Carbon Re

Creating order from chaos:

How AI can optimise alternative fuel use in cement production.

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, October 2024

Summerbell, D and Wong, K (2024) Creating order from chaos: how AI can optimise alternative fuel use in cement production. Carbon Re, London.

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

Contents

Foreword

Seeing alternative fuel handled in a cement plant is to watch a constant battle to make order out of chaos. The sour-sweet silage smell lingering in the hangar-like reception yards reminds you of the biogenic component of the piles of refuse that, while unpleasant to handle, gives this fuel a lower carbon footprint than the coal or petcoke it replaces.

While large new facilities are built to sort, grade and store plastic and other waste, most plants have to pay close attention to the weather, as humidity or rainfall can play havoc with how the fuel burns. It makes you wonder if, in fifty years, we'll look back on burning waste as a lasting success, as it has been in Germany over the last decade, or if advances in plastic reuse and reduction will keep us shifting towards new alternative fuel sources.

Amid all the media hype about how artificial intelligence is going to change the world, efficient waste disposal is not the first topic that springs to mind. But creating order from chaos is what AI does best. Just as Large Language Models can turn vast amounts of text into a coherent response to your question, machine learning can take the enormous datasets painstakingly gathered by the cement industry into new insights about the process. We can apply these insights to answer questions like: How can we optimise cement production despite a constantly changing fuel supply? And, how can we practically leverage AI to reduce carbon emissions? These are critical questions, especially as the European Union begins phasing out free carbon allowances over the next decade.

In this whitepaper, we outline how AI is being used to optimise alternative fuel usage today, illustrated by the publication of our recent case study results. We also highlight the scope for future development of AI applications in this space.

We welcome your thoughts and experiences. As people working in the cement industry, you have the best insight into our most pressing question: What problem should AI tackle next?

Daniel Summerbell & Kenny Wong

To discuss this paper, its findings and implications, please get in touch. carbonre.com/contact-us/

Dr Daniel Summerbell, Chief Solutions Officer
daniel@carbonre.com

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.

Kenny Wong, Head of Product
kenny@carbonre.com

Kenny is an operational improvement specialist with over nine years of experience delivering high-value ($1-15M) process improvement projects in various manufacturing environments globally. He has built and deployed software products for tech and manufacturing start-ups, and combines this expertise to deliver results for Carbon Re and its customers.

Alternative fuels are challenging for production managers and operators

Cement production faces immense pressure to reduce its carbon footprint, and perhaps the most widely adopted solution so far is the use of alternative fuels (AF), commonly derived from a variety of waste streams, including commercial, industrial, municipal and agricultural. While these fuels present a promising alternative to fossil fuels, they bring a host of challenges that complicate the cement production process.

A key challenge with burning AF is that each cement plant operates under unique circumstances. Variables such as kiln design, fuel mix, and local waste availability make it impossible to apply a one-size-fits-all approach, and centralised strategies are very difficult to design and implement. Moreover, it's never a one-time effort. The dynamic nature of fuel availability, composition, quality, and pricing, require flexible operational strategies and tailored solutions to optimise plant performance and cost-efficiency. Local production managers are constantly having to adapt.

Even with optimal sourcing, storing, and handling, AF pose significant challenges for production managers and operators as they seek to balance a new set of constraints on plant operation, such as:

Blockages. The formation of buildups in preheater cyclones can disrupt material flow and force costly shutdowns and escalate operational costs.

Hitting the temperature sweet spot. Inconsistent burning properties make it challenging to sustain optimal temperature profiles in the kiln. In the calciner fluctuations in particle size and calorific value cause havoc with conventional loop control systems.

Making trade-offs between TSR and SHC. Increasing the thermal substitution rate (TSR) leads to a rise in specific heat consumption (SHC). Operators must navigate a delicate balance between maximising AF use and maintaining energy efficiency, as these metrics frequently conflict.

Complexity of alternative fuel markets

One of the key challenges with AF is the highly volatile markets. This can lead to rapidly fluctuating prices, varying availability of certain fuel types or grades, or even contracts where plants are obliged to purchase a certain volume of fuel in order to keep access to the supply.

This adds an additional complexity to the operational and procurement strategy of the plant. Should a plant over-produce clinker in order to secure its supply of low-cost fuel? What trade-offs should be made in procurement to ensure a consistent supply of quality fuel at the minimum cost? It is a truism in the industry that a gigajoule of energy from coal does not have the same impact on the process as a gigajoule from refuse-derived fuel.

Using AI we can determine the cost per **useful** gigajoule of fuel.

Running a plant with two competing metrics

The efficiency of kiln operation is measured by SHC, which is the thermal energy used per unit of clinker production (expressed in gigajoules per tonne or kilocalories per kilogram of clinker). When AF is used, another key metric becomes important: the TSR, which indicates the proportion of the kiln's total thermal energy sourced from AF.

Typically, increasing TSR leads to higher SHC for several reasons: larger fuel particle sizes may require more combustion air and cause greater heat loss, production limits can increase fixed losses per tonne, fuel stored outside may contain more moisture than expected, and differences in flame shape and temperature may demand more fuel to maintain the proper sintering temperatures. Additionally, the variability in fuel composition often requires the kiln to operate at higher temperatures to avoid dropping below acceptable levels.

Despite the increase in SHC, the lower cost of AF usually reduces overall fuel costs as TSR rises. This forces operators to balance two competing metrics: TSR and SHC. Changes in fuel mix, necessary to maintain TSR in a volatile market, can obscure the effects of energy efficiency measures, requiring operators to make tough decisions on trade-offs between managing difficult fuels and maintaining optimal plant operations.

To address this, some plants introduce a third metric: specific fuel cost, which tracks the cost of fuel per tonne of production. While this eliminates the conflict between TSR and SHC, it makes plant performance highly sensitive to fluctuating fuel prices and the availability of AF, limiting the usefulness of specific fuel cost as a performance measure since it is influenced by external factors beyond the operators' control.

The use of AF is limited by various constraints. On the supply side, these constraints include the availability and fluctuating prices of fuels. On the process side, they involve factors such as airflow, blockage risks, and the impact of fuels on flame shape and temperature, which can ultimately affect product quality.

As process variables and fuel properties change over time, plant conditions can sometimes breach constraints even without changes in control settings. To minimise the need for constant operator intervention, plants are typically operated with a significant buffer between the normal operating zone and the plant's constraints.

> **To minimise the need for constant operator intervention, plants are typically operated with a significant buffer between the normal operating zone and the plant's constraints.**

The role of AI in increasing TSR

AI can increase TSR by simulating the impact of potential fuel mix combinations and learning the relationship between performance and blending of different AF. These tools can maximise the use of AF by predicting potential constraint breaches and adjusting plant controls to avoid them. This allows the plant to be run with a smaller buffer zone, increasing the utilisation of AF without increasing the risk of constraint breaches.

In addition, AI can optimise the overall financial performance of the plant by simulating the impact of different fuels on performance, and then making the appropriate trade-offs between differently priced fuels and plant performance metrics such as heat consumption and production. This eliminates the impact of trying to run the plant according to multiple competing metrics.

Real World Results

Carbon Re's AI platform, implemented in 2024 at Heidelberg Material's Mokra cement plant, integrates real-time data from chemical analyses, pressure and temperature sensors, and gas analyser data, with AI.

By analysing trends across the full pyroprocess system to proactively adjust targets for the existing ABB Ability^{TM} Expert Optimizer, Carbon Re increased the TSR at the plant from 83% to 85% in month-to-month trials.

Download the case study here to read in full.

Optimising the fuel-to-air ratio is challenging

Ensuring the best ratio of fuel-to-air is a particular challenge when it comes to burning AF. This is due to the variation in the types, particle sizes and chemical composition of AF, which makes them more difficult to burn than traditional fossil fuels. To optimise combustion, a key control parameter for cement plant operators is airflow. Airflow determines the oxygen supply through the system and the fuel-to-air ratio.

In a perfect system, the relationship for fuel-efficient combustion is clear and fixed over time. As oxygen decreases, higher levels of carbon monoxide and organic carbon are produced asymptotically. In a kiln, this relationship is continually changing due to numerous complex factors influencing combustion, which is exacerbated by AF use.

The image below shows real data of a single kiln using AF over a week from the Carbon Re AI platform. Two lines have been drawn to show the change in the combustion efficiency curve between day 1 to day 2 vs day 3 to day 5, due to an evolving fuel mix. To maintain optimal combustion, the kiln's oxygen target needs to be adjusted.

> Operators don't have the time or tools to adjust this target as frequently as required due to the complexity of the kiln environment.

Traditionally, control room operators haven't needed to adjust their oxygen targets often because the stability of conventional fuels didn't require it. However, the increased variability introduced by AF means that oxygen targets must be adjusted regularly to maintain optimal combustion efficiency. Operators don't have the time or tools to adjust this target as frequently as required due to the complexity of the kiln environment.

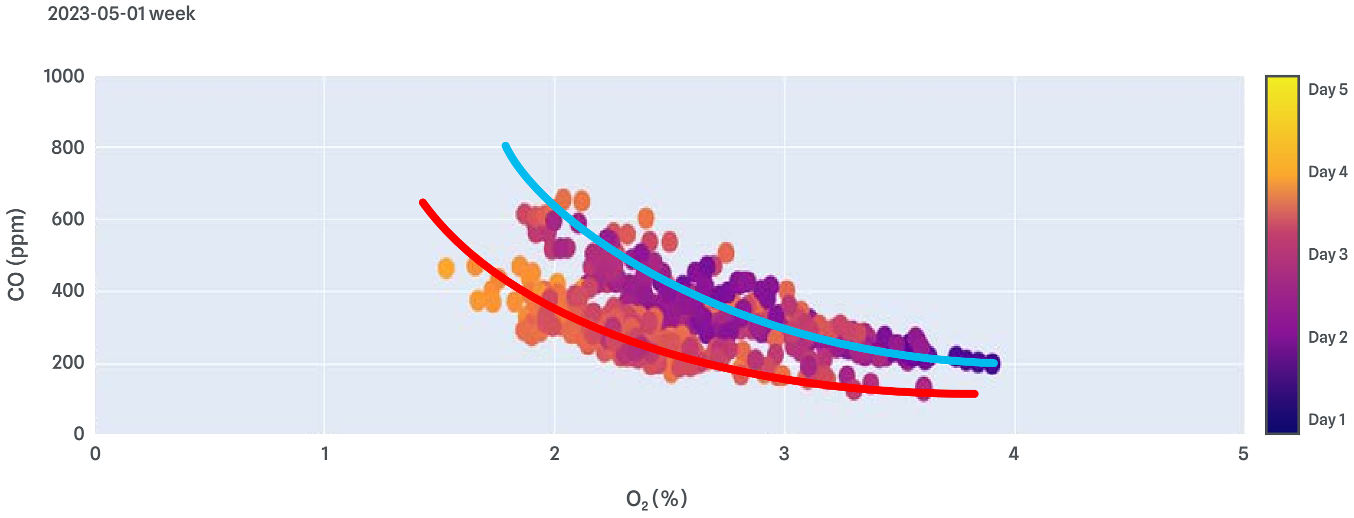

Figure 1: Scatter plot of 15-minute averages of oxygen (O2) vs carbon monoxide (CO) at the top of the preheater tower over 1 week of running at a cement plant. The colour grading indicates the evolution in time. Two lines have been drawn showing the change in the combustion efficiency curve between day 1 to day 2 vs day 3 to day 5, originating from an evolving fuel mix. To maintain optimal combustion, the O2 target should be adjusted.

This complexity is particularly well-suited to machine learning solutions that can comprehend and identify patterns in data with speed and accuracy beyond human capabilities. When applied to kiln combustion, AI can not only identify the ideal target for a given parameter, such as oxygen, but it can also compare that target with the current kiln conditions to make appropriate changes to ensure continuous stability while adjusting parameters towards optimal combustion. Using machine learning solutions for combustion also allows operators and process engineers to spend more time on higher-value projects, away from the minutiae of kiln optimisation.

The importance of optimising combustion

Incomplete combustion, which occurs if there is not enough oxygen available, generates carbon monoxide and, in some cases, solid particles of organic carbon. Both of these pollutants have harmful effects on both human health and the environment and are subject to strict regulatory requirements. Incomplete combustion is much less efficient than complete combustion because the fuels do not burn completely, leading to wasted fuel. If unburned particles are reaching the gas exit to the system, it is an indication that the fuel is not combusting in an optimal burning zone. As a result, more fuel must be burned to achieve the same chemical changes required for calcination or clinkering.

At Heidelberg Material's Mokra Plant, Carbon Re's oxygen target recommender reduced post combustion oxygen by 0.2% while maintaining production capacity and without increasing carbon monoxide or organic carbon limit breaches.

Maintaining consistent quality with inconsistent fuels

Another major constraint is, of course, clinker quality. Maintaining kiln stability, and hence quality is the primary concern of control room operators. The varying performance of AF makes this particularly challenging.

The nature of clinker sampling means that the kiln is being controlled using information that can be hours out of date. It typically takes approximately 45 minutes between the clinker dropping into the cooler, and a control decision being made. This leads to both under- and overburning of clinker as the system is slow to react to changes in quality, and therefore has to make more drastic changes to burning conditions.

From reactive to predictive control

A key component of the Carbon Re platform is the live quality forecast soft sensors that predict the C3S and free lime content of the clinker every 15 minutes, at the moment at which the clinker quenches in the cooler. This continuously updated value enables smoother control of the burning zone, reducing the dependence on fossil fuels for rapid adjustments and increasing TSR further.

In live deployments, the Carbon Re platform has delivered a 30-33% reduction in clinker quality variation.

Alternative fuels increase the likelihood of blockages

Alternative fuels generally have more complex chemical and physical properties than traditional fuels, such as higher levels of chloride and sulphur or lower ash melting points. These characteristics can lead to more frequent and severe blockages and ring formations within the equipment. Chloride is a particular issue: condensing at around 800°C, it forms sticky liquid layers in the preheater tower, which facilitates the accumulation of particles, further intensifying blockages and ring formations. What's more, when AF does not burn completely or uniformly, the risk of material buildup increases, potentially leading to operational disruptions and efficiency losses.

Frequent buildups require shutdowns for manual cleaning and maintenance, leading to significant downtime and reduced throughput. Moreover, these stoppages strain the maintenance resources and escalate operational costs.

AI-powered Blockage detection

Carbon Re's blockage detector is a significant industry advancement resulting from novel research into AI for increasing AF use. Utilising high-resolution data from pressure, temperature sensors, and gas analysers, our AI can detect spikes in these parameters up to twelve hours prior to a shutdown caused by a blockage, preventing unplanned outages.

Additionally, by optimising fuel mixing and feeding strategies, our models ensure more complete combustion and reduce the likelihood of buildups. Machine learning models can continually adjust operational parameters through the live data feeds on fuel characteristics and operational conditions. The platform aids the control room operator, maintaining optimal combustion conditions and minimising the formation of blockages.

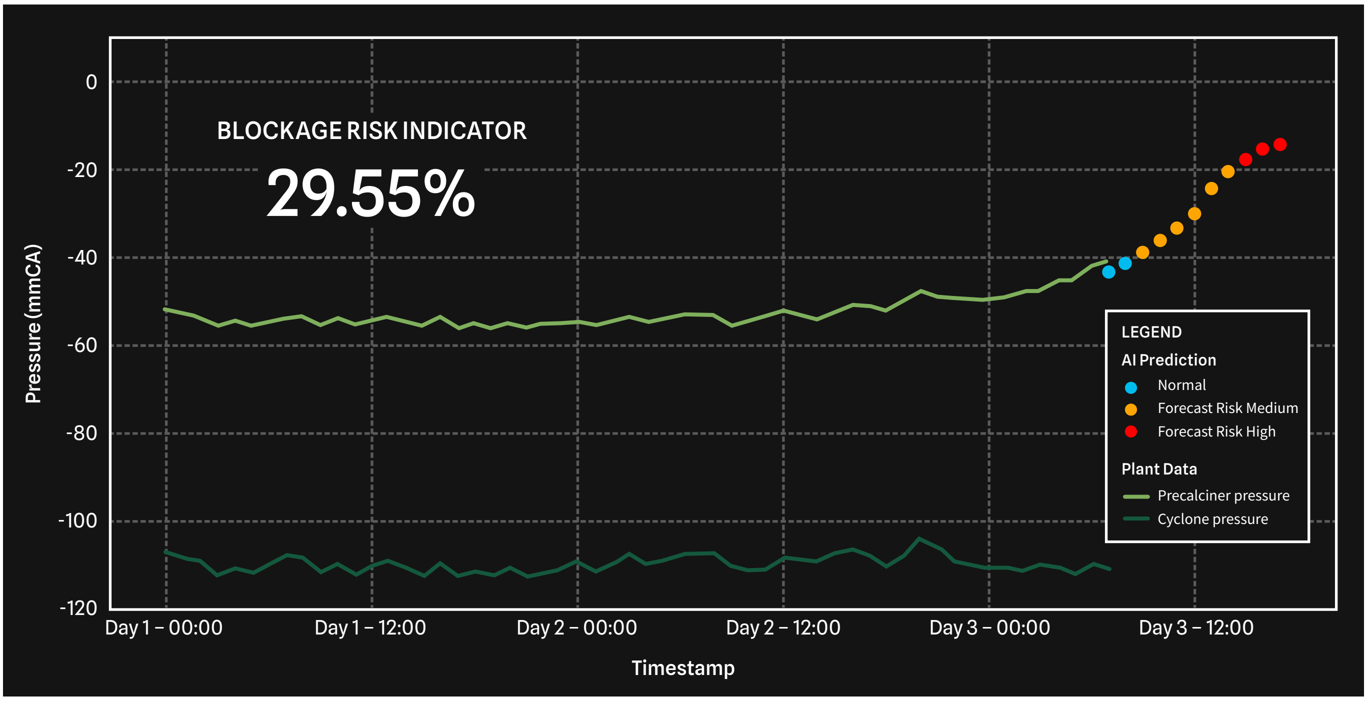

Figure 2 illustrates how Carbon Re's AI utilises high-resolution data to detect the spikes in a range of sensors, including pressure that precedes a blockage, and predict the severity of the risk based on real-time conditions. This prediction gives operators hours of warning to adjust operational parameters to minimise the formation of a blockage and prevent unplanned outages.

A tangible return on investment

It is challenging to assess the direct quantifiable benefit of any new technology in the complex and noisy environment of a cement kiln. However, benefits can be definitively established in a number of areas, as outlined below. In each case, the financial benefits we quantify will vary by plant, as both the AI models will be customised for different strategic needs and opportunities to improve production control settings according to the plant.

Optimising Combustion: By dynamically adjusting airflow targets, AI enables more efficient combustion, minimises fuel waste, reduces specific heat consumption and saves on fan power and electricity costs. On/off trials indicate a reduction in fuel consumption in the calciner of ~2%.

Enhancing Process Control: Advanced AI systems enhance process control improving process stability. During live kiln operation, Carbon Re's software can reduce quality variation by 25-33%, implying improved process stability.

Increased TSR: Improvements in combustion performance and stability enable the plant to increase its use of AF. At the Heidelberg Material's Mokra plant, for example, TSR increased from 83% to 85% during month-to-month comparisons.

Reduced carbon footprint: Carbon emission reductions have an additional financial benefit, depending on the location of the plant. With Carbon Re's AI, plants typically save 4.5kg/tonne of clinker in fuel-derived carbon emissions. In the European Union, UK and Canada, the unused carbon credits allocated have a real tradable financial value.

Stopping Blockages: AI models can analyse operational data to detect early signs of blockages and suggest preventive actions. By monitoring key parameters like temperature and pressure, AI systems can proactively adjust process conditions to prevent material build-up.

Example: Cost of shutdowns

While the financial cost of a shutdown varies from plant to plant, it is commonly estimated at around $300,000 per day, or $12,500 per hour at a 1MT/yr capacity kiln. These losses stem from lost production and the associated labour costs, as well as energy costs resulting from the need to reheat the kiln system after a lengthy shutdown. In addition, unplanned shutdowns increase wear and tear on refractory linings and other plant equipment and come with safety risks.

The high cost of such shutdowns is one reason why plants are often run with a considerable operational buffer to avoid blockages by reducing the consumption of high-sulphur AF. Our platform provides early warning of blockages, which not only allows for an increase in TSR but also reduces the significant downside risk associated with blockages.

> [Carbon Re's] platform provides early warning of blockages, which not only allows for an increase in TSR but also reduces the significant downside risk associated with blockages.

What next for AI in Cement: optimising the fuel mix for cost and clinker quality

Advances in AI will continue to transform how cement plants optimise critical operational parameters.

Fuel split

Further benefits can be brought to plant operation by analysing the fuel split between the calciner and main burner. As the fuel mix changes, the optimal split in both fuel and, subsequently, airflow will also change. Machine learning models can predict the optimal proportions of fuel and the resulting requirements for airflow. Adjusting fuel flow and airflow setpoints both reduces the overall specific heat consumption and maximises the use of the most cost-effective AF, minimising overall fuel cost.

Fuel mix optimisation

The volatility of fuel prices and varying quality of AF also leads to a complex optimisation problem that is well suited to a machine learning type approach. The models that Carbon Re has created are designed to take into account fuel chemistry, measured calorific values, moisture content etc, as well as current fuel prices. The performance of the overall fuel mix can then be assessed based on live plant data. The contributions to that performance from each of the fuel types in the mix are estimated from a combination of observed performance changes as the fuel mix varies and predicted changes based on models derived from historical data and the measured fuel characteristics.

Even over a relatively short period of time, reactive control adjustments will lead to natural variation in the fuel mix, allowing for real-time estimation of the impact of individual fuels as they are fired. Based on this simulation, the true cost of each fuel per unit of useful heat energy delivered to the process will then be estimated, and the fuel mix will be adjusted to minimise the overall fuel spend while maintaining quality parameters and levels of production.

Integration of carbon capture systems

The cement industry’s roadmap for future decarbonisation is heavily dependent on carbon capture and storage technology. There are two main technologies being considered for the decarbonisation of the cement industry: conventional carbon capture and oxyfuel technology. As the use of AF has implications for both, the cement industry must either adapt to the complications caused by combining AF with carbon capture or face the dual cost impact of both the energy requirements of carbon capture processes as well as returning to more expensive heterogeneous fossil fuels.

Carbon Capture

Conventional carbon capture requires the separation of carbon dioxide from flue gases, using amine or similar chemicals to dissolve the carbon dioxide, and then applying heat to release the carbon dioxide from the solvent. Excess oxygen or carbon monoxide in the flue gas causes degradation of the solvent. Therefore, the cost associated with poor control of the fuel-air ratio is considerably increased when using a carbon capture system of this type. The value of AI-driven optimisation of airflows within the kilns is increased accordingly.

Oxyfuel uses pure oxygen in lieu of ambient air during combustion. This reduces the volume of gas flow through the cement process, and means that the exhaust gas is very concentrated carbon dioxide. This eliminates the need for capture using a solvent. However, the energy required to separate the pure oxygen is considerable. Any imbalance in the fuel-to-oxygen ratio leads to either expensive purified oxygen being wasted by ending up in the flue gas or incomplete combustion leading to wasted fuel and the presence of carbon monoxide in the flue gas. Again, management of the process using AI will mitigate the potential costs.

Conclusion: Creating order from chaos, now and in the future

Optimising the burning of AF in cement production is a multifaceted challenge requiring an integrated approach. By improving waste segregation, adapting market strategies, and utilising advanced technologies like AI for process optimisation, cement producers can enhance TSR while minimising environmental impact. The journey toward sustainable cement production is complex, but with informed decision-making, the industry can make significant strides in reducing its carbon footprint.

Is this a long-term plan?

While burning AF is an effective way to reduce carbon emissions in the short term, it cannot completely decarbonise the process. In addition, questions remain over the long-term supply of the best quality fuels; Germany, a leader in AF use, is already exploring ways to reduce the initial use of plastics and promote the reuse of materials rather than incineration. As a result, the industry will need to remain flexible and adaptable as it integrates new materials and technologies into its process.

The great thing about AI models: adaptability

AI models provide flexibility and adaptability, regardless of the future direction of cement production. Even if the industry shifts away from burning waste by 2030, AI tools can help the industry adapt to the changing reality. They can accelerate the adoption of new fuels, process changes, and operational strategies, ensuring that investments in AI technology not only deliver a return immediately but also future-proof the process. This is a key step towards a more sustainable and resilient cement industry.

If you'd like to know more about the models deployed, setup process, and benefits seen to date, please get in touch. Our AI modelling team is based in London to access the best global talent. The plants we deploy to are global, and we are already operating in plants in Asia, America, and across Europe.

Further Reading

> “Versatile and long-lasting, concrete buildings and structures are in many ways ideal for climate-resilient construction. But concrete has a colossal carbon footprint — at least 8% of global emissions caused by humans come from the cement industry alone. We must decarbonise its production.

Source: Nature, editorial, 28 September 2021

Cement is a crucial material for modern infrastructure. Used extensively in constructing homes, schools, roads, offices, and hospitals, Cement is responsible for 8% of global carbon dioxide emissions: greater than deforestation, global shipping and aviation combined. Approximately 40% of these emissions result from fuel required to generate extremely high temperatures, around 1500°C, and more than half from the chemical reactions involved in the manufacturing process.

Currently, much of the demand for cement is from China, with over 50% of the global production occurring there. In the UK, over 8 million tonnes of cement are produced annually, while India, the world’s second-largest producer after China, manufactures over 323 million tonnes each year. While China’s cement use is expected to decline over the coming decades, cement production globally is expected to increase by 45% by 2050 due to soaring urbanisation as well as the civil infrastructure needed to protect against rising sea levels and extreme weather events.

As carbon credit and offset costs rise, environmental regulations tighten and pressure to meet sustainability goals grows, the industry - previously labelled ‘hard to abate’ - has been pushed to explore more and more ways to reduce its carbon footprint. Today, the long-term roadmap to decarbonise cement production still depends on costly Carbon Capture & Storage (CCUS) technologies that have yet to be proven at scale and are not yet commercially viable. Carbon Re’s whitepaper review of technologies to reduce carbon emissions from cement production highlights three technologies with the potential to make a significant impact on carbon emissions by 2030:

1. Substitute Cementitious Materials including “LC3” Cement
2. AI for energy efficiency and SCM blending
3. Alternative fuels (including biomass and waste derived fuels)

To find out more, read our whitepaper:carbonre.com/three-technologies-to-reduce-climate-change

Contact us

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

Dr Daniel Summerbell, Chief Solutions Officer

daniel@carbonre.com

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.

Kenny Wong, Head of Product

kenny@carbonre.com

Kenny is an operational improvement specialist with over nine years of experience delivering high-value ($1-15M) process improvement projects in various manufacturing environments globally. He has built and deployed software products for tech and manufacturing start-ups, and combines this expertise to deliver results for Carbon Re and its customers.

carbonre.com/contact-us/

Join the

Carbon Revolution

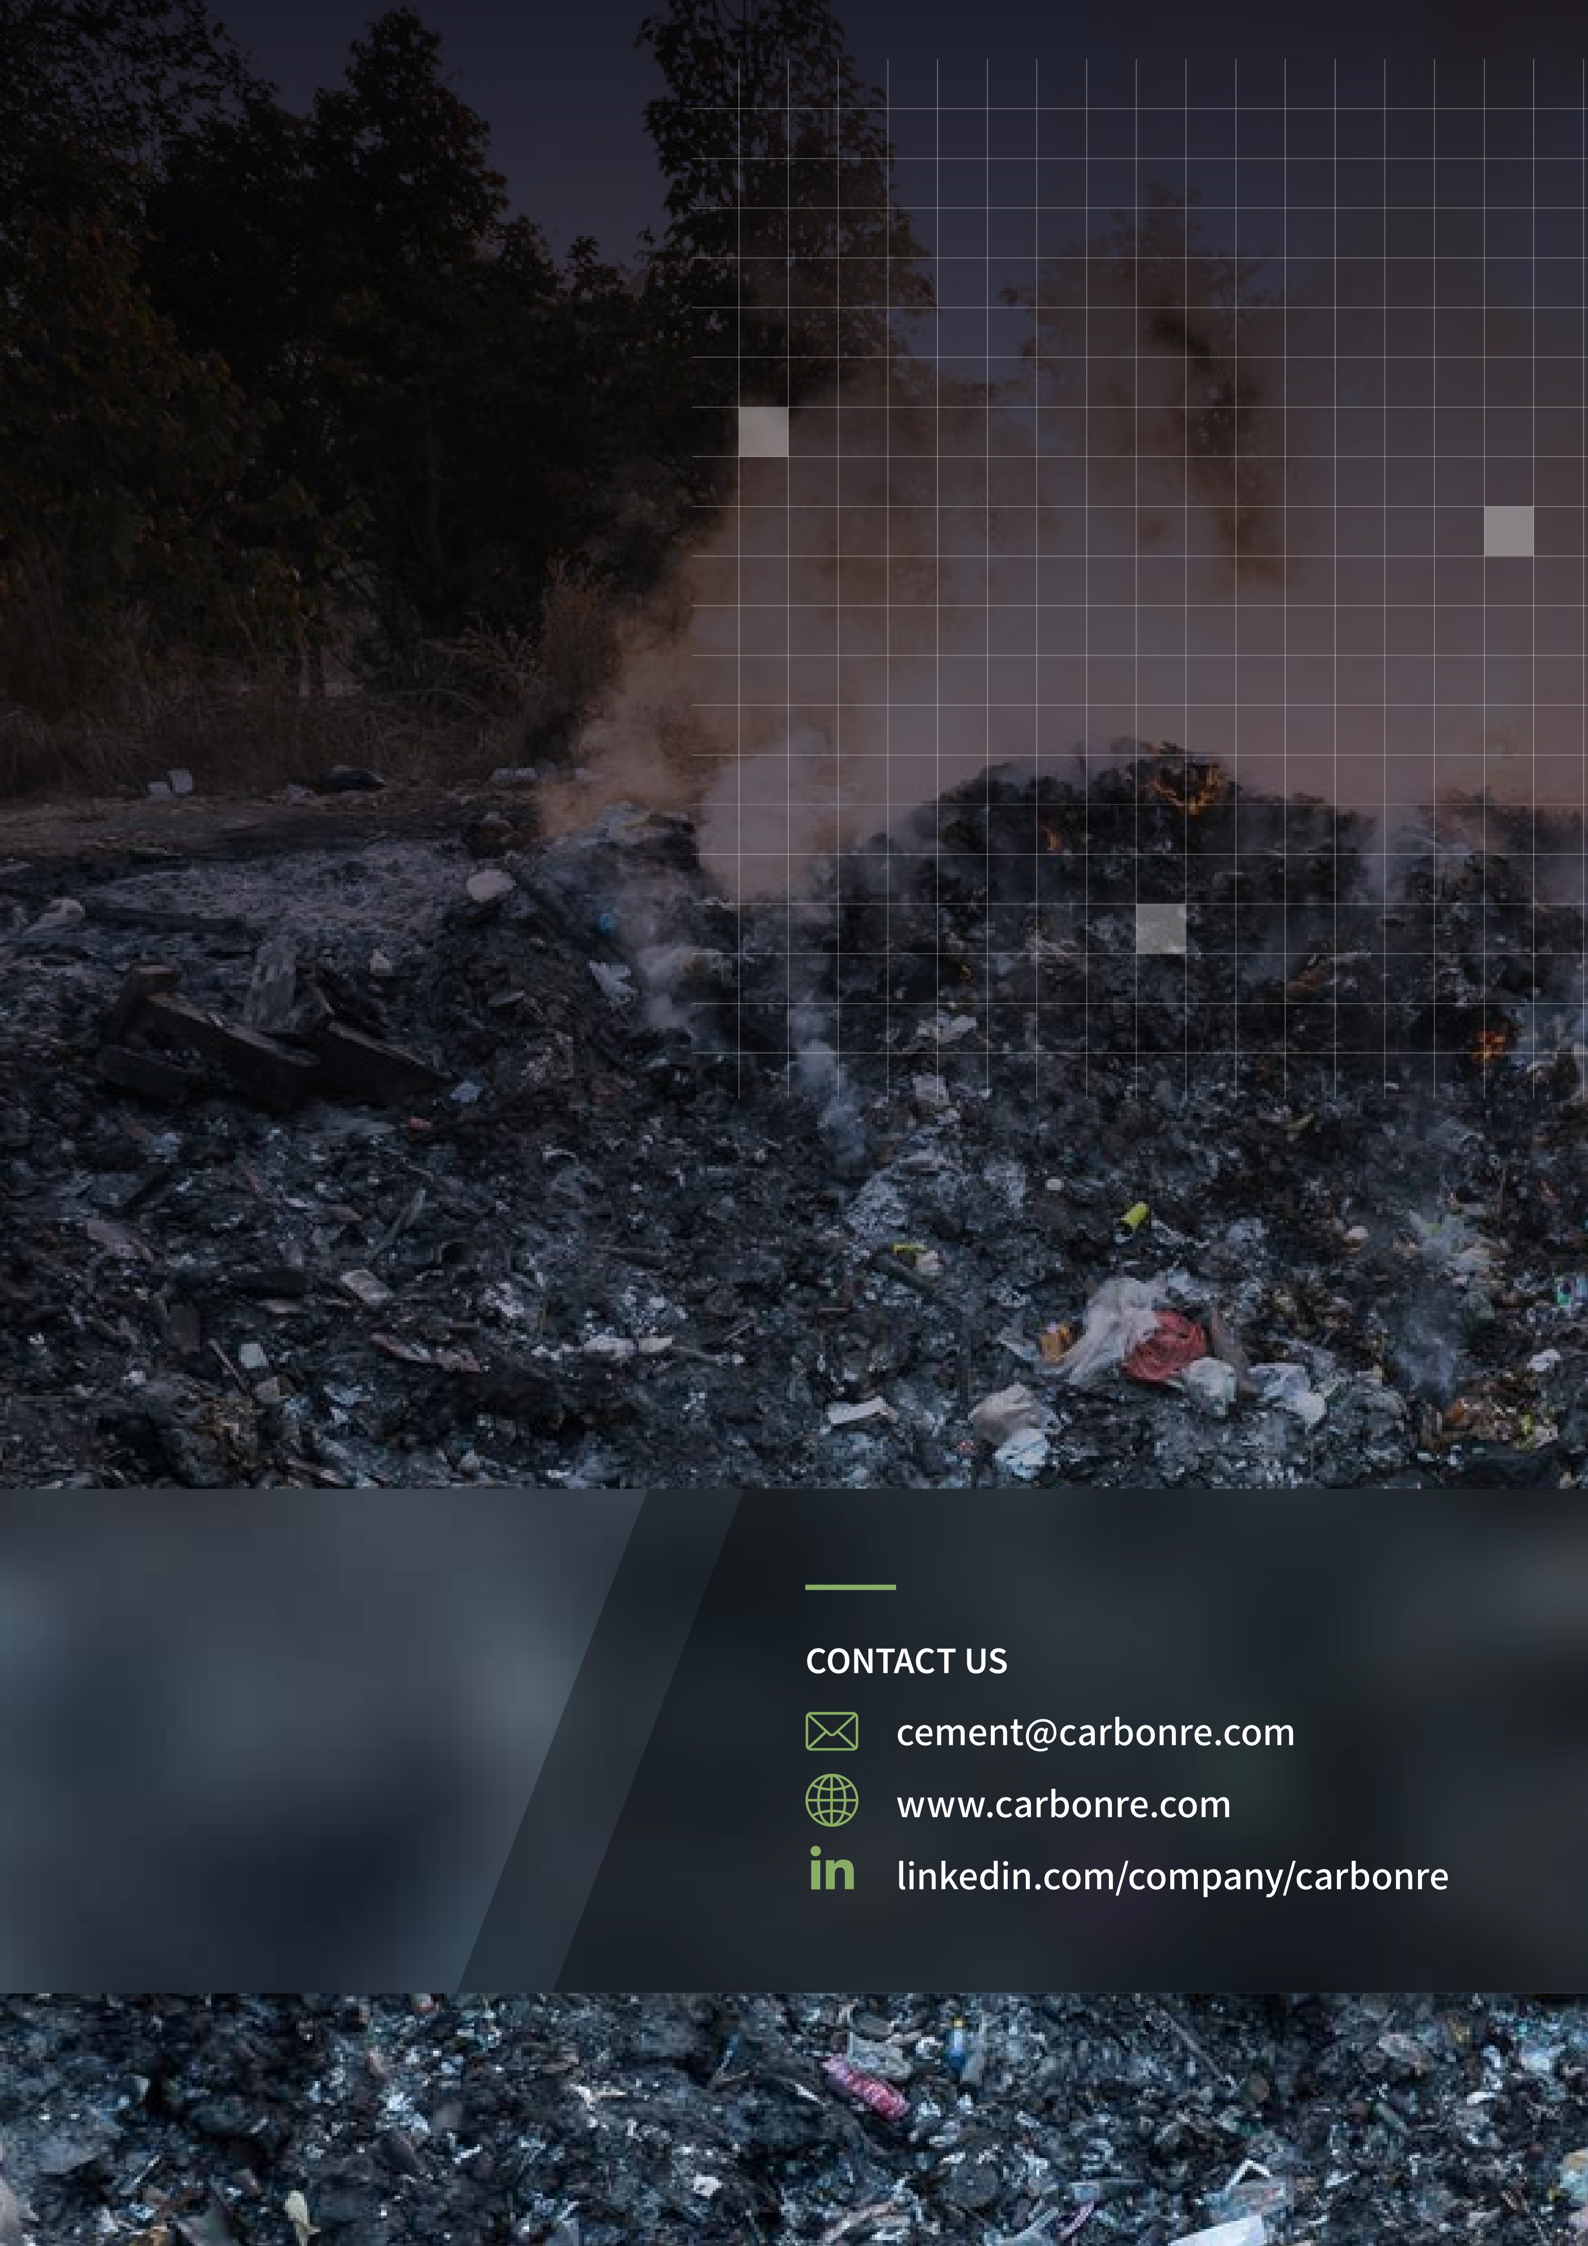

CONTACT US

✉ cement@carbonre.com
🌐 www.carbonre.com
in linkedin.com/company/carbonre