

How AI is helping cement plant operators reduce energy consumption and carbon emissions

Insights from optimising production in cement plants by applying the latest AI technology and simulating 100 years of training

Who we are

We are a company pushing the boundaries of artificial intelligence to accelerate the decarbonisation of cement and other foundational materials. Carbon Re's innovative AI-powered software optimises cement production, specifically targeting the pyroprocess stage to reduce fuel-derived carbon emissions by up to 5%.

Our solution integrates seamlessly with plant Advanced Process Control (APC) systems like ABB Ability[™] and FLSmidth ECS/ProcessExpert[®], using AI models that continuously adjust in closed-loop control to optimise fuel use and manage fuel-mix variability.

By integrating with Carbon Re, plants can utilise a broader set of their process, laboratory and chemical data. Our advanced machine learning models enable real-time, dynamic optimisation of process targets, automating repetitive manual tasks and allowing process engineers to focus on more impactful work.

Carbon Re requires no capital investment, new equipment, or plant shutdowns. It supports ongoing plant optimisation, adapting to changing inputs and external pressures such as volatile fuel costs and emissions regulations. The result is significant energy savings, allowing operators to run plants at peak efficiency and realise substantial cost reductions.

Find out more at carbonre.com/product/

Published by Carbon Re, February 2022

Summerbell, D, O'Sullivan, A and Iftikhar, M. (2022) How AI is helping cement plant operators reduce evergy consumption and carbon emissions: insights from optimising production in cement plants by applying the latest AI technology and simulating 100 years of training. Carbon Re, London.

To see other publications by Carbon Re, please visit carbonre.com/resources/

Contents

Executive summary	2
A 'digital twin' to optimise cement production	3
AI & reinforcement learning	5
100 years of training in a day	5
The complexity of alternative fuels	6
Simplifying the introduction of new technology	6
Emerging secondary benefits and future direction	
FAQs	9
Appendix 1	10
Appendix 2	12

Executive summary

Cement production is a remarkably complex process with ever-changing inputs (fuels, raw materials), conditions (state of equipment, shift changes), and competing priorities (throughput, control limits). Great cement plant operators take many years to learn how to get the best out of their plants through controlling interrelated systems and conflicting requirements to deliver product quality, whilst maximising throughput.

Advances in artificial intelligence (AI) have had a transformative impact on many industries, including finance, healthcare, automotive manufacturing and technology, bringing new capabilities and opportunities for innovation. However, adoption in the cement industry has been slower due to a lack of tools specifically tailored to the needs of plant operators.

While generic modelling tools have been available to the industry for some time, AI technology can now develop a 'digital twin' of an individual plant by analysing the sensor data gathered on a minute-by-minute basis by Distributed Control Systems. This creates a high-resolution digital model that captures the properties of a given kiln (the thermodynamics, and the physical processes of clinker production) and delivers an accurate representation of the real world performance in that specific plant. This digital twin provides a platform for AI 'agents' to learn complex process control in a sophisticated simulation. The outputs are insights on how to maximise the efficiency of the plant that enable operators to adapt to the continually changing plant conditions.

The branch of AI that powers these insights is called 'deep reinforcement learning'. This is an area of AI that has seen considerable breakthroughs in recent years, demonstrating the capability to exceed human performance in highly complex strategic planning tasks.

Applying digital twins and reinforcement learning to fuel use in cement production has enabled us to train AI agents to explore and find the optimal parameters of the production system concerning one or more goals (e.g. minimal fuel cost, maximal throughput), which are then applied to the real plant.

Carbon Re's software platform builds 'digital twins' to support operators in managing the complexities of cement production by providing them with clear, actionable recommendations for process setpoints. This leads to significantly reduced energy consumption and CO₂ emissions.

We know that the specific heat consumption of one ton of clinker can significantly vary from one day to the next and is affected by the interactions of different raw materials, operational settings, conventional and alternative fuels, and Delta specific kiln properties.

Carbon Re's AI platform enables the 'digital twin' technology to capture these complex relationships, and the platform helps our operators by recommending optimal plant parameters to optimise fuel use while maintaining throughput and clinker quality.

A 'digital twin' to optimise cement production

Cement production is a complex process that requires balancing enormous input parameters affecting cost, quality, and reliability. Operators have to manage these conflicting requirements, which is why their job is both highly skilled and demanding. To date, software solutions supporting operators to deal with this complexity have relied on model predictive control methods and rule-based Expert Systems. While these systems can maintain operations within desired limits, they lack the dynamic qualities needed to deal with changing plant parameters or improve performance beyond a narrow operating window.

Combining the research on cement production, carried out by Dr. Daniel Summerbell at the Institute for Manufacturing at the University of Cambridge, with the expertise in artificial intelligence and machine learning from Professor Aidan O'Sullivan, at UCL, Carbon Re has developed, an AI-powered software platform that can optimise production control parameters in the clinker manufacturing stage from the preheater to the clinker cooler. This stage is central to the operating efficiency and costs of the plant, accounting for 100% of thermal fuel use in a typical plant.

The AI system learns by finding relationships between the plant operating data in the same way that a human being does. These data and insights are displayed so that the plant operator can visualise the information being used by the software to make the optimisation decision and build confidence that "the system thinks as we think."

The recommendations, provided by the AI system allow the operator to consider optimisation decisions simultaneously with production targets and quality targets. The system acts as a co-pilot: it recommends settings, but the operator is always in control, and it is the operator who sets limits and authorises changes. Live performance reviews produce recommendations that plant operators can accept or reject.



Preheater to Clinker Conveyor: Three units in sequence make up the burning process central to the production of clinker that is then processed into cement. In a modern plant, the first unit is a multi-stage cyclone preheater with calciner, to preheat the input raw meal to 900°C using kiln exhaust gases. The second stage is a rotary kiln where sintering at 1450°C drives the chemical changes into clinker. Finally, the clinker minerals are stabilised in a clinker cooler, capturing the waste heat to feed into the prior kiln combustion process.

Fuel consumption reductions are primarily achieved by optimising excess air ratios to minimise energy loss to the stack and tailoring heat input to the specific raw materials. In addition, performance improvements are achieved through changes to fuel mixes to achieve the same energy content with lower fuel-derived carbon emissions.

Delta Zero can enable cement plant operators to achieve a 10% reduction in fuel consumption and up to 20% in fuel-derived carbon emissions.

Optimisation is an hourly and daily process, not a one-off adjustment in control settings. Achieving the 10% savings in energy requires regular minor adjustments of parameters by a few percentage points. As the input parameters are constantly changing, so are the recommendations. These benefits are unlocked in a way that couldn't be

> achieved by expanding the size of a process engineering function with additional Chemical Engineers – they result from machine learning algorithms powered by powerful computers working through the massive plant datasets, finding optimisations each time that will work best for your daily or hourly operating parameters.^{1,2}



- 1 Journal of Cleaner Production, June 2016, Daniel Summerbell, Jonathan Cullen, Claire Barlow. *Potential reduction of carbon emissions by performance improvement: A cement industry case study*. https://www.researchgate.net/publication/304532399_Potential_reduction_of_ carbon_emissions_by_performance_improvement_A_cement_industry_case_study
- 2 Applied Energy, July 2017, Daniel Summerbell, Diana Khripko, Claire Barlow, Jens Hesselbach. *Cost and carbon reductions from industrial demand-side management: Study of potential savings at a cement plant*. https://www.researchgate.net/publication/316005154_Cost_and_carbon_reductions_from_industrial_demand-side_management_Study_of_potential_savings_at_a_cement_plant

AI & reinforcement learning

Al has generated massive excitement in recent years. This is now being supported by an ever-growing list of applications where AI has delivered tangible benefits and improvements, in a variety of industries – from healthcare to finance and logistics.



The state of the art in AI is still 'narrow' AI – i.e. an AI that has learned to perform a particular task rather than a general level of intelligence. This narrow AI works best on problems represented as 'closed systems' where the number of components or variables stays consistent over time. Still, there is the opportunity to exceed human performance levels within that framework, particularly in complex optimisation problems. Narrow AI means it is essential to marry domain expertise to the expertise in AI as one without the other is far less powerful. There is no 'one AI to rule them all,' and specific problems require specific capabilities. Reinforcement learning is a powerful approach to train narrow AI that's being used to solve a variety

of problems, from autonomous navigation systems to power grid management.

Al terminology frequently references 'training'. in reinforcement learning this is the process by which the AI learns through repeated interaction with an 'environment' what actions to take in different contexts. The learning in this context is the gradual improvement of a neural network which can be thought of as the brain of an AI. The AI's choice of action is influenced by the 'reward

function' which can be tailored to induce the desired behaviour in the problem solution.

For example, when optimising performance in a cement plant, the choice of reward function could incorporate a carbon cost to encourage decisions that reduce emissions while maintaining production. It can also include specific constraints such as NOx emissions or torque limits to ensure safe operation.

100 years of training in a day

The training process described above can be considered an extremely brute force way of learning. The agent repeatedly experiences a multitude of scenarios testing out actions to learn their impact and discover the optimal step to maximise reward. Over the training process, this enables the AI to make strategic decisions that leverage complex interactions between processes and over time, for example, adjusting the fuel split between calciner and kiln based on changes in the burnability of raw feed.

This training process is enabled by advances in computing and high-speed processors, which means the agent can experience the equivalent of

100 years of cement plant operation in simulation during a day's training.

As a result, Carbon Re can do things to help the most experienced operator out of reach of the most advanced 'Expert System.' For example, it can predict output clinker quality from feed rate, fuel choices, and other kiln performance raw material quality parameters. The predictions are instantaneous and in advance of the changes being implemented in the plant control system. Output clinker quality can then be understood immediately when the input control decisions are made, avoiding the time lag before clinker quality tests can be conducted at the output of the kiln.

The complexity of alternative fuels

In the last decade, being a cement plant operator has meant facing enormous and ever-increasing challenges. To reduce the use of fossil fuels (coal, petroleum coke, and natural gas) and take advantage of lower prices, cement producers are increasingly switching thermal energy for the kiln

Not all gigajoules are created equal...

High flame temperature is critical for cement production: one gigajoule of dried waste could burn as a 'long' flame at 1,000°C rather than the 'short' flame at nearer the 1,500°C temperature needed for cement calcification in the kiln.⁶

to a blend of alternative fuels such as biomass, industrial waste, domestic waste, or dried sewage.³

The composition of the alternative fuels will also affect the production process in different ways, from moisture levels to volatile contents and

> particulate size.⁴ All these different interactions from alternative fuels add complexity to decision-making on plant control parameters.⁵ AI can simulate the impact of potential fuel mix combinations, and learn the relationship between performance and blending of different alternative fuels.

Simplifying the introduction of new technology

So far, improvements in cement plant technology have increased the complexity of operating systems: simple electronic controls have evolved into complex and interdependent systems. Each iteration of technology has also entailed investment in onsite plant equipment. Legacy technology providers to the cement industry have been slow to learn from the broader technology industry and the latest software practices.

No investment in equipment needed

Carbon Re brings these technological advances to cement producers. It works with no systems integration, no investment in equipment, and no investment in IT hardware. It can operate as a standalone solution, as it exploits a massive multi-dimensional simulation of the plant, a 'digital twin' of the actual cement plant.⁷ Carbon Re uses the digital twin to evaluate and understand current plant performance, using the power of cloud computing systems to consider the impact of all possible parameter changes. Operators are provided the resulting prioritised set of optimisation parameters to implement in the existing plant control systems.

The system can ingest plant data from any data source, including the major distributed control systems widely used in the industry. Where desired, Carbon Re can then be directly integrated with an existing distributed control system or an 'Expert System' to set optimum efficiency parameters.

https://pdf.sciencedirectassets.com/278653/1-s2.0-S1877705813X00074/1-s2.0-S187770581300492X/main.pdf

³ Environment Agency (UK), David Baird, Sarah Horrocks, Jenny Kirton, Roland Woodbridge. Science Report: SCO30168 *The use of substitute fuels in the UK cement and lime industries.* February 2008.

<sup>https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/291698/scho1207bnna-e-e.pdf
Procedia Engineering 56 (2013) 393 – 400, Azad Rahman, M.G. Rasul, M.M.K. Khan, S. Sharma. Impact of alternative fuels on the cement manufacturing plant performance: an overview.</sup>

⁵ International Finance Corporation, 2017. Increasing The Use Of Alternative Fuels At Cement Plants: International Best Practice. https://www.ifc.org/wps/wcm/connect/33180042-b8c1-4797-ac82-cd5167689d39/Alternative_Fuels_08+04.pdf?MOD=AJPERES&CVID=IT3Bm3Z

⁶ Mineral Products, Association, October 2019. Options for switching UK cement production sites to near zero CO2 emission fuel: Technical and financial feasibility. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/866365/ Phase_2_-_MPA_-_Cement_Production_Fuel_Switching.pdf

⁷ Forbes. What Is Digital Twin Technology – And Why Is It So Important? https://www.forbes.com/sites/bernardmarr/2017/03/06/what-is-digital-twin-technology-and-why-is-it-so-important/

Every cement plant is unique

Locally available fuel supplies, raw materials, and market conditions are key factors to consider when designing a cement plant for both the material properties of the clinker and production costs. Rules-based 'Expert Systems' typically use industry-standard relationships to set production parameters that are not necessarily optimal for individual plants. The advantage of machine learning systems is that they naturally adapt to the specific properties of each plant. Carbon Re adapts to different kiln sizes, fuel types and mixes, raw materials, and production requirements.

Simplifying plant control decisions

From our development and platform roll-out, we have seen how current plant management systems have reached the complexity limit for any human operator: sitting with plant operators, we observed process control alerts going off every 15 seconds. Other 'helpful' software overwhelms with information. Using feedback from plant operators, recommendations from Carbon Re have been designed to reduce complexity and allow plant operators to deal with what matters most.



Emerging secondary benefits and future direction

Further secondary benefits have been identified with the continuous development of Carbon Re, such as:

- Keeping the plant within operating parameters: Carbon Re enables operators to keep the kiln operating within recommended torque settings.
- Controlling NOx emissions: keeping emissions within control limits.

Furthermore, the 'digital twin' offers future exciting opportunities to better understand and address maintenance issues with the plant, such as kiln blockages, as well as evaluate how best to make capital investments: Avoiding kiln blockages: Finding the optimum ranges for silica, alumina, sulphur, and chlorides for the right amount of coating inside the kiln whilst optimising clinker quality and throughput is a challenging balancing act. Enough coating is required not to damage the refractory lining and shorten the life of the kiln, and not too much, which would create a blockage. We are currently working with cement plant operators to predict and manage coating levels – avoiding damage to the critical refractory lining and minimising the risk of a costly and time-consuming shut-down.



FAQs

- 1. Does Carbon Re control production operations? No, we have deliberately designed Carbon Re as a human-in-the-loop solution, meaning that the plant operator maintains control over production processes. Carbon Re provides specific and quantified recommendations to be actioned by the operator.
- 2. Do we need new equipment to run the system? Carbon Re's system is based in the cloud, so there is very little, if any, 'infrastructure' to install at the plant. Carbon Re analyses incoming data identifies recommendations, and presents them to operators as a priority list on an intuitive browser-based dashboard.
- 3. Why is your approach better than others?

Carbon Re provides much more value than existing solutions: Compared to expert systems, our solution uses the latest AI technologies, which can significantly outperform traditional computing approaches.

Traditionally, process improvement has been undertaken through expensive consulting, where the expertise is lost when the consultants leave. Our approach encodes this expertise into powerful software to provide continuous and increasing value over time.

4. If these savings are possible, why has it not been done before?

With Industry 4.0, the new 'Internet of Things' and sensors have provided continuous data streams from manufacturing processes, but taking advantage of these vast amounts of data and underlying relationships requires powerful analytical tools. We have built Carbon Re to be just that: an enterprise-scale software solution powered by Deep Reinforcement Learning, one of the most powerful – and complex – branches of Artificial Intelligence, with world-class software engineering, design, and features.

Carbon Re is based on cutting-edge research at the Institute for Manufacturing at Cambridge University. Find out more here. https://www. researchgate.net/publication/304532399_ Potential_reduction_of_carbon_emissions_ by_performance_improvement_A_cement_ industry_case_study



Appendix 1

Comparison with alternatives

As technology has evolved, different approaches have been taken to address the opportunities and challenges in managing a cement plant efficiently: improving quality and throughput, reducing cost, and shutdowns. AI is already delivering significant advantages over both the prior types of technology, 'Expert Systems' and 'Predictive Process Control & Automation Systems':

Approach	Year introduced	Technological solution	Key factors
Expert Systems	From 1990	Codification of predefined 'rules' and relationships using historical data. Typically hard-coded into a programming language, e.g. 'C' or 'FORTRAN'.	 Time to implement of 3-5 years, with low project completion rates. Systems mothballed as changes in production processes, input materials or output targets require re-design & re-engineering of the rules-based system. Difficult to maintain and update with new rules Use of theoretical datasets and relationships lead to sub-optimal recommendations. Limited number of inputs and output factors considered by models meant not all target factors for plant being considered in recommendations e.g energy costs, NOx emissions, plant maintenance & stability. Very low tolerance to uncertainty in input data Cannot learn on its own, and update existing knowledge base
Predictive Process Control & Automation Systems	From 2000	Extrapolation of Expert Systems, continuing to use pre-defined rules based on interpretation of historical data. Limited key parameters defined, each independently forecast to predict process control settings. Direct connections with key plant systems for automated control. Enhanced plant visualisation using detailed graphical interfaces.	 Upgrades and investment in hardware to allow direct automated controls. Limited modelling looking at production parameters independently, rather than the interplay of parameters Significant cost and time required to upgrade to (a) incorporate any new requirements such as exploiting alternative fuels, reducing input material costs or reducing environmental emissions, or (b) adapt to any changes in plant process or hardware. Increased system complexity, requiring specialised training & skills for plant operators.

Approach	Year introduced	Technological solution	Key factors
AI-based reinforcement machine learning	From 2020	'Digital Twin' of the plant is created using all historic datasets available. Establishes the real life operating characteristics of each specific cement plant. Using the learned experience of the entire recorded history of the plant, the system makes specific optimisation recommendations to the plant operator (or direct into an 'Expert System') based on daily or minute by minute requirements.	 Ability to operate standalone reduces the set- up time to a 1-3-month phase to capture the plant data history and train the AI models. No systems integration, hardware investment or plant operator training required. Changes to plant process or hardware are reflected in the AI models with simple recalibration. New operating requirements easily incorporated, for example optimising costs through balancing the interplay of raw material prices, utility prices and emissions based carbon taxes.

Appendix 2

Preparing for the impact of carbon taxes

The cement industry accounts for 8% of global greenhouse gas emissions, meaning that carbon tax changes coming in the next 24 months will have a significant economic impact on the cost of cement production. Changes to carbon tax mechanisms will have a global effect on the supply chain, rewarding those plants with the most efficient operations.

In the European Union, carbon taxes have not been applied to industries where international trade would affect the ability of local producers to compete with imports through the use of free allowances. These allowances will start to be removed soon, and be entirely removed in the next 6-8 years. The change is driven by 'Carbon Border Adjustment Mechanism' (CBAM) legislation, starting in the European Union before becoming global. The legislation was formally agreed in March 2022 specifically targeting the cement and steel industries first. Canada, Japan and the broader G20 are coordinating parallel schemes.

The time when we make carbon emissions, and when we take carbon out of the atmosphere, creates a pivotal difference in the temperature of the planet, called the "radiative forcing effect". One tonne of carbon saved today could help us meet global temperature targets as much as two tonnes saved in 2050. As this 'time value of carbon emissions' becomes more widely recognised, governments will accelerate rather than delay changes to carbon tax structures.

Using the worldwide averages in 2022 for a typical cement plant, a plant with an annual production of 1.8 million tonnes of cement will incur annual carbon taxes of US\$63 million. Put into context, the average yearly energy costs are under US\$30 million. This will require roughly a 50% increase in the price of each tonne of cement sold just to break even.

These estimates conservatively use a historic average carbon price at €58/tonne, below current market pricing of €70-€80/tonne and projected increases to over €100/tonne.

Contact us

To discuss this paper, its findings and implications, please get in touch.



Dr Daniel Summerbell, Chief Solutions Officer

Daniel is a pioneering researcher in sustainability in cement. He is a Visiting Industrial Fellow at the Institute for Manufacturing, Cambridge University. With a background in process engineering across a range of industries in the US, Canada & Europe, he began a programme of research into the cement industry in 2014.



Aidan O'Sullivan, Co-founder & CTO

Pioneer in application of AI to energy systems; Professor, AI & Energy at UCL; Fellow, Alan Turing Institute; PhD Imperial College; Postdoc MIT



Muhammad Iftikhar, Senior Solutions Engineer

Muhammed is a professional Chemical Engineer with a Masters Degree in Chemical Engineering. He has over 20 years experience as a Process Engineer/Manager, Plant/ Energy Manager in Soda Ash, Lime and Cement, and Waste Management.

carbonre.com/contact-us/

Join the Carbon Revolution



CONTACT US

- ☑ cement@carbonre.com
- www.carbonre.com
- in linkedin.com/company/carbonre