AI: NOW UP AND RUNNING

alcemy, Carbon Re, and CemAI, discuss how AI/ML optimisation improves reliability, product quality, and energy efficiency at cement plants.

The digitalisation of cement manufacturing in Industry 4.0 has brought about a sea change at the plant. The science of AI/ML applied to operations through digital solution modules can be implemented seamlessly so end users of the technologies achieve tangible benefits in the short and long-term.

The following article relates the history, integration, and current status of AI controls in cement production through the accounts of three industry-leading actors. CemAI, Carbon Re, and alcemy head a cascade effect of digitalisation in the industry; their combined experiences shed light on this ongoing transition.

Complex analysis of the operating parameters of all equipment and processes, coupled with the ability to analyse differences in real-time operational data and build scale with use, define the current standard for 'smart factories'. Available digital solutions produce the most effective and economic use of plant equipment, increasing reliability, maximising productivity, and generating savings. A recent online Zipdo Education Report entitled 'Digital Transformation in the Cement Industry – Statistics' published in May of 2025, states, "Sixty-seven percent of cement companies reported increased operational efficiency due to digital transformation initiatives."[1]

Digitalisation/Industry 4.0

Digitalisation has been a movement in cement production for well over a decade. Using the terabytes of data generated annually by a cement manufacturing plant, the evolution of acting on real-time information began in earnest. An average cement plant produces over one terabyte of data over a year, the equivalent of 6.5 million pages of documents.

By the advent of Industry 4.0 in cement production in the 2010s, digitalisation of the data produced was revamping control rooms and the industry standards for process control. Rapidly, the sensors improved – wireless, sensitive devices – capturing more subtleties in temperature, vibration, and other indices of operational status. The digital data permitted managers to preserve and analyse results in real time, resulting in better process control through reactive responses to data changes. The ability to store and analyse data quickly led to the creation of digital 'twins' (models). The smallest delta in the results could now be recognised by the new AI/ML system software's analytic capacities.

The ROI for the 'smart factory' software solutions is compelling. The AI blueprint for the gold standard cement plant exists now and is gaining impressive scale through use by early adopters. Investment in digital transformation in the cement sector is projected to reach US$5 billion globally by the end of 2025.[2]

The low-hanging fruit for commercial applications came in AI/ML systems impacting plant maintenance. The enhanced data streams and analytics quickly addressed deviations from the models and revealed a new correlation between functions heretofore neither appreciated nor understood (Figure 1). A single case of early maintenance intervention from AI analytics could easily produce six-figure savings in downtime and reduced machine life. Consequently, 45% of cement manufacturers have implemented IoT solutions to monitor real-time equipment performance and adoption of AI-driven predictive maintenance in cement plants has increased by 30% over the last three years.[3, 4]

Digitalisation capacity has scored dynamic growth of late. Better data collection, more powerful processing and storage, and scale-building from informed use with AI/ML optimisation modules flipped a light switch on the 'smart factory' concept in cement-making.

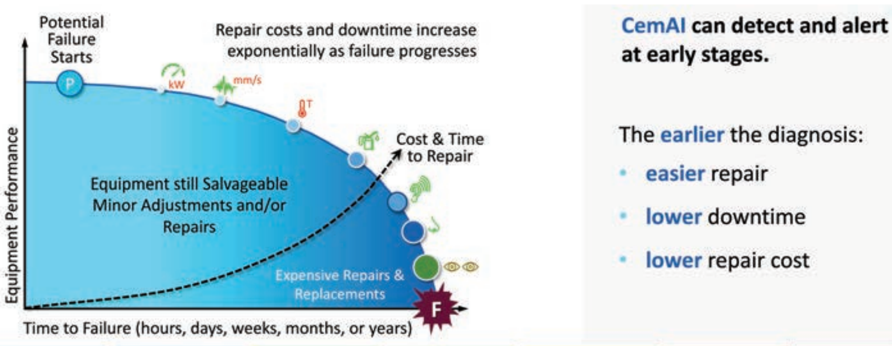

Figure 1. Equipment 'signals' when performance starts deteriorating.

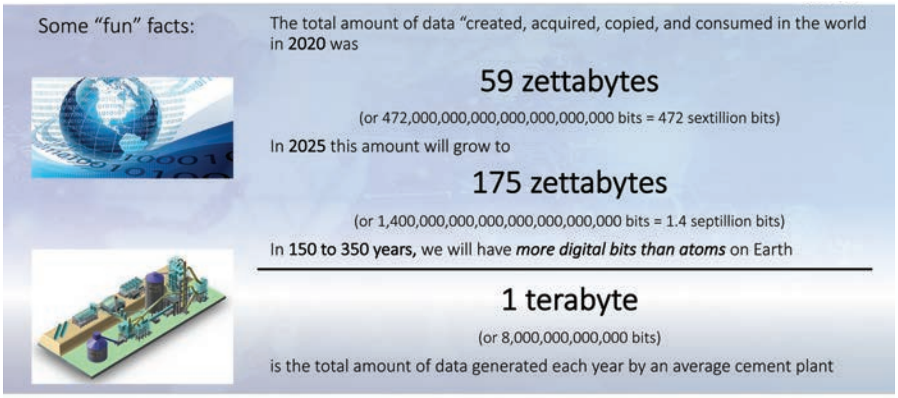

Figure 2. 90% of the world's data today has been created in the last 10 years.

The digitalisation revolution has enabled a never before available understanding of the manufacturing process. Operating data in terms of sensor recordings of temperature, vibrational status, lubrication parameters, and other plant conditions is now widely obtainable and capable of preservation and study.

If the plant team considers one terabyte of data each year (or two million individual computer files) to be the challenge, then the question becomes how the plant exploits the full value of that data (Figure 2). Optimisation in cement making means intelligent control of a complex system.

Optimising software enables plant goals to become control directives. Engineers feed the AI software objectives – fuel cost, or alternative fuel (AF) use, for example. Using the digital model 'twin' to simulate the process permits evaluation of all possible setpoints against the objective. The controller recommends the target which best achieves the priority. Constraints and engineering judgments offer further fine-tuning.

Pyroprocess optimisation

Carbon Re focuses on the pyroprocess – an extremely complex and carbon-intensive area of manufacturing. Improved control at the kiln delivers real gains for producers with significant cost, quality, and carbon emission benefits. Carbon Re's AI tames pyroprocess complexity by modelling the process and training models on high-quality operational data.

Training an AI model requires large volumes of high-quality data to learn the relationship between process inputs and outputs. In the pyroprocess, however, data can be inherently unreliable: the environment is high temperature and often dusty, and sensors can drift or fall. For this reason, the first step is rigorous data cleaning. Incorrect readings can confuse the model, so outliers must be removed. For example, in gas analyser data, cleaning cycles are filtered out to ensure the model is trained only on true measured values.

A high-quality model also requires deep embedded domain knowledge. This expertise is used to select which signals should be used as model inputs – for example, which temperature measurements are most relevant – and to determine the most appropriate type of AI model for the application. In pyroprocess modelling, the most important inputs typically come from chemistry samples and key process readings. Carbon Re also incorporates control data, including plant control limits and labels informed by local experts, ensuring that the models reflect real operational constraints. Put simply, the models are used to generate predictions for desired process outcomes or to recommend setpoint values. When the models are first up and running, Carbon Re analytics offers automatic oversight as needed. In this fashion, the models adapt to changing plant conditions seamlessly.

In the pyroprocess, many factors are difficult to measure continuously but are very important to controlling the process. Software sensors, or 'soft sensors', provide real-time estimates of these physical variables. In the clinker quality area, soft sensors are used to predict key parameters such as free lime, litre weight, and C3S in real time. By providing predicted values, soft sensors enable proactive intervention rather

quality view demo example gives a screenshot example of the precision of the simulation and the degree of control achieved (Figure 6).

Confidence is grown through the steering adherence screen/view of setpoint results. With prediction accuracy significantly exceeding expected correlation to actual results, the move to closed loop is made. The graphics provide compelling proof of reduced variability from strength targets through the alcemy digital processes. On average, the alcemy cloud optimising system provides a 30% reduction in deviation from target values and that translates to higher throughput – in the range of up to 2.3% – and energy savings from grinding coarser.

"The cement grinding process determines the cement's 'genoma': its strength development, workability, durability, and ultimately its carbon footprint," says Oliver Kanders, General Manager of alcemy. "It is the cement grinding processes that enable CO_2-reducing strategies like clinker substitution to be realised."

In essence, alcemy closes the transparency and traceability gap between cement and concrete, aligning both sides around complete data sets, predictive strength models, and clinker-optimised recipes. While this abstract focuses on the cement system, there is a distinct advantage in extending this approach across both product lines: optimising performance, cost, quality, throughput, and CO_2 emissions – collectively defined as key optimisation parameters. Every tonne of cement leaving the plant is verified against performance predictions, supports higher SCM substitution, and positions concrete plants for reliable, low-carbon outcomes.

Figure 5. Control room operator view.

Figure 6. Control room quality view.

Plants start up quickly with AI control systems

CemAI can introduce its AI software modules – 'CemAI Predictive Maintenance' and 'CemAI Optimize' – into any cement plant's control room and start the processes based upon whatever sensors and data currently exists.

The control room serves as the central integration hub of the digitalised cement plant for the entire production process. From a vast network of sensors, instruments, and subsystems across the entire plant, the control room enables operators to monitor all operations from raw material preparation to final product. It has evolved from a simple monitoring station to an intelligent, data-rich environment, critical for a cement plant to realise the benefits of digitalisation, including enhanced reliability, profitability, and safety.

Predictive maintenance seeks to recognise a potential failure before it fully manifests itself, leading to greater impacts on the production process and on the repair or replacement of equipment. Process optimisation using AI/ML drives the process toward maximised throughput and improved energy efficiency.

Complex analysis of the operating parameters of all equipment and processes, along with the ability to analyse differences and dependencies in real-time operational data, underpin both CemAI solutions in the cement manufacturing plant.

CemAI solutions – web or cloud-based – are housed separately and securely. It is preferred to keep the data on virtual modules in a secure cloud, but for customers who believe security is best on premises, CemAI can easily accommodate their choice.

The start-up – including periodic audits of available sensors and recommendations, installation of new sensor process, and equipment mapping – generally occurs over three months. Even legacy plants with a minimum number of sensors can get started. CemAI implementation will point operators to their blind spots and call out the best places to put the next 50 sensors, for example, to gradually build up their digital footprint.

All equipment operates on the potential and functional failure curve, and it is CemAI's predictive maintenance solution that increases the window between a potential failure and a functional failure. Combining vast amounts of information and condition monitoring, the highly complex world of digital predictive maintenance becomes an easy-to-use, day-to-day tool for maintenance, inspectors, and planning.

The predictive maintenance solution covers the entire plant process from primary crusher all the way through to cement loadout and packaging. The CemAI team creates a full end-to-end model of the cement plant, including all processes, equipment, and analogue signals. The model is trained with the historical data of the machine. Learning begins immediately. From the moment the system is online, live data is sent to the algorithms for analysis. The process is periodically audited to determine what additional sensor can be brought online.

The full plant model is critical to the operation of the CemAI solution, as each process may affect another. This end-to-end set up of dependencies and hierarchies helps alert to the smallest anomalies. As the system sees more real-time data, the algorithms will use the new data to create new ranges of operation.

The initial mapping is created by CemAI's engineering and data science teams and the validation of the model is carried out in conjunction with the plant operation team. The alerts generated are shared directly with plant personnel and the CemAI service team. The CemAI service team makes recommendations on what may be causing the alarms with potential course of action for plant personnel. The mapping creates hierarchies for digital plant models, defines relationships, and signal interactions. This process takes about 8 to 12 weeks.

End-to-end optimisation, through the CemAI Real-Time Optimizer module, enables production processes to be directed toward achieving a selected objective while balancing the constraints required for other correlated operations. This installation also takes approximately 8 to 12 weeks. CemAI's service includes 24/7 expert consultations that coordinate with local teams for insights and remedies to any digitally perceived anomalies anywhere along the cement making process.

A big part of the success CemAI has seen with its end users, in over 27 predictive maintenance and 32 real-time optimiser installations, is the buy-in that has almost uniformly taken place. "The process works best when the client has a 'champion,' a manager or executive of the plant who sees the value and urgency of installing our AI systems and what they can accomplish," says Scott Ziegler, CEO of CemAI.

There should be little debate on starting down the path of AI/ML software tools being introduced to the control room. The spectre of older workers fearing being made redundant is a red herring. On the contrary, seasoned staff bring invaluable, deeply embedded knowledge that is essential to building effective models, and as a result, local teams typically welcome the interactive CemAI installation process.

At a time when plant workforces consist of both younger employees and long-serving personnel, this is an ideal opportunity to preserve a wealth of operational knowledge. AI systems are user-friendly in the control room or when accessed remotely, and robust security controls are in place to prevent unauthorised access.

Conclusion

Having existed since the time of the Roman Empire, cement and concrete manufacturing continues to have bright outlook in the 21st century due to AI. The digitalisation of cement manufacturing and the analytics provided by AI optimisation systems – end-to-end in operations – are yielding transformative results in efficiency, energy use, and product quality. ■

References

1. ESER, A., 'Digital Transformation In The Cement Industry Statistics', (2025).
2. Ibid.
3. Ibid.
4. Ibid.

Additional information

alcemy, a Berlin-based deep tech company, provides AI-powered predictive quality control software & sensors for cement and concrete production. Its system reliably supports real-time optimisations of clinker factor, throughput, and CO_2 emissions through fineness and recipe adjustments, despite changing raw materials and the increased use of AFs.

Carbon Re, London, UK, is an industrial AI company, spun out of two of the UK's leading universities, Cambridge and University College London (UCL). Focused exclusively on cement, Carbon Re's AI operates in closed-loop control to optimise the pyroprocess cutting costs and carbon emissions.

CemAI, Norfolk, VA, is a cement-centric concern featuring digital analytic products, e.g. Predictive Maintenance and Process Optimiser, covering cement plants end-to-end. An affiliate company of TITAN Cement Group, CemAI combines deep cement manufacturing knowledge with advanced software and IT infrastructure technical expertise.