Q&A: More-sustainable concrete with machine learning

MIT-IBM Watson AI Lab researchers aim to design concrete mixtures that use AI to shrink environmental footprint and cost, while recycling byproducts and increasing performance.

Lauren Hinkel | MIT-IBM Watson AI Lab
December 7, 2021

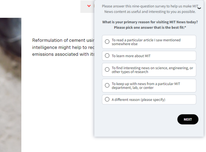

As a building material, concrete withstands the test of time. Its use dates back to early civilization, and today it is the most popular composite choice in the world. However, it's not without its faults. Production of its key ingredient, cement, contributes 8-9 percent of the global anthropogenic CO_2 emissions and 2-3 percent of energy consumption, which is only projected to increase in the coming years. With aging United States infrastructure, the federal government recently passed a milestone bill to revitalize and upgrade it, along with a push to reduce greenhouse gas emissions where possible, putting concrete in the crosshairs for modernization, too.

Elsa Olivetti, the Esther and Harold E. Edgerton Associate Professor in the MIT Department of Materials Science and Engineering, and Jie Chen, MIT-IBM Watson AI Lab research scientist and manager, think artificial intelligence can help meet this need by designing and formulating new, more sustainable concrete mixtures, with lower costs and carbon dioxide emissions, while improving material performance and reusing manufacturing byproducts in the material itself. Olivetti's research improves environmental and economic sustainability of materials, and Chen develops and optimizes machine learning and computational techniques, which he can apply to materials reformulation. Olivetti and Chen, along with their collaborators, have recently teamed up for an MIT-IBM Watson AI Lab project to make concrete more sustainable for the benefit of society, the climate, and the economy.

Q: What applications does concrete have, and what properties make it a preferred building material?

Olivetti: Concrete is the dominant building material globally with an annual consumption of 30 billion metric tons. That is over 20 times the next most produced material, steel, and the scale of its use leads to considerable environmental impact, approximately 5-8 percent of global greenhouse gas (GHG) emissions. It can be made locally, has a broad range of structural applications, and is cost-effective. Concrete is a mixture of fine and coarse aggregate, water, cement binder (the glue), and other additives.

Q: Why isn't it sustainable, and what research problems are you trying to tackle with this project?

Olivetti: The community is working on several ways to reduce the impact of this material, including alternative fuels use for heating the cement mixture, increasing energy and materials efficiency and carbon sequestration at production facilities, but one important opportunity is to develop an alternative to the cement binder.

While cement is 10 percent of the concrete mass, it accounts for 80 percent of the GHG footprint. This impact is derived from the fuel burned to heat and run the chemical reaction required in manufacturing, but also the chemical reaction itself releases CO_2 from the calcination of limestone. Therefore, partially replacing the input ingredients to cement (traditionally ordinary Portland cement or OPC) with alternative materials from waste and byproducts can reduce the GHG footprint. But use of these alternatives is not inherently more sustainable because wastes might have to travel long distances, which adds to fuel emissions and cost, or might require pretreatment processes. The optimal way to make use of these alternate materials will be situation-dependent. But because of the vast scale, we also need solutions that account for the huge volumes of concrete needed. This project is trying to develop novel concrete mixtures that will decrease the GHG impact of the cement and concrete, moving away from the trial-and-error processes towards those that are more predictive.

Chen: If we want to fight climate change and make our environment better, are there alternative ingredients or a reformulation we could use so that less greenhouse gas is emitted? We hope that through this project using machine learning we'll be able to find a good answer.

Q: Why is this problem important to address now, at this point in history?

Olivetti: There is urgent need to address greenhouse gas emissions as aggressively as possible, and the road to doing so isn't necessarily straightforward for all areas of industry. For transportation and electricity generation, there are paths that have been identified to decarbonize those sectors. We need to move much more aggressively to achieve those in the time needed; further, the technological approaches to achieve that are more clear. However, for tough-to-decarbonize sectors, such as industrial materials production, the pathways to decarbonization are not as mapped out.

Q: How are you planning to address this problem to produce better concrete?

Olivetti: The goal is to predict mixtures that will both meet performance criteria, such as strength and durability, with those that also balance economic and environmental impact. A key to this is to use industrial wastes in blended cements and concretes. To do this, we need to understand the glass and mineral reactivity of constituent materials. This reactivity not only determines the limit of the possible use in cement systems but also controls concrete processing, and the development of strength and pore structure, which ultimately control concrete durability and life-cycle CO_2 emissions.

Chen: We investigate using waste materials to replace part of the cement component. This is something that we've hypothesized would be more sustainable and economic — actually waste materials are common, and they cost less. Because of the reduction in the use of cement, the final concrete product would be responsible for much less carbon dioxide production. Figuring out the right concrete mixture proportion that makes endurable concretes while achieving other goals is a very challenging problem. Machine learning is giving us an opportunity to explore the advancement of predictive modeling, uncertainty quantification, and optimization to solve the issue. What we are doing is exploring options using deep learning as well as multi-objective optimization techniques to find an answer. These efforts are now more feasible to carry out, and they will produce results with reliability estimates that we need to understand what makes a good concrete.

Q: What kinds of AI and computational techniques are you employing for this?

Olivetti: We use AI techniques to collect data on individual concrete ingredients, mix proportions, and concrete performance from the literature through natural language processing. We also add data obtained from industry and/or high throughput atomistic modeling and experiments to optimize the design of concrete mixtures. Then we use this information to develop insight into the reactivity of possible waste and byproduct materials as alternatives to cement materials for low-CO_2 concrete. By incorporating generic information on concrete ingredients, the resulting concrete performance predictors are expected to be more reliable and transformative than existing AI models.

Chen: The final objective is to figure out what constituents, and how much of each, to put into the recipe for producing the concrete that optimizes the various factors: strength, cost, environmental impact, performance, etc. For each of the objectives, we need certain models: We need a model to predict the performance of the concrete (like, how long does it last and how much weight does it sustain?), a model to estimate the cost, and a model to estimate how much carbon dioxide is generated. We will need to build these models by using data from literature, from industry, and from lab experiments.

We are exploring Gaussian process models to predict the concrete strength, going forward into days and weeks. This model can give us an uncertainty estimate of the prediction as well. Such a model needs specification of parameters, for which we will use another model to calculate. At the same time, we also explore neural network models because we can inject domain knowledge from human experience into them. Some models are as simple as multi-layer perceptions, while some are more complex, like graph neural networks. The goal here is that we want to have a model that is not only accurate but also robust — the input data is noisy, and the model must embrace the noise, so that its prediction is still accurate and reliable for the multi-objective optimization.

Once we have built models that we are confident with, we will inject their predictions and uncertainty estimates into the optimization of multiple objectives, under constraints and under uncertainties.

Q: How do you balance cost-benefit trade-offs?

Chen: The multiple objectives we consider are not necessarily consistent, and sometimes they are at odds with each other. The goal is to identify scenarios where the values for our objectives cannot be further pushed simultaneously without compromising one or a few. For example, if you want to further reduce the cost, you probably have to suffer the performance or suffer the environmental impact. Eventually, we will give the results to policymakers and they will look into the results and weigh the options. For example, they may be able to tolerate a slight further cost under a significant reduction in greenhouse gas. Alternatively, if the cost varies little but the concrete performance changes drastically, say, doubles or triples, then this is definitely a favorable outcome.

Q: What kinds of challenges do you face in this work?

Chen: The data we get either from industry or from literature are very noisy; the concrete measurements can vary a lot, depending on where and when they are taken. There are also substantial missing data when we integrate them from different sources, so, we need to spend a lot of effort to organize and make the data usable for building and training machine learning models. We also explore imputation techniques that substitute missing features, as well as models that tolerate missing features, in our predictive modeling and uncertainty estimate.

Q: What do you hope to achieve through this work?

Chen: In the end, we are suggesting either one or a few concrete recipes, or a continuum of recipes, to manufacturers and policymakers. We hope that this will provide invaluable information for both the construction industry and for the effort of protecting our beloved Earth.

Olivetti: We'd like to develop a robust way to design cements that make use of waste materials to lower their CO_2 footprint. Nobody is trying to make waste, so we can't rely on one stream as a feedstock if we want this to be massively scalable. We have to be flexible and robust to shift with feedstocks changes, and for that we need improved understanding. Our approach to develop local, dynamic, and flexible alternatives is to learn what makes these wastes reactive, so we know how to optimize their use and do so as broadly as possible. We do that through predictive model development through software we have developed in my group to automatically extract data from literature on over 5 million texts and patents on various topics. We link this to the creative capabilities of our IBM collaborators to design methods that predict the final impact of new cements. If we are successful, we can lower the emissions of this ubiquitous material and play our part in achieving carbon emissions mitigation goals.

Other researchers involved with this project include Stefania Jegelka, the X-Window Consortium Career Development Associate Professor in the MIT Department of Electrical Engineering and Computer Science; Richard Goodwin, IBM principal researcher; Soumya Ghosh, MIT-IBM Watson AI Lab research staff member; and Kristen Severson, former research staff member. Collaborators included Nghia Hoang, former research staff member with MIT-IBM Watson AI Lab and IBM Research, and Executive Director of MIT Climate & Sustainability Consortium Jeremy Gregory.

This research is supported by the MIT-IBM Watson AI Lab.

SHARE

RELATED LINKS
- Jie Chen
- Elsa Olivetti
- MIT-IBM Watson AI Lab
- School of Engineering
- MIT Schwarzman College of Computing

RELATED TOPICS
Interview, Concrete, Artificial intelligence, Infrastructure, Construction, Civil and environmental engineering, Building, Industry, Emissions, Carbon dioxide, Greenhouse gases, Pollution, Sustainability, Climate change, DMSE, Concrete Sustainability Hub, Electrical Engineering & Computer Science (eecs), Global Warming, Faculty, Research, MIT-IBM Watson AI Lab, Machine learning, Natural language processing, Data, Computer science and technology, School of Engineering, MIT Schwarzman College of Computing

RELATED ARTICLES

AI UN climate change conference, trying to "keep 1.5 alive"

MIT makes strides on climate action plan

2,200-year-old Roman tomb offers insights on ancient concrete resilience

Making roadway spending more sustainable

Predicting building emissions across the US

More MIT News

Q&A: A conversation with Helen Elaine Lee about her novel, "Pomegranate"
The MIT professor's new book explores the world of a woman set free from prison and redefining herself in society.

Researchers teach an AI to write better chart captions
A new dataset can help scientists develop automatic systems that generate richer, more descriptive captions for online charts.

Transatlantic connections make the difference for MIT Portugal
The international partnership focuses on climate and sustainability.

Summer 2023 recommended reading from MIT
Enjoy these recent titles from Institute faculty and staff.

Studies at the intersection of equity, computing, and education
"The work I'm doing is deeply rooted in the belief that you can plant seeds in people," says graduate student Cecilé Sachter.

Computer vision system marries image recognition and generation
MAGE merges the two key tasks of image generation and recognition, typically trained separately, into a single system.

More news on MIT News homepage →