Machine-Learning-based sorting for nuclear fuel waste

Stéphane Puydarrieux [1], Jean-Louis Deparis [2], Mickaël Picq [3], Jean-Jacques Dupont [4]

1 Orano Recyclage La Hague, 2 Siléane, 3 PICQ Consulting, 4 Orano Projets

Abstract

In order to securely store old nuclear waste from UNGG reactors out of the Silo 130 at Orano La Hague (France), precise sorting means were investigated with the goal of regulating magnesium levels in each waste container.

Such a project is meant to prevent the high level release of Hydrogen in order to avoid the consequences of this Hydrogen release during ulterior storage. To sort the waste from the Silo and quantify the magnesium, Siléane and Orano have come up with a robotized solution provided with Vision technology (3D Vision system + Laser) and AI (Machine Learning algorithms). After detection, extraction and 3D reconstitution, each part of the waste is analyzed by a predetermined tree-based classifier (5 years of development and runs using nuclear waste simulants). These algorithms allow the kind of waste that is taken from the Silo to be indentified, whether it is Aluminum, Magnesium or Graphite. After a long period of development, runs and inactive tests to validate the solution, this project entered a new phase in summer 2019 : after obtaining the ASN authorization mid-2019, the industrial plant started with active waste at the end of 2019 and has produced 5 containers in a fully supervised modus to test the algorithms' behaviour on real waste. Now, a thorough review of the ML-based classifier is currently underway to optimize and qualify the process as active. As of April 2021, the machine is due to work with the upgraded AI-based Classifiers in fully autonomous modus. Considered as the latest AI case study in the nuclear industry, applications of AI-based recognition technology to material sorting and waste conditioning looks promising.

Keywords: *nuclear waste management, machine learning, 3D images*

1 Context

ORANO La Hague is a leader in recycling nuclear materials. The first facilities at ORANO La Hague site were designed to recycle the first generation of nuclear fuel burnt in French gas-cooled reactors (Uranium Naturel Graphite Gas, UNGG fuel). The structural parts of the fuel were separated from the uranium inner parts by a mechanical process, and the waste was stored in Silos. The Silo 130 mainly contains UNGG structural parts, which include graphite (90%) magnesium and aluminum waste. It received the first UNGG waste in 1976 until the end of the 1980s. Due to the evolution of safety standards, the extraction of this legacy waste is an important task for the Orano Group.

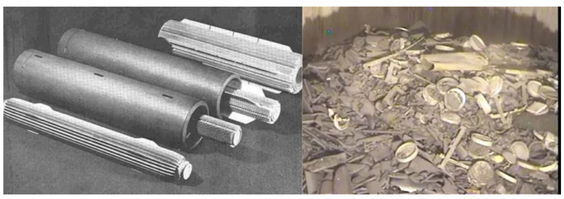

Figure 1 : Original UNGG Fuel over view, and bulk waste in the Silo

The difficulties inherent in the recovery and conditioning of the waste stored in Silo 130 are linked in particular to the lack of traceability of the objects stored on the one hand, and the changing

nature of this waste over time on the other. The characterization prior to retrieval is based on records of its production and calculation. Due to a fire-accident in 1981, the storage in Silo 130 can be divided in 3 main layers of waste. The lower layer has been flooded since this fire-accident, the 2nd layer was temporarily submerged, and the upper layer has never been in contact with water. The corrosion kinetic of magnesium is quite different for each layer.
The aim of the retrieval program is to empty Silo 130 and store the waste in specific barrels within optimal safety conditions. The waste would be submerged in water within the barrels, so H_2 production would take place within the barrel and must be evacuated. The barrels are covered by a lid which carries gasexhaust orifices which have PORAL®. sintered metal porous filter In the accidental situation of a barrel falling down in storage, the filter media would be flooded, thus the gas disposal requirements cannot be met until the barrel is well stored and the filter media dried. To prevent accidental situations due to H_2 production, the safety requirements are laid down to limit the H_2 production rate within the barrels by limiting the quantity of magnesium and aluminum.
There are two ways of H_2 production: the radiolysis of water by radio-nuclides and the corrosion of magnesium and aluminum stored underwater. As the radiolysis of the water cannot be avoided, the magnesium materials must be identified throughout the production, and quantified in each barrel to reduce the H_2 production rate as much as possible. The aluminum material must be identified and taken off from the production line. Thus, the system as a whole must ensure optimal performance of the correct product recognition rate, and even up to 100% for aluminum detection.

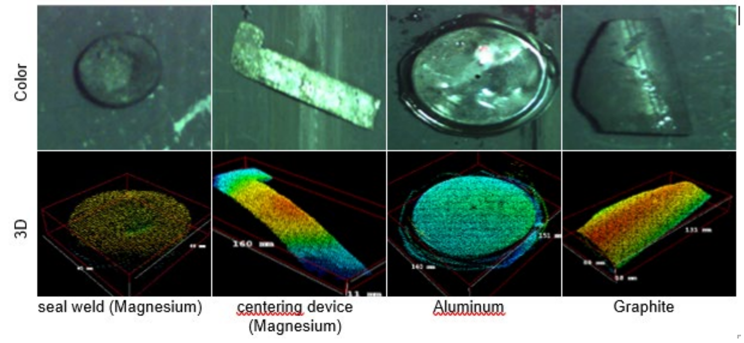

Figure 2: type of waste

Due to the high radiation dose rate, a human handling process cannot be considered. Each type of waste has specific geometric characteristics, both in terms of shape and size: magnesium parts are thin and long or with cylindrical portions; aluminum parts are plate-shaped; and graphite parts have random geometric properties due to the grinding mechanism in the retrieval grapple (Figure 2). Thus, a machine learning approach, relying on tree- based classifiers, was developed based on geometric recognition.

2 Description of the chosen solution

An effective identification requires a process which can pick out each piece individually throughout the production line. So that robotic arms took place upstream of the optical process of recognition which is programmed with “pick & place” artificial intelligence-based proprietary algorithm named Kamido® by Siléane (Figure 3). Each individual waste part is led under the optical AI based recognition process. As the piece passes through a laser ray, a combination of optical devices scans the laser ray deformation. A large range of geometric parameters are drawn from the 3D reconstruction of each single piece and then the data is analyzed by a tree-based classifier. The results are given as a recognition ratio in the range [0 – 1] for each kind of waste. The higher ratio corresponds to the higher probability of the waste material family being recognized.

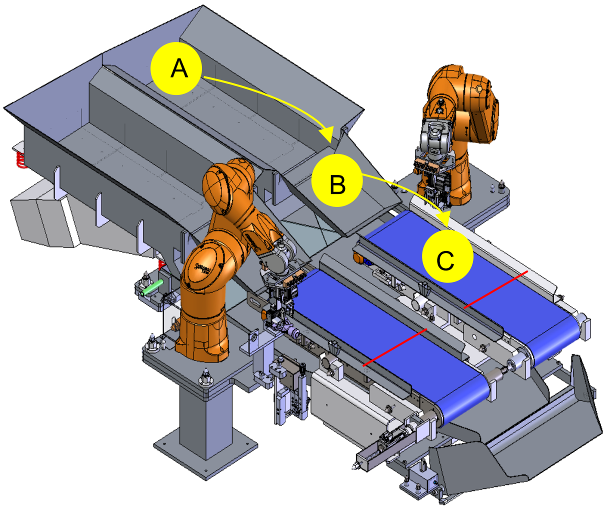

Figure 3: Kamido® System overview

Explanation: The successive vibrations of the bin (A) ensure that the feeding spout (B) of each robot is continuously fed. Optical cameras identify the waste to be picked up and send the coordinates to the robot, which will automatically pick up the waste one by one. The system ensures the choice of the most appropriate gripper tool according to the geometry of the identified waste: suction cup or gripper tool. The robots place each piece of waste individually on the conveyor belt (C) of its own quantification line. The isolated object then passes under the laser triangulation system. A dedicated algorithm reconstructs the 3D point cloud of the object (Figure 4).

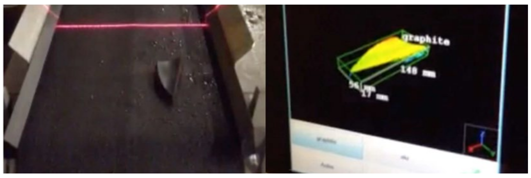

Figure 4 : 3D reconstruction of a graphite piece after passing through the laser ray

Machine learning (ML) is a set of statistical or geometric tools and computer algorithms that automate the construction of a prediction function *f* from a set of observations called the training set. A machine learning model is a specific algorithmic process that builds a prediction function *f* from a training database. The construction of *f* constitutes the learning or training of the model. A prediction corresponds to the evaluation *f* (x) of the prediction function *f* on the predictor variables of an observation x.
Supervised learning algorithms make predictions based on a set of examples. With supervised learning, you have an input variable that consists of labelled training data and a desired output variable. You use an algorithm to analyze the training data to learn the function that maps the input to the output. This inferred function maps new, unknown examples by generalizing from the training data to anticipate results in novel situations.
When the data is being used to predict a categorical variable, supervised learning is also called classification. This is the case when assigning a label or indicator, for example either dog or cat, to an image. When there are only two labels, this is called binary classification. When there are more than two categories, the problems are called multi-class classification. (Answers complex questions with multiple possible answers, such as: is this A or B or C or D?)
The choice of the algorithm is crucial. At the start of the project, in addition to the raw performance of the system, the need for explicability of the algorithm's operation was very important. A first classification algorithm had been designed to test the feasibility of the project (Figure 5: Refer to PATH 1). This algorithm was a simple decision tree only based on an estimation of the shape of an object (Does it look more like a rectangle or a circle?) from a few shape parameters (width, height, area...) and empirically chosen threshold values.
As expected, this algorithm obtained correct results on the batch of samples that had been used to choose the threshold values, but it was not satisfactory when new objects were presented: 2% of "false positive" (graphite identified as magnesium) and 8% of "false negative" (magnesium identified as graphite) … In order to improve these results, a random forest algorithm was

implemented. (Figure 5 : refer to PATH 2).
Decision trees, random forest and gradient boosting are all algorithms based on decision trees. There are many variants of decision trees, but they all do the same thing – subdivide the feature space into regions with mostly the same label. Decision trees are easy to understand and implement. However, they tend to over-fit data when we exhaust the branches and go very deep with the trees. Random Forrest and gradient boosting are two popular ways to use tree algorithms to achieve good accuracy as well as overcoming the over-fitting problem. Random forest or Random Decision Forest is a method that operates by constructing multiple Decision Trees during training phase. The Decision of the majority of the trees is chosen by random forest as the final decision.

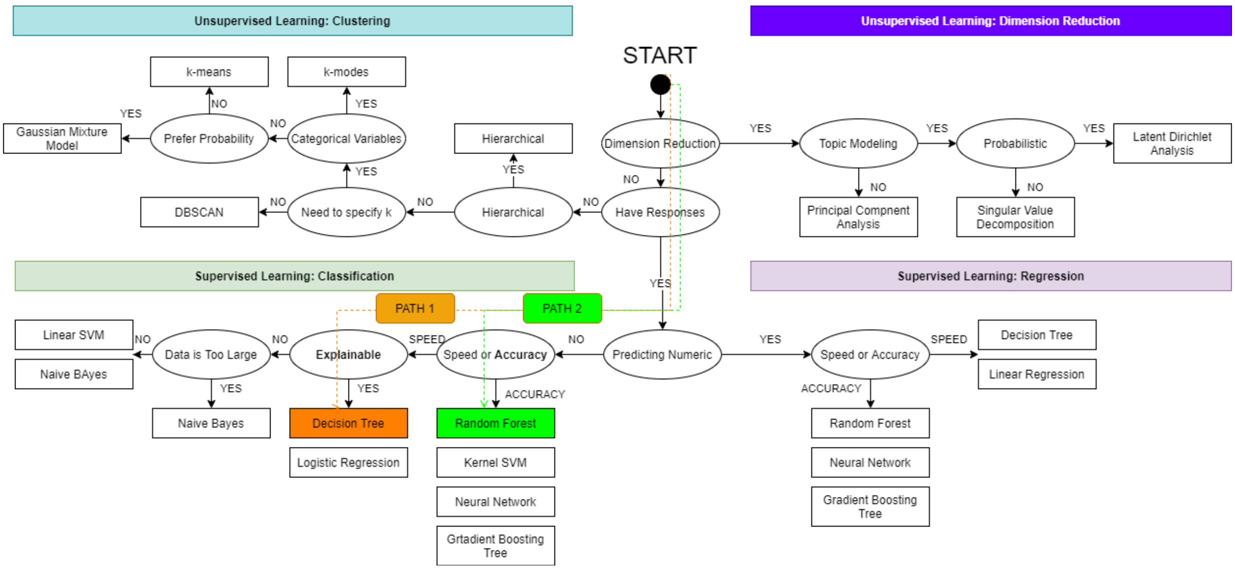

Figure 5: Which machine learning algorithm should I use? (Hui, 2020)

In summary, the advantages and disadvantages of a random forest algorithm are summarised in the table below (Figure 6):

Pros	Cons
✓ The overfitting problem is less pronounced ✓ Can be used for feature engineering i.e. for identifying the most important features among the all available features in the training dataset ✓ Runs very well on large databases ✓ Extremely flexible and have very high accuracy ✓ No need for preparation of the input data	○ Complexity ○ Requires a lot of computational resources ○ Time-consuming ○ Need to choose the number of trees

Figure 6: pros & cons of random forest (BestMachineLearningAlgorithmsforClassification, 2018)

The effectiveness of that kind of machine learning algorithm depends on a good representative range of waste samples taught to the classifier. Normally it takes several hundreds or even thousands of images to obtain a suitable algorithm. The input products consist of real waste from Silo 130. At the start of the project, no samples were taken from the Silo to characterise each family of waste (Graphite, Magnesium, Aluminum, etc.). The composition of the first batches and therefore of the real waste that will be the subject of this active start-up is therefore not known a priori.

A selected range of samples was defined to train the algorithm to identify the objects. These samples were defined by taking into account all available information: production records, UNGG fuel fabrication data sheets, consequences of more than 30 years of underwater storage by analysing the corrosion kinetic in laboratory, etc. Some samples were drawn from the Silo. Simulants of each category were produced. These raw objects were mechanically modified and surface treated to obtain the most representative baseline possible.

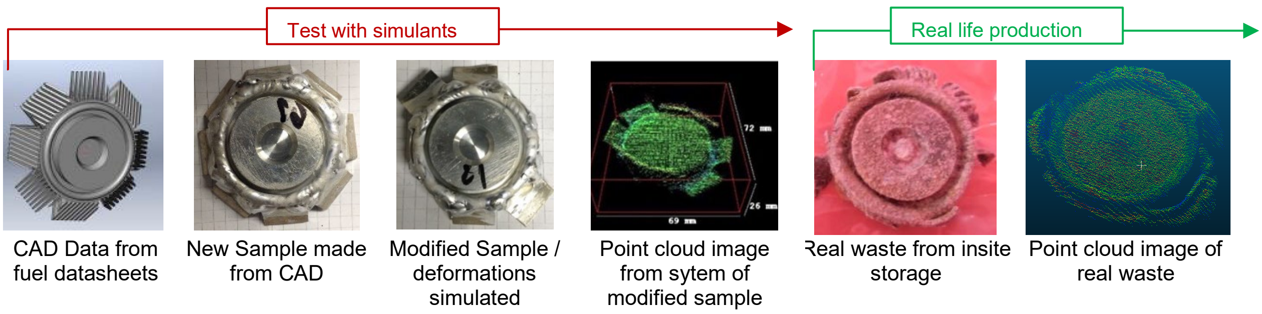

Figure 7 : Figures of seal weld, from simulants CAD Data to real waste point cloud

For the development and qualification of the machine more than 575 magnesium products and 67 Aluminum items were produced and scanned with the system among a total volume of $1m^3$ graphite representative simulants.
Some families contain objects of the same shape but different sizes. A data augmentation could have been done to take this scale factor into account. We preferred to build simulants of different sizes to qualify the other modules of the system (bin-picking subsystem for example).

3 Real life

All the training was done with waste samples designed from production records, drawn samples and estimated effects of corrosion on magnesium parts. So even if these samples are as similar as possible to real waste, it is not possible to be sure that the built random forest algorithm decision software will be able to correctly identify the waste that will be extracted from the Silo. It is agreed that at the start of production with real waste, the performance will not be sufficient. In summary, the initial performance is as follows.

	Drum#1 (C20885)		
Items count	6844		
	failure	OK	DVO
aluminum	25	11	20
centering device	5	32	5
graphite	276	5979	427
other	29	0	7
seal weld	8	11	9
Critical failure	38		
% DVO			6,84%
Overall total	**343**	**6033**	**468**

Figure 8 : performance of classifier#1 / objects of drum#1 (real waste)

As expected, the initial training, i.e. carried out on the simulants, did not give sufficient results for the automatic production of the cell from the start. The performance of the first version of Random Forest is not sufficient for stand-alone production. (See Figure 8). A significant number of aluminum covers (25) are not recognised by the system. The detailed analysis of the images gives a logical explanation for this state. In the figure, we can see a cover as it was learned by the system

and in the Figure 10, Three aluminum covers. The first one (A) looks like the images in the training set, and the confidence rate is good. The second (B) and third (C) correspond to a crushed lid. This type of object had never been seen by the algorithm.
Then it is necessary to anticipate the production launch and the possible problems that may arise. A strategy was developed prior to going into active production. This strategy makes it possible to start production while continuing to optimise the classification algorithm.

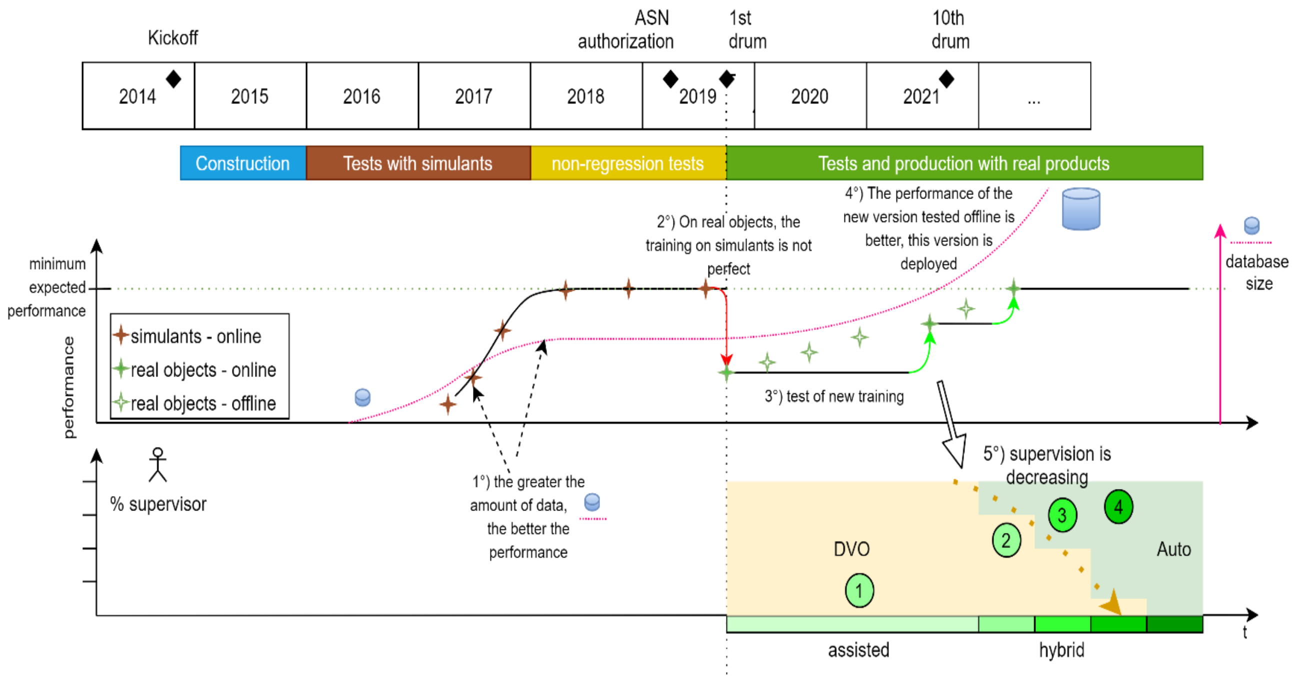

Figure 9 : Project plan and strategy from 1st drum.

As a result, the production of the first dozen drums is carried out in a fully assisted mode. All products are individually observed by the operator on the recognition conveyor. If necessary, the conveyor stops to give the operator time to indicate the correct class. Methodical tests have been carried out in order to define a well calibrated classifier which is able to recognize magnesium parts, avoiding "false-negatives" (magnesium identified as graphite) and "false-positives" (graphite identified as magnesium). This phase ① makes it possible to build up a large database of annotated images, i.e. with the truth about the nature of the object. (Figure 9: database size)
In a second phase (Figure 9: step 2 ②), the production is in a hybrid mode. The operator is only called (a procedure called DVO) upon for products for which the confidence level of the classification is not sufficient. The confidence level ΔS of the system is a numerical value derived from the classification algorithm. In the decision forest algorithm, each tree votes for the class that it considers most suitable. Each tree is different because it has only been trained on a subset of the global training base. Thus, the output of the decision forest algorithm is a table indicating the probability of the current object belonging to each class (Figure 10, column "Distribution of probabilities"). Confidence ΔS is then defined as the difference between the probabilities of the two most probable classes.
In the example of the Figure 10, object *A* has a good confidence rate while the confidence of the object *B* is rather low. The information given by the operator during the execution of the DVO procedure for the object B can confirm the correct choice made by the algorithm. On the other hand, the category of object C is not correctly determined by the algorithm which classifies the object as graphite. The calculated confidence rate is very low, as the algorithm proposes the classes aluminum and graphite with similar probabilities. The DVO procedure is called and the ground

truth is correctly entered in the database. It is reasonable to assume that after further training, this type of object will be correctly classified in the future.
The question arises as to how much product is needed to validate the system. The determination of the minimum sample size of the populations of each category" (seal weld, aluminum, centering device) was calculated using the formulas from COCHRAN's methods (G.COCHRAN, 1977), assuming that the sampling is carried out by polling.
Again, it is customary to split the base into two non-equal parts: 80% for training and 20% for testing. On the other hand, the distribution of products is not uniform, and the critical products are the least present (Figure 11), aluminum objects represent less than 0.18% of the objects extracted from the first 6 drums. Initial projections envisage a number of drums between 10 and 40 to have a stabilised system.

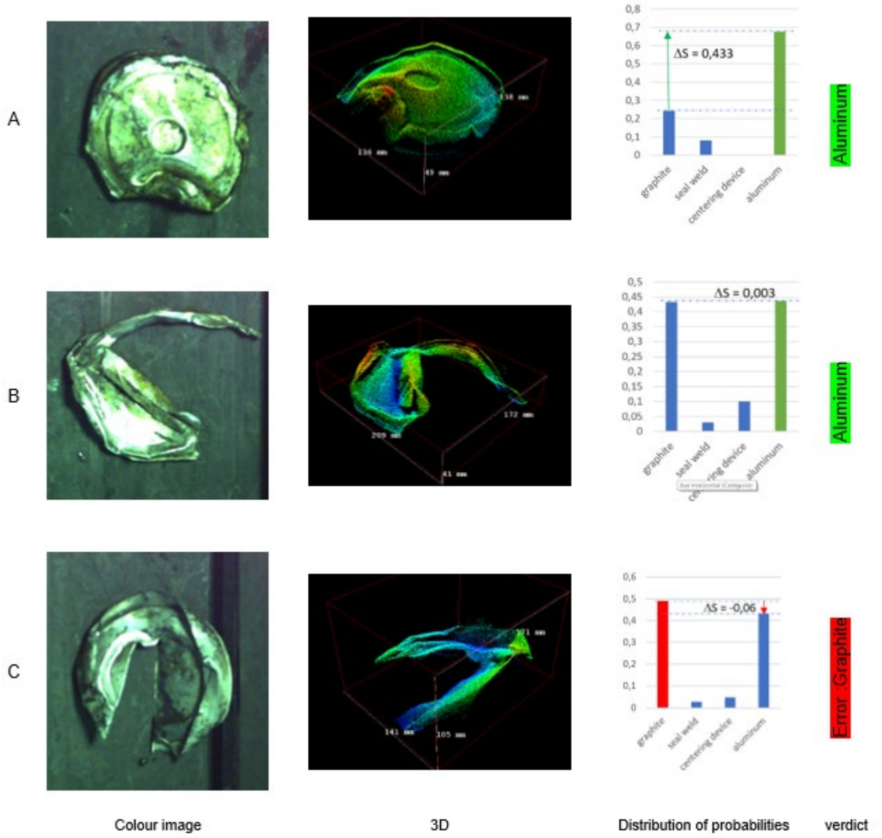

Figure 10 : colour photo, 3D point cloud and classification for 3 different objects. the verdict is that given by the trained classifier on the simulants

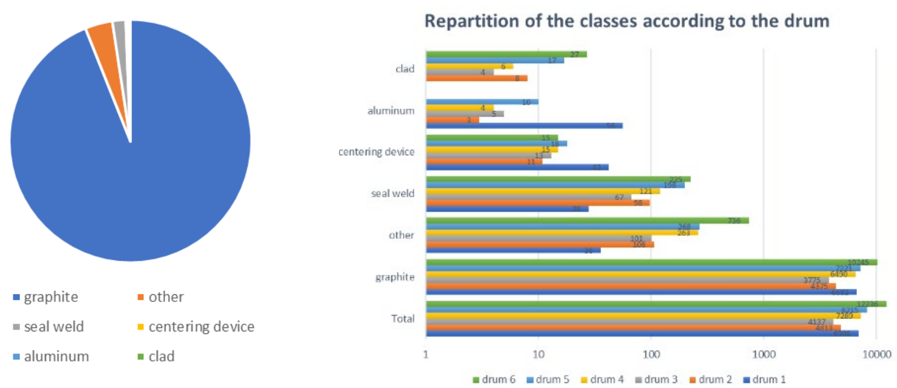

Figure 11 : Repartition of the classes. The number of products in the graph on the right is in logarithmic scale

In order to prepare the training of the algorithm, a specific software has been developed (Figure 12). The first feature is the visualisation of the data. This data is divided into two groups: training data and test data. The software also serves as a parameterisation interface for the creation of training. The software also allows the algorithm to be run with the training created on the basis of the tests. The performance per material type can be visualised. This software runs on a computer not connected to the system. Only a new version of Random Forest algorithm that performs better than the previous version is deployed on the machine (+ Figure 9). All these processes are very rigorously qualified.

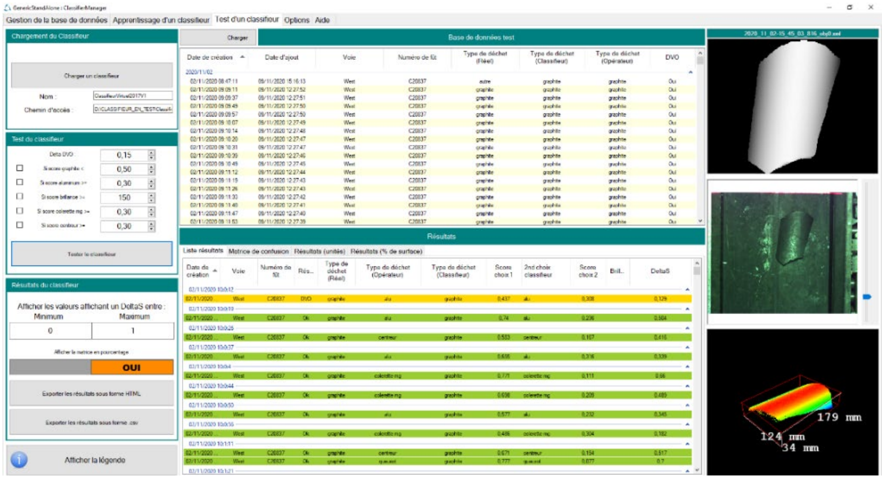

Figure 12: MMI of the software developped to create algorithms

Using this software, new training was performed using the data from the first 6 barrels. As a reminder, this data is divided into training data and test data. Significant improvements (Figure 13) have already been observed, although the desired number of annotated images has not yet been reached.

	drum#5 (C31929)		
Items count	7732		
	failure	OK	DVO
aluminum	2	4	4
centering device	1	14	3
graphite	272	6362	587
other	242	0	26
seal weld	66	78	71
Critical failure	69		
% DVO			8,94%
Overall total	583	6458	691

Initial classifieur

	drum#5 (C31929)		
Items count	7732		
	failure	OK	DVO
aluminum	1	7	2
centering device	0	18	0
graphite	587	6159	475
other	200	0	68
seal weld	7	202	6
Echecs Critiques	8		
% DVO			7,13%
Total général	795	6386	551

optimised classifier

Figure 13: Performance of an improved version of the classification algorithm trained on real images compared to the performance of the algorithm trained only on simulants

For the system to be optimal and in order to allow the operator to use the system with greater confidence, various changes have been made since it was installed on the La Hague site.

The acquisition system operates in the environment of the hot cell, which has its own white lighting. This lighting is dedicated to maintenance operations but also to monitoring the process in general during production. It turns out that this white lighting, which is quite bright, constitutes noise for the 3D acquisition system. In order to improve the accuracy of the system, it is desirable to improve the SNR of the laser line seen by the camera. Therefore, optical filters were added in front of the illuminators to minimise the light power at the laser wavelength. This modification

improved the quality of the 3D reconstructions in general, to the detriment of the visual rendering of the colour cameras dedicated to observation and installed in the cell. A white balance was able to partially compensate for this degradation.
Initially, the operator only had the 3D point cloud of the object displayed on the supervision software to indicate the class of the current object. A colour camera filming the waste stream was added to the system. The operator now has a colour photo of the object to indicate the class. Operation is faster and more accurate.

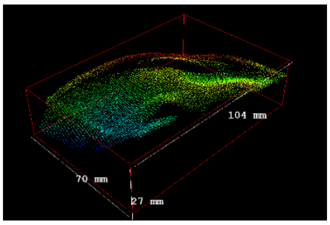

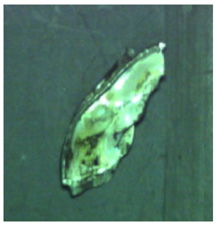

Figure 14 : Determining the class of the object is complicated for a human with only point cloud information (left). It is easier with the product picture (right).

In order to improve the reliability of the designations made by the operator, it is now possible to activate a two-step validation. The addition of this procedure allows the system to limit the number of wrong designations due to a handling error.
Finally, statistical tools have been added to the supervision software, allowing for the retrieval of indicators that previously required an offline post pressing.

4 Conclusion

There is a real desire to manage the deployment of the algorithm in this particularly sensitive area from a security point of view. No self-learning is foreseen as the classifier must be supervised at all times but all tools are ready to learn a new classifier or to improve an existing one with the possibility to check its results on a large database of 3D reconstructions.
The success factor of the implementation of a waste sorting system based on artificial intelligence lies in the ability of the industrialist to precisely define the inventory of waste present in the storage. A rigorous analysis of the storage conditions and the transformations induced by the mechanical processes must be carried out in order to best define the expected shape of the waste simulants used for inactive qualification.
The start of the process with the real waste cannot be done directly in automatic mode. An intermediate step-by-step mode is necessary to check the consistency of the training carried out on the basis of the waste simulants with respect to the real waste.
The time and volume of waste required for the active qualification of the sorting process based on artificial intelligence is directly linked to the representativeness of the simulants on the one hand and to the risk of being confronted with new disturbing elements that cannot be controlled (quality of lighting on site, operating conditions, etc.) on the other.
Despite all the precautions taken to ensure that the simulants were as representative as possible, the storage conditions in the Silo brought an expected uncertainty in the shape and surface condition of the real waste. A re-training or even an optimisation of the waste sorting algorithms was therefore an option considered when starting up the installation with real waste.
Similarly, as the Silo continues to be emptied, it is expected that the shape of the waste will change as a result of being stored underwater, and a stratum phenomenon is identified which leads to variable corrosion kinetics of the magnesium and aluminum waste depending on its depth and the time spent underwater. This phenomenon leads to a periodic verification of the non-regression of the sorting system and, if necessary, a re-training of the system in order to integrate a possible notable evolution of the geometric characteristics of the waste to be sorted.

Before creating a new learning, the data must be prepared. The objective of data preparation is to create a homogeneous quality dataset with a consistent and well-defined structure and formats. Dealing with missing values, removing outliers, redefining the scales of the parameters to homogenise their variability are the main actions carried out during this phase.
Applied to our problem, this principle consists of removing from the database all images that are not of interest. The software automatically saves all images, including those taken during maintenance test periods. These images are not of significant interest for the learning or testing phase. It is wise to remove them from the database.
By analysing the data, we realise that the distribution follows a normal law. An optimisation to be evaluated is the possibility of calling the DVO procedure for objects that are too far from the mean. This strategy will make it possible to quickly isolate the out-of-norm objects.

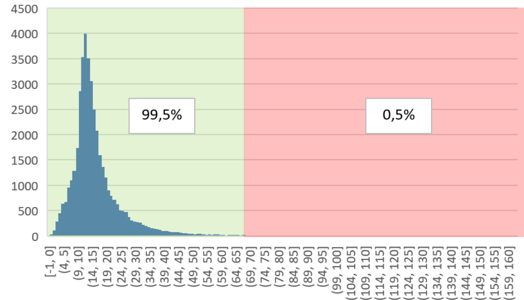

Figure 15: Distribution of the average height value (mm) for the product set

With a proper and well-designed AI system, it is possible today to bring robotics systems an interesting capacity of agility and autonomy when the input data of the process to automatize is uncertain or even unavailable. It is the purpose of Siléane to develop machines associating robotics with various sensors and AI in order to take up the challenges faced by the operating industrials. This specific project with Orano La Hague on Silo 130 is an exemplary project in that matter, as the automation has allowed the project to be launched and operated in a short time frame and within a reduced space, when a mechanical and human-led solution would have had much bigger costs regarding the available financial resources and most of all regarding the implied dosimetry costs for the operators (ALARA principle). The worst case being the most probable without automation : leaving the project unmanaged as a legacy to the next generation. Research and Development to make autonomous robotics safe and sound for operations in the nuclear industry is and stays one of Siléane's most focal project of the last and next decade and the partnership with Orano brings concrete applications and pragmatism to make essential technological advances in various fields. And now, along with autonomous waste sorting solutions, autonomous dismantling coupled with radiological measurement is to become operational.

5 Bibliographie

Andra. (n.d.). Retrieved from Déchets de structure des combustibles UNGG (Orano/La Hague): https://inventaire.andra.fr/families/dechets-de-structure-des-combustibles-ungg-oranola-hague
BestMachineLearningAlgorithmsforClassification. (2018). Retrieved from BestMachineLearningAlgorithmsforClassification: http://intellspot.com/wp-content/uploads/2018/03/5-Best-Machine-Learning-Algorithms-for-Classification-infographic.pdf
G.COCHRAN, W. (1977). *Sampling Techniques third edition.* John Wiley & Sons,Inc.
Hui, L. (2020, 12 9). *Which machine learning algorithm should I use?* Retrieved from The SAS Data Science Blog: https://blogs.sas.com/content/subconsciousmusings/2020/12/09/machine-learning-algorithm-use/