fib Congress 2026
Lisbon - Portugal

AI-powered concrete recipe optimization regarding material parameters, cost and environmental footprint

Jan Suchorzewski[1*], Erik Enecker[2], Claes Holmström[2], Katarina Malaga[1],

1. Department of Infrastructure and Concrete Technology, Research Institutes of Sweden, Gothenburg, Sweden
2. Ecometrix, Drottninggatan 94, 111 36 Stockholm, Sweden

[*]Corresponding author email: jan.suchorzewski@ri.se

Abstract

Concrete production faces a critical need for innovation as traditional trial-and-error mix design is slow, costly, and suboptimal in addressing performance, cost, and environmental objectives. This paper presents development and validation of ACORN, an AI-driven concrete recipe generation system. ACORN integrates large-scale data collection and machine learning to accelerate the development of high-performance, low-carbon concrete mixes. A comprehensive database of over 45,000 concrete recipes was built by aggregating industry data and published research using a large language model for data extraction and anomaly detection for quality assurance. On this foundation, deep learning predictive models were trained to simulate concrete properties (e.g. consistency, strength), production cost, and carbon footprint from a given mix composition. The system features a user-friendly interface for practitioners to obtain instant performance and sustainability evaluations of proposed mixes. Initial validation shows that ACORN can predict key properties within ~15% accuracy of laboratory results, enabling significant reduction in physical testing. By simultaneously optimizing mechanical performance and environmental impact, ACORN demonstrates the potential to cut development time by 50–60%, reduce associated CO_2 emissions, and explore innovative mix designs beyond the reach of manual methods. The paper details the system architecture, the underlying concrete recipe database, validation against experimental data, and discusses the implications for industry practice and sustainability. Future developments toward generative recipe design and broader industry adoption are also outlined.

Keywords: *Concrete, Mix design, Machine learning, Sustainability, Generative AI*

1. Introduction

1.1. Introduction

Concrete is the world's most utilised construction material, yet its development process remains largely manual and empirical. Engineers traditionally refine concrete mix designs through iterative experiments, trying different ingredient proportions and process parameters to meet targets for strength, durability, cost, and carbon footprint. This approach is inefficient and expensive. Minor improvements require extensive laboratory testing, and exploring unconventional mix compositions (e.g. alternative binders to cement and new cement types or recycled aggregates) is often impractical due to the vast design space and non-linear interactions among components. Moreover, cement, the key component in concrete, is responsible for roughly 8% of global CO_2 emissions, so there is urgent need to develop low-carbon concrete. The industry's efforts, such as using supplementary cementitious materials to replace clinker, have grown in recent years, but overall progress in carbon reduction has been modest. New digital tools are needed to break this impasse by guiding concrete design more intelligently and efficiently.

Recent advances in artificial intelligence (AI) and data analytics offer a transformative opportunity for concrete technology. In other material domains (e.g. metallurgy), AI-based design systems have reduced development times from years to months and cut the number of physical experiments by over 60%, while achieving superior performance outcomes. A similar approach in concrete could enable innovative mix design, where optimal formulations are identified via computer simulation before any batching and testing. However, applying AI to concrete design presents challenges. Concrete performance depends on diverse factors (local raw materials properties, mix ratios, curing process) and multi-objective criteria (structural performance, durability, cost, environmental impact) that are difficult to predict simultaneously. Previous studies using machine-learning in concrete research, e.g. to predict compressive strength, have yielded promising results for narrowly defined targets. However, they typically depend on limited, domain-specific datasets and therefore struggle to generalize to broader performance indicators or to emerging alternative cementitious materials.

ACORN (AI-powered COncrete Recipe geNerator) is a project that addresses these challenges by combining data-first strategy with state-of-the-art AI. Instead of limiting the scope to fit available data, ACORN aggregates a wide range of concrete data to build a robust foundation for learning. The goal was to develop a comprehensive platform that could predict and eventually generate concrete mix designs meeting user-defined performance, cost, and sustainability targets. The ACORN system is based on deep domain knowledge in concrete materials and conducted experimental validation, as well as data management, sustainability metrics, and software development of the AI platform. This paper describes the ACORN system's architecture and functionality (Section 2), the development of its underlying database (Section 3), the validation of its predictive models (Section 4), and conclusions including outlook for future work (Section 5). The work is positioned in the context of advancing sustainable concrete innovation, illustrating how AI and big data can revolutionise an industry rooted in empirical tradition.

2. System Architecture and Functionality

2.1. Data Collection Module – LLM-based Parser

ACORN begins by gathering heterogeneous data from multiple sources: environmental product declarations (EPDs) from concrete producers, concrete mixture specifications and test results from research publications, and legacy databases of mix designs and properties accumulated by RISE. To interpret and unify this information, ACORN employs a fine-tuned large language model (LLM) capable of parsing text, tables, and figures. The LLM automatically extracts relevant parameters (e.g. cement type, water-cement ratio, admixtures, curing regime) and outcomes (e.g. 28-day compressive strength, workability, carbon footprint) from documents. Unlike manual data entry, this approach scales to thousands of reports with minimal human intervention. A proof-of-concept demonstrated that a GPT-based model can reliably map concrete data from diverse formats by following natural language prompts.

2.2. Concrete Recipe Database

The output of the data collection and cleaning process is a unified database housing tens of thousands of concrete recipes with their associated attributes. Each data entry typically includes mix composition (all ingredients and proportions), processing details (mixing method, curing time/temperature), mechanical performance metrics (compressive strength, tensile strength, elastic modulus at various ages), durability indicators (e.g. carbonation depth, chloride diffusion), and lifecycle metrics (CO_2 emissions per m^3, cost per m^3). By mid-2025, ACORN's database contained over 45,000 curated mix recipes, making it one of the most extensive concrete datasets ever assembled. This scale is orders of magnitude larger than data used in most prior studies, enabling ACORN's predictive models to leverage deep learning effectively. The database is designed to evolve new data from ongoing R&D projects or future literature can be continuously added. The system's modular architecture allows it to ingest novel data formats with minimal reconfiguration, ensuring longevity and relevance as standards and materials change.

2.3. Predictive Modelling Module

Using the rich database, ACORN trains machine learning models to predict concrete performance. Initially, several algorithms were benchmarked, including tree-based ensemble methods and neural networks. Given the volume and complexity of the data (many features, non-linear interactions), deep learning models were chosen for their representational power. A neural network model forms the core of ACORN's predictor, mapping an input recipe to multiple outputs simultaneously (e.g. 28-day strength, 56-day strength, slump, carbon footprint, and cost). This multi-target approach ensures that trade-offs between different objectives are implicitly learned. The model was trained on a large subset of the database and tested on recipes not seen during training to evaluate its generalisation. The predictive module was iteratively improved by benchmarking different network architectures and hyperparameters as the database grew. In later stages, physics-informed neural network methods were also explored – embedding known formulae (such as water-cement ratio vs. strength relationships) into the model's structure to guide learning.

2.4. User Interface and Integration

ACORN's predictions are accessible through a Graphical User Interface (GUI) integrated a web platform. The target users are concrete technologists in industry and researchers who can input a candidate mix and quickly evaluate its expected performance and impacts. The interface allows the user to specify a mix either fully (exact proportions) or partially (e.g. desired compressive strength class and slump – in the future, ACORN will suggest a recipe that meets those targets). For each query, the system outputs predicted values of all key properties, such as a strength development curve, an estimated CO_2 footprint (e.g. kg CO_2 per m^3), and cost estimation. Results are presented alongside confidence ranges, reflecting prediction uncertainty. If a prediction falls outside the model's validated range (for example, an extrapolation on an unusual mix composition), the system alerts the user. Early versions of the GUI were tested by a subset of Swedish concrete producers (members of Svensk Betong, the Swedish Concrete Association) to gather feedback on usability and presentation of results. Their input informed refinements in the interface and ensured that ACORN's output aligns with practical decision-making needs in mix design. The final prototype allows seamless cloud-based access to ACORN's capabilities, meaning a user can log into the platform, input mix parameters, and obtain results within seconds, without needing to install any software locally.

The transition from a predictive interface to an autonomous generative framework is underpinned by an Agentic "Plan-and-Execute" Architecture. This architectural shift moves beyond static inference by utilizing an autonomous agent to navigate the high-dimensional trade-offs between concrete material kinetics, economic constraints, and environmental impact. At the entry point of this workflow, a Task Generator functions as a semantic parser, analyzing the user's performance intent to construct a dynamic, non-linear task graph.

This system utilizes a specialized classification layer that categorizes every incoming request to determine the appropriate execution path. For instance, when a user provides a specific concrete formulation, the agent initiates an Extraction and Augmentation sequence; here, the agent identifies constituent materials and maps them to an Integrated Material Database to pull missing metadata, such as current market costs and CO_2 intensity values (EPD). Alternatively, if a technologist "locks" certain chemical admixtures or aggregates but leaves proportions open, the agent shifts into a Strict Refinement mode, focusing exclusively on adjusting dosages to meet performance targets within the specified material palette. In cases where the user provides only performance targets—such as a specific compressive strength class or slump limit—the agent enters an Autonomous Discovery phase. In this mode, it searches historical datasets for viable compositions and proposes a novel blend designed to satisfy the requested KPIs.

The execution of these paths is managed by an Iterative Task Executor which functions as a high-throughput digital laboratory. Upon receiving the classified task, the agent formulates a candidate concrete recipe and subjects it to a suite of Machine Learning Property Predictors. The resulting performance vector is evaluated against the objective function; if non-conformance is detected, the agent recalibrates the material ratios and re-initiates the simulation.

This cycle of simulation and refinement can repeat for up to 30 discrete iterations within a standard compute cycle. By conducting these rapid, non-linear simulations, the agent identifies the precise "sweet spot" where cementitious content is minimized to reduce the carbon footprint without falling below the required structural limits. This autonomous navigation of the "Plan-and-Execute" graph allows for a level of design precision and speed that exceeds traditional manual optimization, providing a scalable foundation for the decarbonization of concrete production.

Overall, ACORN’s architecture is modular and scalable. Each component communicates via well-defined interfaces, allowing updates or improvements to be made to one part (for example, swapping in a more advanced algorithm for data parsing or a new predictive model) without overhauling the entire system. At the core of ACORN is an AI-driven pipeline that links data ingestion, model prediction, and user interaction (Fig. 1). The system comprises several integrated components.

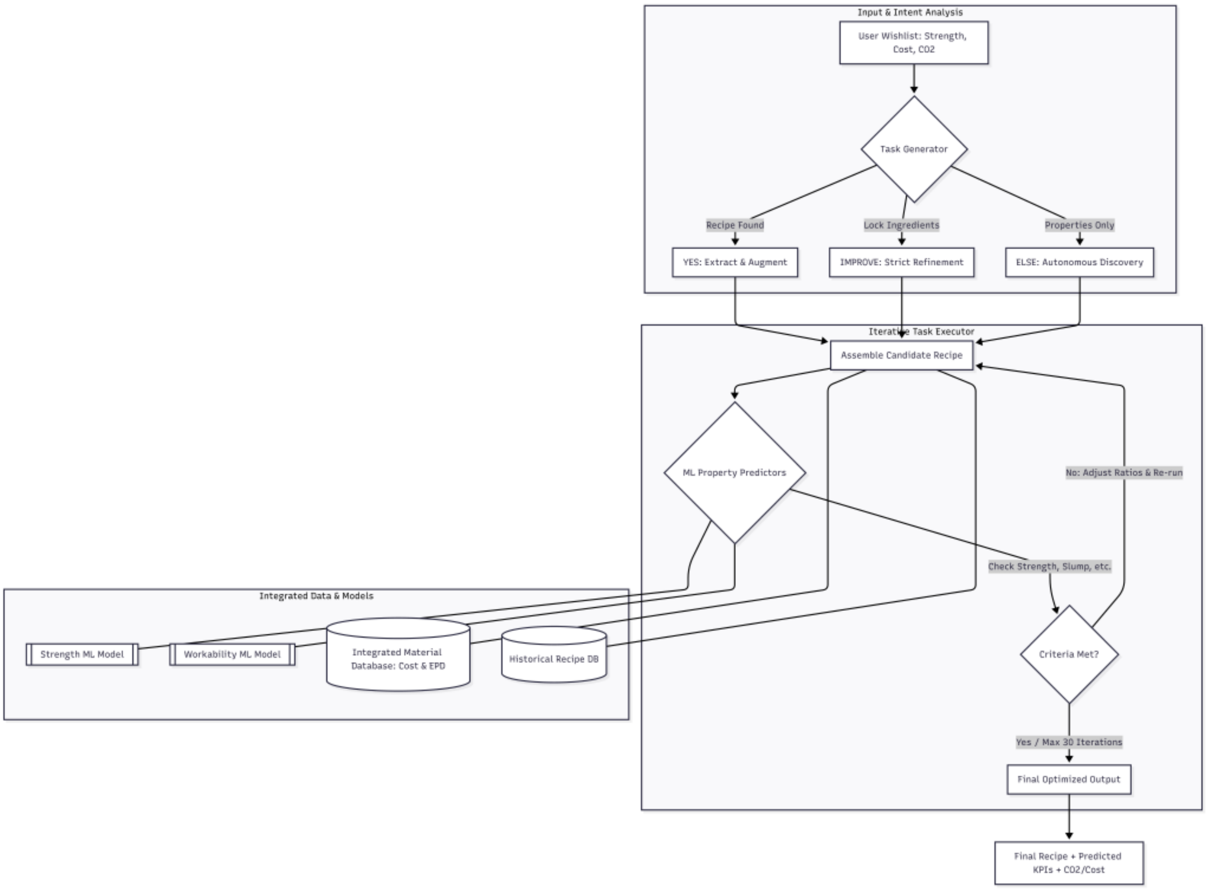

Figure 1. General concept/structure/workflow of ACORN

3. Data Collection and Database Development

One of the distinguishing aspects of ACORN is its emphasis on data-driven development. Building a high-quality, large-scale concrete data repository was a primary project objective, as it underpins the performance of the AI models. This section describes how the ACORN concrete recipe database was constructed, the types of data it includes, and how it uses laboratory data to validate the machine-learning development and data gathering.

3.1. Data Sources

In this study complementary data sources were brought together. On the other hand, a vast collection of concrete research data, including experimental results from publicly funded projects and published literature was implemented for model training. This covers mechanical properties (compressive and tensile strength results, consistency), durability performance (freeze-thaw resistance, chloride penetration, etc.), and novel material trials (e.g. mixes with alternative binders like fly ash, slag, calcined clays). Additionally, synthetic data from physical models was incorporated where available – for example, outputs from cement hydration modelling software that predicts strength gain based on chemistry and curing conditions. By merging industry data with academic and theoretical data, ACORN's database achieves an unprecedented breadth: it spans ordinary structural concretes to high-performance mixes and special concretes (self-compacting, fibre-reinforced, etc.), along with environmental and cost assessments.

3.2. Data Extraction with AI

The challenge in utilising these sources was their disparate formats. Industry data came largely in structured form (spreadsheets and EPD databases), whereas research results were often buried in PDFs of journal articles or reports. The ACORN project applied natural language processing to bridge this divide. As described in Section 2, a custom-tuned LLM was used to read and interpret unstructured documents. To maximise accuracy, a portion of the literature data was annotated manually and used to fine-tune the model — essentially teaching it how to recognize concrete mix information in text/tables by example. For instance, the model learned that a phrase like "w/c = 0.45" indicates water-cement ratio, or that "CEM II/B-V 42.5N" refers to a specific cement type. It could then extract such details from new documents at scale. The LLM approach dramatically accelerated data aggregation: in tests, it could process a 10-page research paper and pull out all relevant numerical data in under a minute, a task that would take hours for a human. By project's end, ACORN had processed hundreds of publications and technical reports, adding thousands of data points to the repository that would have been impractical to curate manually.

3.3. Data Cleaning and Structuring

After extraction, all data were standardized and imported into a structured database. A unified schema was designed to accommodate the varied data fields while maintaining consistency. Key fields include: mix identifier, component quantities (cement, water, each aggregate fraction, admixtures, additions, etc.), component properties (e.g. cement strength class, aggregate type), mixing process parameters (mixing energy or time, curing regime), test conditions (age, specimen size, curing environment for each result), and results (strength, slump, density, etc. along with units). Sustainability metrics like CO_2 equivalent emissions and cost were attached when available or calculated using known emission factors. Each entry in the database is traceable back to its source, allowing verification if something appears questionable. The anomaly detection routine was run on the compiled dataset to catch any outliers introduced during data merge. For example, if a certain compressive strength value was far lower than expected for the given w/c ratio and cement content, it would be flagged for review – it might indicate a typo in the source or an unusual test (like intentionally under-curing the sample). These safeguards ensured the final database was both comprehensive and reliable.

3.4. Database Platform

The resulting database was implemented using a NoSQL database system with a web-service layer on Ecometrix's cloud platform. This allowed the predictive model training code to query data efficiently and the user interface to retrieve relevant records on demand (for instance, to show similar past recipes when a user is exploring a new mix). Data security and access control were built in: proprietary data from companies were anonymised and aggregated when used in the public model, and the platform supports partitioning such that a company could opt to use ACORN with its own private data supplementing the base database (via Identity and Access Management, prepared for future deployment). This is important for industry adoption, as companies need assurance that their mix formulations — often considered trade secrets — are not exposed to competitors. By structuring the system to keep each user's contributions compartmentalised (unless they choose to share them), ACORN encourages data sharing and collaboration without compromising confidentiality. The database, as a living asset, will continue growing. Plans are in place to incorporate international data (through partners in EU and beyond) and to update entries as new cements or admixtures come to market in the drive for carbon-neutral concrete by 2045. In summary, the rigorous approach to building the ACORN database – leveraging AI to gather data and expert knowledge to vet it – provides a solid bedrock for the subsequent prediction and optimization capabilities.

4. Validation of Predictive Models and Results

A crucial aspect of ACORN's development was the validation of its predictive accuracy. Since the ultimate aim is to trust AI recommendations for concrete mixes, the models must be rigorously tested against real-world data. The validation focused on two levels: (1) validating that the AI's predictions for a given recipe match experimental results within acceptable error margins, and (2) validating the overall benefit of the system in reducing development iterations.

4.1. Predictive Accuracy Testing

The ACORN model's performance was first evaluated using a reserved test set from the database (comprising data points never seen by the model during training). On this test set, the model achieved high accuracy. For example, the mean absolute error in 28-day compressive strength prediction was around 3–5 MPa across a 20–60 MPa strength range, which is within ~10% of typical values – a level of precision comparable to the variability of lab tests themselves. Similarly, the model could estimate CO_2 emissions of mixes with a small error (within 5–10% of detailed EPD calculations). To further validate these data-driven findings physical experiments on a set of concrete mixes that were deliberately not part of the training data was performed. Several mix designs were selected, covering a spectrum from conventional (e.g. C30/37 strength class with Portland cement) to innovative (e.g. a mix with 50% cement replacement by slag and a new superplasticizer). These mixes were prepared in laboratory and tested for workability after mixing and compressive strength at 7, 28, and 56 days, and their mix details were fed into ACORN. The predictions were then compared to the actual measured results. An example outcome: for a high-slag concrete, ACORN predicted 28-day strength of 48 MPa, and the lab test returned 46 MPa – a difference of only +4.3%. In another case, a high-strength mix (target >80 MPa) was predicted to reach 82 MPa at 28 days; the lab test was 78 MPa, a 5% difference. Table 1 summarizes a subset of the validation results.

Table 1. Example validation results for ACORN's predictive model against laboratory tests. Positive error means over-prediction.

Mix ID	**Description**	**Measured 28d Strength** (MPa)	**Predicted 28d Strength** (MPa)	**Error** (%)	**Measured Slump** (mm)	**Predicted Slump** (mm)
A1	CEM I, *w/c*=0.55 (standard C30/37)	37.0	38.5	+4.1	65	70
B3	CEM III/B (65% GGBS), *w/c*=0.45	46.0	48.0	+4.3	55	50
C7	CEM I + silica fume, *w/c*=0.33 (high strength)	78.0	82.0	+5.1	210 (flow)	190
D2	CEM II/B-V, *w/c*=0.50 + Limestone filler	42.5	41.9	–1.6	68	65

w/c – is water to cement ratio, GGBS – Granulated Blast Furnace Slag, CEM I – Ordinary Portland Cement, CEM III/B – blended cement with GGBS, CEM II/B-V – pozzolan-blended cement with Coal Fly Ash

As seen, the predictive model showed strong agreement with experimental outcomes for both strength and workability in these trials. All errors were within about ±5%, which is generally acceptable for engineering use (given normal variability in concrete testing). It is worth noting that ACORN slightly over-predicted some high-strength results; this is a known tendency of data-driven models when extrapolating beyond the core of training data. The team is addressing this by incorporating more high-strength data and applying model bias correction based on these tests. Overall, these validation tests gave confidence that ACORN's predictions are reliable enough for practical decision support. They also helped identify areas for improvement, such as better modelling of special mix types like fibre-reinforced or self-compacting concretes, which can be added as the database expands.

4.2. System-Level Validation – Efficiency Gains

Beyond point predictions, ACORN's value lies in reducing the effort to arrive at an optimal mix. To gauge this a comparative study of a mix design optimisation with and without ACORN was performed. The task was to develop a concrete mix meeting a specified strength (50 MPa at 28 days) and a stringent carbon footprint limit (e.g. <250 kg CO_2/m^3, significantly lower than typical for that strength). Traditionally, an experienced concrete researcher might approach this by educated guesses – e.g. trying a mix with 50% slag replacement, then adjusting if the strength falls short – possibly requiring numerous batches. With ACORN, the researcher instead used the tool to simulate various combinations in silico. ACORN was used to screen about 20 candidate recipes in a few hours, identifying two promising formulations that met the CO_2 criterion. These two were then produced in the lab; one achieved the required 50 MPa (the other fell slightly short at 47 MPa). In contrast, a completely manual approach could have easily taken 8–10 trial mixes to find a compliant recipe, especially because the carbon limit was pushing the use of cement alternatives into unfamiliar territory. While this is just one case study, it illustrates ACORN's potential to cut down the "design loop" dramatically. In general, internal project tracking showed that even in early use by project partners, ACORN reduced the number of physical test iterations by approximately 50% for reaching a target formulation. This proportion is expected to improve further as the predictive accuracy and generative guidance features mature.

4.3. Limitations and Ongoing Validation

ACORN's current limitations were also examined during validation. One important consideration is extrapolation: the model is only as good as the data it has seen. If a user in the future queries a very novel mix (say, a completely new binder or an ultra-high-performance mix far outside the training data range), ACORN might respond with higher uncertainty or potentially misleading predictions. Users must be aware that the tool is meant to assist, not fully replace, expert judgement – especially outside the validated envelope. Accordingly, the system includes warnings and encourages follow-up testing for high-risk predictions. Another aspect is speed and scalability. Originally, the project targeted an instant (few seconds) response for predictions to facilitate interactive use. The initial deployment meets this for prediction mode. However, for the planned generative mode (where the system iteratively searches for an optimal recipe), the current implementation using an agent-based AI approach can take a few minutes to converge on a solution. This was observed in the first beta of the recipe generation feature: by employing an AI agent that iteratively adjusts the mix and evaluates guardrail conditions, the search produces excellent results but at the cost of computation time (several minutes per query). While still much faster than days or weeks of lab work, this is an area for improvement. The development team is exploring optimisations and more efficient search algorithms so that even complex generative queries can yield answers in under 30 seconds, aligning with user expectations for a seamless design experience.

In summary, the validation phase confirmed that ACORN's predictive models are technically sound and sufficiently accurate for guiding concrete development. The collaboration between AI predictions and human expertise (in planning which predictions to trust and how to act on them) proved effective. ACORN effectively serves as an ``AI assistant'' for concrete technologists, rapidly sifting through possibilities and pointing toward viable solutions — much like a seasoned advisor that learns from a vast repository of past concrete knowledge. The next section will discuss the broader implications of these results and how ACORN can be deployed moving forward.

Conclusions

The ACORN project has demonstrated a successful integration of modern AI techniques into the domain of concrete materials, yielding a tool that can substantially enhance the way concrete mixes are designed. Here we conclude the paper by highlighting the key outcomes, implications for practice, and future developments.

Presented work validated the concept that data-driven, AI-assisted concrete mix design is feasible and valuable. By constructing one of the largest curated concrete recipe databases to date and leveraging it in deep learning models, the project achieved reliable predictions of concrete performance from mix composition. The AI can balance multiple objectives, providing a more holistic optimization than traditional methods that often focus on strength alone.

The system architecture – particularly the use of LLMs for data collection – showcases how AI can unlock and unify existing knowledge scattered in technical documents. This method could be extended to continuously harvest new data, keeping the system's knowledge up-to-date.

The advent of generative AI could mark a paradigm shift in the concrete industry. Concrete producers and designers can drastically cut the time and cost needed to develop new mixes. This means faster introduction of low-carbon concretes and the ability to respond to project-specific requirements with bespoke mixes without prohibitive development effort. As the case studies indicated, ACORN can facilitate mix optimisations that lower cement content or incorporate more waste-derived materials while still meeting performance requirements. At scale, such optimisations can yield significant carbon reductions. For instance, if major concrete producers achieve even a 10% CO_2 reduction in their mix designs through AI-guided optimisation, the cumulative effect on global emissions would be very substantial (consider that one partner company in India with interest in ACORN produces ~10 million m^3 of concrete annually – improvements there alone could save hundreds of thousands of tonnes of CO_2 per year).

The model prototype is a stepping stone towards a more expansive vision of AI-driven concrete design. One immediate next step is the implementation of the generative recipe design feature. The concept is to invert the predictive model into a prescriptive tool: users would input desired properties (e.g. 50 MPa strength, 500 mm flow, <200 kg CO_2/m^3) and ACORN would intelligently search for a recipe that meets those criteria. Preliminary work has started using evolutionary algorithms and advanced optimisation techniques in combination with the predictive model to achieve this.

Acknowledgements

The ACORN project was funded by Vinnova (Swedish Innovation Agency) under grant no. 2024-00521. The authors gratefully acknowledge this support. We also thank the Swedish Concrete Association (Svensk Betong) and its member companies for providing data and valuable feedback during the development.

References

Elhishi, S. et al. (2023). "Unboxing machine learning models for concrete strength prediction using XAI." Scientific Reports, 13, 19892.

Völker, C., et al. (2023). "Presenting SLAMD – A Sequential Learning based Software for the Inverse Design of Sustainable Cementitious Materials." NanoWorld Journal, 9(S2), 180–187.

fib Bulletin 102 (2022). Guide for Protection and Repair of Concrete Structures: Guide for Good Practice. Fédération Internationale du Béton, Lausanne, Switzerland.

Neville, A.M. (1991). Properties of Concrete, 3rd ed. Longman Scientific & Technical, Essex, UK.