data
- Satellite imagery
- Aerial imagery
- Surface magnetometry
- Ground resistivity tomography
- Gama ray spectrometry
- Global Navigation Satellite System location data
- Region boundaries
- Paper topographic maps
- Paper geology maps
- Paths of roads
- Routes of waterways

Data sources for LETTERPRESS

Some data types not used for the IFE14 environment were used in the LETTERPRESS environment, so we include them for completeness. The data types were terrestrial laser scans of buildings and a computer-aided design (CAD) model of one of the physical props used.

Data preparation

Most data needed some form of preprocessing before being imported into GeoVisionary. The type and complexity of preprocessing varied between data formats and specific data sets during this work. First, some data required significant processing to change it from its initial form into a visual form suitable for 3D display (the topic of this paper). Second, some data needed to be converted from one file format to another, or required a standard processing task; this paper doesn't discuss data preparation of that type.

Processes for Creating 3D Visual Forms to Display Data

The capability to display data in an interactive 3D environment offers new opportunities that require conventional methods for displaying data to be revisited and makes possible totally new ways to display some data types. Careful thought is needed to choose the correct way to display data so the information it contains is conveyed to the viewer.

In this work, we identified a general process to help decide how to display particular data types given specific criteria in terms of the information to convey to a specific audience. The process included the following steps:

- Assess the standard information, variables, and data products produced for the data type or source of data to be displayed

- Assess which of those pieces, or types, of information the audience for the visualization needs to know, and then list them
- Assess the standard way that those variables/information are presented
- List the variables/information to be encoded in the visual form(s) to be designed for that data type
- Against each list item, list options for encoding that item (e.g., magnitude information can be encoded as a color)
- Iterate through design options and test them in isolation, and with other data sets, to decide on the optimal solution for the given constraints

Applying this general process revealed several subsets of it that applied to particular data types or situations:

- Abstracting complex data types into a visual form
- Converting complex data to a simple visual form
- Converting a simple data type to a simple visual form
- Converting a simple data type to a complex visual form
- Combining multiple diverse data sets into a visually uncluttered form with minimal obscuration between data

Here, we define a complex data type either as one that has too many dimensions to be immediately displayed in a 3D environment, or one that is not inherently geospatial in nature but was collected during the IFE14. We define abstracting as the process of designing a method for deciding how to display a complex data type in 3D so that it conveys multiple variables to viewers in an easy-to-understand visual form, without obscuring the other data in an environment, where possible. Examples of complex data that can be abstracted are seismic data and single-sample gamma spectroscopy data.

Additional factors to be considered were the technical details of the measurement technique, the sensor or instrument used, the capabilities and constraints of the visualization tool, and the overall purpose of the visualization. The data available for both IFE14 and LETTERPRESS presented a number of test cases for the process, and the outcome of applying it to them is given in the Results section.

Results

The results are presented by stating the data used for the test case, a short discussion of the case and applying the process to it, a table listing the variables/data to be displayed and how each one was to be visually encoded, and, finally, a figure showing how the visual form is presented. At the end of the section, all of the vi-

sual forms used both in isolation and in the geographical context are presented in tabular form.

Abstracting Complex Data: Seismic Aftershocks

Data used for the test case: spreadsheet from the IFE14 data store

The Seismic Aftershock Monitoring System⁸ (SAMS) was designed by the CTBTO for use in an on-site inspection to detect, locate, and characterize seismic events within and close to the inspection area. Some of those events could be directly related to an underground nuclear explosion. The system consists of a number of seismic stations deployed around the inspection area. Any signals indicative of a seismic event are analyzed to locate the epicenter of the event and characterize it. There is no obvi-

ous way to represent the raw SAMS data in a 3D environment because it has the form of multiple traces of seismic data for each SAMS station. Consequently, information pertinent to the inspection must be abstracted from these raw traces and visualized in a way that adds value to the entire inspection team.

We assessed that several of the data products derived from the initial SAMS data and contained the information (variables) the inspectors actually needed to know. Those variables are listed in Table 3 along with the chosen method to encode them. The visual form chosen to encode those variables did not cause excessive visual clutter when displayed with other results. An example of that form is shown in Figure 1. Additionally, it was determined that information about the SAMS stations themselves was not relevant to most inspectors, so only the location

Table 3. Variables from SAMS to be encoded into its visual form

Variable	Encode form
Seismic event location	A marker at the event location, its form defined by the following variable
Confidence in the calculated parameters of that event	A multipointed star with a number of points equal to the number of SAMS stations (or another proxy for confidence)
Magnitude	The color of the location marker (using a normalized or absolute scale as required)
Associated uncertainty ellipse	An ellipse with the same dimensions and orientation as the calculated spatial uncertainty. A solid line outline with a semitransparent area color (this color could encode event depth).
Depth of event	A visually narrow rod dropping vertically from the marker to the event depth and terminated by a sphere

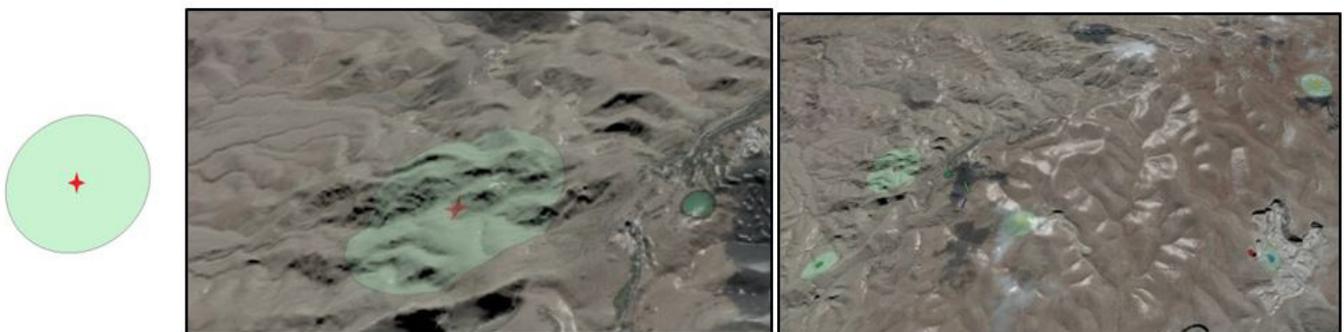


Figure 1. Abstracting complex data: the symbol for the seismic aftershocks. Left: The symbol designed to represent the seismic aftershocks is visually simple but encodes five relevant variables: (1) the multipointed star is placed at the event epicentre, (2) the number of points on the star is equivalent to the number of stations used to reconstruct that event and gives a measure of the confidence, (3) the color of the star represents the magnitude of the event, (4) the shape and orientation of the ellipse is the same as the calculated spatial uncertainty for the event location, and (5) the depth of the event is indicated by a vertical rod (terminated by a sphere) dropping vertically down from the star symbol to the event depth (not shown). Center: The symbol placed at the correct location in a 3D reconstruction of the region. Right: Symbols for multiple events placed on the landscape.



needed to be represented (using a different visual form than the seismic events).

Complex Data to Complex Visual Form: Terrestrial Laser Scan Data

Data used for the test case: text files in standard .las format

Several types of bunkers were used during LETTERPRESS, and terrestrial laser scans of each type were commissioned, which resulted in colored point cloud data for the interior and exterior of each bunker type. Each point cloud stores the spatial location of every point scanned and, when available, its color. This results in files containing data for hundreds of thousands or millions of measured points. These can be displayed as point clouds within GeoVisionary. Point clouds are problematic for visualization tools because each point is unique but, in general, all points cannot be rendered at all times. Visualization tools have to use various optimisation routines, custom data formats, and selective rendering routines (i.e., only displaying a subset of the points that are actually within the field of view of the user) to display them and not slow down the rendering speed.

The terrestrial laser scanner data was displayed in the environment by first converting the .las format file into the format used by GeoVisionary using the conversion tool provided with GeoVisionary. The data was not geolocated, so it was positioned

manually using the terrain and imagery already loaded into the environment as fiducials. Most of the points had color information associated with them, so they were displayed with those colors, which resulted in the point cloud having a photorealistic appearance. The scanner data and the terrain data overlapped in areas, so the terrain data had to be switched off to see and fully explore all of the scanner data. Figure 2 shows an example of a point cloud from the scanning of a bunker, compared to a real-world photograph of the same bunker.

Complex Data to Simple Visual Form: Magnetometry Data

Data used for the test case: a text-based data file from the IFE14 data store. We decided to display the variables listed in Table 4.

The magnetometry data collected during IFE14 is considered a complex data set for this work, but in contrast to the seismic data, it has an inherent geospatial nature. Measurements were taken as the sensor was moved over an area, and one measured variable was the local magnetic field gradient. The data points form an irregular trail over the inspected area of the ground. We did not have access to specialized software to read and display the data available from the IFE14. Such software could produce data products suitable for 3D visualization, including geolocating them, with fewer preprocessing steps than were required for this work. The chosen visual form for the data was a plot of the magnitude of the local field gradient, color-coded according to magnitude, for each measurement location, and it was wrapped over the terrain data. Some conventional magnetometry analysis software could geolocate such a plot and overlay it onto either 2D mapping data or 2D aerial imagery.

The results for some surveys during IFE14 are shown both in isolation and geolocated within GeoVisionary in both Figure 3 and in row three of Table 9. Even a nonspecialist can easily identify where the values diverge from background values and the physical location of those anomalies, which is exactly the type of information inspectors need to know.



Figure 2. Complex data to complex form: terrestrial laser scanning data. **Top left:** Photograph of the entrance of one of the bunkers used during LETTERPRESS. **Top right:** Laser scan data from the same bunker displayed in GeoVisionary and replicating the photograph. **Bottom:** View from inside that bunker.

Table 4. Magnetometry variables to be encoded into its visual form

Variable	Encoded form
Measurement location	A circular marker at the measurement location; its color is defined by the local field gradient
Local field gradient	Color of marker; a red-yellow-green color scale was used

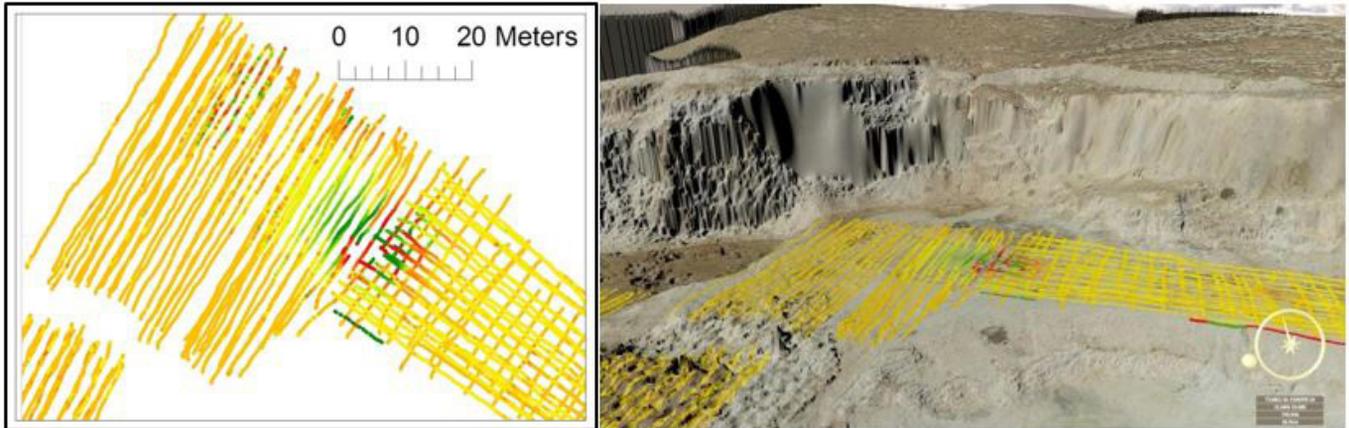


Figure 3. Magnetometry data. **Left:** The visual form chosen for this data in isolation. The tracks follow the path of the sensor, and each measurement point is colored according to the magnitude of the magnetic field point at that location. This is a standard way to display this type of data. It clearly shows large regions of uniform values with two regions of different values and two anomalies, but the actual location of the data and anomalies is not shown. **Right:** That data applied to the terrain at its correct geographical location, in a GeoVisionary environment. Very high-resolution terrain and imagery data give the magnetometry data its context such that the actual location of the visible anomaly is instantly identified.

Complex Data to Simple Visualization: Ground Resistivity Tomography Data

Data used for the test case: a spreadsheet from the IFE14 data store. The variables that we decided to encode for this data are listed in Table 5.

In this technique, ground resistivity measurements are made at intervals along a survey line, and then computer tomography techniques are used to reconstruct the resistivity of a cross-section of the ground below the survey line. This sort of data did not have a default option for display suitable for 3D visualization, so the full process was applied. Applying the process to this data set was more onerous than for the magnetometry data and required three iterations of the process to discover the optimal visual form to use. The data extended both along and below the ground, and it could have a width associated with it, so there were several options available to encode the variables.

The first option was to create a suitably sized virtual poster board for each survey line and wrap the image of the standard data product produced by the analysis software onto it; the data product is a plot of the data values. This was the simplest option, but it caused a problem because the whole of the plane was opaque, so the image from one survey line would totally obscure any other data behind it. Also, a significant amount of nondata white space was also displayed using this method. Thus, this method was rejected.

The second option was to display the data as a point cloud, with each point colored according to the magnitude of the variable to be displayed. This removed the problem of obscuring other data sets but lost the visual representation of the volumetric nature of the data. Thus, this method was rejected.

The third option was to render the data as voxels. Each data point was represented by a voxel that occupied the approximate volume related to the data point and was colored according to the magnitude of the data point (see Figure 4). The voxels were rectangular prism geometric solids. An advantage of using voxels was that those below a particular value could be made to disappear, so that all voxels containing background values were not displayed. Only relevant data were displayed, and no data were obscured by the display of background values. This method was accepted.

Table 5. Ground resistivity tomography variables to be encoded into its visual form

Variable	Encoded form
Measurement location	A marker at the measurement location; its color is defined by the resistivity at that point
Resistivity	Color of marker; we used a red-yellow-green color scale

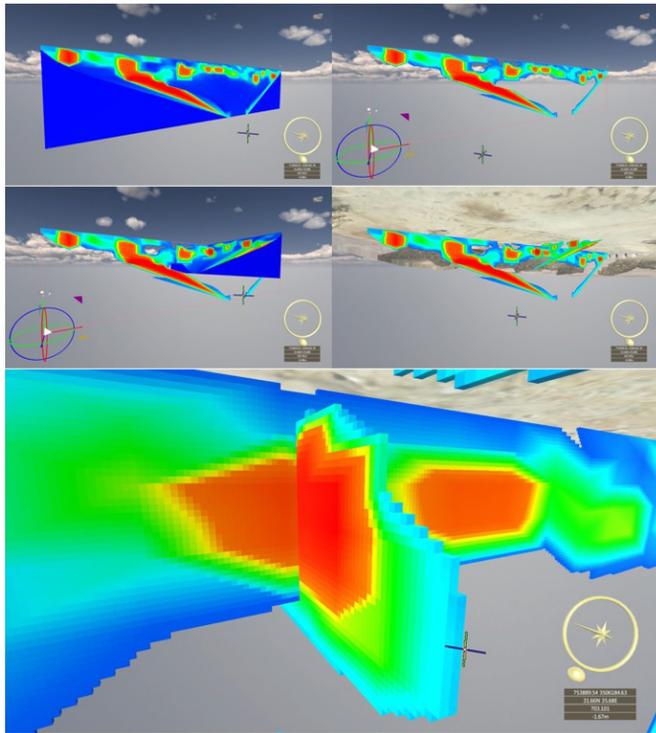


Figure 4. Complex to simple test case: ground resistivity tomography data. **Top left:** Data from one survey line (92 m long) displayed as voxels with a color scheme linked to resistivity magnitude in its correct geographical location. **Top right:** The same data, but with the background values made to disappear so they do not obscure other data in the environment. **Middle left:** Data for a second survey (44 m long) added at its correct geographical location, with background data visible. **Middle right:** The same data, but with all background data values made to disappear. **Bottom:** A close-up of the intersection between two survey lines showing the anomaly (red voxels) present in both data sets. Note that with the chosen display form and the background values made invisible, there is no ambiguity about which features are common to both survey lines.

Complex Data to Simple Form: 3D CAD Models

Data used for the test case: a model of the casing of a decommissioned nuclear weapon, supplied in the STEP (.stp) format. There are no variables to encode because the appearance of the model in the 3D environment should be the same as it was when viewed in the CAD program used to create it.

Models are complex data types because they contain many components and can include one or more image files that are used to color the model when it is displayed. We considered the displayed model to be a simple visual form because it replicates the visual form from the CAD program. We used a CAD model of the empty casing of a decommissioned nuclear weapon; such casings were used as props used during LETTERPRESS.

GeoVisionary has readers import some 3D model file formats natively, but other formats may require extra licenses. The file supplied to the authors was not in a compatible format (STEP format) and had been converted using a separate program to a compatible format (COLLADA). The model was then imported into the environment, and multiple copies of it were positioned to replicate the placement of the actual casings, as shown in Figure 5.

Simple to Simple: Handheld Photography 1

Data used for the test case: photographs from IFE14 data store in JPEG (.Jpg) format and a separate shapefile containing the geo-location data for the photographs. The variables to encode for some of the handheld camera photographs are listed in Table 6.

During IFE14, some photographs were taken using handheld cameras of a general scene or object. The chosen visual form for such photographs was as an image wrapped onto a virtual poster



Figure 5. Complex to simple test case: 3D CAD model. **Left:** Model of weapon casing (approximately 2 m long) used as a prop during LETTERPRESS imported into an environment. **Right:** The same image, but with laser scan data of the bunker where the casing was placed during LETTERPRESS turned on to provide context.

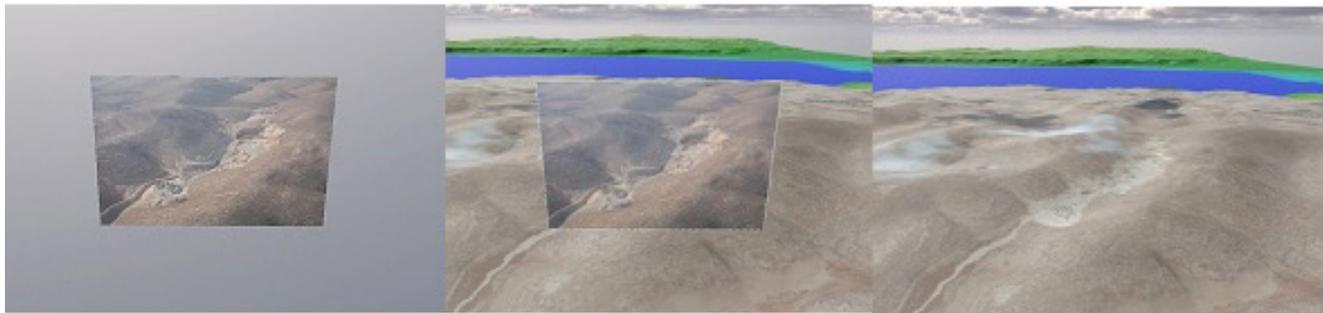


Figure 6. Simple to simple data display: poster boards for photographs. **Left:** A photograph taken during an overflight by an inspector has been displayed on a virtual poster board (43 by 30 m) and positioned and orientated at the approximate location where it was taken. The topography and appearance of the local landscape and the quarry (approximately 700 m long) can be seen. No other data is shown in the environment. **Center:** The terrain data and imagery from satellites (90 m and 0.5 m resolution, respectively) has now been turned on to show the context of the photograph. These data were available to the inspectors prior to the start of IFE14, but the visualization tool was not. **Right:** The poster board has been switched off to reveal the view of the quarry according to the satellite data. Note how the presence of the photograph in its correct geospatial context informs the viewer of detail that is missing from the view with only the satellite data, even though the photograph is flat, but the satellite data is rendered in 3D.

board created within the environment. An example is shown in Figure 6 for a photograph taken by an inspector using a handheld camera during an overflight during IFE14. The photograph is of an abandoned quarry used during IFE14. The poster board was placed at the approximate location where the photograph was taken. Comparing the actual scene captured in the photograph with the virtual one using satellite data and aerial imagery (see Figure 6) shows that they match well.

Simple to Simple: Handheld Photography 2

Data used for the test case: photographs from IFE14 data store in JPEG (.Jpg) format. The variables to encode for some of the handheld camera photographs are listed in Table 7.

During IFE14, some photographs were taken from a height, looking down onto the ground or an object. For small numbers of such photographs, the conventional option was taken; that is, they were visualized as scaled and geolocated images directly wrapped over the terrain. To do so, the manual geolocation

function in ArcGIS was applied using the satellite imagery as the geolocation reference (see Figure 7).

Simple to Complex: Aerial Photography

Data used for the test case: the aerial photographs from the IFE14 data store in JPEG (.jpg) format. The variables to encode are given in Table 8.

Aerial photography using survey cameras was used during IFE14 and produced individual images numbering in the thousands. It was neither sensible nor practical to process them individually for visualization. This test case required us to decide the best method for displaying the information contained in more than 2,000 photographs. Assessing the data set led to the conclusion that converting it into two more complex data types, orthoimages and terrain models, was the optimal way to display the information the photographs contained. When displayed together in 3D, these data sets would provide the inspectors with a high-fidelity reproduction of the landscape contained in the

Table 6. Variables from handheld photography to be encoded into its poster board visual form

Variable	Encoded form
Location where the photograph was taken	The location of the poster board
Direction of photograph	Match the orientation of the poster board to the approximate orientation of the camera
The image itself (the rectangular raster array of the color values for each of the three color channels in the image: red, green, and blue)	The color texture map for the poster board

Table 7. Variables from handheld photography to be encoded into its geolocated form

Variable	Encoded form
Location of the scene or object of interest captured by the photograph	Photograph geolocated onto the scene it captures or the object of interest
The image itself (the rectangular raster array of the color values for each of the three color channels in the image: red, green, and blue)	A geolocated image to be wrapped onto the terrain data



Figure 7. Example of wrapping a single aerial photograph over the ground captured in it.

Left: Location of IFE14 base of operations (BoO) before the exercise was rendered using satellite-derived terrain data (90 m resolution) and satellite imagery (50 cm resolution). **Center:** Photograph of the BoO taken with a handheld camera during an overflight manually scaled to 240 m by 213 m, then geolocated before being loaded into this GeoVisionary environment and set to render on top of the satellite imagery. **Right:** A close-up of the aerial image showing its detail. Wrapping this single photograph over the location of the scene it captures rapidly informs the viewer what the BoO's layout, relative scale, and wider geographical context is within the inspection area.

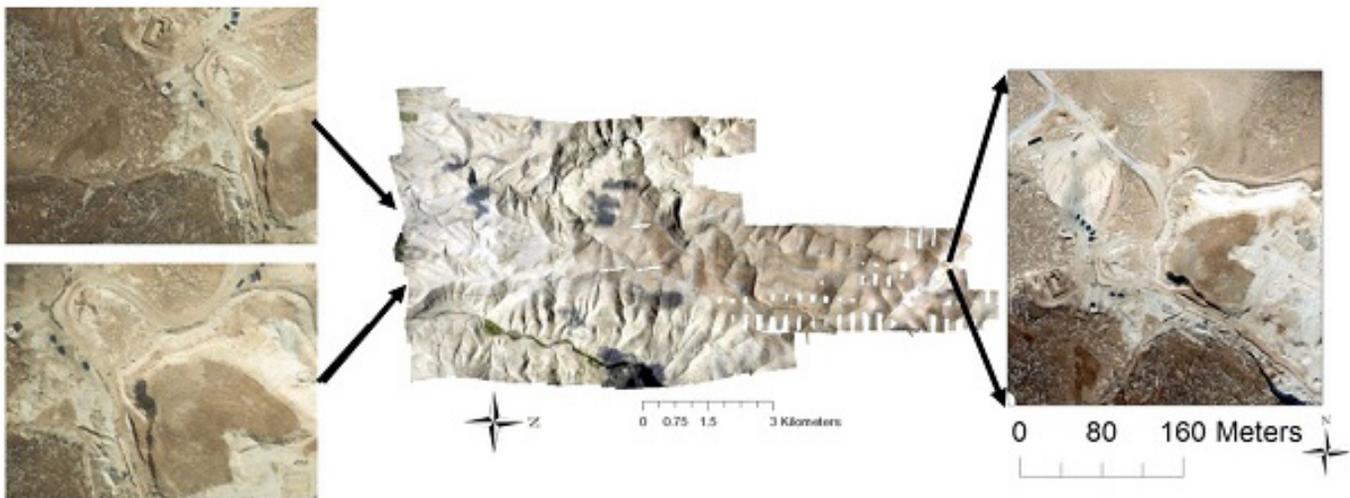


Figure 8. An orthoimage used to display some of the information contained within 800 aerial photographs captured by a survey camera during IFE14. **Left:** Two of the individual aerial photographs. Each contains only a small part of the inspection area (approximately 263 by 197 m) but with very fine detail (a few centimeters' resolution). **Center:** The orthoimage (approximately 20 km² in area) was produced using a structure from motion software and is an example of complex data due to its large file size. The orthoimage corrects some of the unwanted optical effects from the imaging process and retains the high resolution of the original images. **Right:** A close-up of a portion of the orthoimage that contains the information from the two photographs shown on the left of the figure. The area captured in those photographs can now be seen in its wider geographical context. The effect of improving the visible detail of the imagery wrapped over the terrain can be seen in Figure 9. It drastically improves the visible detail in the scene when the view point is close to ground level.

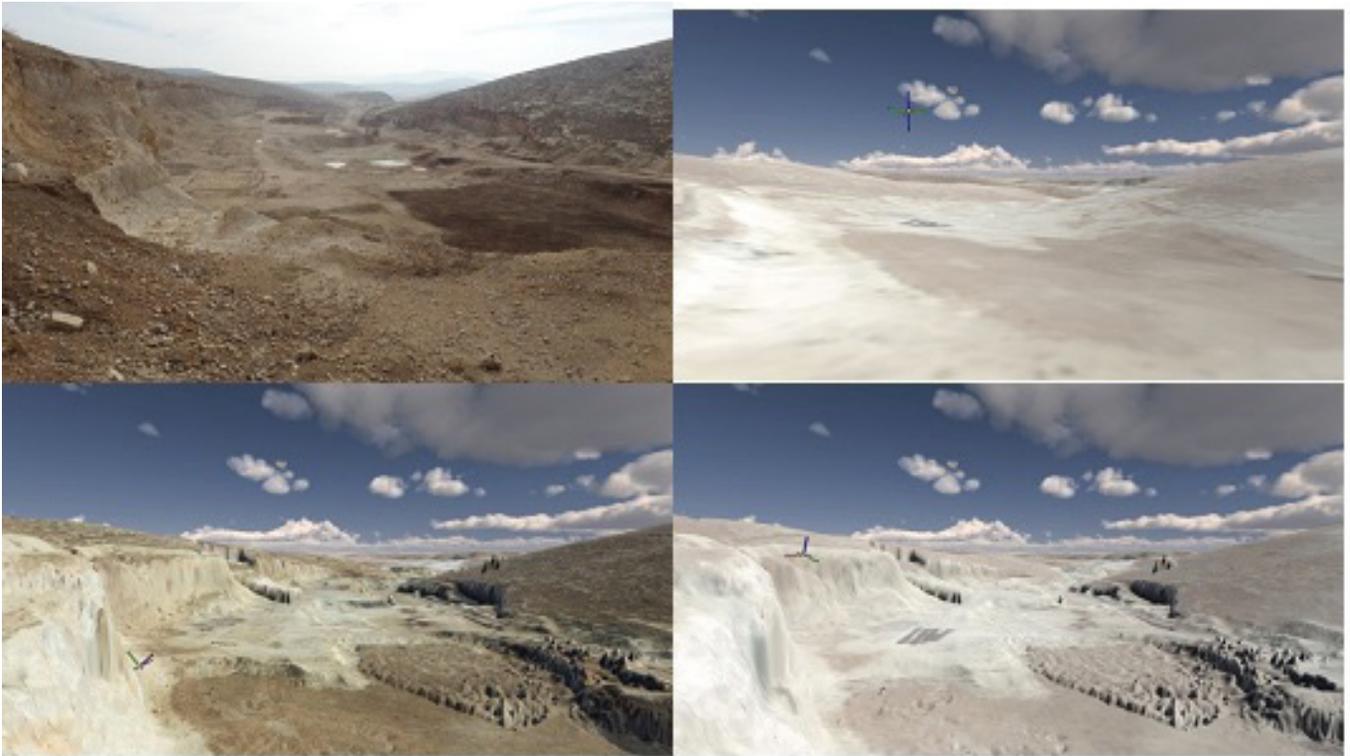


Figure 9. The effect of using the elevation data generated from the aerial imagery. **Top left:** A photograph taken during IFE14 from the head of a disused quarry used in the exercise, approximately 180 m at the widest point seen. **Top right:** The same view recreated in GeoVisionary, using the 90 m resolution terrain data and 50 cm imagery data available before the exercise. The shape of the quarry has been smoothed out by the low resolution of the terrain data. **Bottom right:** The high-resolution terrain data (approximately 7 cm) created from the aerial photography is applied, and the imagery is the same. The shape of the quarry can now be seen. **Bottom left:** The orthoimage (approximately 4 cm resolution) created from the aerial photography has replaced the satellite imagery. Significantly more detail can now be seen within the quarry, and this image now closely resembles the photograph, despite the presence of some artifacts in the terrain data. That scene could have been used to plan field missions in that quarry, if these data had been available during IFE14.

photographs. These complex data types were created using a structure from the motion software PhotoScan Pro. Although that process required significant effort and computation processing time, it is within the capabilities of the systems that were used by the inspectors during IFE14.

Table 8. Variables from aerial photography to be encoded into its visual form

Variable	Encoded form
Location of the contents captured by each photograph	Geolocate and scale the contents of every photograph
The actual contents of each photograph itself (the rectangular raster array of the color values for each of the three color channels in the photograph: red, green, and blue)	A geolocated image to be wrapped onto the terrain data
The elevation value to every pixel in every photograph	The elevation value for each pixel

Orthoimages: Orthoimages (and elevation data) for 40 km² of the IFE14 inspection area were produced. Each orthoimage was a mosaic of multiple aerial photographs merged together, which was also corrected for some of the distortion effects from optical imaging. The orthoimages had a ground resolution of approximately 4 or 7 cm compared to the satellite imagery's 50 cm resolution. An example is shown in Figure 8.

Terrain Data: The second data product produced from the aerial photography was a terrain model of the photographed ground. The photography was not originally captured with the purpose of creating terrain data, so this reconstruction contains some processing artifacts that would not be present in an optimized data set. To display the results, attempts were made to replicate the image in a photograph of a disused quarry used as part of IFE4. That photograph is shown in Figure 9, and it is compared to the derived aerial photography data and the satellite data showing



the same viewpoint. The aerial photography data is significantly closer in appearance to the photograph than the satellite data.

Facilitating Cross-Comparison of Different Data Sets: The Disused Quarry

The results for the IFE14 data discussed above produced a large data set of information for the location of the disused quarry. All of that data was added to a single 3D environment so it could be examined and cross-compared simultaneously. Some of screenshots of that work are shown in Figure 10. The figure includes the use of a dynamic section — a plane where everything on one side of the plane was not rendered. It could encompass the whole scene or one of the objects within the scene. Where data above and below ground needed to be seen simultaneously, a dynamic section might be the solution, as shown in Figure 10. Once the environment was assembled, we explored it using the same search logic the inspectors would use. Refer to the Discussions section for the details of this data exploration.

Discussion

Designing a process to visualize many different data types in a 3D environment was complex and required significant foresight and planning. Several of the visual forms we designed required only that the output from conventional analysis software to be reparsed and saved in a file structure and a file format compatible with the visualization software. Either no new analysis was required, or it was simple analysis, such as normalizing the magnitude of a variable. This could be automated by modifying the analysis software to produce the new files, or a simple translation program could be written to provide that translation function. Conversely, there were some data sets that needed more thought and more manipulation before we devised a data product that made sense. The symbols used to denote the location and magnitude of a seismic event, derived from analysis of the SAMS data, are one such example of this.

Our experience using the symbol designed for SAMS was that it enabled us to assess a cluster of such objects easily, without having to extract all the information in each symbol, even if the symbols were displayed with other information. When we saw an interesting pattern, we could carefully study it and extract the relevant information. During an inspection, this could allow an inspector to rapidly assess whether a particular event was associated with the original triggering event.

The reproduction of the topography and visual appearance of the location to be investigated in 3D was the most important

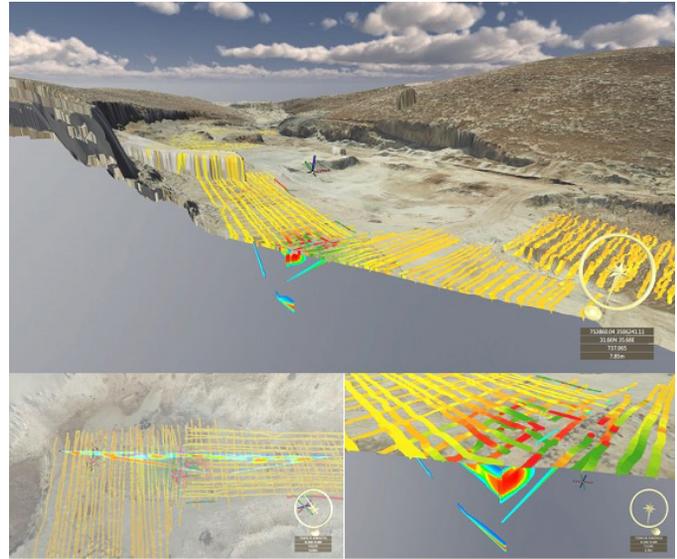


Figure 10. Assessing multiple data sets simultaneously. **Top:** View of the quarry, showing the terrain and orthoimage derived from aerial photography, magnetometry data, and ground resistivity tomography data. A dynamic section has been added so subsurface data can be seen with surface data. For scale, the distance between the walls indicated by the gray dashed line is approximately 180 m. **Bottom left:** A closer view of data from above, with the terrain made partially transparent so subsurface data can be seen and the dynamic section is turned off. **Bottom right:** Close-up of anomaly seen in both magnetometry and ground resistivity data. (Above ground is magnetometry: red/green area; underground is ground resistivity: red region.)

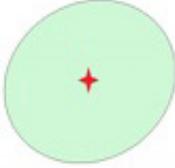
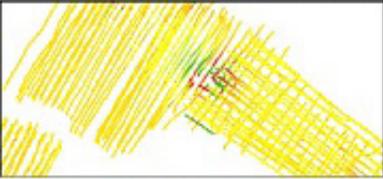
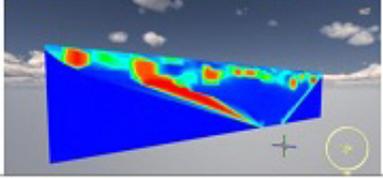
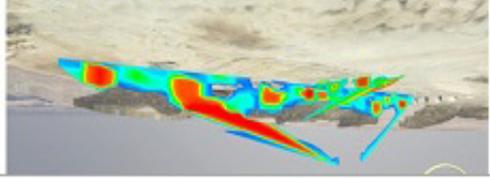
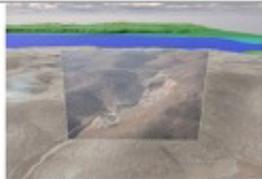
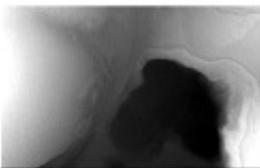
data set, as it gave context to all of the other data sets. Our opinion was that detail in the terrain imagery had a higher priority than detail in the terrain elevation data, as visual cues regarding the topography could be extracted from the imagery, but fewer cues could be extracted from just high-resolution terrain data.

An unexpected feature of exploring data in 3D was that the local ground level changed according to the terrain elevation data being used. Data such as ground resistivity is measured relative to the local ground level, not the height above sea level, but the visual form of the data is a discrete 3D object that is positioned using absolute coordinates. Within the 3D environment, the vertical offset between the high- and low-resolution elevation data at the site of the ground resistivity surveys was 9.5 m.

GeoVisionary allows the user to move freely within the 3D environment using five degrees of freedom: up/down, left/right, forward/backward, rotate left/right, and look up/down. When using a 3DConnexion SpaceMouse to move within the environment, the sensation to the user is similar to flying; therefore, we described movement within the environment as “flying” below.

Once the environment for the disused quarry contained

Table 9. A matrix of data types displayed with and without geographical context

		Visual form for data type.	Example shown in context.
SAMS data. (Seismic waves)	Complex to simple		
Terrestrial laser scanner	Complex to complex		
Magnetometry	Complex to simple		
Ground resistivity tomography	Complex to simple		
3D CAD model	Complex to simple		
Handheld photography	Simple to simple		
Handheld photography	Simple to simple		
Aerial photography	Simple to complex	Orthoimage 	
Aerial photography	Simple to complex	Elevation data 	



multiple data sets, we could explore that location and extract relevant information from it. First, we displayed only the data derived from the aerial photography: the high-resolution terrain data and an orthoimage. We flew around the environment to familiarize ourselves with it, which was enlightening because one author (Dr. Palmer) had visited the site in reality during IFE14, and the other (Dr. Stevanović) had not, but is currently a trainee surrogate inspector for the CTBTO. Our experience of examining the data together was similar to a real inspection, where one inspector may have visited a site and the other has not yet visited it but needs to become familiar with it. As multiple data types were added to an environment, there was a need to reassess how each was displayed due to the interaction between the data sets (e.g., one obscuring or clashing with another unnecessarily). Such issues were resolved by reapplying our process with a new constraint: for example, “minimize the extent that data X obscures data Y.”

After a short time (approximately 10 minutes), we were familiar with the topography of the site and so turned on the display of the other data. To begin, we turned on the magnetometry data, which was set to be visible above the terrain imagery. It was immediately clear that most of the region surveyed contained background readings for the color scale used. There were a few areas with readings outside of the background range (anomalies) that were clearly visible. There were also some areas showing high readings, which were instantly dismissed from immediate investigation because they clearly indicated instrument artifacts (e.g., single sections of high readings at the start of a survey line that were surrounded by background readings). We could immediately see the exact locations of anomalies and their place in their local and wider geographical context without the need for additional information.

We flew closer to the anomalies to examine them. We turned the data on and off to see if there were any surface indications for the nature of the anomalies; there were surface features over the anomalies that were different from their surroundings so they would be easy to locate in reality. We turned the data back on and changed its opacity so that it was partly transparent and so we could see the imagery and the data simultaneously over the whole area. From the context of an inspection, we had rapidly assessed a large quantity of magnetometry data and identified the areas in which the data required further investigation and eliminated those that did not.

We then turned on the data for the two ground resistivity survey lines. We could not immediately see the survey lines because they were underground, and our viewpoint was above

the ground. The terrain was therefore made partially transparent, and we could then see the data. We also flew below the ground level to view the data close up from below. To make the data clearer, the background readings were made to disappear; several anomalous features were immediately seen.

Both survey lines contained a feature in which the lines crossed, so it was the same feature detected by both surveys; the magnetometry data also detected anomalous features at the same location. Even with the terrain made partially transparent, there was still some difficulty in examining features, so a dynamic section was added and positioned so that it sectioned the whole environment through the anomalous feature (see Figure 10).

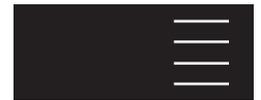
With the dynamic section in place it was easier to examine the anomalies and compare how it was detected by the different techniques; techniques that measured totally different physical properties. In an inspection context we were able to rapidly assess multiple data sets, cross compare them, identify anomalies requiring further investigation and produce reporting materials that could be used to justify that investigation to the State being inspected. This rapid assessment was made possible by our work in preparing the visual forms for the data and displaying it in a 3D environment.

The environment built for LETTERPRESS was informed by the experience of building the IFE14 environments and had the benefit of GeoVisionary being upgraded to use virtual reality systems for displaying and interacting with environments. The environment has been used for briefings given to attendees of the event and those who did not. Those who have explored the environment in virtual reality were able to gain a comprehensive understanding of the location. The environment was rendered at a scale of 1:1 and the VR headset gave depth perception, so the true size and appearance of the bunkers and weapon casing could be appreciated.

Conclusion

The use of 3D software packages to display data of any type is becoming more and more prevalent across different industries. For the display and analysis of geospatial data sets, we have found that a default visual form to use to display a data type in 3D did not generally exist. We have presented herein a process to find an optimal visual form for any data set that has a geospatial nature which is to be displayed in a 3D environment.

The process we suggest is a forked decision tree whose end points give a classification for the initial data and a choice of visual forms to use to display it in 3D.



We present two case studies of real-world events: the CTBTO's IFE14, held in Jordan in 2014, and LETTERPRESS, a quad exercise by Norway, Sweden, the United States, and the UK for nuclear arms control verification research, held in the UK in 2017. Both occurred within a defined geographical area and generated a variety of technical data sets that were either geospatial in nature or had a geospatial aspect.

We used a specialist geospatial/geotechnical 3D visualization tool, GeoVisionary, for this work. It was deemed to have benefits over similar available tools:

- It was able to load very large terrain and imagery files, yet retain and display their original resolution
- It was compatible with standard GIS data types, including shapefiles
- Environments could be viewed either using a standard computer monitor or using an immersive virtual reality system
- Depth perception of the features within an environment was possible when using a suitable monitor or projector
- It had an in-built capability to capture screenshots and create videos from within an environment. The screenshot capability was used to produce all of the images for this report, and the video capability was used to create the presentations in references 9 and 10.
- It did not need to be connected to the Internet

We found that the data types produced by conventional analysis tools did not generally produce output files suitable for loading straight into our 3D visualization tool. The data generally needed to be processed into a suitable form before it could be displayed in 3D. Once a visual form for a particular data type has been specified then, with relatively little effort, the conventional analysis software used could be modified to produce a file compatible with the visualization tool, or a simple translation program could be written for the same purpose.

We believe that the work required to design optimal visual forms is justified due to the benefits we experienced when investigating the environments we built. We were able to rapidly become familiar with the topography and visible appearance of a given location we had not previously visited using a standard computer monitor and immersive virtual reality systems. We also were able to rapidly assess data scattered over a large area, identify spatial patterns in the data, and locate areas of interest; cross-compare data sets collected from the same location using different geophysics techniques; and locate features of common

interest between them. We could also capture visual records of the environments and our data analysis in the form of screenshots and produce informative briefing materials to convey our findings to other audiences.

Acknowledgments

Satellite terrain data was from the shuttle radar topography mission (Jarvis, A., Reuter, H. I., Nelson, A. & Guevara, E. 2008. Hole-filled seamless SRTM data V4. International Centre for Tropical Agriculture. Retrieved from <http://srtm.csi.cgiar.org>). The satellite imagery data was provided by the CTBTO and is copyright DigitalGlobe, 2014.

Keywords

On-site inspection, 3D data visualization, geospatial, virtual reality

Author Biographies

Dr. James B. A. Palmer is a senior scientist at the Atomic Weapons Establishment (AWE) and has been the visual imaging specialist in the Arms Control Verification Research Group for 5 years. He has a remit to understand imaging remote sensing technologies that can be used from the ground, air, and space and how they could be used for arms control purposes. As part of that role, he has worked extensively with the CTBTO in Vienna on their use of imaging technologies during an on-site inspection and has attended field tests, workshops, and conferences hosted by the CTBTO. He attended the CTBTO's Integrated Field Exercise 2014 in Jordan as one of the UK's resident observers and witnessed first-hand how such technologies were used to collect data and how the data were then used and displayed. That led to his work on 3D data visualization presented in this paper. He received a PhD in experimental Z-pinch physics from Imperial College London, UK, and a degree in physics and optical science from Reading University, UK.

Dr. Jennifer Stevanović is a scientist and forensic seismologist at the AWE. She specializes in seismic monitoring of unusual events and site investigations using near-surface geophysical techniques. She has worked with the CTBTO on their use of near-surface geophysics techniques during an on-site inspection and is being trained as a CTBT surrogate inspector. She received a PhD in planetary seismology from the University of Bristol, where she worked with NASA on the InSight lander mission to Mars. That work included modeling of potential seismic sources for the lander and analogue experiments on coupling mechanisms for



seismic signals from the ground into the seismometer. She has a degree in earth sciences from the University of Oxford.

References

1. Comprehensive Nuclear Test Ban Treaty CTBT/MSS/RES/1. (1996, November 19).
2. CTBTO. (2013). *The CTBT and on-site inspections*. Retrieved from www.ctbto.org/fileadmin/user_upload/pdf/IFE-14_Final_corrected_SinglePages_WEB.pdf
3. (2018). CTBTO. Retrieved from www.ctbto.org/specials/integrated-field-exercise-2014
4. (2015). CTBTO. *IFE14: The movie*. Retrieved from www.youtube.com/watch?v=XfCjslo_nUc
5. (2018). AWE plc. Why dismantle a perfectly good warhead? *New Discovery Magazine*, Spring/Summer Edition.
6. *GeoVisionary v3.0*. Virtualis Limited. Retrieved from www.geovisionary.com
7. Virtualis Limited. Retrieved from www.virtualis.com
8. Joswig, M. (2013). Seismic aftershock monitoring system (SAMS) of the CTBTO: Status and experience from field measurements. *Verhandlungen der Deutschen Physikalischen Gesellschaft*, Dresden 2013 issue, 1.
9. Palmer, J. & Stevanović, J. (2017). Integrating OSI data into an interactive 3D environment to aid inspection activities. *The CTBT: Science and Technology 2017 Conference*. Video available at www.youtube.com/watch?v=cWONFkqrZCA&list=PLMRJ0gX0j4Tno_FL7ebzBO5qnrFF2OMuP&index=6 at 20 minutes and 16 seconds from the start of the video.
10. Stevanović, J. & Palmer, J. (2017). Investigating options to visualize the technical data generated during an OSI. *The CTBT: Science and Technology 2017 Conference*. Video available at www.youtube.com/watch?v=EGpJEVYbFQA-&list=PLMRJ0gX0j4Tno_FL7ebzBO5qnrFF2OMuP&index=0 at 59 minutes and 13 seconds from the start of the video.



Geo-based Data Integration (GDI) in the IAEA Department of Safeguards

M. Barletta, A. Yeoh, A. Marques-Patricio, N. Sumarokov, S. Robb, M. Fowler, M. Carey
International Atomic Energy Agency, Vienna, Austria

Abstract

The International Atomic Energy Agency (IAEA) Department of Safeguards has deployed the Geo-based Data Integration (GDI) platform for information integration, analysis, and activity planning involving geospatially related information used for nuclear safeguards verification. GDI provides interactive, layered maps in a user-friendly collaborative environment for IAEA inspectors, analysts, and managers to access, utilize, and share geospatially attributable information regarding nuclear facilities, sites, and other locations and activities relevant to the implementation of states' safeguards agreements. GDI operates in the secure Integrated Safeguards Environment (ISE), and access to information in GDI is limited and controlled via the Safeguards Authorization Management system, which is implemented as part of the Modernization of Safeguards Information Technology (MOSAIC) project.

Because the IAEA is an on-site inspection agency that verifies nuclear materials and activities in physical locations in states in accordance with their safeguards agreements, nearly all safeguards-relevant information has essential geospatial attributes. This includes state-declared information, information collected by IAEA inspectors and instruments, and open source information, including commercial satellite imagery. The GDI platform allows integration of information from multiple sources to better leverage analysis, evaluation, and activity planning in IAEA headquarters, in support of the in-field activities of IAEA inspectors around the world. This paper describes the technical and organizational context for the development and deployment of GDI. It explains safeguards activities for which GDI is designed and presents an initial experience in the use of GDI. The paper also outlines how emerging information technologies can further advance the utility of GDI for nuclear safeguards verification.

Introduction

This article is the first published description of the development and deployment of the International Atomic Energy Agency (IAEA) Geo-based Data Integration (GDI) platform for nuclear safeguards

verification. It is presented in four parts. The first provides background information on the Modernization of Safeguards Information Technology (MOSAIC) project and two IAEA projects and associated applications that complement GDI: the Geospatial Exploitation System (GES) and the Safeguards Multimedia Information Analysis and Integration Project (SG-MM). The second section offers a detailed account of the developmental process of GDI, details its capabilities and features as currently deployed, and provides notional visual illustrations of selected GDI features. The third section outlines the work processes that GDI has been designed to support — design information verification and complementary access — and notes how GDI is also expected to contribute to other processes in the Department of Safeguards. This section also reviews initial experiences in the use of GDI, including advances and challenges identified to date. The final section of this report outlines plans for further development. It concludes by identifying emerging information technologies that may enhance GDI's contributions to fulfilling the IAEA's nuclear verification mission.

GDI Context

The deployment of GDI reflects ongoing efforts to modernize information technology (IT) in IAEA safeguards, and concurrent and complementary efforts to enhance capabilities for expert analysis of commercial satellite imagery and multimedia information in the Department of Safeguards.

Modernization of Safeguards Information Technology

The MOSAIC project was a 3-year, €41 million effort involving 150 IAEA staff and IT contractors, which concluded successfully in May 2018.¹ Based on continuous engagement of empowered users with teams of developers in an agile, iterative approach, MOSAIC delivered more than 20 bespoke, in-house developed solutions to improve the structure and integration of all safeguards information in a modernized IT system, as well as strengthened



protections for information security. These new integrated applications spanned information collection, processing, analysis, and management and the planning, conduct, and reporting of IAEA in-field verification activities. As described below, GDI was one of the applications developed in MOSAIC, and it provides a user-friendly environment for work with geospatially attributed information.

With the protection of confidential information being of the highest priority for the Department of Safeguards, GDI has been designed to tightly integrate with the Department's Authorization Management (AM) system, which is the department's single system for requesting, managing, approving, and auditing access rights to safeguards information assets. GDI's integration with AM guarantees that the information that is available to users through GDI is appropriately filtered based on users' defined enterprise roles and specified work needs. For this reason, by design GDI does not replicate data that is managed by other applications; all communication with other systems is performed via secure services that return only the information that the logged-in user is authorized to see.

Geospatial Exploitation System

The GES is a data repository and enterprise-wide, collaborative platform that uses geographical information system technology to exploit and disseminate commercial satellite imagery (CSI) and geospatial data within the Department of Safeguards.² In 2012, GES was deployed into the Integrated Safeguards Environment (ISE) to provide authorized users across the department with up-to-date access to CSI, geospatial data, and analytical products. The application uses commercial off-the-shelf (COTS) software integrated with customized expert geographic information system (GIS) tools to ensure secure and efficient storage, management, analysis, and dissemination of safeguards-relevant data to authorized users.

CSI has become an increasingly valuable information resource for the IAEA in identifying and remotely monitoring infrastructure changes and activities at nuclear fuel cycle (NFC) sites around the globe. In recent years, the IAEA has acquired thousands of satellite images from commercial vendors, derived hundreds of geospatially attributable site plans from this CSI, and generated a substantial number of detailed satellite imagery analysis reports — all of which are distributed and accessible through GES. Ensuring that authorized users have timely access to this valuable and unique resource is crucial for up-to-date integration and analysis with other safeguards-relevant geospatially

attributable information across the department. A recent significant upgrade of the GES application will further improve interoperability with GDI and other MOSAIC applications and will ensure that users have access to a diverse range of CSI sensors, including Synthetic Aperture Radar and other emerging CSI technologies.

Safeguards Multimedia Information Analysis and Integration Project

Based on a multiyear effort to specify needs and identify opportunities for optimization of the use of multimedia information in the Department of Safeguards,³ SG-MM was launched in 2017. The project aims to optimize the use of multimedia information (photographs, videotapes and audio soundtracks, site maps, floorplans, diagrams, schematics, process flowsheets, etc.) in the Department of Safeguards analytic activities in headquarters and in planning and assessment in support of nuclear verification activities in the field. The project is implemented by a multidisciplinary team leveraging the complementary expertise of open source, trade, and satellite imagery analysts and nuclear technology and safeguards experts. Based on extensive experience in state evaluation and through engagement with IAEA inspectors, the project team aims to specify the full range of multimedia information use cases; identify gaps in existing IT tools, work processes, and multimedia analytic expertise; and realize synergies with existing in-house tools and expertise. The project is currently focused on procurement, deployment, and training for a targeted set of COTS software for photo and video analysis and digital asset management and on augmenting existing systems and tools to optimize their use with multimedia information. The project includes ongoing IT surveys to assess whether emerging capabilities are commercially mature for effective adoption for continued improvement in multimedia capabilities.

GDI Development Process

The initial releases of GDI in 2017–2018 were developed as part of the Department of Safeguards' MOSAIC project. The implementation process for GDI followed the best practices that had proven effective in the realization of other MOSAIC projects: an adapted PRINCE2 (Projects in Controlled Environments) methodology for project management and a Scrum Agile framework for software delivery.

PRINCE2 helps program managers control project initiation, review progress and project closure, and monitor human resources, budget, risks, and dependencies among projects. The Agile development framework ensures that future users of an



IT system are closely involved at all stages of the development process to guarantee that development is truly driven by and effectively meets users' needs. In contrast to the traditional waterfall method, Agile development encourages incremental delivery of product features. Product owners, who represent the user community, work closely with business analysts and the technical team to define the functionality to be delivered within a predetermined time period. Each development iteration results in a usable product for immediate review by users, which can be quickly modified or enhanced with additional functionality if needed. This flexible approach, which encourages learning from past experiences and continuous improvement, enables the development team to promptly and appropriately tackle problems and address changing requirements, including as new capabilities are identified that could enable improved work processes.

Building a geospatial system in a highly secure IT environment is a major challenge, so the availability of developers and testers with the appropriate skills and experience is a key factor for the project's success. To be effective, technical staff should not only possess a deep knowledge of the underlying technology and available state-of-the-art tools but should also have a good understanding of the business processes and the ways that end users will interact with the future system. Moreover, development efforts need to be supported by infrastructure specialists, including server and database administrators and information security engineers. For GDI, this kind of joint effort was required for the deployment and configuration of various ArcGIS components and OpenStreetMap, pregeneration of user interface tiles, data migration and transformation, and setting up authentication and authorization mechanisms.

Another important factor is the involvement of users in populating the IT system with data, including the definition of the boundaries of objects in GDI, linking objects to various information sources, and performing quality control. This is a major undertaking, and the required efforts should not be underestimated, especially taking into account the volume and sensitive nature of the relevant data, the need to train pilot users in advance, and the need to assure that quality standards are met. In the case of GDI, it was recognized early on that there are significant efficiencies in setting up automated procedures for integrating information from existing systems and data sets, such as GES and lists of locations maintained by the Department of Safeguards Operations Divisions (for the geospatial coordinates of facilities, sites, and other places of interest).

GDI Current Capabilities

From the outset, the goal of the GDI project was to fill the geospatial capability gap for the widest possible range of authorized users in the department's secure ISE. GDI is a lightweight supplement to highly specialized satellite imagery and geospatial analysis tools in GES. In its simplest application, GDI offers functionality similar to that of Google Maps, Bing Maps, Apple Maps, and other web mapping services, but for use only in the ISE with no access to the Internet. This includes the ability to use a variety of base maps (street map, satellite imagery, topographic map, etc.), search by address and zoom to locations, measure distances and areas, create temporary drawings and annotations, and export maps into various file formats for internal reports and briefings. Since the users interact with the maps in a secure environment, sensitive information — such as state-declared information and the location of the next inspection — remains protected.

GDI also provides authorized users with tools that allow them to record the coordinates and draw the shapes of various entities, such as state-declared sites and facilities, locations of interest, material balance areas (MBAs), nuclear material key measurement points, and individual buildings and rooms. As its name suggests, integration with existing and future applications in Department of Safeguards is one of GDI's most important functions. GDI allows the association of objects on the map with various entities maintained in other applications. For example, sites, facilities, and MBAs depicted in GDI are associated with the corresponding records in the department's information system that maintains the current and historical records for these entities; buildings and sites can be linked to the relevant entries in the Additional Protocol (AP) declarations. Thus, GDI provides a map-based geospatial entry point to information stored in ISE, allowing better visualization, analysis, and other use of that information. This integration is planned to be bidirectional: for instance, clicking on a building on the map in GDI can display an extract from the AP declaration describing that building, whereas clicking on a hyperlink in the IAEA's AP system (which makes the content of AP declarations available to authorized users and supports the process of reviewing these state declarations) could open the GDI-provided map and indicate the building's location.

Through the use of placemarks and by allowing the association of objects with documents, GDI enables users to geo-tag various documents stored in the department's document repositories in ISE. Photographs recorded during in-field verification activities can be linked to a visualization of the exact locations

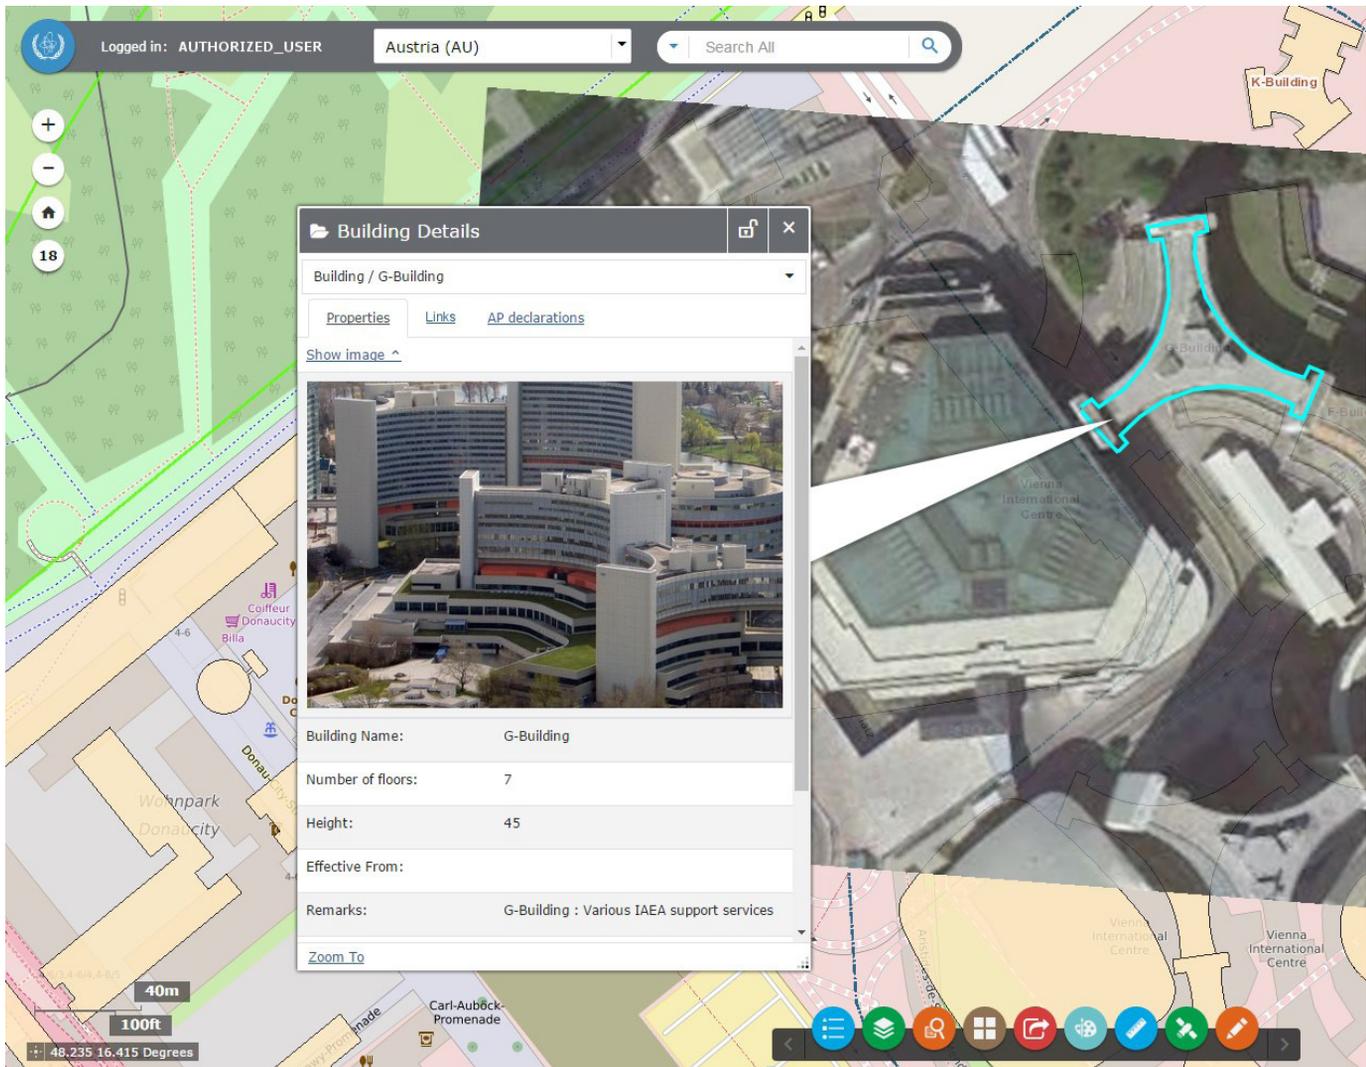


Figure 1. The GDI platform provides authorized users with access to interactive, layered maps and geospatially attributed information in the secure Integrated Safeguards Environment. (Copyright 2018, Google Earth/DigitalGlobe)

where the photographs were taken; environmental sampling (ES) reports can be linked to layouts showing where the samples were collected; open source documents can be associated with the locations that they refer to. GDI also enables links to multimedia documents, such as videos and 3D panoramas, and will in the future link to inspectors' measurement data sets integrated through the mobile Instrument Records Integrator for Safeguards (IRIS).⁴ The placemarks can be grouped by category, allowing the users to display only those categories of interest relevant to their specific tasks. This functionality helps improve visualization of available safeguards information and contribute to better knowledge management.

GDI Current Look and Functions

Information in GDI is presently represented by layers showing polygons, icons, and placemarks positioned over up-to-date maps overlaying satellite imagery. A layer may present features such as polygons that represent site and facility boundaries and buildings on sites or show placemarks to indicate the location of activities such as indicating a building in which verification activity has been conducted. For all layer types, it is the attribution of metadata of the individual layer features that is used to define the feature as well as links to relevant information about it. The contextualized visualization of this information in GDI may enable new or more accurate understanding of information about the site

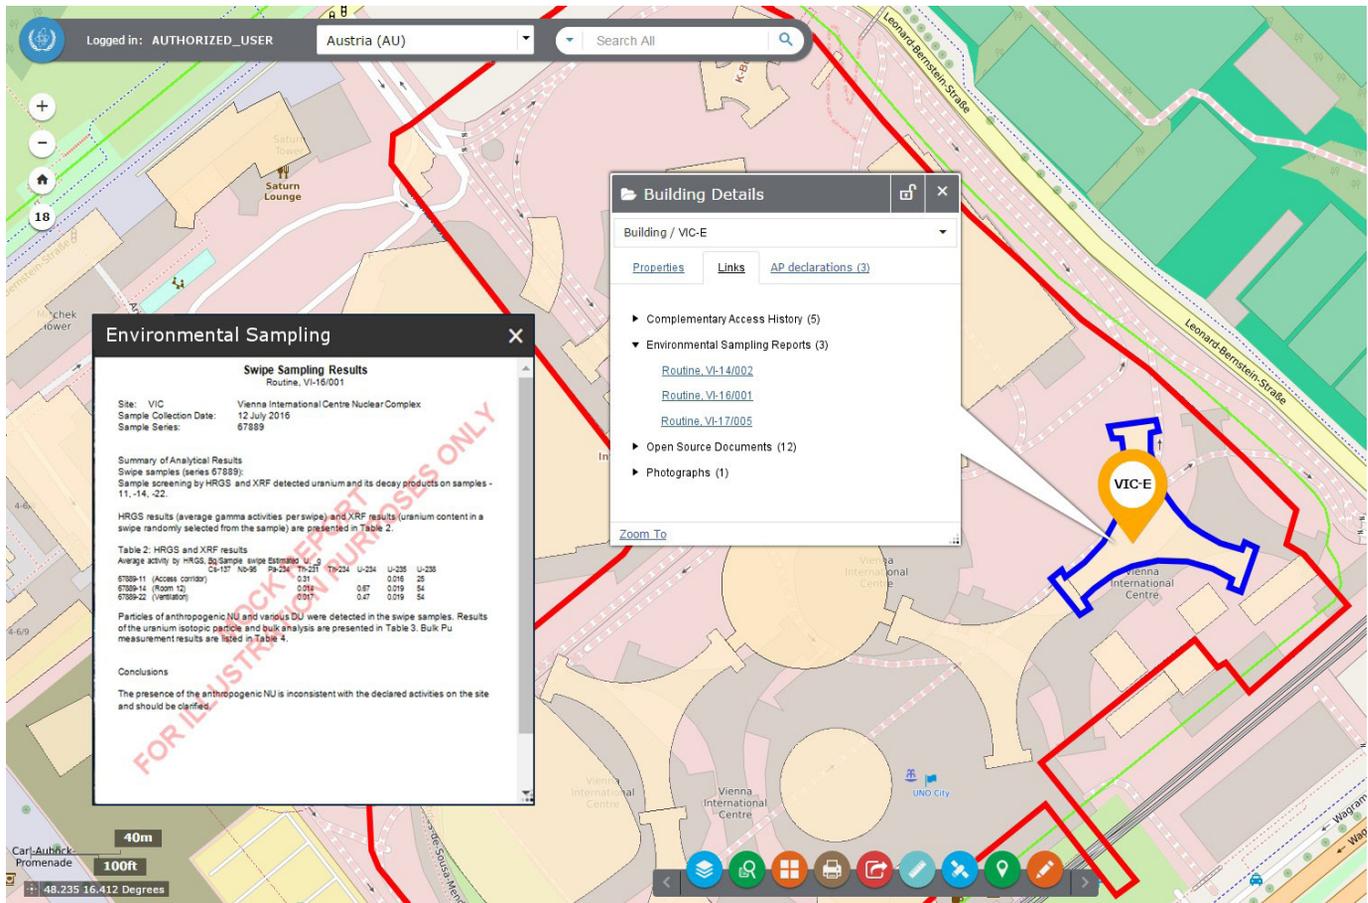


Figure 2. GDI can show site boundaries and building outlines and provide access to information from Additional Protocol declarations, environmental sampling analysis, and complementary access previously conducted on site.

and/or state, which may also lead to new lines of investigation.

However, it is important to note that GDI alone is no panacea; it is only one of a set of functionally complementary applications and data sources within ISE that work in coordination, in part through sharing common metadata attribution. So, while GDI can visually represent the buildings in which ES was taken, it will not store information specific to the taking of ES or the results of ES analysis. It is the unique attributes of a placemark depicting the location of the activity on a map that is the “glue” that allows analysts to deeply examine the specifics of an activity or data source. When a user clicks on a placemark on the map, the unique attributes of that placemark will be used to interrogate data sources and applications within ISE, which in turn will determine how the resulting layer or pop-up window is populated. As for the case of ES, one would see links to data sources related directly to the

activity, such as relevant planning/supporting documents, inspection reports, ES analysis, and, perhaps, again through attribution, the ability to open another software tool specifically designed to represent the linked data, such as a visual ES system.⁵ A visual ES system may show the exact location within a room or part of a building where ES was taken and may provide a greater level of detail (quantitative values, data graphs, and qualitative assessments).

Because GDI is accessible only to authorized users in ISE and integrates information about real nuclear facilities — sites and other locations that must remain confidential — Figure 1 and Figure 2 depict features of the IAEA headquarters in Vienna to illustrate selected features of GDI. These static illustrations, however, cannot convey the fluidity with which users are now able to access and use safeguards information through GDI.



GDI Applications for Safeguards

GDI is designed to improve the efficiency and effectiveness of design information verification (DIV) and complementary access (CA) activities in relation to states' Comprehensive Safeguards Agreements (CSAs) and APs, respectively. This section summarizes how GDI can contribute at a working level to these and related safeguards activities.

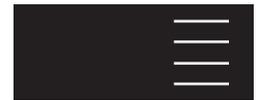
For each nuclear facility declared in relation to a CSA, a safeguards inspector is assigned as the Facility Officer (FO). The FO's responsibilities include conducting the design information examination (DIE) — that is, to review and assess the state's declarations for each facility submitted according to Design Information Questionnaires (DIQs); preparing teams of inspectors who will perform DIV and inspections related to nuclear material verification for the facility; and generally maintaining knowledge of all safeguards-relevant information relevant to the facility. GDI provides a working platform to assist in all of these tasks. By providing access to all types of safeguards information that is geospatially relevant to the nuclear facility, GDI aids the FO in working with other inspectors and analysts to conduct DIE for the initial declaration, and for any subsequent DIQ updates to reflect changes in a facility's capacity or operation. FOs can engage and coordinate any additional needed technical expertise — such as satellite imagery analysis, open source information collection and analysis, or ES planning — by using the common geospatial environment provided by GDI. By presenting information geospatially, FOs can also seek guidance from Senior Inspectors and other management with regard to safeguards implementation in the facility. Any question that may arise through DIE (or any other safeguards activity) regarding a facility can be presented to the state for clarification and followed up on by inspectors in the next DIV.

In preparing teams of inspectors in IAEA headquarters in Vienna, the FO provides logistical and technical guidance to facilitate smooth and effective implementation of the activity in the field. GDI provides an environment for the FO to convey key information, such as geospatial orientation of features of a facility (e.g., process lines, storage and flows of nuclear materials, location and coverage of IAEA containment and surveillance measures such as seals and cameras). This geospatial orientation for verification activities in the field — especially where multimedia information, such as photographs, can be coupled with other visual information, like a floor plan from the DIQ — can aid the FO in rapidly and reliably conveying necessary information. By

reviewing the briefing materials prepared by the FO for the facility in GDI, inspectors can familiarize themselves with relevant information and then seek any needed clarification from the FO. For example, by providing inspectors with photographs and technical information regarding essential nuclear-related equipment used in a facility, they can more quickly and reliably conduct visual observation or other verification activities to confirm that all such essential equipment is present during DIV and being utilized in accordance with the state's declaration. To aid in facility orientation and equipment identification, photo captions or other relevant open source information that is not in English can be translated in advance so that inspectors can focus on technical observations, use of instruments, and other safeguards procedures. Effective planning through GDI can enable smooth implementation in the field, thus minimizing the time required for inspectors to perform their activities and the time required by state and regional safeguards authorities and facility operators to support DIV or other verification activity.

It should be emphasized that nuclear facilities are typically safeguarded for decades — usually longer than the career of any IAEA staff member. An inspector serving as FO may leave the service of the IAEA, or for development and other reasons be transferred to a different position inside the organization. Hence, every safeguarded facility will be the responsibility of a series of FOs, each of whom acts in that capacity for a limited period. GDI offers a platform by which new FOs will be able to quickly assimilate the cumulative knowledge about a facility that has been identified as important by their predecessors. They can contribute to this cumulative process of knowledge management by adding their insights as facilities are renovated or otherwise change in capacity, applications, or operations. The IAEA holds extensive and detailed information electronically for each declared facility, and the secure ISE environment and the integrated set of applications in MOSAIC have given inspectors and analysts much easier access to all of that information. The geospatial integration of that facility-relevant information through GDI will enable far more efficient long-term knowledge management over the full life cycle of nuclear facilities, from initial construction through operation and any subsequent modifications, and on to eventual decommissioning and dismantlement.

With regard to AP implementation, for each site declared by the state, an inspector is assigned as the Site Officer (SO), whose responsibilities include reviewing and assessing the state's initial AP declaration and updates for each building on the site, coordinating in-field activities with FOs for facilities on the site, and



coordinating CAs to be conducted at other buildings on the site. A site is a geospatially defined area that may include one or more nuclear facilities or locations outside facilities where nuclear material is customarily used; other installations or activities relevant to the AP, such as laboratories conducting nuclear-related research and development activities not involving nuclear material; and buildings with no functional relevance to the NFC. Some sites are large and complex, with a diverse range of facilities spanning several NFC stages, and may include many buildings that have no nuclear-related function. Like facilities, sites are typically operational for decades.

In providing maps and layers regarding sites that link to relevant safeguards information sets, GDI offers powerful advantages for SOs and their colleagues implementing verification of the state's AP declarations. The geospatial disposition of buildings and their physical connections (such as process piping, including waste flows) are important in assessing capacities, possible changes in operations, and the declared purposes of activities on a site. Basic features and constraints, such as enclosure of a set of buildings within security fencing, can be understood and assessed through site layouts displayed in GDI. Changes in site features — such as site expansion to develop new areas; construction, refurbishing, and demolition of buildings; and alterations in access roads or security gates — can be observed through visual analysis of CSI of the site. SOs can request new acquisition of CSI and analytical reports of NFC sites. This imagery, derived site plans, and the analytical reports can then be served through GES to GDI for use by inspectors, analysts, and managers.

In providing map-based visualizations and access to information in ISE, GDI enables SOs and their colleagues to conduct a much more efficient review of annual AP declarations regarding buildings on a site. The geospatial specificity of attributing state-declared and open source information to particular locations, complexes, buildings, and laboratories enables more rapid differentiation among the possibly numerous buildings on a site that have no nuclear-related function, so that inspectors and analysts can prioritize their limited time and attention on relevant buildings and activities. GDI's support for geospatial attribution and information access and visualization is also very important in enabling IAEA staff to filter through the potentially very large volume of open source information about nonrelevant buildings on a site — again, to concentrate attention on only the information that is relevant for safeguards implementation.

As in other in-field verification activities, conducting CA has inherent logistical dimensions that typically pose greater

challenges than conducting, for example, DIV in a declared facility in which agency inspectors have been physically present many times in the past. In conducting CA in a building or set of buildings on a site, safeguards managers require clear logistical planning that relates conducting the CA to the achievement of specific technical objectives. GDI facilitates this logistical planning and its presentation for approval to managers by enabling inspectors to map out the exact sequence of buildings, floors, and rooms to be visited by inspectors and the activities to be conducted at each step. GDI will also enable more effective knowledge management in providing map-based layers showing the physical sequence, date(s), and other information related to past CAs that have been conducted on a site. This aids inspectors in implementing any follow-up actions that may have been recommended after prior CAs and in outlining the scope and objectives of the new CA in light of all the agency has learned in the past about the buildings on the site.

In addition to supporting DIV and CA activities, the GDI environment also offers enhanced capabilities that can contribute to designing the safeguards approach for a new facility. As for DIV, the state's initial declaration can be geospatially integrated with other relevant information from CSI and with other open source information. These types of information can inform and complement the technical observations of inspectors who conduct the initial DIV before the new facility is operational. Visualizing hypothetical routes for nuclear material diversion or facility misuse is a necessary part of the process of designing safeguards measures that are both sufficient to cover possible diversion paths and also as efficient as possible in use of resources (such as the deployment of cameras to maintain visual observation of safeguards-relevant areas when inspectors are not present in the facility).

Because GDI is used for preparation of subsequent DIVs (or in CAs), the use of this common map-based geospatial visualization environment will facilitate tacit knowledge sharing among inspectors and their managers in briefings and debriefings related to in-field verification activities. GDI can also aid in facilitating accurate and efficient communication about logistical and other physical details of facilities and sites among staff who work and interact in English, although for many of whom that is not their native language. For these and other safeguards tasks and the associated processes that support them, GDI offers a working platform for accessing state-provided, inspector-collected, and open source multimedia information, such as photographs, video footage, site maps, floorplans, equipment schematics and process flowsheets related to NFC activities. These types of information are especially



powerful when displayed geospatially and in combination, such as through the attribution of photographs and videos of equipment and operations at each stage of a complex process plant. For example, the FO may select a key set of illustrative photos to show inspection teams where to conduct specific activities, a technology analyst may annotate the photographs to explain key process parameters and how to observe key indicators of the status of operations, and an environmental sample analyst may indicate with a visual marker where samples should be taken in the future. In this integrated collaborative analytic and preparatory work to support in-field activities, GDI will enable more effective exploitation of the IAEA's extensive historical collections of data-rich multimedia information that are relevant to safeguards implementation.

GDI Initial Experience

Because GDI is a new system recently deployed concurrently with other new applications in MOSAIC, the IAEA has limited experience to date in the use of GDI. As the next section describes, the system's features and functionalities also are under further development. This section summarizes advances and challenges found in the initial use of GDI.

The initial experience focused on integrating DIQ information with all other information about facilities for access via GDI. It is evident that the shared platform can better leverage the combined assets and expertise of the FO with satellite imagery and open source analysts, especially through the integration of visually oriented information. Inspectors can commission satellite imagery and analyses from imagery analysts that can be viewed as layers in GDI. Open source analysts can contribute ground-level external and internal photographs, videos, and other information that can offer more granular-level detail regarding external observables and how building features reflect processes inside a facility. Inspectors are then in better position to plan and brief colleagues to conduct DIV to directly acquire additional safeguards-relevant information, confirm whether possibly relevant open source information (e.g., photographs) accurately reflect the physical reality inside the facility, and otherwise verify the state's declarations for the facility. Based on the assessment of DIV findings, follow-up actions can be tasked to imagery and open source analysts and to the next team of inspectors who will conduct DIV. Experience with GDI indicates that it provides the IT basis for a tighter information collection-analysis-activity cycle, in which activities in headquarters inform field activities, the results of which in turn focus analytic work, and so on, with improved and

cumulative consolidation of knowledge about the facility.

In compiling concise, synthetic overviews of information in preparation for DIV and to inform and orient analysts and managers about key safeguards-relevant aspects of a facility, the volume of information held by the IAEA constitutes a significant challenge. The inspector or analyst may need to sift through information sets, which may include multiple DIQs, each with a corresponding DIE report and, if questions arose, communications with state authorities on clarifications, as well as inspection, satellite imagery, ES analysis, and other reports. The IAEA also has collected open source information on declared facilities since the late 1990s, and as annual reports and other historical technical documents have been digitalized and made publicly available, it has collected relevant information on the full lifespan of facilities. All of this information that is in text form is indexed and text-searchable, structured by state and fuel cycle step and in part by facility, and accessible to authorized staff in ISE. The IAEA's extensive and data-rich collections of multimedia information are not yet as readily accessible as its quantitative data sets and text-based information collections, but multimedia data frequently provides even more important contextual or other information that can aid in DIE and preparation for DIV. All state declarations and inspection records, and most relevant analytic reports, are directly accessible in ISE by facility. With the exception of CSI, however, much of the relevant information in the IAEA's open source collections is not yet structured geospatially (whether by geocoordinates in the file metadata or in structured databases). Review, geospatial attribution, selection, analysis, and organization of all of this information to provide a complete and up-to-date picture of the facility via GDI can require significant staff time.

We should emphasize that GDI — like other MOSAIC and IAEA Department of Safeguards applications — is an enabling technology. Information analysis requires expert analysts and inspectors with deep knowledge of the state's NFC in which a nuclear facility operates. Information evaluation often requires additional NFCs and other technical expertise and linguistic abilities, and, for quality assurance, also peer review, revisions, and approval. Given these and other requirements for effective information evaluation, a prioritized approach to integrating information for access via GDI is required. The obvious basis for this prioritization is the schedule for conducting DIVs and CAs, in accordance with the Annual Implementation Plan for each state with a safeguards agreement. Realizing the full integrative potential of GDI will require significant allocation of staff time. But, as noted above, IAEA safeguards are typically implemented for

decades along the course of a facility or site's lifecycle. Investing time once in organizing and synthesizing key information from the Agency's extensive historical collections for display via GDI will make all future preparations for DIV and CA more efficient, and will likewise aid any analyses related to the facility that may be conducted in the future.

Future Development of GDI

Realization of the full potential utility of GDI will require efforts in three areas: internal work processes; continued agile development of bespoke functionalities for GDI users and continued improvement in the integration of GDI with other IT systems in ISE; and incorporation of emerging IT that could complement GDI. This paper discusses the latter two areas.

For further internal IT development and to effectively support the goals for GDI noted above, GDI needs to be equipped with tools that allow simple and user-friendly interaction with the system without sacrificing its functionality. It is an important requirement that the system be usable, with little or no training, by inspectors, managers, state evaluation group members, and information analysts who are not trained in expert geospatial tools or systems. Some of the capabilities that the system must possess — and which are currently at different stages of analysis, planning, and development — are the following:

- Positioning GDI as a map service for other applications operating in ISE. This involves the development of services and application programming interfaces that would allow other Department of Safeguards applications to seamlessly integrate GDI maps into their user interfaces and utilize these GDI maps to visualize the information these applications are managing. As a pilot project, the GDI development team has helped the team of the Integrated Scheduling and Planning system to integrate a visual representation of inspector locations at any point in time using an interactive world map.
- The ability to upload images (state-provided or collected from open sources), superimpose them over a map, and control their size, rotation angle, visibility, and transparency. This feature is especially useful to simplify and accelerate the process of creating geometry, since the users will be able to draw vector objects by tracing the boundaries of the objects (such as specific rooms) shown on the images.
- A “smart” search function that would allow searching for different types of objects using a single search box.

For more complex scenarios, a more comprehensive advanced search feature can be developed.

- An “object explorer” feature that would show users the list of all the objects available in the current map's viewport. As users move the viewport and zoom in or out, the list will automatically update itself. Therefore, users will be able to select objects by clicking on their boundaries on the map as well as by selecting them from the list.
- A possibility to easily import and display geotagged information (pictures, footpaths, etc.) collected from different sources, such as various equipment used by inspectors during verification activities.
- The capability to add geospatial metadata to information items and collections held in folders in repositories in ISE through a “drag-and-drop” feature of information items and folders “into” buildings and other polygon objects displayed in GDI.
- A way to maintain the history of changes to various objects, including their shapes and attributes. The changes can then be visualized on the map using “time sliders”; the system should also help users compare maps of the same location at different points in time.

Looking beyond internal IAEA development efforts and toward the future, there are several areas of emerging information technologies that could complement GDI. These include the application of machine learning and artificial intelligence in automated processing and categorization of image-based data. For image data that is not accessible in structured databases or which otherwise lacks geospatial attributes or metadata, emerging IT capabilities could aid in initial geospatial attribution, for subsequent review and confirmation by analysts and inspectors. This could contribute to the collection of new or additional historical open source image data, and also in more complete exploitation of the Agency's internal holdings of multimedia information. Other potentially promising areas include IT developments in 3D modeling and in augmented and virtual reality. Although these technologies may further enhance facility and site visualizations in GDI, commercially available applications are currently too expensive and skilled-labour-intensive to justify their use for normal applications in the Department of Safeguards. As costs drop and user-friendliness increases, however, these information technologies may offer capabilities applicable across a wide range of safeguards activities.⁶ The IAEA is exploring these and other areas of IT development and will continue through internal efforts and external assistance from Member States to modernize the



information technologies that it employs for nuclear safeguards verification.

Conclusion

In sum, the deployment of GDI offers a significant advance for the IAEA. The implementation of nuclear safeguards is an inherently multidisciplinary enterprise in which the geospatial context of safeguards-relevant information and nuclear verification activities are essential dimensions. GDI offers user-friendly geospatial visualizations to engage the wide range of expertise required to analyze diverse information streams and prepare for and assess information collected through in-field verification activities. This shared and secure geospatial working environment encourages a culture of collaboration and enables better integration of work in IAEA headquarters in Vienna with in-field verification activities worldwide. GDI will also support long-term, contextualized knowledge management over the lifetimes of nuclear facilities, sites, and locations. GDI will thus enable present-day IAEA inspectors and analysts to consolidate corporate knowledge to prepare their successors to sustain IAEA implementation of states' respective safeguards agreements in the future.

Keywords

Geospatial data integration, knowledge management, multimedia information, agile application development, nuclear safeguards

Author Biographies

M. Barletta is a staff member in the Division of Operations B, Department of Safeguards, at the International Atomic Energy Agency

A. Yeoh is a staff member in the Division of Operations B, Department of Safeguards, at the International Atomic Energy Agency

Marques-Patricio is a staff member in the Division of Operations B, Department of Safeguards, at the International Atomic Energy Agency

N. Sumarokov is a staff member in the Office of Information and Communications Systems, Department of Safeguards, at the International Atomic Energy Agency.

S. Robb is a staff member in the Office for Verification in Iran in the Department of Safeguards, at the International Atomic Energy Agency.

M. Fowler and M. Carey are staff members in the Division of Information Management, Department of Safeguards, at the International Atomic Energy Agency.

References

1. See Abazi, F., Coyne, J., Partee, S., & Mutluer, A. (2017). MO-SAIC — a major upgrade for safeguards implementation. *Proceedings of the 58th Annual Meeting of the Institute of Nuclear Materials Management*. Indian Wells, CA: IAEA. 2017. See also MOSAIC — the modernization of safeguards information technology: Completing the picture. (2017, January 24). Retrieved from www.iaea.org/sites/default/files/17/01/mosaic.pdf
2. Keskinen, A., Baute, J., Carey, M., Ng, J., & Ujkani, J. (2018). Enhancing the geospatial exploitation system within the IAEA Department of Safeguards. *Journal of Nuclear Materials Management*, 46(3).
3. Barletta, M., Fowler, M., & Khaled, J. (2016). Integrating multimedia information in IAEA Safeguards. *Advances in Nuclear Nonproliferation Technology and Policy Conference (ANT-PC)*. Santa Fe, NM: American Nuclear Society.
4. IAEA. (2016). New indoor positioning technology to improve efficiency and accuracy of safeguards inspections. Retrieved from www.iaea.org/newscenter/news/new-indoor-positioning-technology-to-improve-efficiency-and-accuracy-of-safeguards-inspections
5. Vilece, K., Norman, C., Baute, J., Giaveri, G., Kiryu, M., & Pellechi, M. (2012). Visualization of environmental sampling results at inspected facilities. *53rd Annual Meeting of the Institute of Nuclear Materials Management* (pp. 1572–1576). Orlando: INMM.
6. Barletta, M., Crété, J. M., & Pickett, S. (2011). Integrating virtual reality applications in nuclear safeguards. *Proceedings of GLOBAL 2011*. Makuhari: Atomic Energy Society of Japan.



The IAEA's Physical Model: Fine-Tuning Nuclear Fuel Cycle Understanding for Robust State-Level Safeguards

Brian D. Boyer, James Disser, Sebastien Richet, Daniel Gagne, Stephanie Poirier, Jeremy Whitlock, Agatha Walczak-Typke, and Claude Norman
International Atomic Energy Agency, Vienna, Austria

Abstract

The International Atomic Energy Agency (IAEA) has been making use of increased amounts and types of data and reports on states' nuclear and nuclear-related activities. This information includes declarations provided by states (including those made in accordance with safeguards agreements and additional protocols), data collected by the IAEA (including inspectors' findings, environmental sampling, and measurement data), and other safeguards-relevant information available to the IAEA (from open sources, including geospatial data). This information is integral for the IAEA to fulfil the safeguards objectives common to all states with comprehensive safeguards agreements in force, which are to detect (1) any diversion of declared nuclear material, (2) any undeclared production or processing of nuclear material in declared facilities and locations outside of facilities, and (3) any undeclared nuclear material or activities in the state as a whole. The Physical Model describes the main technical activities involved in a nuclear fuel cycle, from source material acquisition to the production of weapon-usable material. The Physical Model identifies and describes indicators of each process, including specially designed and dual-use equipment, nuclear and non-nuclear materials, related technology/training/R&D, environmental signatures, and other observables of the process. These indicators have assigned strengths, depending on whether they have a weak, medium, or strong specificity to the particular process within a state's fuel cycle. The Physical Model is a living encyclopedia of this knowledge as well as a seamless searchable resource. This allows, among other capabilities, objective open source collection of state nuclear fuel cycle-related activities, capabilities, and nuclear material flows against the framework of the Physical Model, as done by state evaluation groups, using IAEA search and visualization tools. For example, Physical Model-based visualizations of declared nuclear material flows will help to determine typical material flows and scales of activities in a state. The IAEA makes use of many disparate sources of data used in a variety of safeguards processes. The Physical Model is the bridge between

these data and processes. Updating, improving, and enhancing the Physical Model benefits multiple stakeholders in the IAEA, including those analysts needing to search for and manage open source and geospatial information. As revised, the Physical Model will also be a stronger resource for visualizing material flows and potential acquisition paths.

Introduction

The International Atomic Energy Agency (IAEA) makes use of a wide variety of information in safeguards implementation. The IAEA collects information from states' declarations, field activities, and open sources such as scientific literature and news articles. The Department of Safeguards uses this information to develop state-level safeguards approaches; plan, conduct, and evaluate the results of safeguards activities; and draw credible safeguards conclusions that a state is abiding by its safeguards obligations.

The Physical Model provides a framework and an ontology that unites the processes and information sources involved in the safeguards implementation process. The Physical Model is a vital analytical tool, with the objective of describing every process and technology in the Nuclear Fuel Cycle (NFC) capable of producing weapon-usable nuclear material. This paper will outline the Physical Model's role in information analysis and collection; the process to update, improve, and enhance the Physical Model; and the use of the Physical Model in the safeguards implementation and evaluation processes, including visualization of data.

The Structure and Content of the Physical Model

The Physical Model is a reference for safeguards inspectors and analysts on fuel cycle activities. The Physical Model provides the building blocks to model a state's nuclear activities and describes detectable "indicators" of each relevant part of the fuel cycle, rated by their relative strength as evidence of that particular activity or process.



The Department of Safeguards and experts from the IAEA member states prepared the original eight volumes of the Physical Model.¹ Primary foreseen uses² of the Physical Model were to provide the following:

- A fundamental technical basis for improved information analysis
- A technical tool for evaluating states' nuclear activities
- Assistance for safeguards inspectors in preparing for routine and ad hoc inspections, design information verification (DIV), and complementary access (CA) under the Additional Protocol (AP)
- An evaluation tool for assessing the consistency of state declarations and activities

The Physical Model comprises volumes on the fundamental fuel cycle processes, starting with mining and ore processing. The Physical Model then progresses through conversion processes; enrichment;

fuel fabrication; reactors, including neutron sources; and reprocessing. The volumes also include heavy water production, spent fuel management and disposal, radioactive waste management, and hot cell activities. Figure 1 illustrates the layout of the latest formulation of the Physical Model. Each volume of the Physical Model describes a number of technologies or processes for that part of the NFC. The following 10 sections contain descriptions of each process or technology:

- General Process Description
- Especially Designed or Prepared Equipment
- Dual-Use Equipment
- Non-Nuclear Material
- Nuclear Material
- Technology – Training – Research and Development
- Other Observables
- By-Products – Effluents
- End Product
- Proliferation Aspects

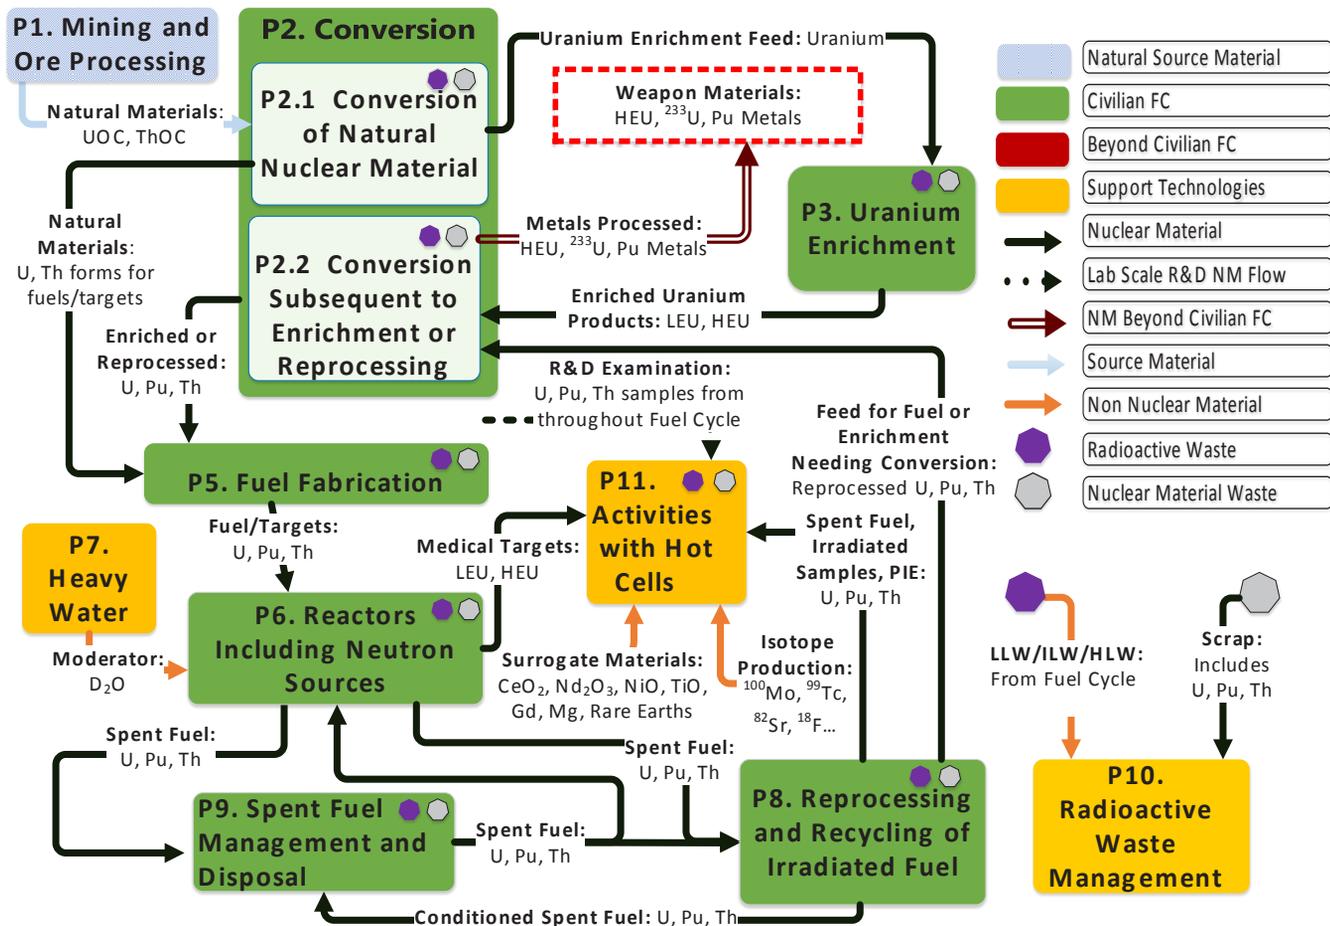


Figure 1. Conceptual framework of the Physical Model volumes



Each section has indicators of the technology or process described, with ratings of weak, medium, and strong to assist in judging the extent to which each indicator truly points to the existence of that technology or material. The inspectors and analysts working on a state evaluation must then consider those indicators as possible evidence of the likelihood of a technology or material existing in a state. They then follow up on possible inconsistencies with a state's declaration, if necessary.

The Physical Model's Role in Information Analysis and Collection

The IAEA makes use of safeguards-relevant open source data.³ The IAEA's definition of open source information is "information available to the public from sources external to the Agency."⁴ The IAEA conducts ongoing reviews of such information to assess the internal consistency of state-declared information and the consistency of state-declared information with that generated and collected by the IAEA.

The Physical Model provides a framework for assigning areas for searching for data for specific fuel cycle indicators in a state. The IAEA has looked for ways to use computer power to increase and enhance the utility of the data available to make analysis more effective by integrating a collection of reference documents, including regulatory documents, technical handbooks, and trade nomenclatures.^{5,6} The documents are stored in a database. They can be searchable by text in a structured way (i.e., on database fields) and by correspondence tables that relate items by their meaning.

The IAEA processes the declarations made under the various safeguards agreements and checks them for correctness and completeness. Comprehensive analysis of all available safeguards-relevant information includes the following steps:

- Matching documents with fuel cycle relevance to the appropriate process step(s) within the Physical Model. This categorization allows the Safeguards Department to perform analysis on the documents that are relevant to specific technical objectives that are part of the state evaluation process.
- Matching documents to specific states based on geo-spatial and entity information in the document.
- Being horizontally scalable to process large volumes of information efficiently.
- Feeding a knowledge repository suitable for analytical inquiry across the body of processed documents to link together data such as nuclear material flows and

balances in the analysis.

- Iterating over feedback from the Safeguards Department to improve results and integration into safeguard processes.

The Physical Model provides a unifying framework to categorize and sort NFC activity. The IAEA has analysis platforms that incorporate the Physical Model as predefined metadata. Users are then able to use the Physical Model to annotate documents, entities, and events to model NFC activity. The fusion of unstructured data with structured data from multiple heterogeneous sources increases the timeliness and comprehensiveness of analytical conclusions. The Physical Model has use across various information collection software packages as a main focal point to integrate disparate data sources and facilitate the access of information to analysts. The use of all-source integrated views of a state's NFC helps state evaluation groups provide credible assurances of the correctness and completeness of a state's declarations. These platforms can meld the Physical Model's NFC information, including the indicators of the processes and activities described, into a multidimensional assessment of a state's nuclear activities.

Updating, Improving, and Enhancing the Physical Model

In recent years, after analyzing the scope and progress of earlier revisions, the IAEA developed a comprehensive program for a thorough and timely revision of the Physical Model, with three prime goals:

- Updating to include new technologies
- Improving the scope of the information, including improved text and detailed illustrations
- Enhancing the usability and accessibility of the information

These three goals of the revision of the Physical Model have the following qualities to assist analysts in understanding the fuel cycle. *Updating* includes adding comprehensive descriptions of new fuel cycle technologies, as well as how plant and nuclear material operations changed over the past two decades. *Improving the scope* provides additional details about the technology, such as in diagrams, descriptions, photos, and references for the reader to dig deeper and timelines for understanding how long it would take to acquire and operate such technology to provide a state with weapons materials. *Enhancing usability* envisions



linking the document references to an IAEA library of references, providing videos of plant processes, incorporating Microsoft Excel worksheets for simple algorithms, and intimately linking the Department's analytical process software and analyst teams to the Physical Model.

These changes enhance the Physical Model's use as a key tool for analysis for the state-level safeguards process. Updating the Physical Model ensures that the text will include changes in fuel cycle technologies, equipment, operations, and nuclear material compositions. Furthermore, this update takes the opportunity to modify the title of some of the volumes to reflect more precisely the content and scope of the revised version. For example, in the past two decades, there have been developments in small modular reactors, Generation IV fuel cycle systems, laser isotope separation, and pyroprocessing that are not sufficiently included in the present volumes. Another aspect of updating is to ensure that the descriptions fit the needs of the state-level safeguards implementation as it has evolved over the years.

These improvements in the Physical Model ensure the addition of essential information not included in the first editions of the text.⁷ These early editions lacked photographs, fuel cycle flow diagrams, and sufficient depth and breadth of descriptions of technologies and equipment. To make an accurate analysis, an inspector would benefit from the ability to visualize and understand what technology is in their acquisition path analysis (APA), in AP declarations, and at a site selected for a visit during a CA.

Enhancing the Physical Model allows integration with databases of the IAEA Safeguards Department, creating interfaces that link it to projects addressing the following areas:

- Information management
- Information analysis
- Technological reference
- Explorative reference⁸

Information management categorizes the safeguards relevance of a document with respect to a certain technology in a state. Information analysis includes information collection strategies and APA. Technological reference provides knowledge about a certain fuel cycle technology.

The IAEA's systems for open source searches and information sorting⁹ will use the enhanced Physical Model to help direct and categorize searches for state activities in the fuel cycle under such categories as described above. The data from such searches and from nuclear material accountancy (NMA) will be grouped and presented in a visual format making it easy to

identify trends and activities in the NFC in a state. The locations of facilities and activities including nuclear material flows will be able to link into Physical Model descriptions.

Physical Model's Role in Safeguards Evaluation by Visualization of Data

The Physical Model is a keystone for state-level safeguards implementation.¹⁰ It provides the framework and structure for state evaluation, consistency analysis of safeguards-relevant information for a state, and open source information collection. The introduction of the state-level safeguards necessitated this view of the state as a whole. One example of how this has been implemented in practice is the integrated approach that the IAEA has taken to Material Balance Evaluation (MBE)¹¹ to assess declared data from verification activities¹² that depend on the scheme of the Physical Model for structure and order. The IAEA's safeguards statistical analysts are using the structure inherent to the Physical Model as a reference framework for fusing data coming from the various sources into a series of visualization solutions to understand the fuel cycle of a state.

NMA data provides information about quantities and types of material in facility inventories and the flows of materials leaving, entering, and under conversion in isotopic, chemical, or physical form in a facility, as seen in Figure 2. IAEA statistical analysts developed a visualization concept and implementation that displays the declared nuclear material flows and inventories organized along the lines of the Physical Model.¹³ This concept utilizes the Sankey diagram technique, a specific type of flow diagram in which the width of the arrows is proportional to the flow quantity. Sankey diagrams graphically represent the flows of materials and quantities in a system, showing succinctly the behavior of a system.¹⁴ IAEA statistical analysts used the Sankey diagram representation to create a fully interactive prototype tool called SNAKEY.¹⁵ The computerized visualization SNAKEY tool received its name from the inspiration found in Sankey diagrams, the sinuous shapes of the flow arrows, and the use of the Python language in the development of the prototype. SNAKEY can illustrate distinctly the fuel cycle material flows and processes in a state. The tool consists of two parts:

- A back end that queries databases containing NMA data and aggregates the data into JavaScript Object Notation (JSON) format
- A richly interactive Data-Driven Documents JavaScript library (D3.js) based front end, which displays the resulting image via a browser

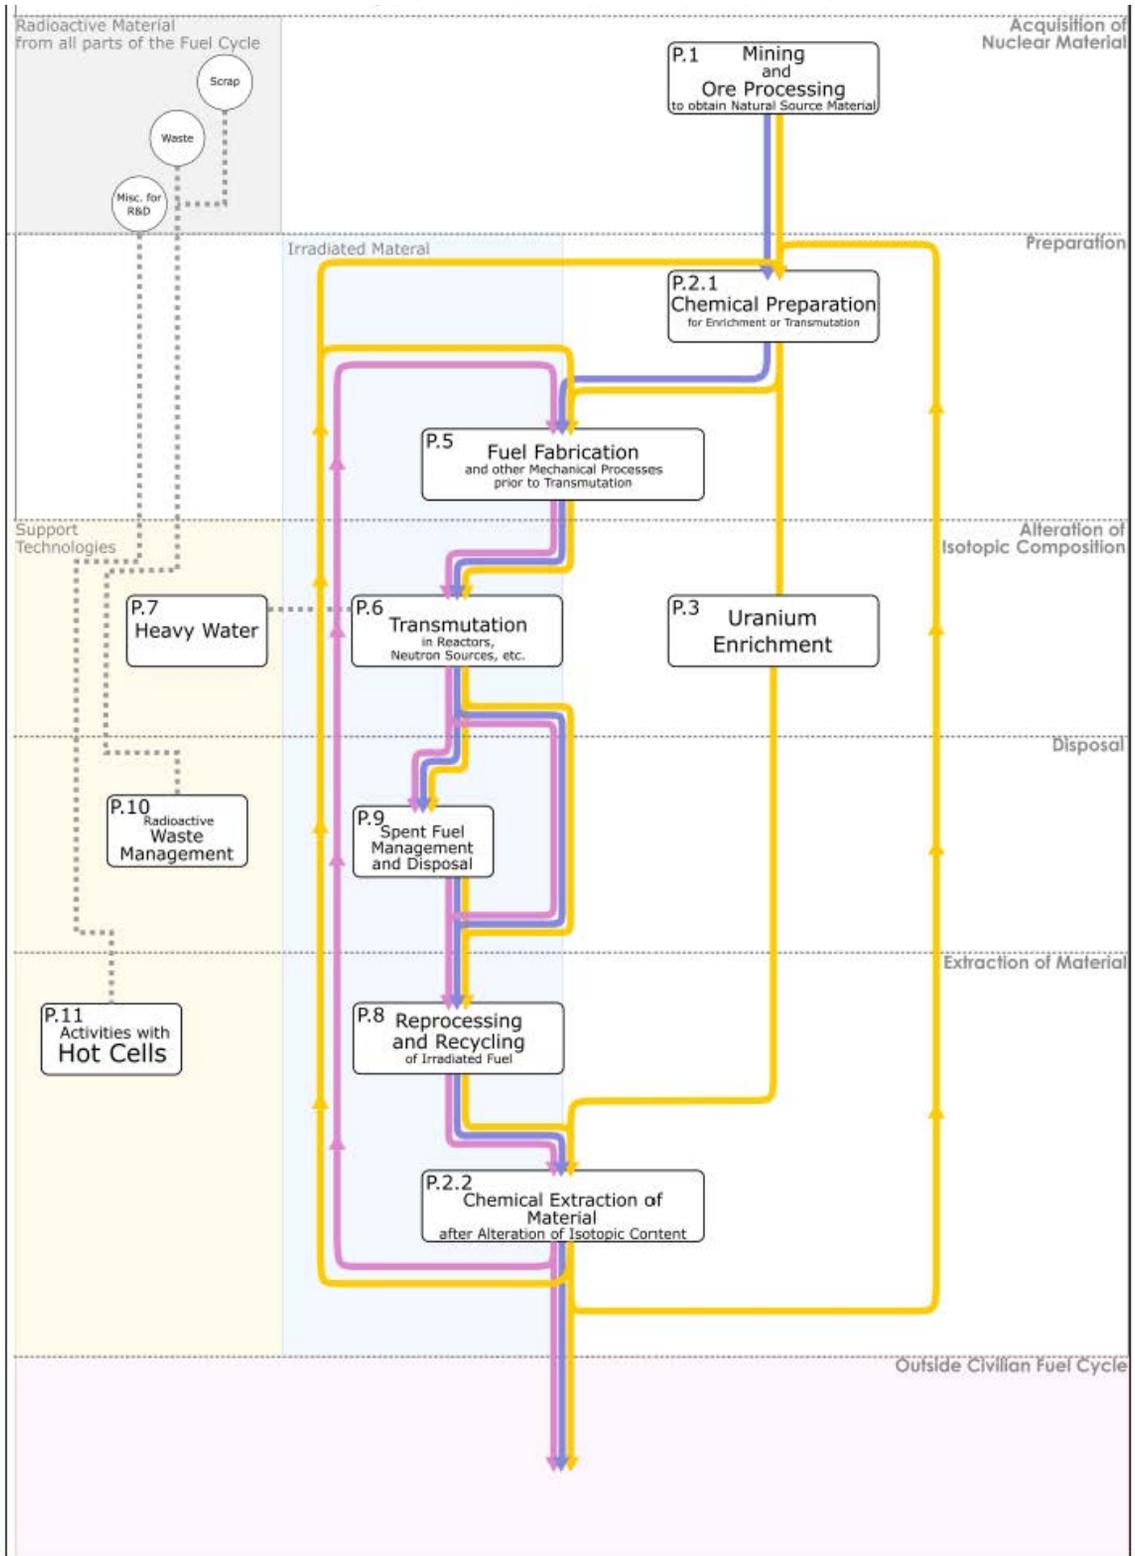


Figure 2. The Physical Model: a process view



Middle Earth Uranium and Plutonium Fuel Cycle Material Flow between Third Age 2941 and Third Age 3019

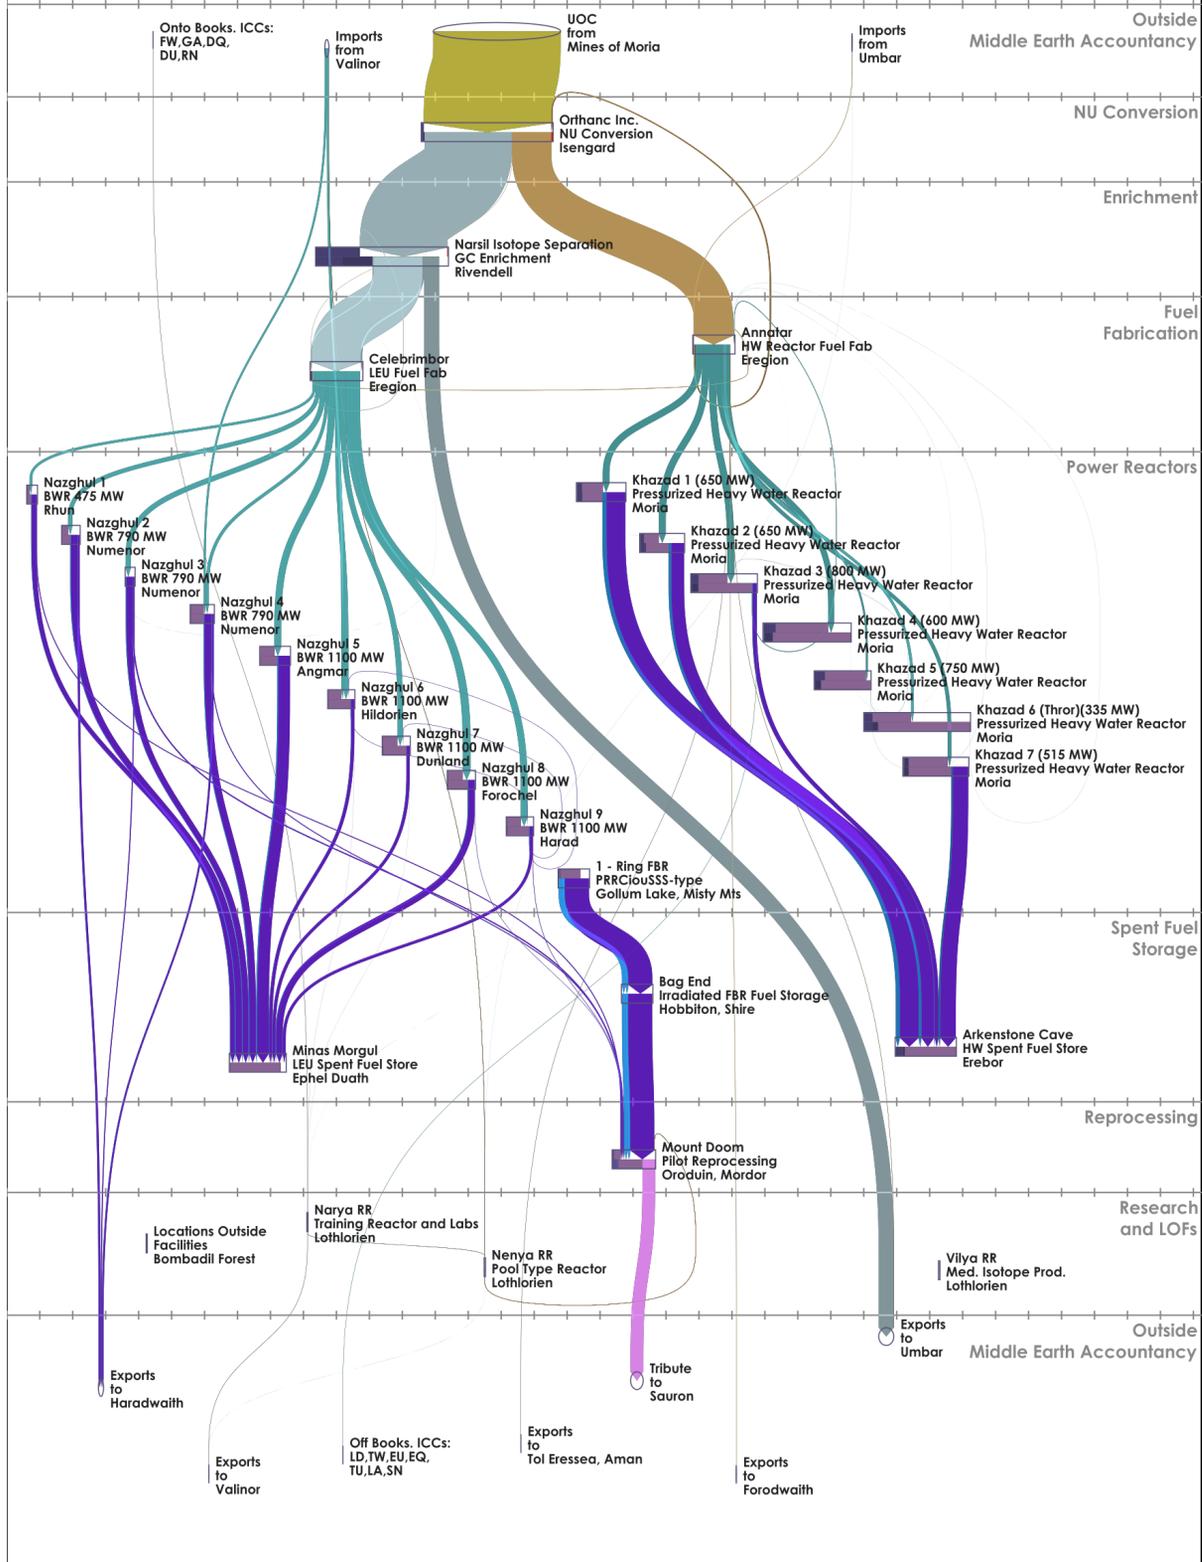


Figure 3. SNAKEY diagram for example hypothetical state of "Middle Earth"

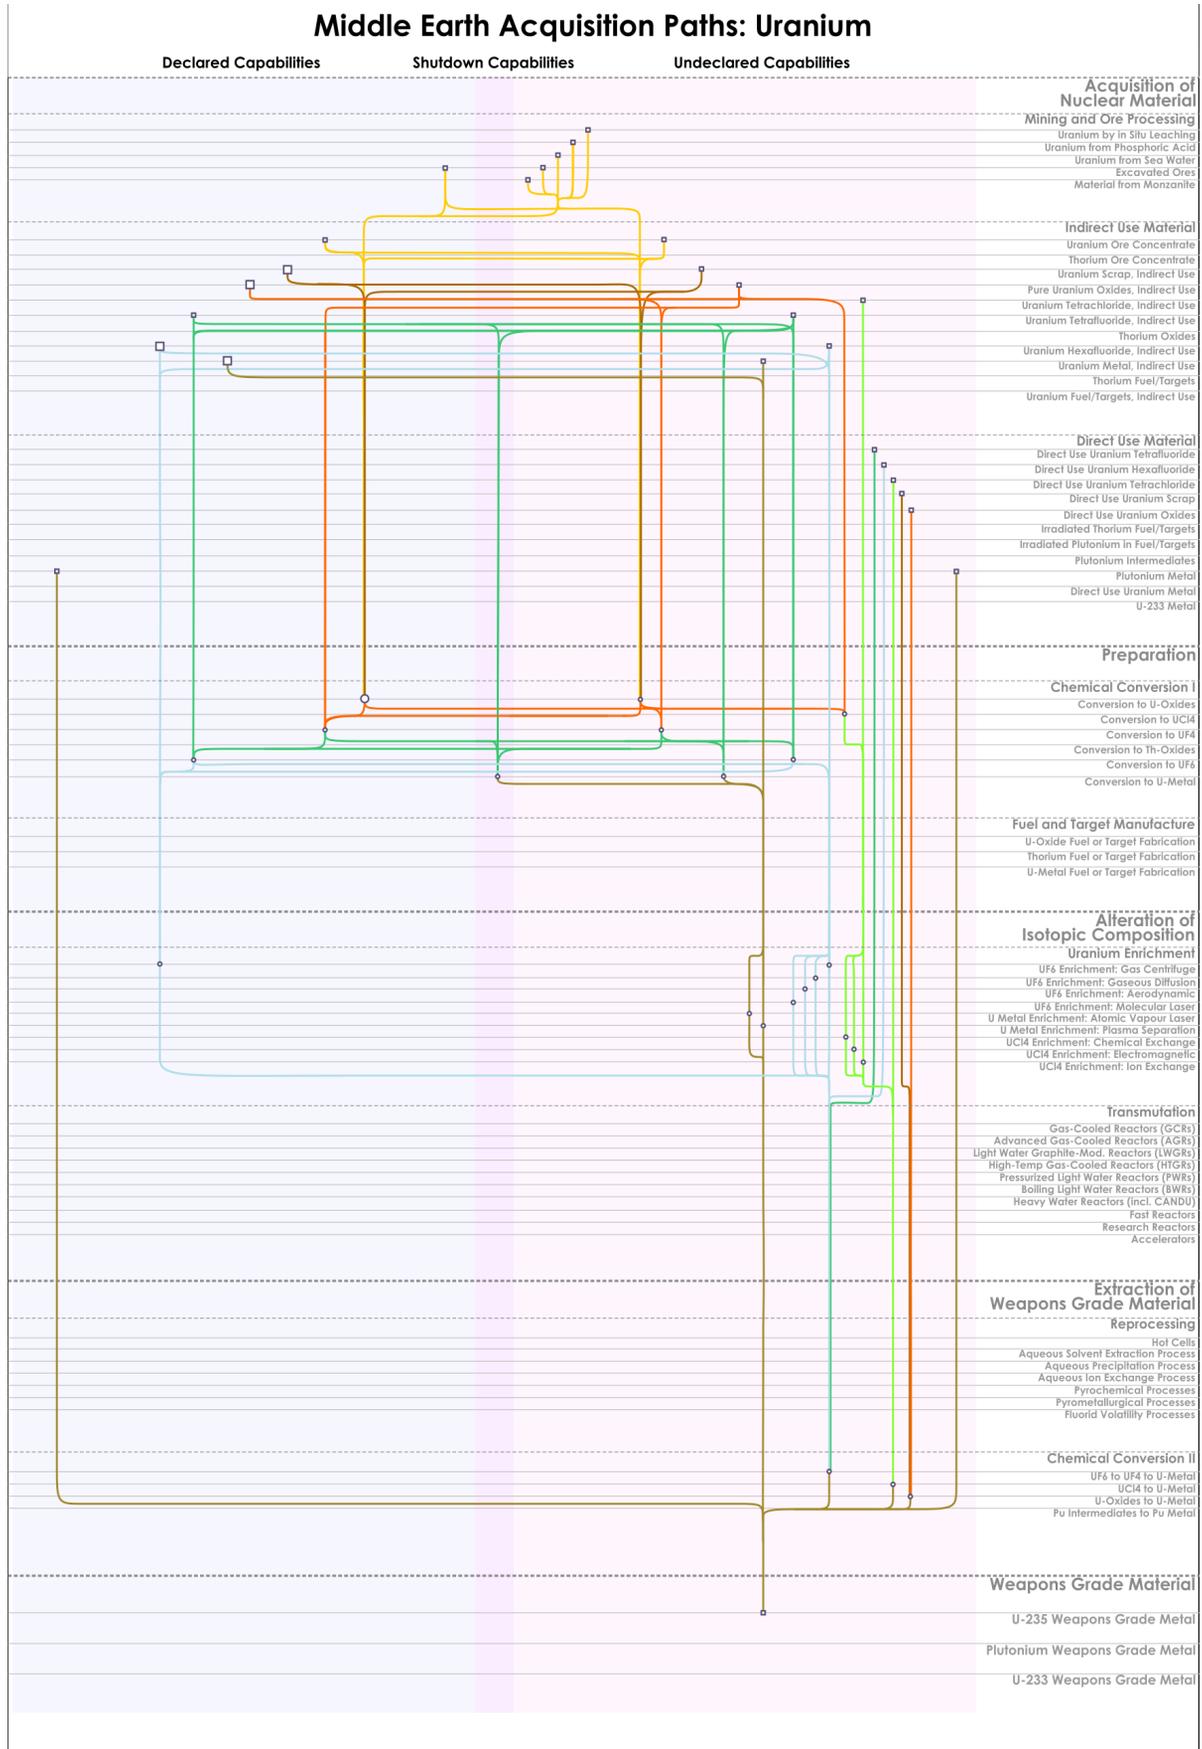


Figure 4. Automatically generated visualization of all possible uranium acquisition paths in a hypothetical state of "Middle Earth"



Figure 3 shows a sample SNAKEY diagram for a hypothetical state.¹⁶ Because this diagram includes the names of the facilities in the state, analysts can link the Physical Model, the material flows, and the geospatial data linking the facilities to photographs, maps, and satellite imagery in its collections.

Work to improve activities such as MBE at the state level suggests that a good way to view the NFC is to visualize materials as they flow through a sequence of transformative processes. This view differs from the traditional facility-centric view, making it ideal for use in developing a state-level approach (SLA), taking into account undeclared materials and activities in APA as well as linking fuel cycle facilities and material flows across the state. In particular, evaluation of material unaccounted for should be done in concert with DIV activities and a thorough understanding of the processes in a bulk handling facility. Such numerical and meticulous analysis is essential to reliably appraise potential diversion of material or misuse of technologies.

A partially automated proof of concept for the visualization of potential acquisition paths evolved from the same front-end technology as the SNAKEY tool. The concept relies on the process view of the Physical Model. For the purpose of the visualization, certain steps (e.g., uranium metal production) need treatment separately based on whether the input material is in direct use. The various processes are sorted into steps that advance a potential acquisition path (e.g., enrichment of indirect use UF_6), and those that do not (e.g., irradiated fuel long-term geological repositories, or irradiation of direct use uranium metal). The nodes on the diagram, seen in Figure 4, represent processes that can advance an acquisition path. The colored lines represent different types of material. The resulting visualization can serve as an interactive safeguards tool.

The visualization helps to organize the different acquisition paths into a Physical Model-based structure that can help reconcile data from different sources. The information aggregated by the SNAKEY tool as well as information from DIV, open sources, MBE activities, and so forth could be used to inform an analysis tool for assessing acquisition paths.

Conclusions

The IAEA's Physical Model describes the main technical activities involved in the NFC, from source material acquisition to the production of weapons-usable material. The present ongoing revision of the Physical Model focuses on updating, improving, and enhancing the documents to have a complete and comprehensive description of the NFC. This effort includes adding to

the document vital photographs, maps, and diagrams of the processes and equipment, and the crucial references, all linked to facilitate more in-depth topical analysis. This ability to plumb the nuances of the fuel cycle as well as review and improve the indicators for each volume provides benefits for the inspectors and analysts at the IAEA to do more detailed and technically informed state evaluations and build better SLAs.¹⁷

With the use of APA as a foundation for developing an SLA, the SEGs rely on the Physical Model as an indispensable resource to provide information on the fuel cycle for relevant technical analysis of a state as a whole. As shown in the statistical analysis SNAKEY project, the Physical Model provides a structure for clustering and qualifying data streams to make the NFC come alive with graphic representations. The IAEA intends the Physical Model to be a living and updatable library of the NFC to function as a seamless searchable resource integrated with IAEA's software for searching and collecting safeguards-relevant information. It is a consolidating architecture for categorizing NFC activity for relative significance in state evaluation including APA. Hence, the convergence of the Physical Model revision, statistical analysis, visual representation of fuel cycle data using the Physical Model, and computational methods for Physical Model data management is an excellent example of the fusion of disciplines and techniques to advance open source analysis with advanced data analysis. Such safeguards-relevant information can be a major component in the effort to understand the time scale it may take a state to develop NFC capabilities.

Keywords

IAEA, physical model, data visualization

Author Biographies

Brian D. Boyer is a safeguards analyst in the IAEA Dept. of Safeguards since 2015. Previously he was a safeguards expert at Los Alamos National Laboratory (2006-2015) and Brookhaven National Laboratory (2002-2004). He was a safeguards inspector and analyst in the IAEA Dept. of Safeguards from 1997 to 2002. He was a nuclear engineer in safety analysis at Brookhaven National Laboratory (1992-1996). He holds a Bachelors, Masters and Ph.D.(with Mech. Engineering Minor) in Nuclear Engineering from the Pennsylvania State University.

Jay Disser has been safeguards analyst in the IAEA Department of Safeguards, Concepts and Approaches Section since 2017. Current projects include safeguards-by-design, safeguards at small

modular reactors, and the revision of the Department's Physical Model. Jay has 11 years of experience in the nuclear field and has spent the last eight years working on nonproliferation and international safeguards issues. While at the IAEA Jay is on-leave from Idaho National Laboratory and has previously worked at Brookhaven National Laboratory. Jay has a Master's in Nuclear Engineering from Pennsylvania State University, a Bachelor's in Mechanical Engineering from the University of Central Florida, and is currently working on a MBA at the University of Pennsylvania.

Sebastien Richet is a Safeguards Analyst at the IAEA. His experience embraces commercial and technical duties, as well as large mathematical and computer science knowledge, that included measurement-related databases and tools production. He has worked for Cogema/Areva, particularly in Russia, and worked in nuclear waste management in Andra. He is president of the French Nuclear Energy Society Chapter in Austria. He holds a *Ingenieur, Nucleaire – Informatique* degree from *Ecole Nationale Supérieure d'Ingénieurs de Caen* and *DEUG, Mathematics* from *Université Scientifique de Lille*.

Daniel Gagne has been an Information Architect at the IAEA for the last 5 years. Previously he was at Pivotal Labs and a Senior Software Systems Engineer at MITRE. He received a Bachelor's Degree in Computer Science from Northeastern University and studied Engineering Systems at the Massachusetts Institute of Technology.

Stephanie Poirier is a Diplom Engineer from the INSA de Rennes, and formerly worked for ALSTOM. Ms Poirier has now been working for 15 years in the department of Safeguards of the IAEA. She spent 8 years deploying surveillance systems and developing ad hoc laser based solutions. She has worked three years for the Concept and Approaches Division before returning to the Technical and Scientific Services Division where she is now occupying the position of Safeguards Equipment Implementation Analyst and in this regard can share her experience related to safeguards by design.

Jeremy Whitlock is the Section Head for Concepts and Approaches, IAEA Dept. of Safeguards since 2017. Prior to that he had 22 years' experience as a reactor physicist and manager in the Canadian nuclear industry, mostly at Chalk River Laboratories. From 2006-2016 he was the manager for Non-proliferation and Safeguards at Chalk River (now operated by Canadian Nu-

clear Laboratories). He has a PhD and Masters in reactor physics from McMaster University, and a B.Sc. in physics from the University of Waterloo.

Agatha C. Walczak-Typke is a Statistical Data Analyst in the IAEA Department of Safeguards, Division of Information Management, Fuel Cycle Analysis Section since 2013. Before joining the IAEA, she worked in academia in Mathematical Logic for 6 years in various positions mainly at the University of Vienna, Austria, and the University of Helsinki, Finland. She also contributed to research in the field of music information retrieval. She then served as COO and then CEO of Theano GmbH, a Germany-based start-up which provides web-based solutions for validation of financial-services-related data. She has a PhD in Mathematical Logic from the University of Leeds, UK and a Magister in Theoretical Mathematics from the University of Wroclaw, Poland.

Claude F. Norman is a scientist manager at the IAEA. After completing a thesis in high-energy physics, Ms Norman joined the Belgian Nuclear Research Centre (SCK-CEN) in Mol Belgium and then the Euratom Safeguards inspectorate in Luxembourg when she served as safeguards inspector and was appointed first as statistical methodology officer and later as Operations Sector Head for the inspector of research reactors, hot cells and storages. In 2004 she joined the IAEA as Unit Head in the Statistical Analysis Section of the Division of Concepts and Planning in the Department of Safeguards and is now the Section Head of the Nuclear Fuel Cycle Information Analysis Section of the Division of Information Management (SGIM-IFC). This Section is in charge of the collection and evaluation of the IAEA verification data from environmental samples, non-destructive assay measurements and destructive analysis samples and their comparison with State declarations for the detection of nuclear material diversion and undeclared material and activities.

References

1. Liu, Z. & Morsy, S. (2007). Development of the physical model. IAEA Report SM-367/13/07. Vienna: IAEA.
2. Liu, Z. (2000). Making use of the physical model. IAEA Report STR-325. Vienna: IAEA.
3. Wallace, R., Anzelon, G., & Essner, J. (2009). Safeguards information from open sources. *Journal of Nuclear Materials Management*, 37(4).
4. IAEA. (2014). Supplementary document to the report on the conceptualization and development of safeguards imple-



mentation at the state level (GOV/2013/38) GOV/2014/41. Vienna: IAEA.

5. Versino, C., El Gebaly, A., Marinova, E., Pasterczyk, C., & Cojazzi, G.G.M. (2011). Integrating IAEA's physical model with JRC's the big table document search tool. *IAEA Safeguards Symposium*. Ispra.
6. Barletta, M., Zarimpas, N., & Zarucki, R. (2012). Open source information acquisition, analysis, and integration in the IAEA department of safeguards. *Nuclear Technology*, 179, 156–159.
7. Reed, J. & Feldman, Y. (2017). Updating the physical model for open source analysis. *Proceedings of the 58th INMM Annual Meeting*. Indian Wells, CA.
8. El Gebaly, A. (2017). Private communication. Vienna: IAEA.
9. Dupuy, G., Feldman, Y., Reed, J., & Gastelum, Z. (2017). Enhancing the use of multimedia information for IAEA safeguards analysis. *Proceedings of the 58th INMM Annual Meeting*. Indian Wells, CA.
10. Cooley, J. (2009.) The state-level approach to international safeguards. *Journal of Nuclear Materials Management*, 37(4).
11. Norman, C., Bevaart, L., & Fuhr, W. (2006). An integrated approach to material balance evaluation. *IAEA Safeguards Symposium*. Vienna.
12. Norman, C., Barletta, M., & Ferguson, M. (2010). Collaborative analysis in support of state evaluation. *IAEA Safeguards Symposium: Preparing for Future Verification Challenges, Poster Session*. Vienna: IAEA-CN-184/268.
13. Norman, C., Zhao, K., & Baute, J. (2013). Nuclear fuel cycle verification data and the state evaluation process challenges and opportunities. *Proceedings of the 54th Annual INMM Meeting*. Palm Desert, CA.
14. Sankey, H. (1898). The thermal efficiency of steam-engines. (Including appendixes). *Minutes of the Proceedings of the Institution of Civil Engineers*, 25(1896), 182–212.
15. Baute, J., Norman, C., Binner, R., Walczak-Typke, A., Caillou, F., Zhao, K., & Bonner, E. (2015). Dynamic exploratory visualization of nuclear fuel cycle verification data in support of the state evaluation process. *Proceedings of the 56th Annual INMM Meeting*. Indian Wells, CA.
16. Norman, C., Baute, J., Binner, R., Nikkinen, M., & Walczak-Typke, A. (2017). Evolution of verification data evaluation under the state-level concept. *ESARDA Symposium*. Düsseldorf.
17. Claude, F., Coyne, J., Calle, D., Gagne, D., Simmons, A., Wilson, B., & Iancu, I. (2017). Data fusion at scale: Strengthening safeguards conclusions through applied analytics. *Proceedings of the 58th INMM Annual Meeting*. Indian Wells, CA.



November 5-8, 2018

Symposium on International Safeguards: Building Future Safeguards Capabilities

Vienna International Centre
Vienna, Austria

March 12-13, 2019

Just Trust Me Workshop

Sandia National Laboratory
Albuquerque, New Mexico, USA

August 4-9, 2019

PATRAM 2019

New Orleans Marriott
New Orleans, Louisiana USA

January 22-24, 2019

INMM Spent Fuel Management Seminar -- XXXIV

Hilton Alexandria Old Town
Alexandria, Virginia USA

July 14-18, 2019

INMM 60th Annual Meeting

JW Marriott Desert Springs
Palm Desert, California USA

July 12-16, 2020

INMM 61st Annual Meeting

Baltimore Marriott Waterfront
Baltimore, Maryland USA

For more information, visit the INMM Events Page.

Author Submission Guidelines

The is the official journal of the Institute of Nuclear Materials Management. It is a peer-reviewed, multidisciplinary journal that publishes articles on new developments, innovations, and trends in safeguards and management of nuclear materials. Specific areas of interest include facility operations, international safeguards, materials control and accountability, nonproliferation and arms control, packaging, transportation and disposition, and physical protection. also publishes book reviews, letters to the editor, and editorials.

Submission of Manuscripts: reviews papers for publication with the understanding

that the work was not previously published and is not being reviewed for publication elsewhere. This restriction includes papers presented at the INMM Annual Meeting. Papers may be of any length. All papers must include an abstract.

The is an English-language publication. We encourage all authors to have their papers reviewed by editors or professional translators for proper English usage prior to submission.

Papers should be submitted as Word or ASCII text files only. Graphic elements must be sent in TIFF, JPEG or GIF formats as separate electronic files.

Submissions may be made via email to Managing Editor Amy Chezem at achezem@inmm.org.

Download an article template for the proper format for articles submitted to *JNMM* for possible peer review.

Papers are acknowledged upon receipt and are submitted promptly for review and evaluation. Generally, the corresponding author is notified within ninety days of submission of the original paper whether the paper is accepted, rejected, or subject to revision.

Format: All papers must include:

- Corresponding author's complete name, telephone number and email address
- Name and address of the organization where the work was performed
- Abstract
- Tables, figures, and photographs in TIFF, JPEG, or GIF formats. Color is preferred.
- Numbered references in the following format:
1. Jones, F. T., and L. K. Chang. 1980. Article Title. 47(No. 2): 112–118. 2. Jones, F. T. 1976. , New York: McMillan Publishing.
- Author(s) biography and photos
- A list of keywords

Download the article template from the INMM website.

The does not print "foot notes." We publish references and/or end notes. If you choose to include both references and notes, you may combine them under the same heading or you may keep them separate, in which case you must use numbers for the References (1., 2., 3., etc.) and letters (A., B., C., etc.) for the End Notes.

is published digitally in full color. Color graphics and images are preferred.

Peer Review: Each paper is reviewed by at least one associate editor and by two or more reviewers. Papers are evaluated according to their relevance and significance to nuclear materials safeguards, degree to which they advance knowledge, quality of presentation, soundness of methodology, and appropriateness of conclusions.

Author Review: Accepted manuscripts become the permanent property of INMM and may not be published elsewhere without permission from the managing editor. Authors are responsible for all statements made in their work.

INSTITUTE OF NUCLEAR MATERIALS MANAGEMENT



60th ANNUAL MEETING

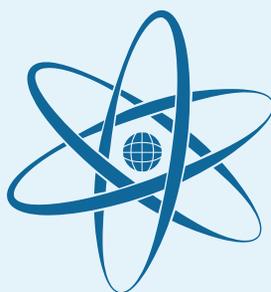
July 14-18, 2019
Palm Desert, CA, USA

JW MARRIOTT DESERT SPRINGS

www.inmm.org/Events/Annual-Meeting

 **#INMM19**

CELEBRATING
60
YEARS



INMM
INSTITUTE OF
NUCLEAR MATERIALS
MANAGEMENT