X-RAY DIFFRACTION

AI-Assisted XRD Quality and Process Control in Cement Production

Application Note 632

This application note describes an AI-based add-on step that classifies an XRD scan by material type and then automatically starts the matching TOPAS method from the Cement Package Advanced. The goal is to make routine quantification faster and more consistent by reducing manual method selection.

Powder X-ray diffraction (XRD) is widely used for quality control in cement production. In routine operation, however, the result quality still depends strongly on choosing the right evaluation method and starting model for the measured sample.

Instead of attempting full phase identification across very large crystal-structure databases, we focus on a practical decision that directly impacts the downstream analysis: **How do we ensure the correct evaluation method is applied?** If the material type is known, the correct TOPAS template can be selected automatically and the standard refinement can run with far less manual input.

Key benefits (at a glance)

- Intelligent workflow guidance ensures consistent method application
- More consistent results across shifts and sites
- Compact AI model for easy deployment
- Fits existing TOPAS BBQ workflows off the shelf availability as Cement GUARD
- Prevents evaluation of invalid measurement data

The ultimate goal is to train a machine-learning model on indexed (i.e. labelled as a certain type of material), measured scans to recognize the material type and then route the scan to a matching TOPAS input file for high-quality quantification. This helps prevent mis-selections at analysis start (e.g., running a clinker scan with a cement method) and makes results more consistent across operators and sites. Throughout this application note, we use the term AI for readability, while noting that the underlying approach is machine learning based on neural networks.

The application target is the **Bruker's Cement Package Advanced**: a set of TOPAS methods for automated quantitative analysis along the cement process chain. It covers typical raw materials, intermediate products (hotmeal/bypass), clinker, and standardized cement types (CEM).

For method selection we use a simple two-step decision:

- **Step 1:** Is the scan a cement or not cement (raw material / hotmeal-bypass / clinker / fly ash / blast furnace slag)?
- **Step 2:** If it is cement, assign the CEM family (CEM I Portland Cement, CEM II Composite Cement, CEM III Blast Furnace Cement) and the Subclasses according to EN 197-1.

The machine-learning workflow was trained using data from European cement plants and therefore uses European cement standards in the examples shown. However, the underlying workflow logic is not tied to a specific norm system and can be transferred to other regional classifications without fundamental limitations.

AI approach

The AI model reads the measured diffraction pattern as a standardized 1D signal. Before prediction, the scan is processed in the same way for every sample (background correction, normalization, and rebinning to a fixed grid). The neuronal-network model uses supervised learning techniques, then outputs a material-type class label (Step 1: cement vs. not cement; Step 2: CEM I/II/III and its respective sub-classes). The final model layer is simply configured to match the number of classes that are required for the intended workflow.

Training set

In total, we analyzed **over 320,000** measured XRD scans from **nine plants**. Labels were available from factory indexing, and we additionally re-labelled the scans based on the TOPAS results to double-check indexing quality and to apply consistent acceptance criteria.

How performance is measured

Performance is reported mainly using the **F1 score**, a standard metric for classification that combines **precision** (how many predicted assignments are correct) and recall (how many true items are found). It ranges from 0 to 1 (or 0–100%); **1.0** means no missed assignments and no false assignments. As a practical guideline, values **>0.95** are usually considered very good for stable, well-separated classes, while values around **0.80–0.90** can already be useful when classes overlap or when training data and labels are noisy.

To assess **robustness**, we repeated the evaluation with multiple random splits of the available scans. In general, the dataset was randomly partitioned into **80/10/10%: teaching (training)** data to fit the model, a **validation** set to tune settings and monitor overfitting during training, and an independent **test** set kept aside for final performance reporting. Potential dependencies on **client-/plant-specific (factory) data** were additionally checked by selectively eliminating or combining data from specific factories and/or classes and confirming that performance remained stable.

In addition, we used **5-fold stratified cross-validation** for key results. Here the dataset is split into five equally sized parts ("folds") while keeping the class distribution similar in each fold (stratified). The model is trained five times, each time holding out one fold for testing and using the remaining folds for training. Using **five** folds is a common compromise: it provides a more stable estimate than a single holdout split, while keeping runtime and compute effort manageable compared to higher fold numbers. Cross-validation therefore gives an estimate of how well the model is expected to generalize to new scans and indicates how sensitive performance is to the particular split (i.e., variability between folds).

We compare the AI predictions against two reference label sets: (i) the **factory indexing** available with the measurements, and (ii) a harmonized **TOPAS-based indexing** from our own re-refinement. This shows both agreement with the production labels and agreement with a consistently refined reference.

Results: screening, quality control, and accuracy

Screening the training data showed that a noticeable share of scans should be removed before training. After screening based on initial classification and the quality of subsequent TOPAS refinement, over 300,000 scans were retained for model training. These removals are important because they are often linked to handling or labeling issues and can otherwise bias the model. We grouped the main reasons for rejection into three practical filter types:

1. **Corrupt scan files:** incomplete or damaged measurement data that cannot be evaluated reliably.

2. **Sample and labeling issues:** preparation/loading problems (e.g., height errors exceeding limits, broken sample pucks, upside-down mounting) and wrong labels in factory indexing (e.g., clinker labeled as cement).

3. **Refinement quality rules:** we used simple thresholds to sort scans into remove, low, and high quality. Criteria included R_{wp}, angular offset, alite crystallite size, alite preferred orientation, and—specifically for cement—total sulfate (gypsum + hemihydrate + anhydrite) above a minimum threshold.

After applying the above 3 filters a distribution of retained scans that establish the final classification model is:

- **Step 1 classes** (ballpark): raw materials (limestone + gypsum + fly ash + blast furnace slag) ~6,500; hotmeal/bypass ~25,000; clinker ~135,000; cement ~150,000

- **Step 2 (cement) classes** (ballpark): CEM I ~53,000; CEM II ~72,000; CEM III ~7,000.

A clear trend was that manually handled measurements showed more outliers and mislabels than scans produced under standardized, automated routines. This supports the need for automated screening and AI-based routing to further reduce operator-dependent variability.

Classification performance

For Step 1 (cement vs. not cement) we achieved F1 = 99.98%, detailed results are shown in table 1 and deviations from class assignement are mapped in Figure 1.

The errors are mainly explained by wrong sample identification at the point of measurement. The most common mix-ups are **hotmeal vs. clinker** and **clinker vs. cement**. The number of AI errors is negligibly small.

Table 1
F1 prediction accuracy for cement vs non-cement classification (Step 1) based on the AI-predicted vs TOPAS ground-truth label.

	F1-score / %	Candidates
Clinker	99.98	134.197
Hotmeal	99.93	25,819
Limestone	99.89	5,893
Gypsum	99.37	632
BFSlag	98.70	38
Flyash	97.56	21
Cements	99.99	149,443

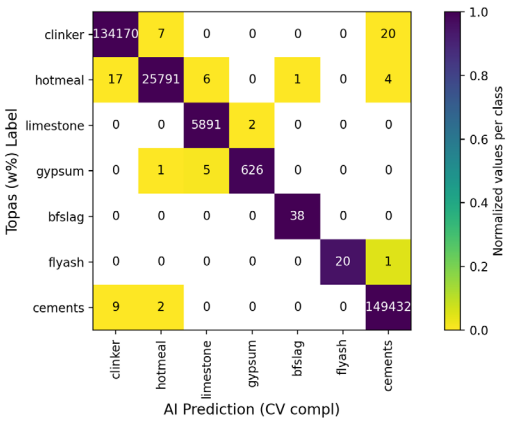

Figure 1
Correlation chart showing the distribution of false assignments at the off-diagonale fields.

Table 2
F1 prediction accuracy for cement classification (Step 2) based on the AI-predicted vs TOPAS ground-truth label and factory labels.

Cement type	Additive / wt%	F1-score / %		Candidates
		AI label	Factory label	
CEM-I (Portland cement)		99.30	91.75	53,496
CEM-II/A-L (Portland Limestone cement)	Limestone 6 … 20	99.40	91.67	51,917
CEM-II/B-L (Portland Limestone cement)	Limestone 21 … 35	97.09	86.72	1,100
CEM-II/A-S (Portland- Blast furnace cement)	BFSlag 6 … 20	92.46	62.61	6,077
CEM-II/B-S (Portland-Blast furnace cement)	BFSlag 21 ... 35	95.18	79.54	10,794
CEM-III/A (Blast furnace cement)	BFSlag 36 … 65	97.24	85.94	6,578
CEM-III-B (Blast furnace cement)	BFSlag 66 … 80	89.01	74.48	466
CEM-II/A-V (Portland Flyash cement)	Flyash 6 … 20	98.26	82.64	11,411
CEM-II/B-V (Portland Flyash cement)	Flyash 21 … 35	97.46	59.09	7,593

Step 2 (CEM I vs. CEM II vs. CEM III) is less accurate than Step 1 (table 2), as expected. Across the overall task, the AI achieved **98.44%** accuracy. For comparison, the factory indexing agrees with our TOPAS--based re-indexing at **87.44%**, indicating a relevant baseline level of labeling inconsistencies in routine operation, and illustrating the performance margin achievable under highly standardized conditions. Figure 2 shows the distribution of false assignments of the AI step vs. the TOPAS ground truth and factory labels.

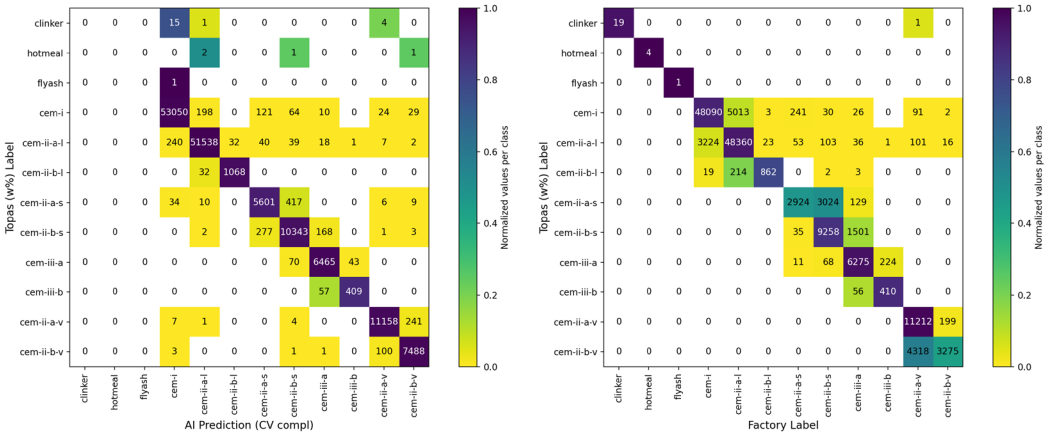

Figure 2 **(a)** **(b)**
Correlation charts showing the distribution of false assignments at the off-diagonale fields for AI prediction vs TOPAS ground truth (a) and factory label vs TOPAS label (b).

Most remaining uncertainties are linked to (i) small sample numbers in some subclasses and (ii) scans close to class borders (Figure 3). For slag-related cements, these borders correspond to typical slag addition limits (about **5**, **20**, **35**, and **65 wt%**), where real production can naturally move along gradual transitions.

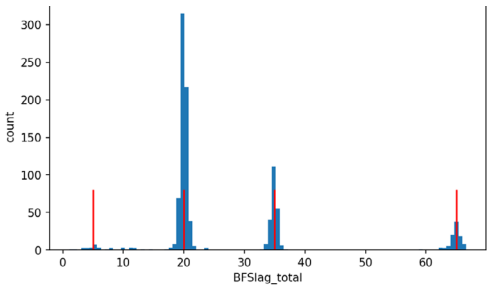

Figure 3
Histogram of Blast Furnace Slag (BFSlag) cement samples, misassigned by the AI as a function of slag concentration. The red lines indicate the class borders according to EN 197-1.

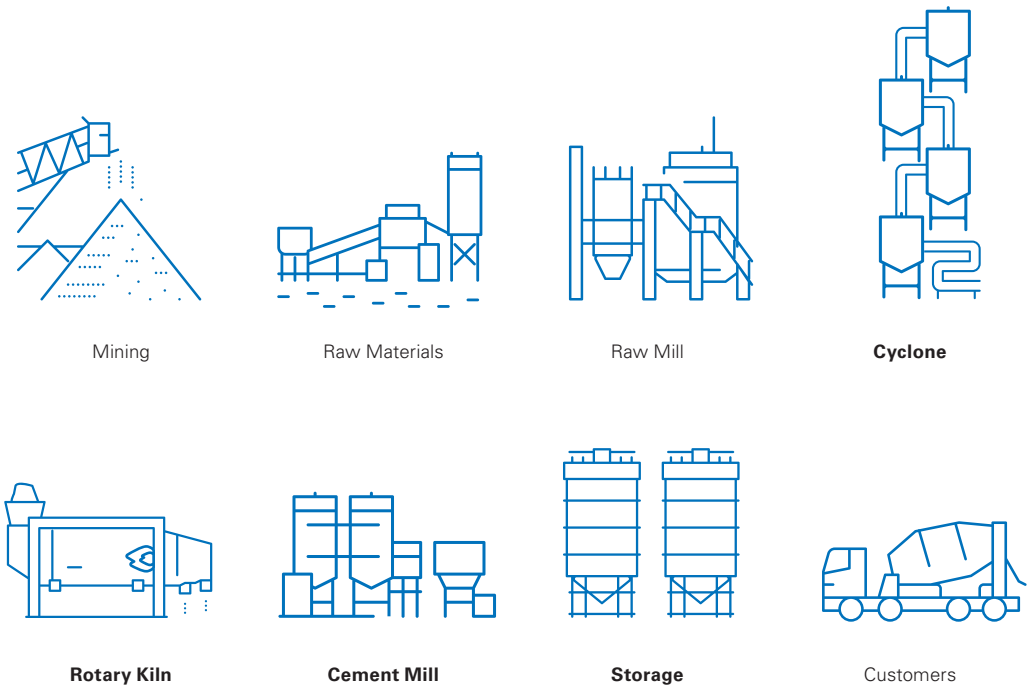

Figure 4
Graphical representation of cement manufacturing. Process steps printed in bold are covered by materials classification with the Cement GUARD package.

Workflow implications and recommendations

The filters established during dataset curation to stabilize the AI training model can now also be applied routinely in **DIFFRAC.BBE** workflows: scans that are corrupted, poorly prepared, or otherwise outside acceptance criteria can be detected and excluded automatically in unattended quality-control runs before starting the automated TOPAS evaluation. This prevents poor or misleading refinements from propagating into downstream reporting—especially avoiding that low-quality evaluations are written into the factory **LIMS** system.

In particular, **Step 1** classification (cement vs. non-cement) proved to be **extremely robust** and reliably separates cement from non-cement scans even in the presence of occasional operator labeling inconsistencies. This enables the AI to automatically select the correct **TOPAS input file** (template) in the **DIFFRAC.BBE** automated workflow, ensuring that each scan is routed to the appropriate evaluation method from the start.

For **Step 2** (classification of **cement classes and subclasses**, i.e., CEM families and DIN 19721 subclasses), the proposed AI-assisted workflow acts as an automatic **labeling and routing** layer before the standard **TOPAS Rietveld refinement**. Compared to factory indexing, this yields an overall improvement of about **~10%** in correct labeling, thereby directly reducing operator-dependent mislabeling and wrong method selection. Importantly, even if the AI occasionally assigns a **blast furnace slag (BFS)** cement candidate to a neighboring cement class close to a class border, the subsequent TOPAS refinement still analyzes the scan with the appropriate quantitative model and typically converges to the correct result.

The above discussed workflow items from filtering to Step 1 and Step 2 classification are implemented in the Cement GUARD addition to the Advanced Cement Application pack. Looking along the production chain of Cement, depicted in Figure 4, Cement GUARD covers material classification at the Cyclone, Kiln, Cement Mill and Storage steps.

For best robustness, the model benefits from broad coverage across plants and operating conditions. Synthetic augmentation by combining real measurements can help in edge cases, but the strongest results come from a sufficiently large and well-curated training set.

We therefore recommend a curated, multi-site training set; in our case the pool spans **over 300,000 scans**. We tested robustness by systematically removing subsets of data and found the model to be stable. As a practical guideline, around **~100 scans per covered class** are useful to integrate new material into an existing class, and around **~500 scans** are recommended when introducing a new class.

Outlook: deployment in the TOPAS BBQ workflow

The main one-time effort is running the training and validation pipeline, which can take **hours to days** depending on the class scope and dataset size. The trained neural network itself is **very compact** and can be integrated into the **TOPAS BBQ** workflow with minimal overhead. Within the **DIFFRAC.BBE/BBQ** architecture, the same pipeline can automatically screen a scan, classify an unknown sample, and launch the matching TOPAS method. This turns the **Advanced Cement Package** quantification into a single, consistent workflow that can run at the customer site directly after measurement (or at any later time) while reducing operator-dependent method selection errors.

The workflow is flexible in how the AI model is deployed and maintained: it can be operated with a **fixed neural network** that is trained once and then used unchanged for routine classification; it can be **upgraded** by periodically retraining the curated **global model** as additional (and newly curated) plant data become available; and it can also be provided as a **customer-specific** solution, trained on **proprietary customer data** to reflect local raw materials, process specifics, and labeling conventions.

Credits

This work was enabled by anonymous donors of training data from leading international cement producing companies. Their contribution is essential to achieving scientifically sound progress by providing FAIR (findable, accessible, interoperable, reusable) data to the community.

This is a joint investigation by Bruker scientists and developers, together with the Karlsruhe Institute of Technology (KIT). We especially mention André Orth (Institute for Automation and Applied Informatics) for his contributions.

Bruker AXS is continually improving its products and reserves the right to change specifications without notice.
Order No. DOC-A88-EXS632 © 2026 Bruker AXS SE.

Bruker AXS
info.baxs@bruker.com

bruker.com

Worldwide offices
bruker.com/baxs-offices

Online information
bruker.com/xrd-software