EU Horizon program: Horizon-CL4-2021-TWIN Transition
Reducing environmental footprint, improving circularity in extractive and processing value chains (IA)
Grant Agreement No 101058310

WP 6 Software

D6.3 Functional user- and machine learning interface

ReSoURCE

Project Reference No	101058310
Deliverable	Functional user- and machine learning interface
Work package	WP 6
Type	R – document, report
Dissemination Level	PU – Public
Date	August 2025
Status	Final
Editor(s)	Cord Fricke-Begemann, Yannick Conin, Sebastian Ulrich
Contributor(s)	ILT, LSA, NEO, RHIM
Reviewers	All partners
Document description	This document describes the interfaces to the ML algorithms for user interaction for training required to adapt the modules' performance in operation.

Document revision history

Version	Date	Modification introduced	
		Modification reason	Author
V0.1	21.07.2025	1st version	ILT/yaco
V0.2	25.07.2025	Add implementation details from LSA	LSA/jm
V0.3	30.07.2025	Add HSI details	ILT/yaco
V0.3.1	01.08.2025	Refine HSI details	NEO/jh
V1	01.08.2025	Final release	ILT/yaco

Table of Contents

1. Executive Summary

The work reported in this deliverables refers to ReSoURCE project (*Resource - Refractory Sorting Using Revolutionizing Classification Equipment*) in the Framework of WP6 – Software and artificial intelligence. This deliverable presents results mainly originating from Task 6.3. It describes the machine learning (ML) algorithms used for classification of refractory material in demonstrator A as well as the interfaces implemented for interaction between the user and the ML models. It outlines how the various sensors get linked into different ML-components and how their outputs get combined to a final sorting decision for the extraction system. A key feature of the software system is the real-time sensor-fusion and classification of objects to provide the Demonstrator A's extraction systems with sorting decisions. The intelligent LIBS control through combination of 3D-, HSI- and LIBS-Sensors allows up to 92 % precision compared to 43-63 % precision from single-sensor operation.

2. Introduction

The ReSoURCE concept for the sorting of refractory materials consists of a set of different sensors which are used to characterize and classify individual pieces of the material which are moving on a conveyor belt. From the sensor data the information for physical separation and fractionation of the pieces is derived. Sensors of three different types are used: 3D camera, HSI camera and LIBS. A schematic drawing of their arrangement at the conveyor belt is displayed in Figure 1 with refractory bricks being moved from left to right.

All sensors are generating data about the material at different positions on the belt. While the cameras are measuring the whole area of the belt, the LIBS sensor performs consecutive single spot measurements with a fixed repetition rate. To make efficient use of the LIBS measurements it is insufficient to simply statistically distribute the measurement locations. Instead, they are actively controlled to be carried out in so-called regions of interest (ROI) on each sample. The ROI are pre-determined by an algorithm based on the combined data from the cameras. Subsequently, the LIBS measurement is carried out in the ROI and the resulting LIBS data is merged with the camera data to determine the classification of the inspected material piece. Therefore, the ROI must be selected for each object in such a way as to optimize the classification success for the specific object. The data evaluation involves algorithms based on machine learning and artificial intelligence as well as classical analytical procedures.

Figure 1 Schematic drawing of the automated refractory brick sorting system with 3D scanning, HSI, and LIBS sensors integrated along a conveyor belt, followed by a sorting unit for material separation.

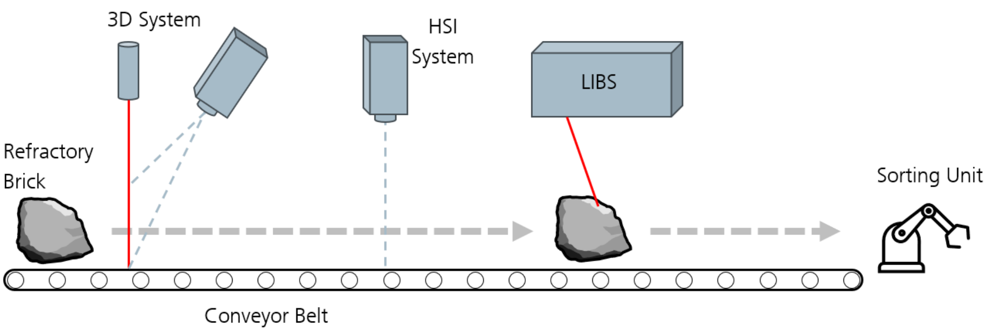

A schematic drawing of the soft- and hardware components is depicted in Figure 2. The following basic steps of data acquisition, treatment and communication are carried out:

1. The 3D camera sensor continuously generates information about the topology of the material stream on the moving belt. The height information is pre-treated by the iMotion software package and transferred to the LIBS ROI Selector as well as the final Sorting Decider.

2. The HSI camera continuously collects the spectral reflectance signal for every pixel of the imaged material stream on the moving belt. The spectral information is classified per pixel and per object by the software package Breeze and transferred to the LIBS ROI Selector as well as the final Sorting Decider.
3. The LIBS ROI Selector algorithm determines a set of measurement locations for the LIBS sensor control for each object to be measured.
4. When the object of interest reaches the LIBS measurement area on the moving belt, the LIBS measurement is positioned and executed at the defined ROI. The resulting LIBS spectrum is classified by the software SpectraAnalyser and transferred to the final Sorting Decider algorithm.
5. The final Sorting Decider algorithm in iMotion combines the LIBS data with the HSI- and topology data and determines the material class for each inspected object. The class information in combination with geometrical information about the object is transferred to the extraction system.

Previously, D6.1 has described the general architecture for the ROI selection algorithm and the real-time data interface concepts. D6.2 has described the target ROI selection algorithm and the material classification interface.

Figure 2 Schematic drawing of the interaction of the individual software and hardware components for material classification.

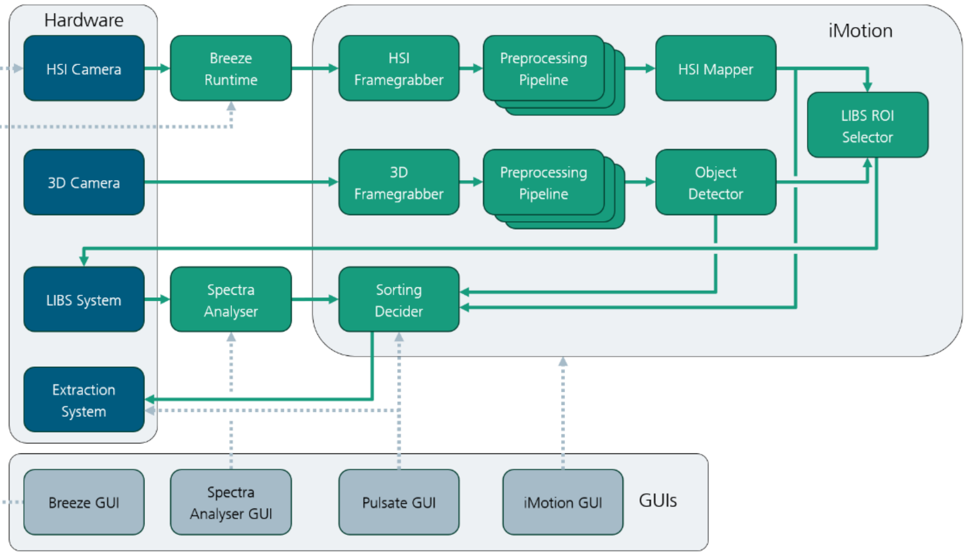

3. Data Acquisition

For the purpose of this document, a demonstration use case has been selected which intends to sort three types of refractory material used in cement rotary kiln (CRK) plants, namely two hercynite variants with either low or high iron content (HERC-lowFe and HERC-highFe, respectively) and one magnesium spinel variant (MgSp). For training of the HSI classifier, for each material group approximately 200 unused, crushed samples were recorded using the VNIR-camera, measuring in the 400 – 1000 nm wavelength range. Class labels were manually annotated to the data. Prior to the data generation two separate reference measurements were taken, one of a calibrated reflectance panel as white reference, the other one with closed shutter as dark reference. The raw sample spectra were calibrated against these references to obtain reflectance values.

Additional data was recorded to train the LIBS classifier. For this, the same material groups were used but included both original and used brick conditions, representing practical industrial applications. Each material type was sampled with approximately 30–60 full spectra per sample, covering the wavelength range from 200 to 900 nm. The spectra were utilized in their raw, uncorrected form to maintain data authenticity. Six distinct material classes were defined: magnesium spinel original (MgSp-orig), magnesium spinel used brick (MgSP-UB), hercynite low iron original (HERC-lowFe-orig), hercynite low iron used brick (HERC-lowFe-UB), hercynite high iron original (HERC-highFe-orig), and hercynite high iron used brick (HERC-highFe-UB). The dataset was intentionally designed to represent conditions found in practical industrial usage. A classifier was subsequently trained to distinguish between all six material classes, with data saved in a structured format and appropriately annotated with class labels during the acquisition process.

Upon completion of training for HSI and LIBS classifiers new data was recorded as foundation for the machine learning models present in iMotion. Again, fresh samples of the same original CRK material classes were used, leading to a dataset of 441 unique measurements after outlier removal. The recorded data consists of topological data of each sample from the 3D camera, the HSI classification as output by Breeze, as well as the LIBS classification as output by SpectraAnalyser. Due to technical limitations at time of recording only a single LIBS ROI per brick sample could be analysed and measured. Depending on conveyor belt speed, up to 12 LIBS measurements per sample will be available for future refinements. LIBS ROI were chosen randomly on a sample but weighted towards areas likely to offer fair measurement conditions for LIBS (ref. section 5.2). Data was saved in a structured format for training, annotating the class labels in the progress. To verify the trained LIBS ROI selector’s performance, new verification data was recorded. This training data consists of a total of 132 measurements of the same objects as used before.

4. Classification Models

To evaluate the data from the various sensors, multiple classification models are employed to classify and fuse the data into a unified sorting decision. These models are outlined in the following sections.

4.1. HSI Sensor – Breeze

For real-time classification of HSI line image data, the Software Breeze is used. When a new image line is received, the spectra get calibrated against the pre-recorded white and dark references to obtain calibrated reflectance values on each image pixel. This approach ensures comparability between recordings under different conditions. 2D arrays are composed from the calibrated line image data to facilitate the classification of objects.

For classification two different models are used. Samples are segmented from the data arrays representing the whole conveyor belt by means of a pre-trained Fast Segment Anything model (FastSAM). For classification, a binary partial least squares discriminant analysis (PLS-DA) is fitted for each material class in a one-versus-rest approach.

Figure 3
Data flow for HSI classification as depicted in Breeze.

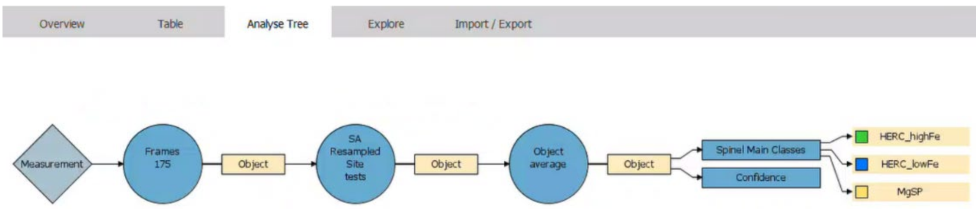

The data output of the classifier is structured into two distinct streams. First, all spectra are averaged over all pixels of an object, then the PLS-DA models classify each object and output the greatest object-wise confidence and the respective class. The schematic drawing for this workflow is depicted in Figure 3. It can be seen how the distinct PLS-DA models to the very right share the same preprocessing steps. For the second data stream, the PLS-DA models operate on a per-pixel basis to create a map of classifications and respective confidences per pixel. While the first method delivers a more accurate classification per object, the second method is needed for the consecutive classification steps in iMotion. It yields important local information about the object structure, which is relevant for optimal LIBS ROI selection.

4.2. LIBS ROI Selector - iMotion

To perform LIBS measurements, it is necessary to determine the area on a sample, known as the region of interest (ROI), where the spot measurement should take place. Previous works have solely used three-dimensional topology data of an object to determine regions that are well suited for LIBS measurements – these are generally areas that have low surface roughness and are angled perpendicular to the optical axis of the LIBS laser beam. To enhance the sorting process in terms of classification accuracy and material throughput, ReSoURCE uses a HSI camera to aid in ROI selection and material classification.

The software package iMotion is responsible for aggregation of data from different sources and outputting a sorting decision. One key component is the LIBS ROI Selector module. It utilizes an object's 3D topology from the 3D camera in conjunction with classified HSI data as output by Breeze to calculate one or more ROI per object. For this, a machine learning based approach is chosen.

The first step in the data pre-processing toolchain is data cleaning where outliers are excluded, for example objects smaller than a certain minimum size. From the topology data two new intermediary maps are calculated for surface roughness and angle between surface normals and optical axis of LIBS. These get combined into a map indicating the surface quality for taking a LIBS measurement. An exemplary surface quality map is displayed in Figure 4 (centre), next to the topologic map in Figure 4 (left) it is derived from.

Figure 4
Data of a sample brick used as inputs for LIBS ROI selection.

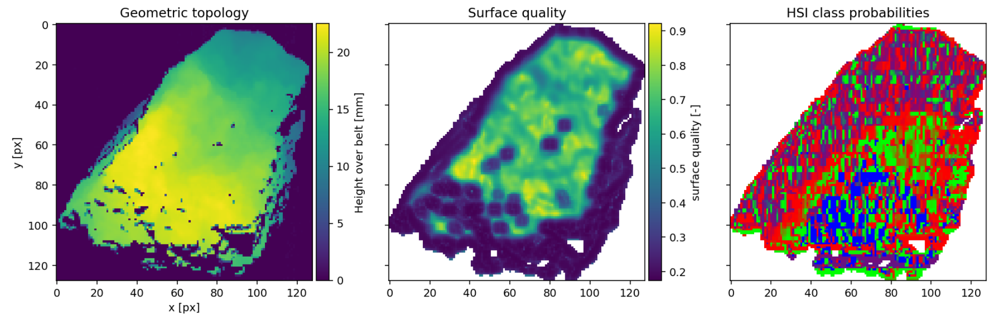

The HSI data – consisting of two pixel-maps where each pixel holds the most likely material class and the classifier's confidence of that material class – will have the missing classes imputed by assuming the non-predicted classes as equally likely. This results in pixel-maps of confidences, one for each material class. An exemplary map of per-pixel class probabilities is displayed in Figure 4 (right), where pixels display mixed colours representing the probability distribution across classes. The classes are colour-coded as HERC-highFe (red), HERC-lowFe (blue), and MgSp (green). These probabilities are noted as $\tilde{p}_{\mathrm{H},x,y}(\hat{c}_{\mathrm{n}}|m_{\mathrm{i}})$, meaning the probability of predicting material class $\hat{c}_n$ at position x, y given a measurement m_{i}. $\tilde{p}_{\mathrm{H}}$ indicates that the data originates from the HSI classifier and has not yet been calibrated.

For training, each object has been annotated with its material class. They are divided into a training and test data set, making sure to balance the class distributions in both sets. From the training data set a probability calibration is performed. This aims to enable the usage of the predicted class probabilities in further statistical calculations. Here matrix scaling is used which performs linear regression on the class probabilities by multiplying them with a matrix $\boldsymbol{W}$ and applying a bias b. This results in calibrated class probabilities $p_{\mathrm{H},x,y,m_i} = \tilde{p}_{\mathrm{H},x,y,m_i} \cdot \boldsymbol{W}_{\mathbf{H}} + b_{\mathrm{H}}$. The parameters of $\boldsymbol{W}$ and b are optimized by minimizing the cross-entropy loss over object-wise class predictions using a gradient descend method. For HSI data first the matrix scaling is applied to all image pixels. Then for object-wise aggregation all class probabilities are averaged over all pixels, making sure to exclude pixels from the background.

As next step in data pre-processing LIBS data will be considered. Capturing of LIBS data is detailed in the following section 4.3. Matrix scaling is also used to calibrate the LIBS class probabilities with $p_{\mathrm{L},x,y,m_i} = \tilde{p}_{\mathrm{L},x,y,m_i} \cdot \boldsymbol{W}_\mathrm{L} + b_\mathrm{L}$. This also allows combination of both sub-classes (used and original) per material. After calculating the scaling parameters used in probability calibration, they get applied to both training and test data of HSI and LIBS data, respectively. The resulting confusion matrices after matrix scaling are displayed in Figures 5 (HSI) and 6 (LIBS). In this example the HSI confusion matrix shows a clear bias towards predicting HERC-highFe. However, it should be noted that the HSI classification model achieves significantly better performance when pixel-wise spectra are aggregated before prediction instead of using aggregated pixel-wise predictions. The final step in the pre-processing chain is scaling the probability maps and the surface quality map to a common size of 128 × 128 px.

Figure 5
Confusion matrix of the HSI classifier's output after applying probability calibration.

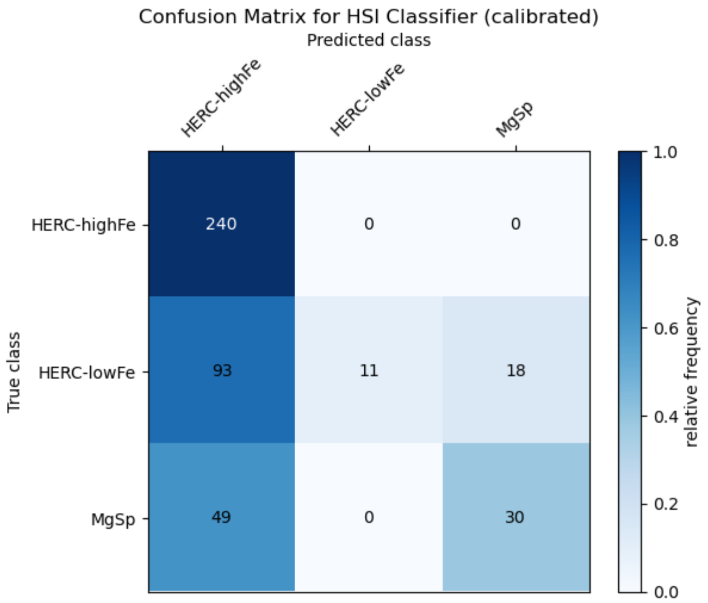

Figure 6
Confusion matrix of the LIBS classifier's output after applying probability calibration.

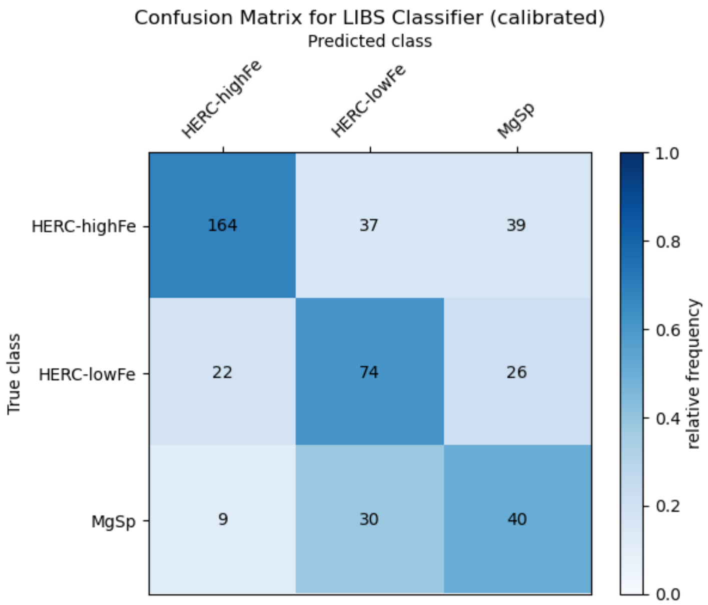

The uncertainty in material classification can be quantified using a concept from information theory, the Shannon entropy H, measured in bits. For this the now calibrated HSI class probabilities are averaged over all pixels of an object. From the resulting object-wise class probability, the entropy can be calculated with $H = -\Sigma_{n=1}^{3} p_{c_n} \log_2(p_{c_n})$, considering all 3 material classes c_n. When the classifier is certain of a prediction this would deliver $H = 0$ bits. In the worst case, when all classes are equally likely, this would deliver $H = \log_2(3) \approx 1.59$ bits. To give an intuition this can also be expressed as a confidence value with $c = 2^{-H}$. An entropy of 1.59 bits gives a confidence of 0.33, an entropy of 1 bit gives a confidence of 0.5, and an entropy of 0 bits gives a confidence of 1. The resulting entropy per object is displayed as histograms in Figure 7. All materials yield an approximately equally high entropy per object on average of 1.39 to 1.49, as indicated by the respective dashed line. This signifies a low classification confidence from HSI-pixels alone.

Figure 7
Histograms of entropy from HSI data, per object, partitioned per material class. The dashed black lines represent the mean entropy for that material class.

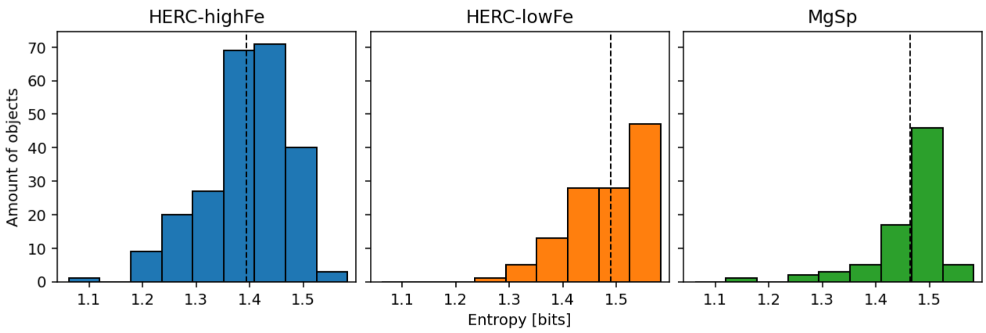

This entropy is the basis for calculating how many LIBS measurements are needed to obtain a classification accuracy as specified by the user. Each new measurement adds new information and therefore reduces the entropy. To quantify the possible information gain per LIBS measurement a supervised convolutional neural network (CNN) model is deployed that predicts from HSI probabilities p_H and surface quality data q_3D how a LIBS measurement at each possible ROI is expected to look like and how much it could reduce the entropy. The CNN's structural composition is displayed in Figure 8. It consists of an input layer, 5 hidden layers to compute the feature map, another 5 hidden layers and the output layer to obtain the predicted LIBS probabilities $\hat{p}_L$. This output layer has the same size as the inputs and represents a map of class probabilities as output by the LIBS sensor.

Figure 8
Structure of CNN used for prediction of pixel-wise LIBS class probabilities.

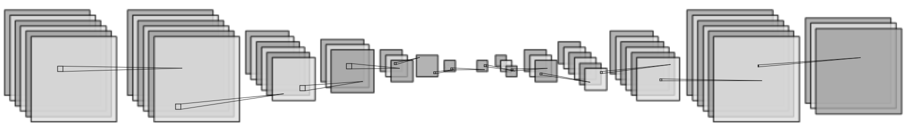

To overcome the challenge of having only a single LIBS spot measurement per object to train the CNN on, an extra convolution is introduced to 'blur' the predicted output using a Gaussian kernel. This way the entire 128x128 pixels area can be trained despite having only very sparse target data. The model is trained to minimize the Brier score (mean-squared error) between its predictions and the LIBS class probabilities as given by SpectraAnalyser.

Figure 9
Iterative choice of LIBS ROI based on most information gain per measurement for a HERC-lowFe sample. Probabilities are represented as mix of colours with red for HERC-highFe, green for HERC-lowFe, and blue for MgSp.

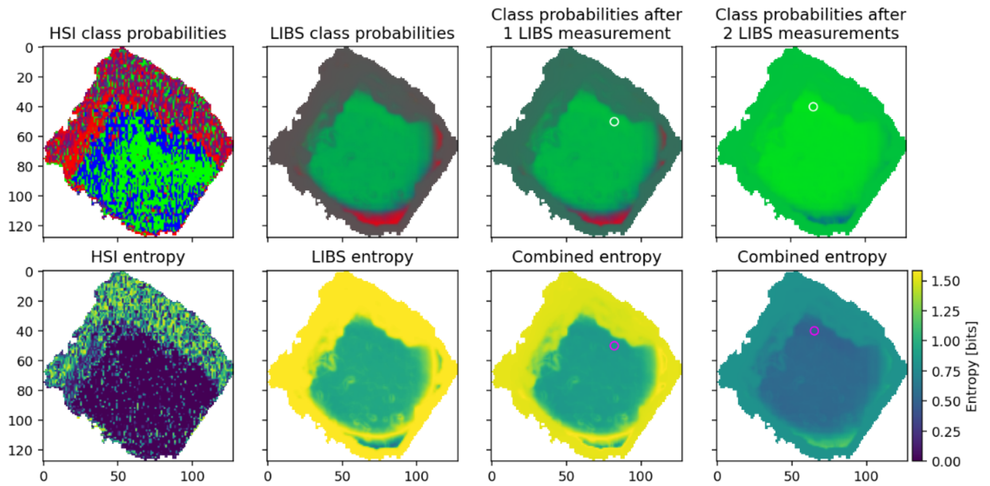

The model's predicted class probabilities for LIBS $\hat{p}_{\mathrm{L}}(p_{\mathrm{H}}, q_{3\mathrm{D}})$ can be seen in Figure 9, row 1 column 2. Here, the colours represent the different class probabilities with a mix of red (HERC-highFe), green (HERC-lowFe), and blue (MgSp). Grey therefore represents an equal distribution of predicted class probabilities. Calculating the pixel-wise entropy from the expected LIBS-predicted classes shows where the greatest information gain regarding object classification is to be expected. This can be seen in Figure 9, row 2 column 2. The yellow-greenish colour corresponds to high entropy. This suggests there is little correlation between these predicted LIBS class probabilities from HSI- and geometry data, on one side, and the actual LIBS classifier's predicted class probabilities on the other side. This is in line with previous analyses conducted in detail on specific samples as part of the project. Still, the pixel-wise entropy can be used as a guide for LIBS ROI choice.

To choose the optimal LIBS ROI spots for further object classification, again the class probabilities for an entire object are calculated from HSI data by averaging over all pixels as $p^{(0)}_{O,m_i} = \frac{1}{128^2}\sum_{x,y} p_{\mathrm{H},x,y,m_i}$, making sure to exclude any pixels of the background. This vector of class probabilities is then updated iteratively in the following steps. It is combined with the LIBS predictor's classification map, resulting in new class probabilities $p_{\mathrm{H,L_1},x,y,m_i} = p^{(0)}_{O,M_i} \cdot \hat{p}_{\mathrm{L},x,y,m_i}$ for each pixel (normalised over all classes), which can be seen in Figure 9 row 1 column 3. In this example, this step mainly reduces the predicted LIBS confidence in red areas. From this the pixel-wise entropy is calculated, displayed just below, corresponding to where the classification of HSI data would benefit most from a single LIBS measurement. Therefore, the position with the lowest entropy $\min_{x,y} H_{\mathrm{H},L_1,x,y}$ is chosen as LIBS ROI, as indicated by the purple circle. The expected class probabilities from the LIBS probability predictor at this position are finally combined with the object's running class probabilities as $p^{(1)}_{O,m_i} = p^{(0)}_{O,m_i} \cdot \hat{p}_{\mathrm{L},x_{min},y_{min},m_i}$ (normalised over all classes). From this class probability vector the object's classification confidence $c^{(1)}$ can be estimated as $c^{(1)} = 2^{-}$. These iteration steps are repeated until a user defined desired accuracy (ref. section 5.4) is achieved. The second iteration is displayed in Figure 9 column 4. Here, a larger reduction in remaining entropy is visible than what was achieved through the first iteration step. The resulting LIBS ROI spots are then sent to the LIBS control.

The predicted LIBS class probabilities can also be used to estimate the amount of LIBS measurements required to achieve an expected confidence. Each consecutive LIBS measurement adds information about the object, therefore resulting in a greater classification confidence. This change in confidence $c^{(n)}$ is visualised in Figure 10 for each sample of the training set. Here the class confidence is only about 35 % – 45 % from HSI data alone. A single LIBS measurement can raise the confidence to up to 86 %. In some cases, when HSI and LIBS disagree, it can even lower the confidence for the first few LIBS measurements, requiring many LIBS measurements to achieve a high expected confidence. The confidence for the brick sample used in Figure 9 is marked by the black dashed line.

Figure 10 Expected classification confidence for increasing amounts of LIBS measurements per object. Colours indicate the true material class. The dashed black line represents the sample detailed in Figure 9.

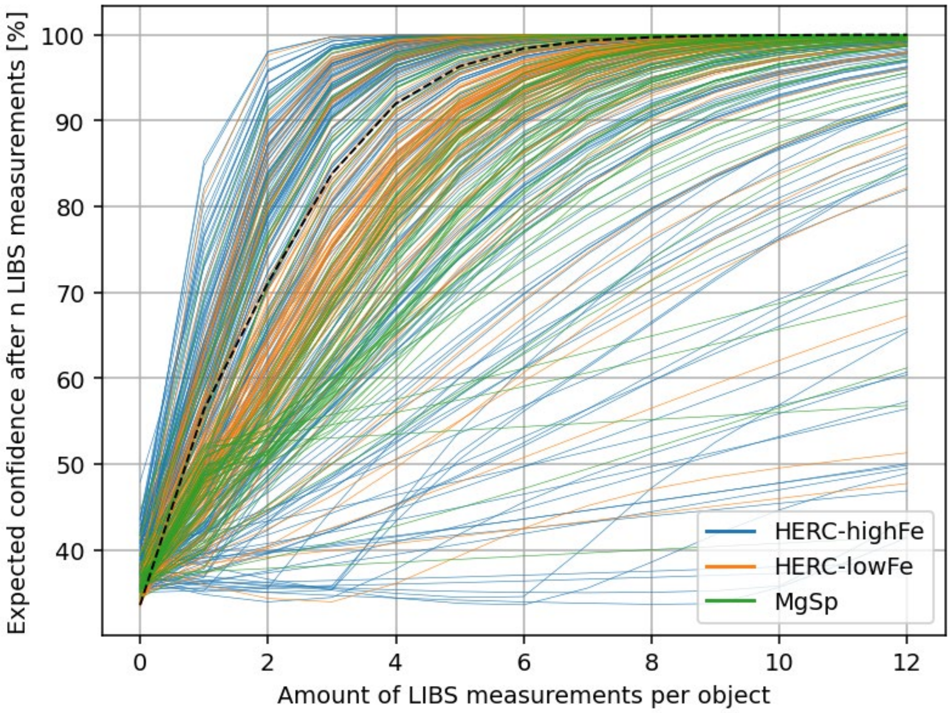

4.3. LIBS Sensor – SpectraAnalyser

LIBS data is acquired by capturing spot measurements on the object surface. To improve signal quality, each spectrum is integrated over multiple laser pulses. Segments within the spectra were defined based on characteristic emission lines of elements relevant to the material classes. These segments are integrated and normalized relative to adjacent background regions to enhance comparability. All segment values are treated as a feature vector, and scalar distances to predefined reference classes are calculated. For each spot on an object, a score is computed for all material classes based on these distances. A tuple of these scores is then passed to iMotion.

4.4. Sorting Decision – iMotion

For making a final sorting decision of an object, all LIBS classification results for that object together with the pixel-wise and object-wise predictions from HSI are being used. Like in training, HSI- and LIBS-classification data are pre-processed to obtain calibrated class probabilities. Then all class probabilities are combined into one object-wise class prediction, containing the determined probabilities for all material classes considered. They are combined by weighted multiplication, where the weights are optimized to minimize cross-entropy loss. Typically, the class with the highest probability is selected as the predicted material class.

To adjust the precision of the classification, the user can specify in Pulsate's system control GUI the desired "class purity" for each material class. This configuration is transferred to iMotion when starting the sorting task. It is used to adjust the classification threshold. If the probability of the selected class is lower than the threshold set for this class, the sample is discarded. This way the user can adjust the balance between recall against precision to achieve a required purity in the resulting sorted material stream. The final sorting decision for each object is finally transmitted to the material extraction system where robots or an air ejection system sort the objects into different bins.

4.5. Performance Validation

Validation data shows that through intelligent choice of LIBS ROI the precision for the final sorting decision can be raised to 92 % through a single LIBS measurement. The confusion matrix, displayed in Figure 11, shows a clear improvement over the previously discussed classification performances from HSI or LIBS data alone (Figures 5 and 6).

Figure 11
Confusion matrix of the final sorting decision for a single LIBS measurement at optimal ROI.

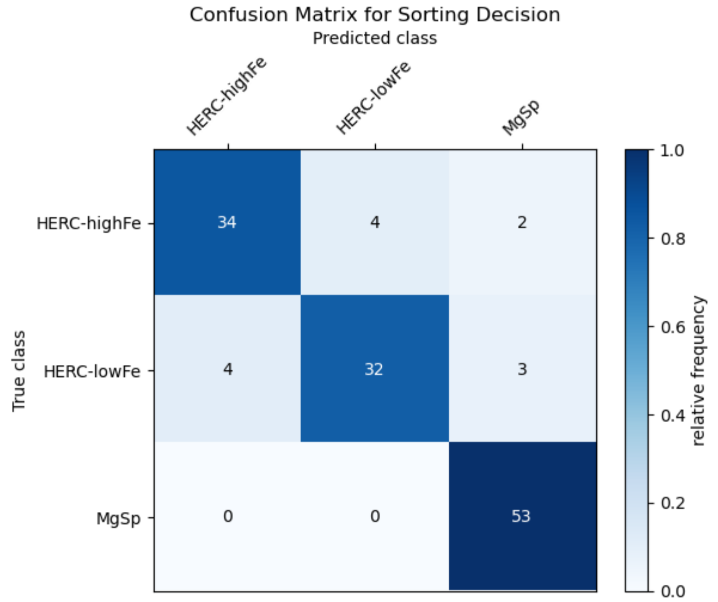

This is in comparison to a precision of 43 % from HSI data alone, 63 % when using a single, randomly positioned LIBS measurement, 70 % when using a single LIBS measurement located at best surface quality position, and 72 % when combining HSI data with a single randomly positioned LIBS measurement. The latter three metrics have been calculated by using test data from the initial data set by weighting each sample's influence on the overall precision with the distance between actual LIBS position and the respective target. Future revisions will show how multiple LIBS measurements influence classification performance.

5. User Interfaces for Classification Models

Adapting the ML models is required every time the input conditions change. Reasons necessitating a refinement of the models could be new material classes, changes in operating conditions, or technical changes for one of the sensors. In some cases, it may be sufficient to adapt the pre-processing of data. In others, re-training the models on new, unseen data may be required. To facilitate these operations there are several interfaces for the user to capture new data and facilitate data handling.

5.1. HSI sensor – Breeze

All HSI data handling is carried out in Breeze. This allows a user to intuitively capture new data and re-train the models on that data. Figure 12 shows Breeze's overview for data management, allowing the user to inspect, label, or analyse recorded data. To create the model, data is selected and divided into a training and test set. The user can then guide the model by selecting spectral segments relevant for the classifier, as shown in Figure 13. The resulting model is exported to Breeze Runtime format, where it can be remotely loaded by the iMotion software package.

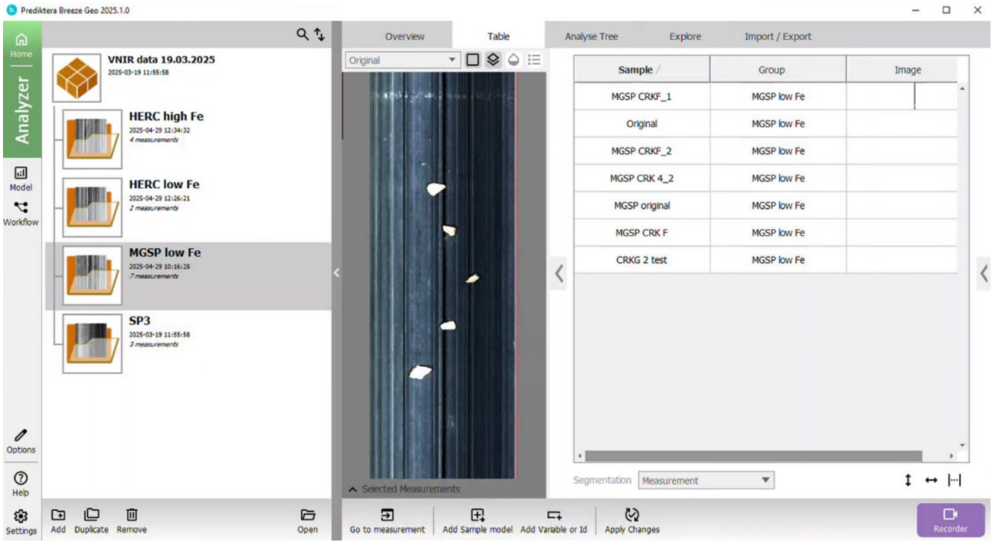

Figure 12
User interface in Breeze for managing recorded data.

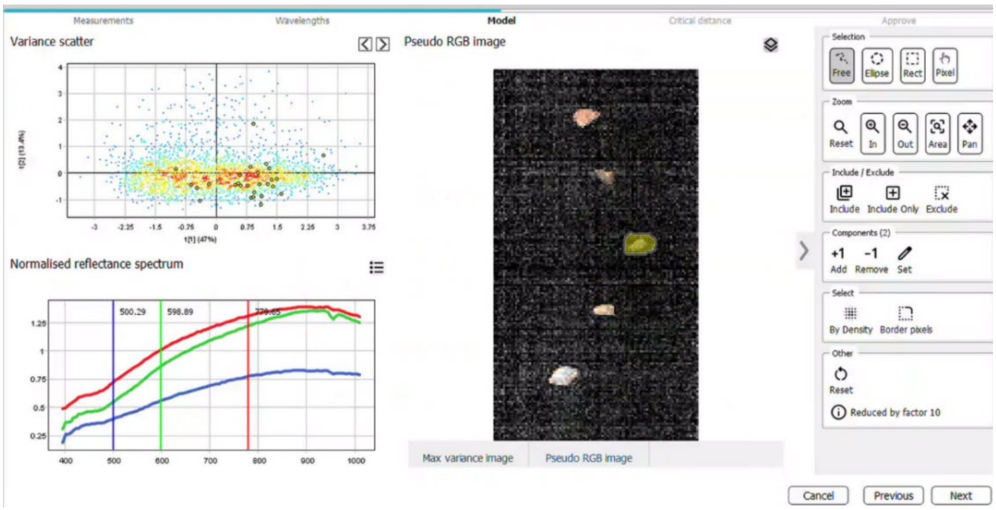

Figure 13
User interface in Breeze for training of the classifier.

5.2. LIBS ROI Selector – iMotion

A special training mode has been implemented in iMotion which allows capturing Breeze's HSI output, the topology for each object, and SpectraAnalyser's classification output from LIBS data. The user interface to select training mode and to specify the current material is displayed in Figure 14. The recorded data will be saved in a structured data store, annotated with the specified material. For generating training data, LIBS ROI are being generated by using the topology only. The object's surface quality for expected LIBS measurements is calculated and LIBS ROI are being sampled from it using the surface quality map as sampling probabilities. Compared to purely random sampling from an object this allows weighting the LIBS ROI to be in favourable areas, whereas random sampling often generates more ROIs near the object's edges, which are typically unfavourable. This ROI generation mode is chosen automatically when selecting to record training data.

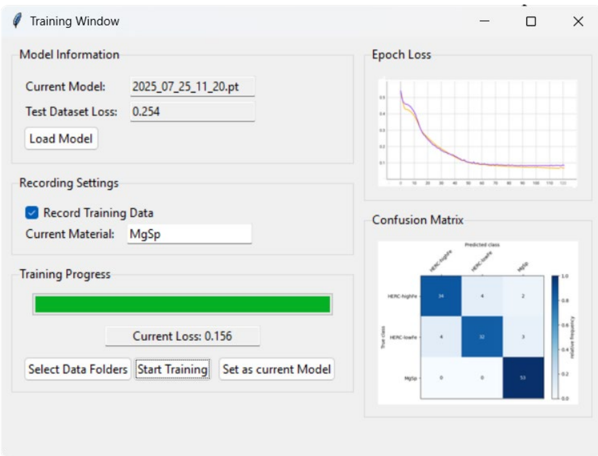

Figure 14
User interface in iMotion for training. Progress and performance can be verified through the loss display and the confusion matrix.

Upon start of training the selected training data is pre-processed as described in section 4.2 and split into a training and a test set. The training will continue until the loss on the test set does no longer improve. During training the epoch loss history is displayed for both training and test set, as depicted in Figure 14 top right. The confusion matrix in the bottom right gives an intuitive overview of classification performance. The resulting trained model is saved alongside the iMotion configuration for use in LIBS ROI selection.

5.3. LIBS sensor – SpectraAnalyser

To set up the classification in SpectraAnalyser, LIBS spectra are acquired with each material class measured independently under controlled conditions. For each class, all recorded spectra are averaged to obtain a representative reference spectrum, minimizing random fluctuations and enhancing signal stability. These class-specific reference spectra are subsequently incorporated into SpectraAnalyser’s “Classificator” component as distinct reference entries, as depicted in Figure 15. This procedure can be systematically repeated for all material classes. Analytical parameters, including the selection of relevant spectral segments and background normalization strategies, are configurable to ensure optimal adaptability to varying material characteristics.

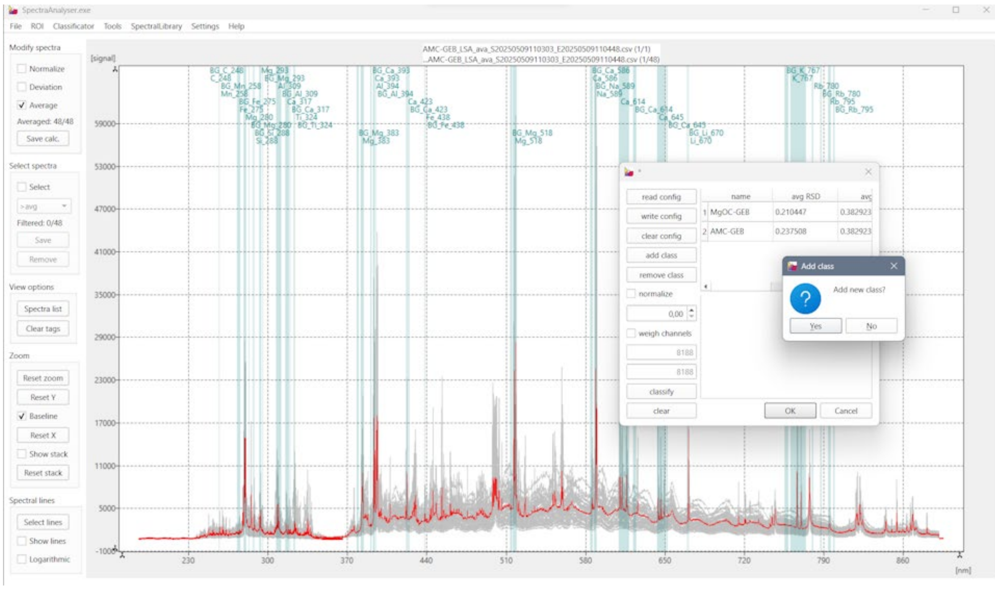

Figure 15
User interface for adding a new material class in SpectraAnalyser software.

5.4. Sorting Decision – iMotion

The user can define in Pulsate’s system control GUI into which physical fraction each material class should be sorted. Additionally, the user can specify the desired class purity for each material class. On machine start this configuration is automatically transferred to iMotion, where it is being used to determine the trade-off between precision and recall for each material class.

6. Conclusions

The development of AI-driven software modules for the combined operation of sensor systems in refractory material sorting has demonstrated significant advancements in real-time classification capabilities. By effectively integrating 3D, HSI, and LIBS sensors, the system enhances classification precision and material throughput, surpassing the limits of single-sensor operations. Successful validation of the software system within the Demonstrator A at the RHIM recycling plant in Mitterdorf underlines its practical applicability and effectiveness in industrial settings.

Future revisions will focus on refining the balance between classification performance and material throughput, ensuring that the system continues to meet evolving user needs. This will be done by acquiring a larger dataset with more material classes and with multiple LIBS ROI per object. Additionally, for HSI classification, it is planned to aggregate spectra of multiple pixels into "superpixels" which are aimed to deliver an adaptable trade-off between classification performance and spatial resolution.