Machine learning for strength prediction and optimal design of sustainable concrete formulas

by

Olivia Pfeiffer

B.S., Mechanical Engineering

University of Massachusetts Amherst (2019)

Submitted to the Institute for Data, Systems, and Society
and the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees
of
Master of Science in Technology and Policy
and
Master of Science in Electrical Engineering and Computer Science
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
May 2022

Author .....
Institute for Data, Systems, and Society
and the Department of Electrical Engineering and Computer Science
May 13, 2022

Certified by .....
Elsa A. Olivetti
Associate Professor, Department of Materials Science and Engineering
Thesis Supervisor

Certified by .....
Stefanie Jegelka
Associate Professor, Department of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by .....
Noelle Eckley Selin
Professor, Institute for Data, Systems, and Society and
Department of Earth, Atmospheric and Planetary Sciences
Director, Technology & Policy Program

Accepted by .....
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Machine learning for strength prediction and optimal design of sustainable concrete formulas

by

Olivia Pfeiffer

Submitted to the Institute for Data, Systems, and Society
and the Department of Electrical Engineering and Computer Science
on May 13, 2022, in partial fulfillment of the
requirements for the degrees of
Master of Science in Technology and Policy
and
Master of Science in Electrical Engineering and Computer Science

Abstract

Given the large environmental impact of the concrete industry, which represents 8-9% of global CO_2 emissions, the design of concrete mixes with low carbon footprints that still meet structural performance requirements will be an essential part of global decarbonization efforts. In this work, we build a concrete performance model, which maps from concrete constituents to compressive strength, a key structural property. Specifically, we leverage the quantity and quality of information provided by our industrial concrete partners (whereas most existing related studies use small, narrow datasets derived from laboratory experiments) to establish an improved concrete performance model that captures the role of several concrete ingredients and a wide variety of formulas. We find that the features which are predicted to be important to concrete strength are compatible with industry knowledge, and that predictions can be improved in the case of small datasets by leveraging information from other larger datasets. Additionally, we integrate our machine learning model into an optimization procedure, and identify mixtures which have minimal cost and minimal climate impact. Lastly, we discuss the trade-offs between these two design parameters, and how these considerations differ by the required strength of the concrete.

Thesis Supervisor: Elsa A. Olivetti
Title: Associate Professor, Department of Materials Science and Engineering

Thesis Supervisor: Stefanie Jegelka
Title: Associate Professor, Department of Electrical Engineering and Computer Science

Acknowledgments

First and foremost, I would like to thank my research advisor, Elsa Olivetti, for her endless support and guidance over these three years. I am extremely grateful that you encouraged me to work on projects most aligned with my interests and gave me the opportunity to do so.

Thank you to my thesis supervisor, Stefanie Jegelka, for being a constant source of knowledge; your comments always pushed me a bit further.

I would like to thank the others that made this thesis possible: Kristen, Jeremy, Kai, and the IBM Team. Thank you for your insightful input at every step of the way.

Thank you to my peers who made my time at MIT so enjoyable. Thank you to the Olivetti Group, for our weekly lunch conversations, movie nights, and overall camaraderie. A special shout out to the Synthesis Project Subgroup members- Zach, Vineeth, Chris, and Elton- for always helping me through whatever coding troubles I was facing. Thank you to the TPP community, especially Ragini and Abhishek, for your support and friendship throughout.

Lastly, I would like to thank my family, who have been by my side on this journey and beyond. I could not have done this without you all.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Problem setup

Concrete is the most widely used building material in the world, with a global annual consumption of approximately 30 billion metric tons [1]. The demand for concrete continues to grow at a rate faster than steel and wood, and the versatility and durability of concrete has made it a fundamental building staple for centuries [1]. In particular, concrete will be necessary to create infrastructure for a climate-resilient future, where buildings, roads, and bridges are long-lasting and resistant to disaster.

Concrete is made up of two basic components: 1) aggregates (rock fragments), and 2) a paste that binds the aggregates together. Aggregates may be coarse (e.g., crushed stone, gravel) or fine (e.g., sand). The paste that is used in concrete is traditionally made by mixing ordinary Portland cement (OPC) and water. The manufacturing of OPC involves heating limestone and clay in a kiln to a sintering temperature of 1450°C. This calcination process causes the limestone ($CaCO_3$) to decompose into lime (CaO) and carbon dioxide (CO_2). When also including the emissions associated with fossil fuel combustion required to generate the heat, it is estimated that one metric ton of CO_2 is released for every one metric ton of cement produced [1]. In fact, the manufacturing of OPC is responsible for 8% of global CO_2 emissions [2].

The partial replacement of OPC with supplementary cementitious materials (SCMs), such as those derived from industrial wastes, can therefore both improve resulting con-

crete properties and reduce the overall carbon footprint of concrete. Commonly used SCMs include fly ash (a by-product of burning coal), slag (a by-product of iron production), and silica fume (a by-product of silicon metal production). The inclusion of SCMs can increase the long-term strength of concrete, improve durability, and mitigate permeability, among other effects. Almost all of produced ground granulated blast furnace slag (GGBFS) is used in concrete; its usage is limited by supply. However, this example is the exception and most SCMs are under-utilized in the concrete industry, due to challenges with understanding the performance of various compositions. The design of environmentally friendly concrete mixtures will be an essential part of global decarbonization efforts. Thus, there is an opportunity to use machine learning to improve the design of concrete formulas, achieving lower environmental impacts while still meeting structural performance requirements.

1.2 Literature and Contribution

Numerous groups in the literature have explored machine learning (ML) as a means to predict the properties of a given concrete formula, and some have extended this to optimal mixture design. In their review of machine learning for the prediction of mechanical properties of concrete, Ben Chaabene *et al.* [3] reviewed nearly 100 instances of machine learning algorithms applied to concrete data. Across all, the average size of the dataset used was about 451 data points, where few (16) had size greater than 1,000 observations, and only 1 had size greater than 2,000 observations. A range of machine learning models have been applied to concrete in the literature, including support vector machines (SVM), artificial neural networks (NN), decision trees, and hybrid models. Many input features are considered, such as type and quantity of constituents.

Young *et al.* [4] predicted compressive strength using data (roughly 10,000 observations) from a vertically integrated cement/concrete industrial producer (VIP), and also evaluated their model performance on a benchmark dataset from the academic literature. Four models were considered: linear regression, NN, random forest (RF),

and SVM. They found that the RF model achieved the lowest root mean squared error (RMSE) and best R^2 for both datasets, and that overall model performance was slightly poorer on the VIP dataset than on the lab benchmark. They believed this to be due to differences in measurement style or ambient conditions at job sites (as opposed to the controlled laboratory setting) and suggested that future work should aim to expand the size of the data and incorporate a wider set of input variables for improved results using job site data. Young *et al.* [4] then integrates their NN model (which had available gradients, unlike the RF) into an optimization procedure to design concrete mixtures which minimize cost while still meeting target strengths. The authors also consider a second optimization which minimizes mixture cost subject to both target strength constraints and target embodied carbon impact.

While the study by Young *et al.* [4] provides many new insights on the application of machine learning to industrial-based datasets, it is limited in the following aspects. First, the concrete mixtures that are considered in the study contain only fly ash as SCM; they do not contain slag nor silica fume, which are two other key SCMs that are often used in concrete. Second, the design strengths of the mixtures range from 30 MPa (~4,300 psi) to 55 MPa (~8,000 psi); information on the minimum costs that can be achieved in the case of high-strength projects (e.g., >10,000 psi) remains unknown. Additionally, the authors consider environmental impact via a constraint in their optimization of cost; it would be beneficial to optimize directly for low carbon footprint. Lastly, when faced with multiple design criteria, it is valuable to understand the trade-offs via a multi-objective optimization, i.e., a visualization of the pareto front with each objective along the axes.

Many researchers have investigated the design of more environmentally friendly concrete, such as those containing recycled aggregates [5–7], those which use industrial waste as supplementary cementitious materials [8–10], and other materials, such as grass for lightweight aggregate [11] and steel fibers [12].

Zhang *et al.* [10] couple an ML model for strength and optimization for design of concrete mixtures. They formulate a tri-objective optimization which is posed as a weighted sum of the following three objective functions: 1) maximize compressive

strength, 2) minimize cost, and 3) minimize embodied CO_2. The compressive strength was estimated using a NN. Input variables included content of cement, silica fume, water, coarse aggregate, fine aggregate, maximum size of coarse aggregate, high-range water reducing admixture, and curing age. They modified the beetle antennae search (BAS) algorithm so that it could be applied in a multi-objective setting. The authors validated the performance of their multi-objective beetle antennae search algorithm (MOBAS) on a number of multi-objective benchmark functions and showed that it performed comparably with other heuristic optimization algorithms, including genetic algorithm (GA), particle swarm optimization (PSO), and differential evolution (DE), with significantly lower computation time. The authors also performed global sensitivity analysis and determined that after age, high-range water reducing admixture content has the most significant impact on compressive strength.

Nunez *et al.* [6] studied the prediction of compressive strength of concrete mixtures which contain recycled aggregates, called recycled aggregate concrete (RAC). They compiled an experimental dataset (n=1,134) from available studies (55 peer-reviewed publications) in the literature. Input variables include water-to-cement ratio (w/c), cement content, sand content, recycled aggregate content, high-range water reducing admixture content, silica fume content, age, specimen type, and 28-day design strength (f'c). The authors explored three different ML models for predicting the compressive strength of the RAC, namely recurrent neural network RNN, gradient-boosted regression tree (GBRT), and Gaussian processes (GP). Separately, the authors applied the PSO algorithm to minimize the cost of a concrete mixture, using four sets of bounds associated with four different target strength classes. Then, the best-performing ML model (in this case, the GBRT) was used in order to verify that the predicted strength of optimal mixtures in each of the four strength classes met or exceeded specified targets. The authors note that future work should integrate SCMs.

Pereira Dias *et al.* [11] use data driven modeling for the design of concrete which uses Miscanthus grass as lightweight aggregate. There are nine inputs considered (Miscanthus, cement, lime, high-range water reducing admixture, $CaCl_2$, water, form

of specimen, curing time, and pre-treatment condition). They create 413 specimens which are evaluated at 7, 14, and 28 days after casting. The authors employ Gaussian process regression. The authors create a useful graphical user interface with knobs which allow a concrete designer to visualize the prediction compressive strengths and sensitivities to selected parameter values.

Using a large dataset provided by an industrial concrete producer, we leverage rich information to build an improved concrete performance model, mapping from concrete mix attributes to key properties and explicitly modeling time as a dimension of the problem. With this model established, we formulate an optimization procedure which aligns with a concrete designer's needs and enables them to better understand the trade-offs between structural performance, environmental impact, and cost. This enables the exploration of various potential mixtures before investing in and carrying out laboratory experimental work, and ultimately this framework can be used to improve utilization of regionally available raw materials at scale.

Chapter 2

Compressive strength prediction on industrial datasets

2.1 Data

Through an initiative launched in partnership with industrial concrete producers, we curated a database from three companies (referred to as C1, C2 and C3) of concrete mixture attributes and their corresponding compressive strength measurements, totaling over 200,000 observations. Although strength is measured at various ages, the most important age is 28 days, when approximately 99% of the strength has developed. After this day, the gain in strength is minimal. Many building specifications are thus based upon the strength achieved at 28 days. In this section, we focus on prediction of the 28-day strength specifically.

As the raw data provided by the companies was in various formats, we first identified a set of standard, or generic, concrete mixture constituents and then processed the data accordingly. The generic mixture constituents include: coarse aggregate, fine aggregate, water, cement, fly ash, slag, silica fume, and five chemical admixtures (high-range water reducing, water reducing, accelerating, air-entraining, and other (e.g., viscosity modifying) admixtures). Additionally, constituent quantities were divided by the total yielded volume (provided by the companies) when necessary in order to obtain mass quantities required to produce one cubic yard (CY) of concrete.

All units are converted to the English system, with masses of admixtures measured in ounces (oz), masses of all other constituents in pounds (lbs), and compressive strengths in pounds per square inch (psi). Lastly, outlier cleaning was performed on an individual company basis in order to remove values more than three standard deviations from the mean of each constituent, as these are believed to be errors. We also include a benchmark dataset (referred to as L) by Yeh [13] which consists of 1,030 strength observations. This dataset is also pre-processed in the same manner as the industrial datasets. The sizes of the subsets of the four datasets which are relevant in the 28-day strength prediction are listed in Table 2.2.

Each mixture can be represented by the 13-dimensional vector $\mathbf{z} = [z_0, ..., z_{12}]$. The first element, z_0, is a derived feature; it is the mass-based ratio of water to cementitious materials (w/cm), a key design attribute of concrete. The remaining elements z_i, with $i > 0$, are each the constituent amount used to produce one CY of concrete respectively as follows: pounds (lbs) of coarse aggregate, fine aggregate, water, cement, slag, fly ash, silica fume, and ounces (oz) of high-range water reducing admixture, water reducing admixture, accelerating admixture, air-entraining admixture, and other admixture.

The kernel density estimation (KDE) plots for the four datasets and the 12 constituent quantities and 28-day compressive strength are shown in Figure 2-1 and 2-2 respectively. The KDE is similar to a histogram, however, it smoothly estimates the probability density function of a random variable. The area under this curve is 1. Note that some features are not available in some datasets. This information is provided in tabular format in Table 2.1.

2.2 Models to predict compressive strength

2.2.1 Random forest

Developed by Breiman [14], a random forest is an ensemble method in which the predictions of its individual components, decision trees, are aggregated in order to

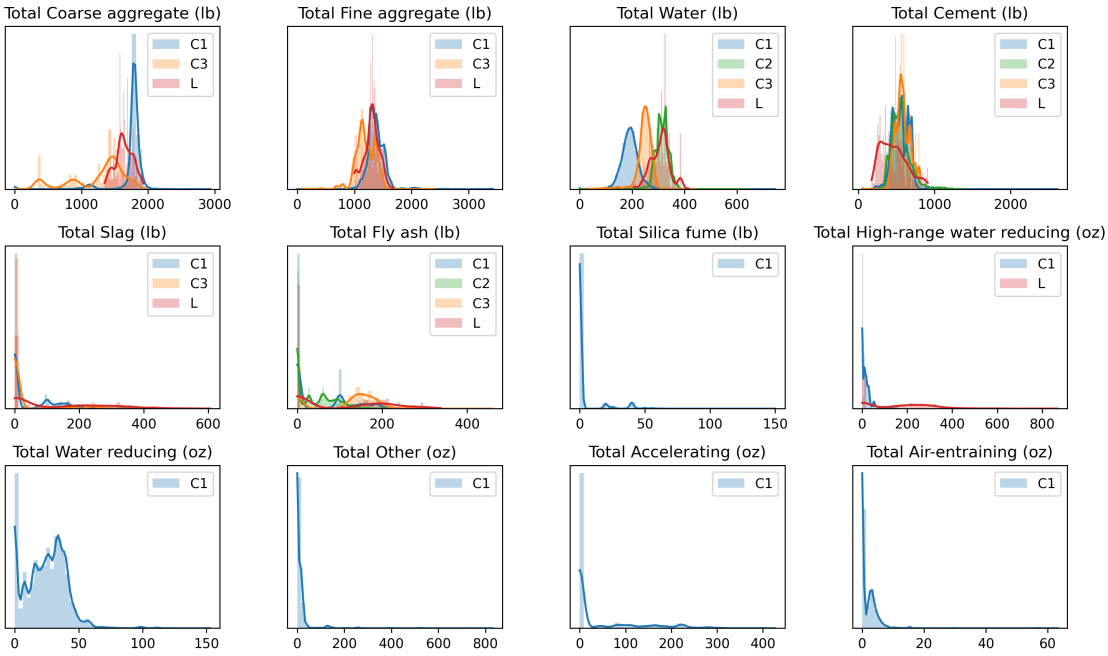

Figure 2-1: Kernel density estimation plots for the constituent quantities across the four datasets.

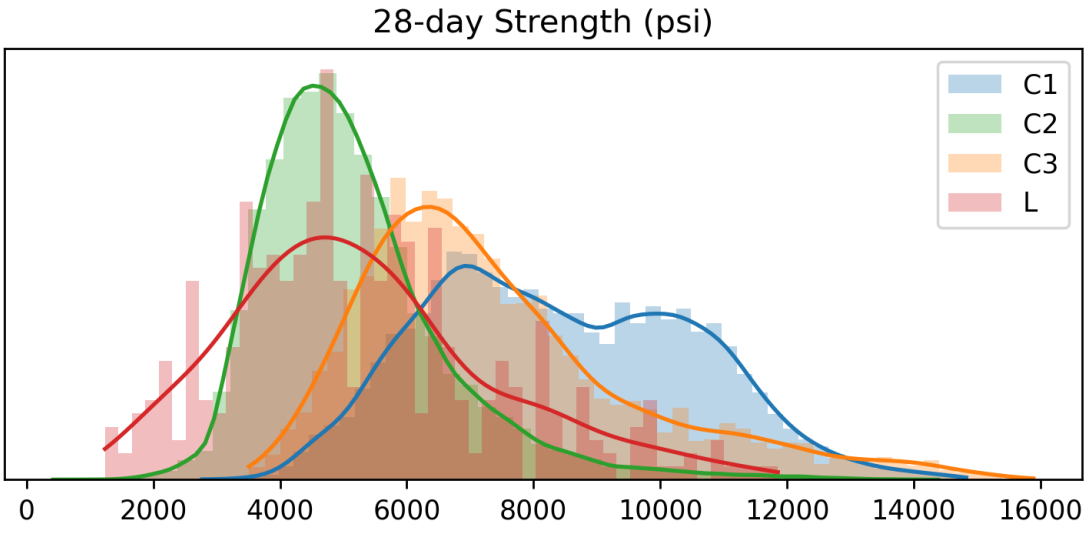

Figure 2-2: Kernel density estimation plots for the 28-day compressive strength across the four datasets.

Table 2.1: Features available for the four datasets.

Feature	C1	C2	C3	L
w/cm	✓	✓	✓	✓
Coarse aggregate	✓		✓	✓
Fine aggregate	✓		✓	✓
Water	✓	✓	✓	✓
Cement	✓	✓	✓	✓
Fly ash	✓	✓	✓	✓
Slag	✓		✓	✓
Silica fume	✓			
High-range water reducing admixture	✓			✓
Water reducing admixture	✓			
Accelerating admixture	✓			
Air-entraining admixture	✓			
Other admixture	✓			

Table 2.2: Sizes of datasets (train and test) for the 28-day strength prediction task.

Dataset	Training size	Testing size
C1	15,294	5,097
C2	42,696	14,231
C3	3,007	1,002
L	313	104

obtain a single output. Decision trees are non-parametric supervised machine learning techniques for both classification and regression tasks. The tree structure involves using rules to repeatedly split the data based on feature values in order to obtain predictions of the target. A random forest fits many decision trees to different subsamples of the dataset, a technique to mitigate overfitting. The random forest models used in this work were implemented using the `scikit-learn` [15] library for machine learning in Python.

Key hyperparameters for training the random forest include the number of decision trees in the forest, the number of features to consider at each split, and the maximum allowed depth of the trees. We found that the best number of features to consider at each split was 3. We train 100 trees with no maximum depth.

2.2.2 Multi-layer perceptron

A multi-layer perceptron (MLP) is the simplest class of artificial neural networks. Generally, a MLP is a feedforward, fully connected neural network with at least three layers (input, hidden layer, output) and nonlinear activation functions. Feedforward indicates that information in the network moves only in the forward direction (from input to output), without any cycles. Fully connected means that all of the nodes in one layer are connected to all of the nodes in the next layer.

The design and training of an MLP involves optimal selection of the number of hidden layers, the hidden layer sizes, activation functions, learning rate, and solver. In this work, the best architecture selected had an input of size 13, six hidden layers of size 50, output of size 1. The activation function is the rectified linear unit (ReLU). The MLP was trained for 151 epochs using a learning rate of 1e-3 using the Adam solver [16]. The MLP models in this study were implemented using the `pytorch` [17] library for deep learning in Python.

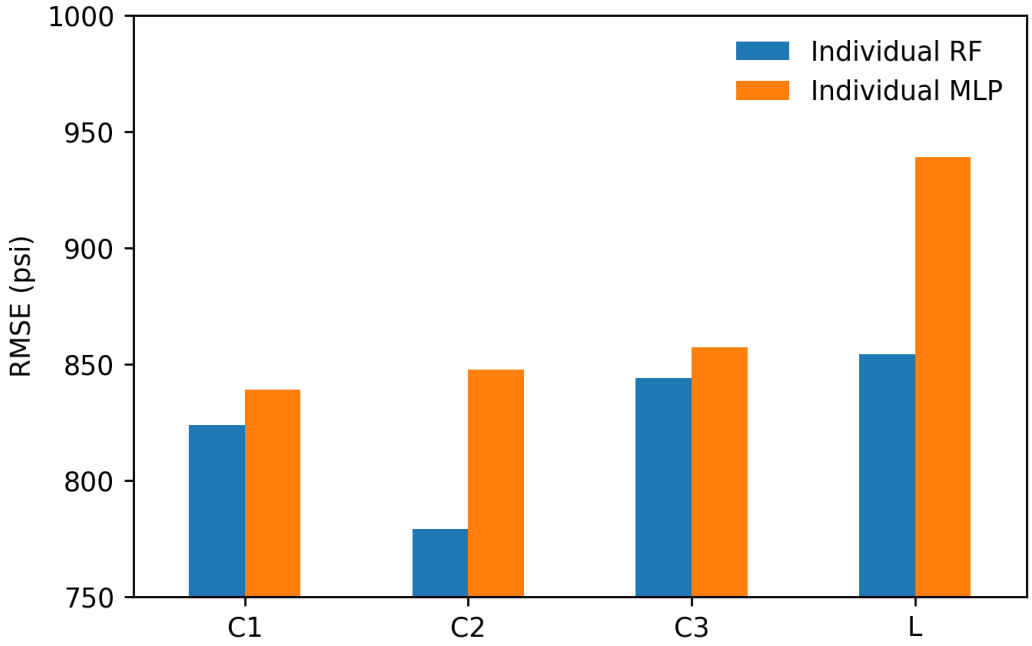

Figure 2-3: Test RMSE for the random forest and multi-layer perceptron on the four datasets.

2.2.3 Results

We find that the RF model has the best performance on all datasets, as shown in the plot of the root-mean-squared-error (RMSE) in Figure 2-3. The lowest RMSE achieved by the RF is on dataset C2, at 779 psi. It is noteworthy that this dataset had the fewest number of features available for predictions, however, it had the highest number of training points and the narrowest distribution of 28-day strength measurements. The RF is only moderately better than the MLP for the other two industrial datasets. In the case of dataset L, which had only 313 points available for training, it is likely that the MLP was not able to accurately infer parameters from this small dataset.

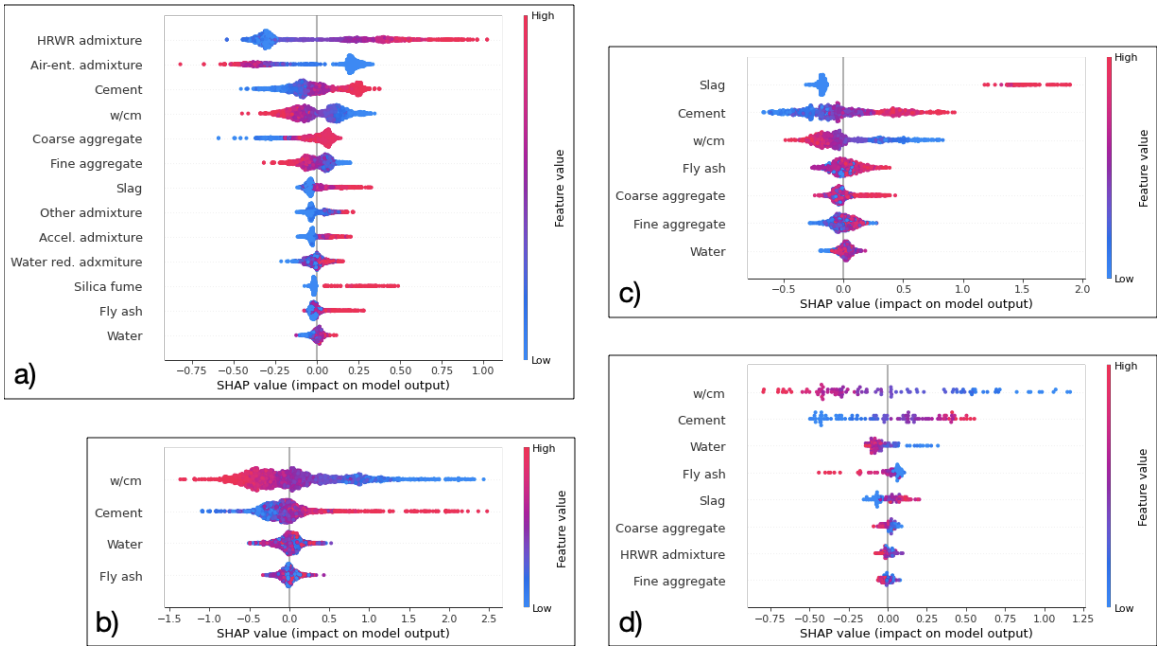

Figure 2-4: SHAP summary plots for the four datasets, with a, b, c, and d corresponding to C1, C2, C3, and L respectively.

2.3 Feature analysis: Comparison with existing domain knowledge

We use SHAP [18, 19], a Python package developed for interpreting model results, in order to explore whether features that are important for ML model predictions align with knowledge in materials science. SHAP stands for SHapley Additive exPlanations and uses a game theoretic approach to obtain additive feature attributions. We explore the feature importance of the RF model using the `TreeExplainer`. The summary plots for testing sets for the four datasets are shown in Figure 2-4, with a, b, c, and d corresponding to C1, C2, C3, and L respectively. In these summary plots, the x-axis represents the impact on the prediction, where values greater than 0 indicate an increase in the predicted 28-day strength, and values less than 0 indicate a decrease in the predicted 28-day strength. Features are marked along the y-axis and are ranked according to importance, with the most important at the top. The colors of the data points in the swarm for each feature indicate the value of the feature, where pink corresponds to high values and blue corresponds to low values.

We find several consistent patterns in the ranking of importance and the directionality of feature values on the predicted 28-day strength. We find that higher quantities of cement are associated with higher predicted 28-day strengths. Cement ranks within the top three important features across all four datasets. In concrete design it is well known that the amount of cement, the hydration of which contributes to strength gain, is a key driver of 28-day strength. Additionally, in our summary plots, lower w/cm ratios are associated with higher predicted 28-day strengths. Across all datasets, the w/cm ratio ranks within the top four most important features. This is once again a well known relation in the domain.

In panels a, c, and d of Figure 2-4, the amount of slag is shown to increase the predicted 28-day strength. In general, SCMs are often used to increase the late term strength of concrete. This effect is observed in our model interpretation with the exception of fly ash in panels b and d, where we observe a mixed effect and negative impact respectively. This may be due to the limited feature set with dataset C2 in panel b, and the low number of data points in dataset L in panel d. Given knowledge that the inclusion of SCMs impacts strength differently based on age, future work should investigate the feature importances along the time dimension in detail.

2.4 Toward generalization: Model architectures for improved prediction

As demonstrated in Section 2.2, strength prediction performance differs depending on the size of the dataset and the number of features available. Thus, it seems possible to leverage information included in richer datasets in order to improve strength predictions on lower-quality datasets. To this end, we explore three architectures which enable higher generalization: 1) Pooled RF, 2) Pooled MLP, and 3) Multi-trunk MLP.

2.4.1 Pooled Architectures

The term 'pooled' refers to a pre-processing step in which all datasets are combined, or pooled, prior to training a single model. The model does not explicitly know which data point belongs to which dataset, however we keep track of this information externally. It is then possible to compute performance scores and errors over the entire single pooled dataset, as well as over the original separate datasets. We explored two versions of this pooled architecture: a pooled random forest and a pooled multi-layer perceptron. In both cases, the same model architectures from Section 2.2 are used, with the pooling applied prior.

2.4.2 Multi-trunk MLP

The multi-trunk MLP is a hierarchical structure with two levels, as shown in Figure 2-5, that allow the model to simultaneously learn dataset-specific transformations and share information across datasets. We first concatenate K datasets, each with features x, targets y, and a dataset indicator variable z in Eq. 2.1, to produce $\mathcal{D}$ of size N. The bottom level of the multi-trunk model has K MLP trunks, f_{θ_k} for $k = 1, \ldots, K$. Each model f_{θ_k} consists of five hidden layers of size 50 and ReLU activation function, ultimately learning parameters θ_k for a 50-dimensional representation for each subset of $\mathcal{D}$ where $z_n = k$. The upper level of the multi-trunk model, $f_{\theta_{shared}}$ is an MLP which takes as input the K 50-dimensional representations and produces a single final output y_n. The parameters θ are chosen to maximize the log likelihood of the data, as in Eq. 2.2 and 2.3.

$$\mathcal{D} = \{x_n, y_n, z_n\}_{n=1}^{N} \tag{2.1}$$

$$\hat{\theta} = \arg\max_{\theta} \sum_{n=1}^{N} \sum_{k=1}^{K} \mathbf{1}_{z_n=k} \log p(y_n \mid f_{\theta_k}(x_n), \theta_{shared}) \tag{2.2}$$

$$\theta = \{\theta_{shared}, \theta_1, \theta_2, \ldots, \theta_K\} \tag{2.3}$$

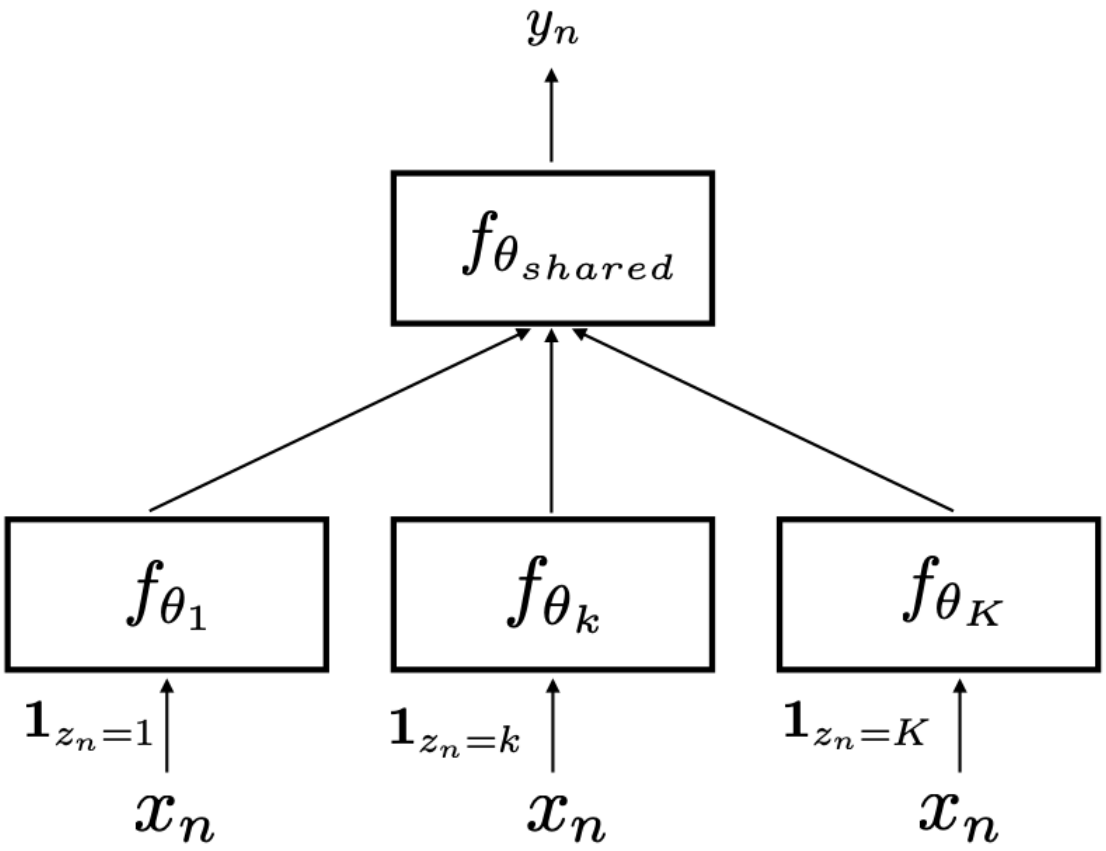

Figure 2-5: Multi-trunk architecture.

Table 2.3: RMSE for the four datasets (C1, C2, C3, L) and five architectures (RF, MLP, pooled RF, pooled MLP, and multi-trunk MLP. Lower values are better.

	Individual		**Pooled**		**Multi-trunk**
RSME (↓)	**RF**	**MLP**	**RF**	**MLP**	**MLP**
C1	823.93	839.13	1266.36	1197.70	871.69
C2	779.05	847.52	780.81	852.03	823.25
C3	843.94	857.30	1029.13	1147.01	843.00
L	854.23	939.20	995.20	1253.62	771.84
Overall	-	-	938.56	967.47	836.3

Table 2.4: R^2 for the four datasets (C1, C2, C3, L) and five architectures (RF, MLP, pooled RF, pooled MLP, and multi-trunk MLP. Higher values are better.

	Individual		**Pooled**		**Multi-trunk**
R^2 (↑)	**RF**	**MLP**	**RF**	**MLP**	**MLP**
C1	0.84	0.83	0.63	0.66	0.82
C2	0.72	0.67	0.73	0.67	0.69
C3	0.85	0.85	0.78	0.72	0.85
L	0.83	0.80	0.75	0.64	0.87
Overall	-	-	0.82	0.81	0.86

2.4.3 Results

We plot the RMSE for the four datasets (C1, C2, C3, L) and five architectures (individual RF and individual MLP from Section 2.2, and the proposed pooled RF, pooled MLP, and multi-trunk MLP) in Figure 2-6. Scatter plots for predicted 28-day compressive strength vs. actual strength for each generalization model can be found in Table A-1 of the Appendix. The multi-trunk MLP model performs consistently well across all four datasets, often comparable with the individually trained RFs and MLPs.

2.5 Discussion

Across the four datasets in this study, we identify common themes in the predictive ability of certain features of a concrete toward compressive strength, and these findings align well with the existing science. Additionally, we find an opportunity to use a hierarchical model, which learns both dataset-specific information and population-

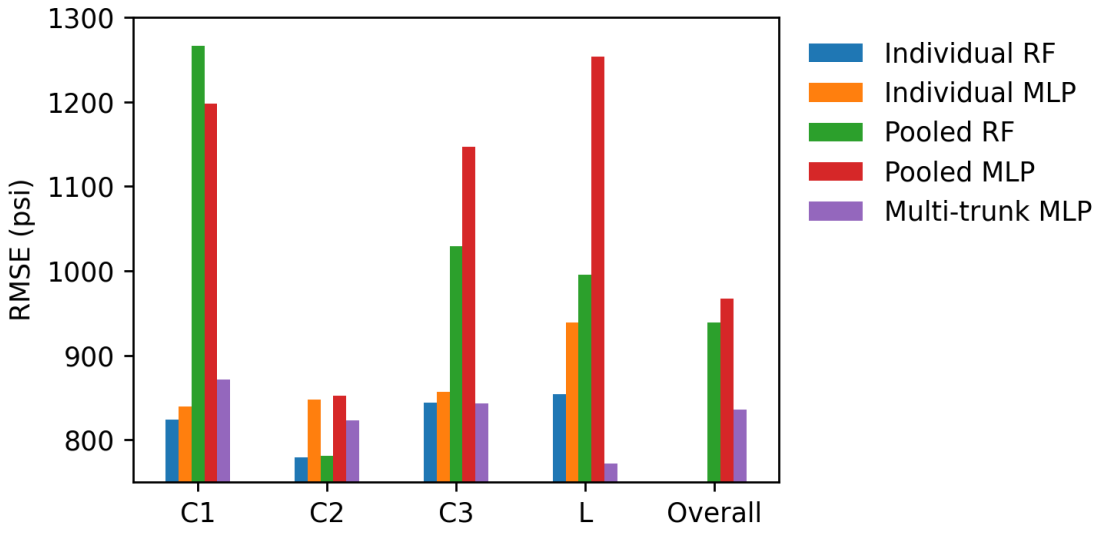

Figure 2-6: Comparison of RMSE across four datasets (C1-L) and overall for the various model architectures (RF, MLP, Pooled, and Multi-trunk).

based parameters, in order to improve prediction accuracy on smaller datasets in particular. Further, while complete information on the attributes of a concrete mixture may be necessary for downstream tasks, it is not required for achieving compressive strength prediction with low error.

Future work should focus on understanding the role of more detailed and specific features (e.g., type of cement, size of coarse aggregate) on the prediction of concrete compressive strength. This will aid in understanding the nuance of the interaction between constituents, and ultimately enable improved design. Additionally, concrete is cast in an environment with a particular temperature, humidity, and other atmospheric conditions. In a laboratory setting, specimens are prepared following ASTM Standards with the ultimate exposure conditions in mind. In practice at an on-site project, conditions may be drastically different than the lab setting, and they may change between measurements. As such, it would be beneficial to incorporate estimates for the measurement uncertainty or account for the environmental conditions at the time of curing in order to improve prediction accuracy using machine learning.

Machine learning can serve as a screening tool for concrete designers to evaluate candidate mixtures before investing in the labor and equipment associated with the

trial-and-error experimental process.

Chapter 3

Optimal concrete design framework

A key performance metric for evaluating a concrete mixture is its compressive strength, a property that increases with time as the concrete sets. Experimentalists typically perform several compressive strength tests on a number of cast cylinders in order to gather strength measurements at 3, 7, 14, 28, and 56 days. This collected data amounts to a time series, and the problem of predicting compressive strength is well-suited to time-series analysis methods that involve curve fitting, where the underlying assumption is that a target y is ordered by x, the time variable. Thus, we first use a novel Gaussian process model which was presented in Severson *et al.* [20] to estimate the complete strength trajectory of a concrete mixture over the entire curing time period. This model is described in detail in Section 3.2. Second, we integrate this strength model into an optimization procedure via a nonlinear constraint, where the predicted strength must exceed a specified design target. This ensures that optimal mixtures meet structural requirements. We optimize the concrete on two dimensions: cost and climate impact (specifically, greenhouse gas (GHG) emissions). We also analyze the trade-offs between these two objectives using a weighted sum method to obtain the Pareto front. The optimization procedure is described in detail in Section 3.5.

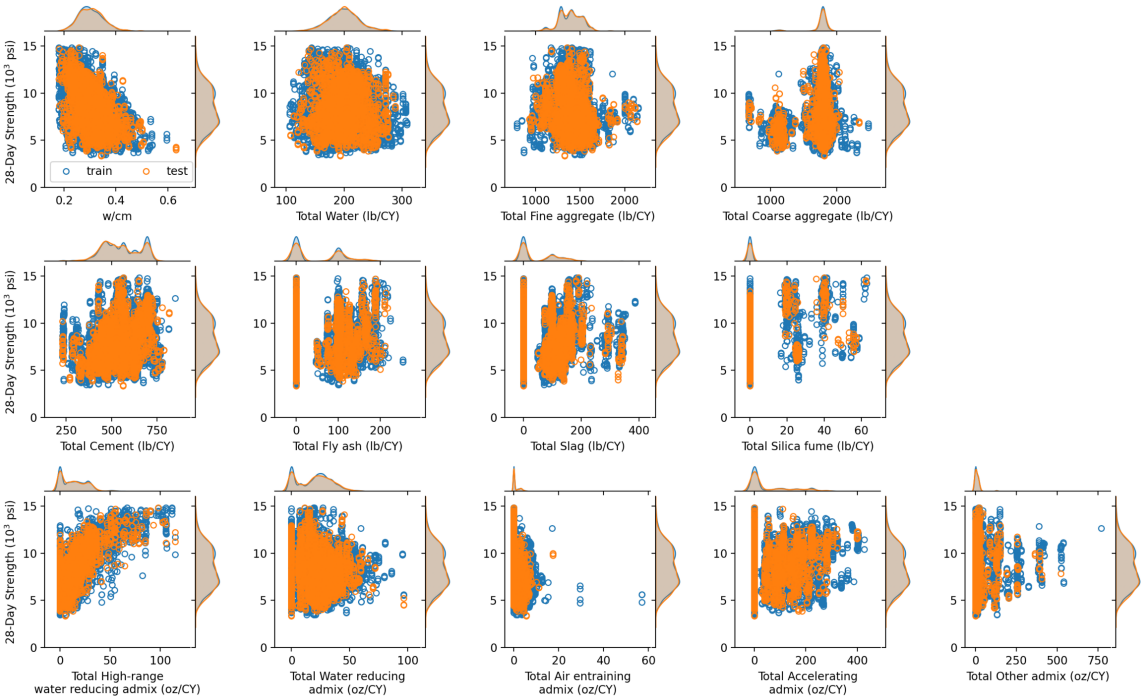

Figure 3-1: Visualization of 28-day compressive strength as a function of each of the 13 measurements.

3.1 Industrial concrete data

The concrete mixture dataset used in this study is provided by an industrial concrete producer. Information is available for 10,796 unique mixes, each with a varying number of strength measurements at different ages (range 1-26, average 4.1 measurements per mix) for a total of 44,490 measurements. As in Section 2.1, each mixture is represented by a 13-dimensional vector, with the first element being the water-to-cementitious material ratio, and the remaining 12 elements being the content of constituents for one cubic yard of concrete. Figure 3-1 displays 28-day compressive strength as a function of each of the 13 measurements.

We use 80% of the data for training the model and the remaining 20% for evaluation. To avoid leakage between the train and test sets, we partition the data based on mixture ID, such that a given mixture with multiple strength observations appears only in the training set or testing set (or vice versa), but not both. All values are z-scored prior to training, by subtracting the mean and dividing by the standard

deviation.

3.2 Gaussian process regression model for prediction of strength trajectories

By definition, a Gaussian process (GP) is a collection of random variables, any finite number of which are jointly Gaussian distributed Rasmussen & Williams [21]. GPs are suited for time series predictions, as motivated by Roberts *et al.* [22], in particular because the use of Bayesian inference enables direct uncertainty quantification.

GPs are distributions over functions, and they are specified by a mean function and covariance function. The mean function is our expectation prior to any observations, and is thus informed by domain knowledge of the expected trend. The covariance function specifies the relation between any pair of outputs (i.e., the function magnitude and time scale). Thus, a GP is governed by the hyperparameters, θ, of these functions. There are two typical approaches for estimating these hyperparameters. In the first, the parameters are optimized individually for each sample by maximizing the log likelihood of the data. In the second, one set of optimal parameters is chosen for the entire population.

In the context of a concrete mixture and its compressive strength measurements, we have a set of static covariates (i.e., the mixture ingredients). Thus, the novel GP model described in Severson *et al.* [20] uses a multi-layer perceptron (MLP) to estimate the hyperparameters of the mean and covariance functions based on the static mixture ingredients ($\boldsymbol{\theta} = MLP(\mathbf{z})$). This is advantageous because our observations per sample, where a sample is the strength trajectory of a particular mixture, are limited in number and thus restrict our ability to infer accurate hyperparameters via the individual approach.

Using domain knowledge on the progression of strength, we specify a linear mean function in logarithmic time (i.e., $\mathbf{t} = \log_{10}(time)$) as in Eq. 3.1. Although there exist readily available empirical equations describing strength in the literature, in practice

we found that these resulted in poorer performance than the linear mean function. We use squared exponential covariance function as in Eq. 3.2, which is a standard choice.

$$m(\mathbf{t}_i; \boldsymbol{\theta}_i) = \boldsymbol{\theta}_{i1}\mathbf{t}_i + \boldsymbol{\theta}_{i2} \tag{3.1}$$

$$k(\mathbf{t}_i, \mathbf{t}_i'; \boldsymbol{\theta}_i) = \boldsymbol{\theta}_{i3} \exp\left(-\frac{\|\mathbf{t}_i - \mathbf{t}_i'\|^2}{2(\boldsymbol{\theta}_{i4})^2}\right) \tag{3.2}$$

3.3 Estimating the climate impact of concrete

We estimate the climate impact of a given concrete mixture on a cradle-to-gate basis, considering the GHG emissions that are associated with the concrete's raw materials extraction, transportation, and production phases (Eq. 3.3). The climate impact factors for the raw materials extraction ($g_{m,i}$) and transportation ($g_{t,i}$) of each concrete constituent i are shown in Table 3.1. These factors are measured in units of CO_2 equivalents (CO_2-eq) per mass of constituent, where the global warming potential (GWP) index is used in order to express the total warming effect of all GHGs in terms of CO_2. For example, 1 kg of methane has about 30 times the warming potential of 1 kg of CO_2 over a fixed time period (e.g., 100 years), thus the GWP of methane is about 30. The climate impact associated with the production of 1 CY of concrete is 5.49 kg CO_2-eq. Since each z_i represents the mass of the constituent i in 1 CY of concrete, the function in Eq. 3.3 results in the climate impact of a mixture in units of kg CO_2-eq/CY of concrete.

We consider two main data sources for the environmental impact analysis: the ecoinvent inventory database for raw materials [23] and the life cycle assessment (LCA) report commissioned by the National Ready Mixed Concrete Association (NRMCA) [24] for transportation and production. Environmental product declarations for chemical admixtures provide materials climate impacts estimates in [25–27].

Table 3.1: Climate impact factors associated with raw materials production ($g_{m,i}$) and transportation ($g_{t,i}$) of each constituent i on a mass basis. The climate impact associated with the production of 1 CY of concrete (g_p) is 5.49 kg CO_2-eq.

i	**Mixture constituent**	**Raw materials ($g_{m,i}$)**	**Transportation ($g_{t,i}$)**
	(Non-admixtures)	kg CO_2-eq/lb	kg CO_2-eq/lb transported
1	Coarse aggregate	0.00228	0.0033
2	Fine aggregate	0.00228	0.00327
3	Water	0.00096	0
4	Cement	0.473	0.0133
5	Fly ash	0	0.00856
6	Slag	0.0667	0.00881
7	Silica fume	0	0
	(Admixtures)	kg CO_2-eq/oz	kg CO_2-eq/oz transported
8	High-range water reducing	0.05341	0
9	Water reducing	0.05341	0
10	Accelerating	0.038	0
11	Air-entraining	0.01497	0
12	Viscocity modifier	0.063	0

$$f_g(\mathbf{z}) = \mathbf{g_m}^\intercal \mathbf{z} + \mathbf{g_t}^\intercal \mathbf{z} + g_p \tag{3.3}$$

3.4 Estimating the cost of concrete

We estimate the cost of a concrete mixture per CY (Eq. 3.4) by considering the materials costs and transportation as in Table 3.2. We do not estimate production costs. We note that our cost estimates do not account for forming, reinforcements, or labor involved with the production of one CY of concrete and are thus underestimates of the true cost. Except for SCMs, material cost values are based upon Young *et al.* [4]. Transportation cost estimates use national average transport distances and quantities reported in [24], and assume costs per ton-mile for truck, rail, ocean, and barge to be \$0.15, \$0.05, \$0.04,and \$0.03 respectively.

It is important to consider how optimal mixtures may change as SCMs become more expensive, for example, as coal plants are retired and fly ash becomes more scarce. Thus, we have developed three cost scenarios in addition to our baseline data,

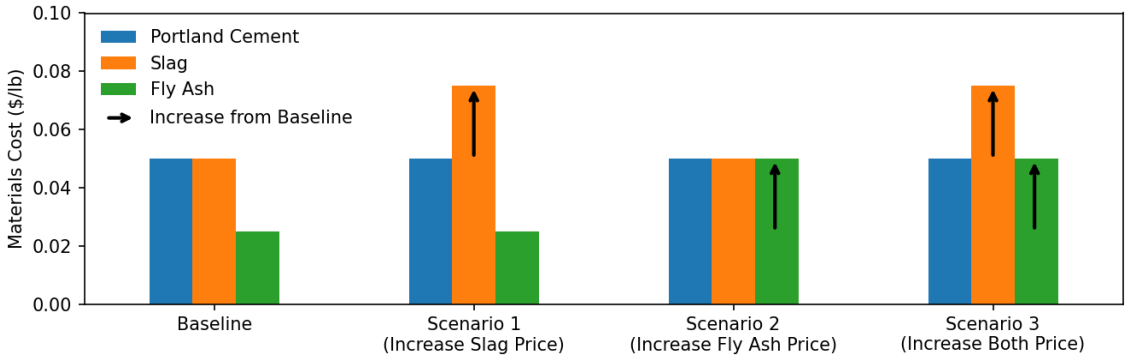

Figure 3-2: Cost scenarios for per-mass prices of cement, slag, and fly ash.

as in Figure 3-2. In the Baseline Scenario, we assume that the cost of slag is equal to the cost of cement ($0.05/lb), and that the cost of fly ash is 50% of that ($0.025/lb). In Scenario 1, we increase slag costs; we set slag costs to be 150% of its Baseline value, totaling $0.075/lb. In Scenario 2, we increase fly ash costs; we set fly ash costs to be 200% of its Baseline value, now equivalent with both cement and slag, or $0.05/lb. In Scenario 3, we simultaneously increase the cost of slag and the cost of fly ash as in the previous two scenarios.

$$f_c(\mathbf{z}) = \mathbf{c_m}^\intercal \mathbf{z} + \mathbf{c_t}^\intercal \mathbf{z} \tag{3.4}$$

3.5 Optimization for concrete mixture design

In Figure 3-3, we show the histograms for estimated costs and climate impacts for all mixtures in the industrial dataset. We note that the distribution of costs is wide, ranging from $50 to $100, and that some of the spread can be accounted for by the achieved strength (Fig. 3-4, left panel). However, there is still a clear opportunity for improved design for cost. Similarly, in Figure 3-4 (right), we observe that there exists a high mode at 390 kg/CO_2-eq for mixtures which achieve compressive strengths of 10,000 to 11,000 psi; however, there exist a number of mixtures with lower climate impacts that still achieve comparable strengths. Therefore, in this section we design

Table 3.2: Cost factors for materials ($c_{m,i}$) and transportation ($c_{t,i}$) for each constituent i on a mass basis.

i	**Mixture constituent**	**Raw Materials** ($c_{m,i}$)	**Transportation** ($c_{t,i}$)
	(Non-admixtures)	\$/lb	\$/lb
1	Coarse aggregate	0.0045	0.002388
2	Fine aggregate	0.0027	0.002302
3	Water	0	0
4	Cement	0.05	0.012379
5	Fly ash	0.025	0.00598
6	Slag	0.05	0.010896
7	Silica fume	0.18	0
	(Admixtures)	\$/oz	\$/oz
8	High-range water reducing	0.083	0
9	Water reducing	0.083	0
10	Accelerating	0.083	0
11	Air-entraining	0.083	0
12	Other	0.083	0

optimization procedures to 1) minimize cost, 2) minimize climate impact, and 3) understand the trade-offs associated with each objective.

3.5.1 Single objective minimization

In order to ensure that optimal mixtures obey the physical requirements of a concrete system, we first define bounds and constraints necessary in Sections 3.5.1.1 and 3.5.1.2 respectively. We then pose the full optimization procedures for the single objective minimization of climate impact and cost in Section 3.5.1.3.

3.5.1.1 Bounds

Bounds for each of the constituent quantities are displayed in Table 3.3 based on the ranges of values of constituents in the industrial dataset.

3.5.1.2 Constraints

Strength With the GP model established in Section 3.2, we can now obtain predictions of strength as a function the mixture vector $\mathbf{z}$ and age t as expressed Eq. 3.6. Note that Eq. 3.6 is equivalent to the mean function of the GP (Eq. 3.1). Thus,

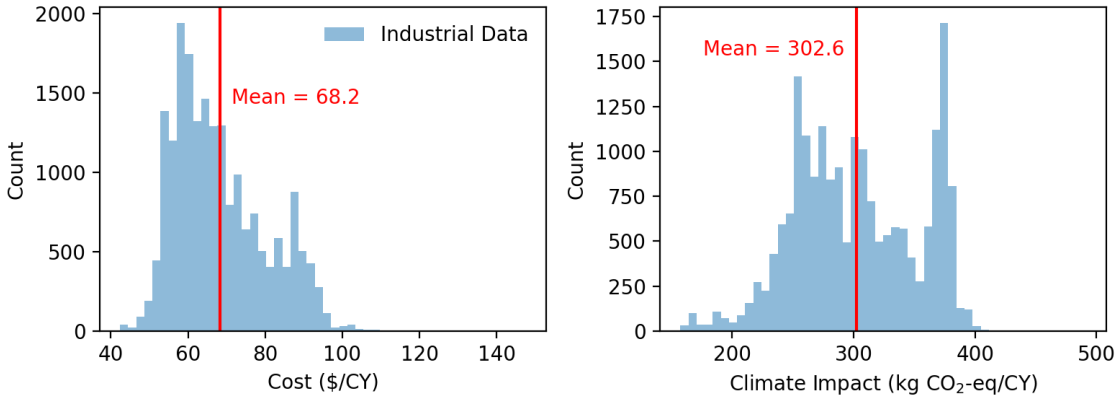

Figure 3-3: Estimated cost (left) and climate impact (right) values for mixtures in industrial dataset. The red line represents the mean.

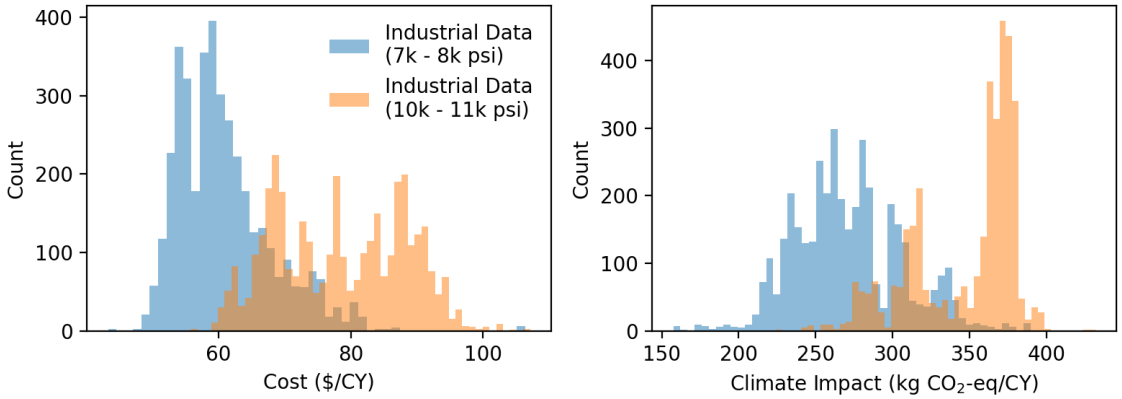

Figure 3-4: Estimated cost (left) and climate impact (right) values for mixtures in industrial dataset, for two distinct strength classes: 7,000 to 8,000 psi and 10,000 to 11,000 psi.

Table 3.3: Densities and bounds for each constituent.

i	**Mixture constituent**	**Density**(ρ_i)	**Lower Bound**	**Upper Bound**
	(Non-admixtures)	lb/CY	lb	lb
1	Coarse aggregate	4214	700	4000
2	Fine aggregate	4467	700	4000
3	Water	1686	30	1000
4	Cement	5310	100	3000
5	Fly ash	4214	0	500
6	Slag	3371	0	500
7	Silica fume	3708	0	100
	(Admixtures)	oz/CY	oz	oz
8	High-range water reducing	29666	0	200
9	Water reducing	32363	0	200
10	Accelerating	37756	0	600
11	Air-entraining	26969	0	100
12	Other	36408	0	1000

the constraint as written in Eq. 3.7 can be interpreted as the following: the mean predicted strength of mixture $\mathbf{z}$ at time t should meet or exceed a minimum specified target, s_{min}. In this work, we evaluate six target strengths, the first starting at 6,000 psi, and incremented by 1,000 through 11,000 psi. Additionally, we use t=28 in our work and recommend future investigations to explore multiple constraints of strength over time.

$$\boldsymbol{\theta} = MLP(\mathbf{z}) \tag{3.5}$$

$$f_s(\mathbf{z}, t) = \theta_1 t + \theta_2 \tag{3.6}$$

$$s_{min} \leq f_s(\mathbf{z}, t) \tag{3.7}$$

Total Volume Constraint Optimal mixture proportions should produce one cubic yard of concrete. Using standard densities, ρ, for constituents as given in Table 3.3, we estimate the total volume of a mixture in Eq. 3.8. We assume concrete mixtures which contain a nonzero amount of air-entraining admixture contain 6% air by volume; this estimate is based on a widespread trend present in the industrial data. Subsequently, when Eq. 3.8 is applied to the industrial dataset we obtain

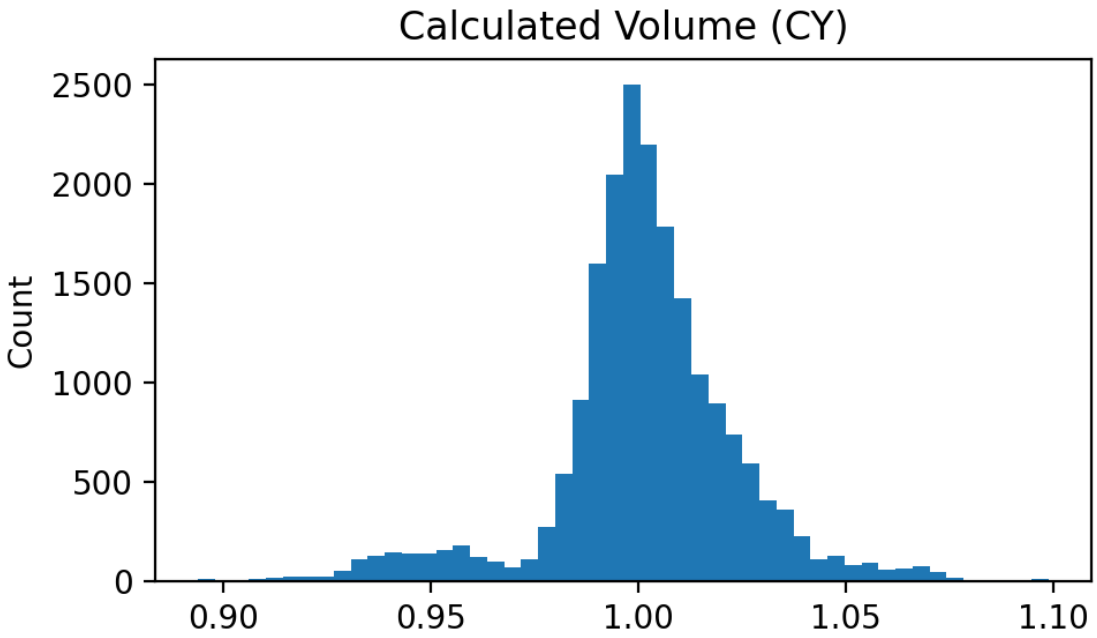

Figure 3-5: Histogram of calculated volumes of mixtures in industrial dataset using constituent densities in Table 3.3 and Eq. 3.8. Units are in cubic yards (CY).

calculated volumes V as shown Figure 3-5. We observe that there is a clear mode at 1 CY, and a majority (51%) of mixtures fall in the range from 0.99 CY to 1.01 CY, suggesting that our volume estimation method is reasonable. Therefore, we formulate the volume constraint as an inequality in the optimization as shown in Eq. 3.9. This form of the constraint is more attainable, as it can be thought of as a relaxed version of an equality constraint.

Aggregate Volume Ratio Constraint Additionally, the range of acceptable volume fraction of aggregates is known in concrete design, and is 0.65 to 0.75. We calculate the volume fraction of aggregates in Eq. 3.10 and its associated constraint is shown in Eq. 3.11.

$$V = \frac{z_1}{\rho_1} + \frac{z_2}{\rho_2} + \frac{z_3}{\rho_3} + \frac{z_4}{\rho_4} + \frac{z_5}{\rho_5} + \frac{z_6}{\rho_6} + \frac{z_7}{\rho_7} + \frac{z_8}{\rho_8} + \frac{z_9}{\rho_9} + \frac{z_{10}}{\rho_{10}} + \frac{z_{11}}{\rho_{11}} + \frac{z_{12}}{\rho_{12}} \tag{3.8}$$

$$0.99 \leq V \leq 1.01 \tag{3.9}$$

Table 3.4: Ratio constraints used in the optimization procedure.

Parameter	**Expression**	**Lower**	**Upper**
Total supplementary cementitious material ratio (R_{SCM})	$\frac{z_5 + z_6 + z_7}{z_4 + z_5 + z_6 + z_7}$	0	0.6
Fly ash ratio (R_{FA})	$\frac{z_5}{z_4 + z_5 + z_6 + z_7}$	0	0.4
Slag ratio (R_{SL})	$\frac{z_6}{z_4 + z_5 + z_6 + z_7}$	0	0.6
Silica fume ratio (R_{SF})	$\frac{z_7}{z_4 + z_5 + z_6 + z_7}$	0	0.1
w/cm (R_{WCM})	$\frac{z_3}{z_4 + z_5 + z_6 + z_7}$	0.2	0.6

$$V_{agg} = \frac{z_1/\rho_1 + z_2/\rho_2}{V} \tag{3.10}$$

$$0.65 \leq V_{agg} \leq 0.75 \tag{3.11}$$

Ratio Constraints There are a number of constraints informed by materials science knowledge that must be imposed in order to ensure the optimal concrete mixtures are physically sound. The water to total cementitious material ratio should be between 0.2 and 0.6 to ensure workability. The total supplementary cementitious material should represent not more than 60% of the total cementitious material of a concrete mixture. Similarly, individual ratios of fly ash, slag, and silica fume to total cementitious material ratio should be limited because addition of SCMs reduce the amount of cement, which is more reactive than SCMs, that is available for hydration. There are subtle differences in terms of the limit of their ratios described as follows: (i) Fly ash and silica fume are, by themselves, unreactive. Both react with $Ca(OH)_2$, which is a product of OPC hydration, in order to provide the strength-giving phase (C-S-H gel). When there is insufficient cement to hydrate to provide $Ca(OH)_2$, the additional fly ash and silica fume act as inert fillers. (ii) Slag is slightly reactive without the presence of $Ca(OH)_2$; hence, relatively higher proportions are acceptable. (iii) Silica fume has a much higher surface area than the other SCMs, and requires larger amounts of water in order to maintain workability. Table 3.4 lists expressions and

bounds for these constraints.

3.5.1.3 Objective functions

There are two objectives of interest when designing a concrete mixture: 1) minimizing the climate impact and 2) minimizing cost.

Algorithm The optimization procedures in this work use the differential evolution (DE) algorithm, a heuristic optimization algorithm developed by Storn & Price [28]. Heuristic methods are typically stochastic direct search methods which involve generation of variations of a candidate solution vector. These methods are able to search large spaces with reasonable computation cost. DE is a highly versatile evolutionary algorithm for solving multidimensional, continuous, minimization problems. Unlike genetic algorithms, DE is specifically designed for real-valued functions. We use the DE implementation available in the SciPy Python package [29]. The algorithm is run 50 times for each weight and target strength pair using the Sobol population initialization method.

Minimize Climate Impact

$$
\begin{aligned}
\min \quad & f_g(\mathbf{z}) \\
\text{subject to} \quad & s_{min} \leq f_s(\mathbf{z}, 28) \\
& 0.99 \leq V \leq 1.01 \\
& 0 \leq R_{SCM} \leq 0.6 \\
& 0 \leq R_{FA} \leq 0.4 \\
& 0 \leq R_{SL} \leq 0.6 \\
& 0 \leq R_{SF} \leq 0.1 \\
& 0.2 \leq R_{WCM} \leq 0.6 \\
& 0.65 \leq V_{agg} \leq 0.75
\end{aligned}
\tag{3.12}
$$

Minimize Cost

$$\begin{aligned} \min \quad & f_c(\mathbf{z}) \\ \text{subject to} \quad & s_{min} \leq f_s(\mathbf{z}, 28) \\ & 0.99 \leq V \leq 1.01 \\ & 0 \leq R_{SCM} \leq 0.6 \\ & 0 \leq R_{FA} \leq 0.4 \\ & 0 \leq R_{SL} \leq 0.6 \\ & 0 \leq R_{SF} \leq 0.1 \\ & 0.2 \leq R_{WCM} \leq 0.6 \\ & 0.65 \leq V_{agg} \leq 0.75 \end{aligned} \tag{3.13}$$

3.5.2 Bi-objective minimization

3.5.2.1 Objective normalization

Following the methods in Marler & Arora [30], we normalize each of the single objective functions in Eq. 3.14 and 3.15, where $\mathbf{z}^*{}_c$ is the formula which which minimizes f_c and $\mathbf{z}^*{}_g$ is the formula which minimizes f_g.

$$f_c^{norm}(\mathbf{z}) = \frac{f_c(\mathbf{z})}{f_c(\mathbf{z}^*{}_c)} \tag{3.14}$$

$$f_g^{norm}(\mathbf{z}) = \frac{f_g(\mathbf{z})}{f_g(\mathbf{z}^*{}_g)} \tag{3.15}$$

3.5.2.2 Weighted sum method

There are several classical methods for multi-objective problem, including ϵ-constraint, goal attainment, and weighted sum. The weighted sum method, among the simplest and most popular to implement, involves converting the multi-objective problem into a weighted sum of all single objectives. We adopt this method in our work, and define

a new function F in Eq. 3.16 representing the bi-objective problem. In this form, we select various values of w in the range $[0, 1]$ in order to vary the prioritization of one objective function over another. As a result, we will be able to understand trade-offs between the two objectives.

$$F = w \cdot f_g^{norm}(\mathbf{z}) + (1 - w) \cdot f_c^{norm}(\mathbf{z}) \tag{3.16}$$

3.6 Results & Discussion

In Fig. 3-6 we plot the optimal costs (left) and climate impacts (right) for the specified target strengths along the x-axis. There is convergence to a common minimum, as shown by the violin shape displaying the distribution of points achieved over all algorithm runs. It is noteworthy that cost is more sensitive to increasing required strength than the climate impact. Specifically, there is a percent change in the median costs over the increasing target strengths of 28%, while only a 3.7% change in median climate impacts.

Unsurprisingly, in Figure 3-7 we observe that the optimal cost increases with the scenarios in which SCMs are made more expensive. A large jump from the obtained Baseline optimal costs occurs in Scenario 2, when fly ash costs are set at 200% of the Baseline value, becoming equal with both cement and slag. This suggests that fly ash is relied upon more heavily than the other SCMs.

It is interesting to analyze how optimal mixtures differ from existing mixtures in the industrial dataset based upon the constituents. In Figure 3-8, we plot the overall constituent distributions using the kernel density estimation for minimum climate impact mixtures, minimum cost mixtures, and the industrial dataset. Note that there exist multiple peaks for some constituents; these peaks typically align with the various required strength targets. For example, observe the bi-modal nature of the amount of high-range water reducing (HRWR) admixture for the minimum climate mixtures (green curve). We can break down this curve by the various levels of target

compressive strengths as in Figure 3-9. In fact, we observe that concrete mixtures which must meet a strength requirement of $>$ 8,000 psi contain increasing amounts of the HRWR admixture (all contributing to the wide spread of the right peak in Figure 3-8), while concrete mixtures which only satisfy lower strength requirements contain negligible amounts of the admixture (contributing to the left peak). The amounts of all other admixtures are negligible. In the minimization of climate impact, nearly all mixtures contain approximately 200 lb of cement. Additionally, the mixtures contain all three SCMs, though the predominant SCM is fly ash.

Within Figure 3-9, there is generally little variation in the constituent quantities based on target strength, with the exception of coarse aggregate, fine aggregate, water, and HRWR admixture. This suggests that there is a relatively small set of mixtures which can achieve minimum climate impact. However, in contrast, there are several constituents within the minimal cost mixtures that vary greatly based on target strength (Figure 3-10). For example, the coarse aggregate content is seen to oscillate with target strength, from roughly 1550 lb up to 1800 lb and then back down. There also appears to be a shift in the type of SCM used depending on the target strength; low-level target strengths use both slag and fly ash, mid-level strengths almost exclusively use fly ash, and the highest target strength uses fly ash and silica fume.

In Figure 3-11 we plot the industrial data with the optimal cost (left) and climate impacts (right). We note that the plotted industrial data is a filtered version of the original dataset; it is the subset which obeys the optimization constraints and bounds within +/- 10%. As a result, there are a few points which lie below the optimal line. These extreme points in particular violate the volume fraction of aggregate ratio. In general, the lower end of the industrial data points are reasonably close to the optimal costs, with many being within \$10/CY. The vast majority of industrial mixtures have costs upwards of \$75, suggesting large room for improvement. However, it is important to note that the industrial data mixtures were likely subject to additional constraints, such as workability and durability requirements.

As discussed previously, climate impact is less sensitive to increasing target strength

than cost. Now plotted in context with the industrial data, the nearly flat slope or 'floor' of the optimal climate impact line becomes more apparent. The gap between optimal climate impact line and the industrial data is larger than the gap for cost. Additionally, there is a cluster of low climate impact industrial data points, nearly parallel to the optimal line, which incorporate greater levels of SCMs. This provides us with confidence that it is indeed possible to achieve the floor of climate impact for a variety of strength requirements.

Lastly, the results of the bi-objective optimization using the weighted sum method are shown in Figure 3-12. In this plot, each color in the series corresponds to a particular target strength. The points in each series are shaded according to the weight w, with $w = 1$ (dark shaded points) effectively minimizing only cost, and $w = 0$ (light shaded points) effectively minimizing only climate impact. Points with values of w within the interval create the Pareto front of the trade-offs between the two objectives.

We observe that for the low-level target strengths (e.g., 6000, 7000, and 8000 psi), the Pareto fronts are nearly vertical, with a slight negative slope. This shape indicates that there is a wide variety in cost of mixtures, however each achieves nearly the same climate impact, with cheaper mixtures being just slightly more carbon intense. For the mid-level to high-level target strengths, we observe that there exists a trade-off between the cheapest and most climate-friendly mixtures. However, for the 9,000 and 10,000 psi strength targets, a less expensive mixture is always available at a nearly constant climate impact. In these cases, the cheapest mixture is only marginally less expensive than the alternatives, but it is significantly more carbon-intense.

3.6.1 Potential GHG abatement

Based on our industrial dataset, the average climate impact of a concrete mixture is 302 kg CO_2-eq. The U.S. alone uses nearly 260 million cubic yards of concrete each year. These two values combine to result in a total annual estimated impact of the U.S. concrete industry of 78.5 million metric tons of CO_2-eq. If all mixtures instead were optimized (achieved an embodied impact of 130 kg CO_2-eq.), this would amount

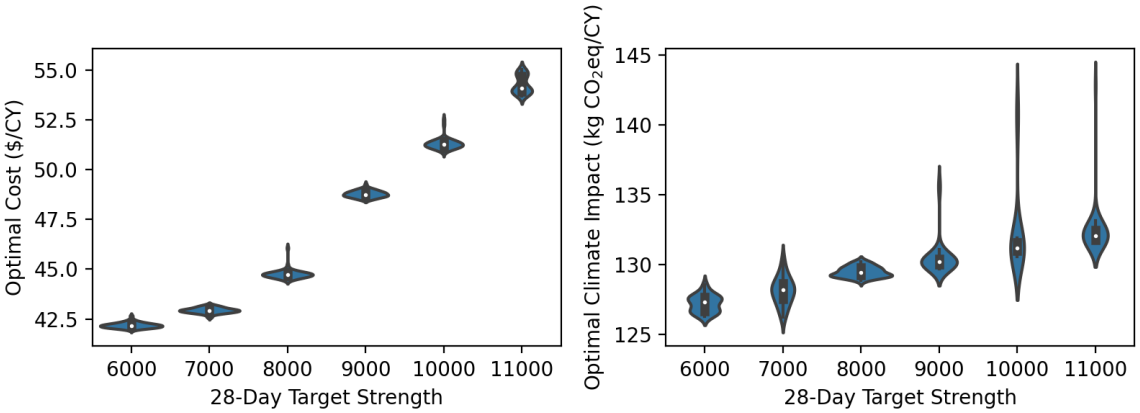

Figure 3-6: Optimal cost and climate impact achieved for specified target strengths.

to an annual savings of 44.7 million metric tons of CO_2-eq.

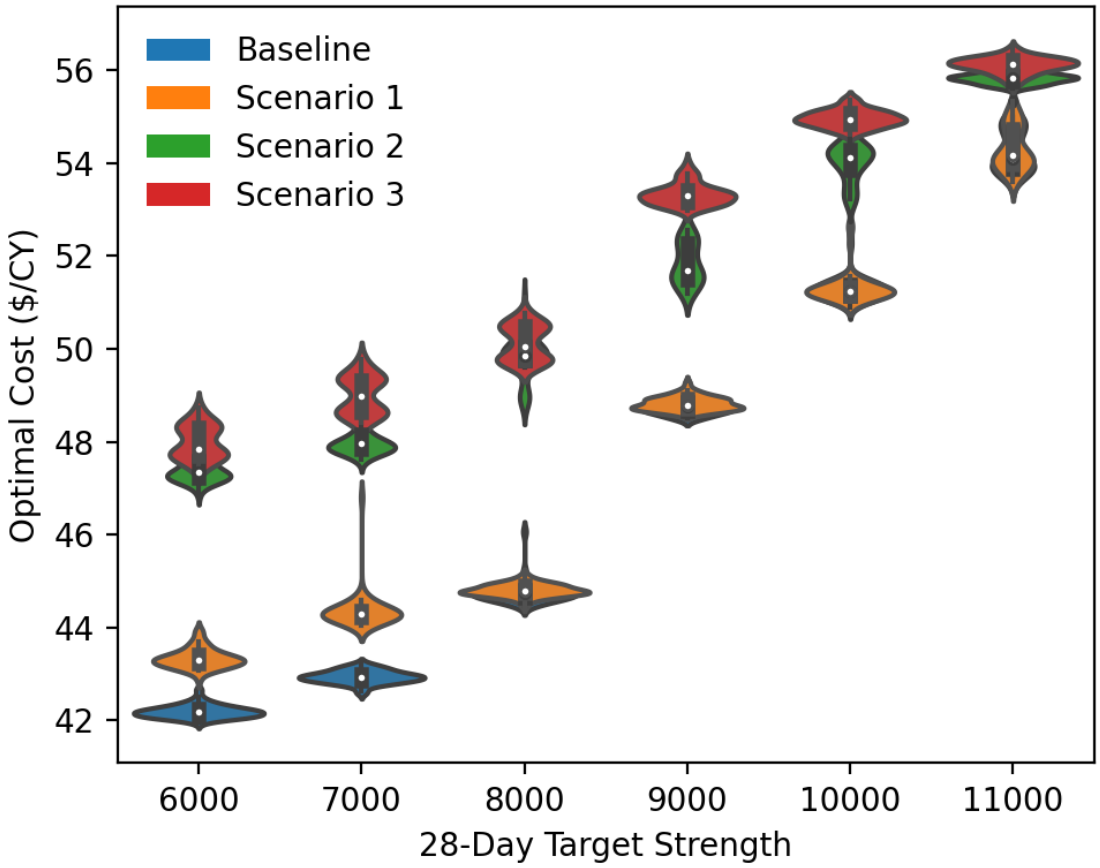

Figure 3-7: Optimal cost achieved for specified target strengths, colored by cost pricing scenario.

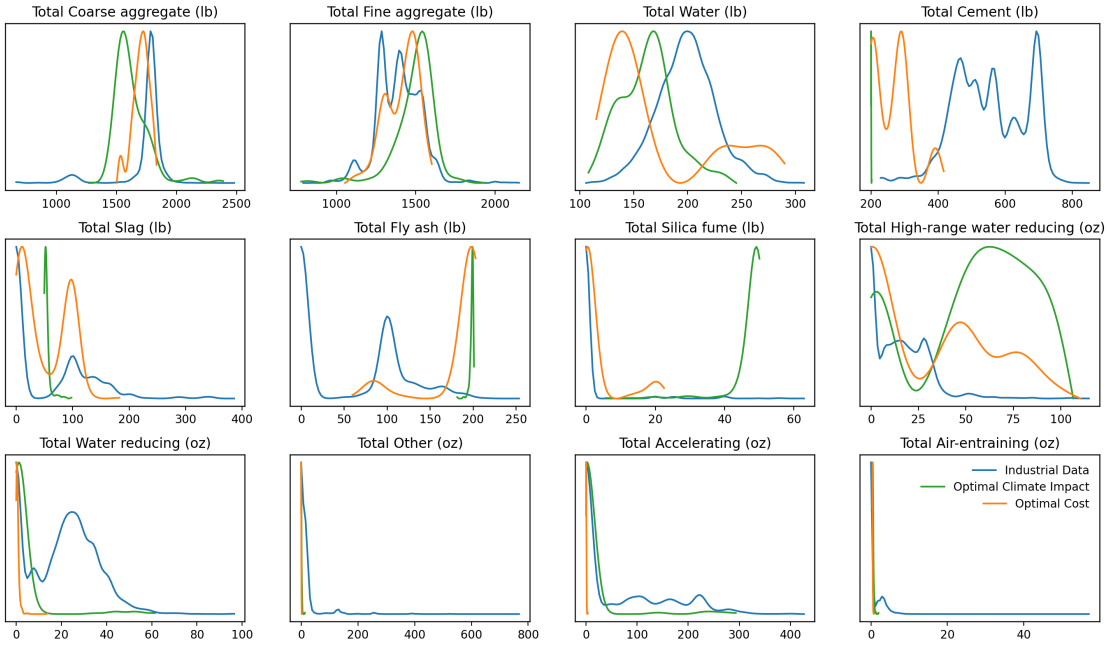

Figure 3-8: Overall constituent distribution comparison between the optimal climate impact mixtures, optimal cost mixtures, and industrial data. Lines represent kernel density estimations along a shared x-axis.

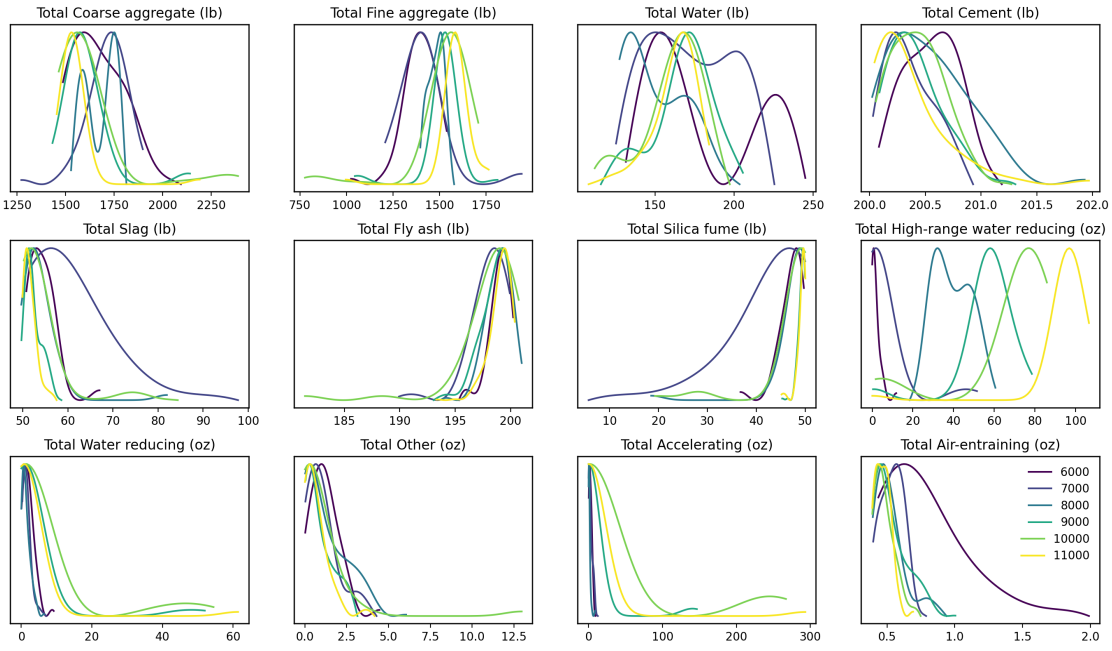

Figure 3-9: Minimum climate impact mixture constituent distribution comparison between various target strengths. Lines represent kernel density estimations along a shared x-axis.

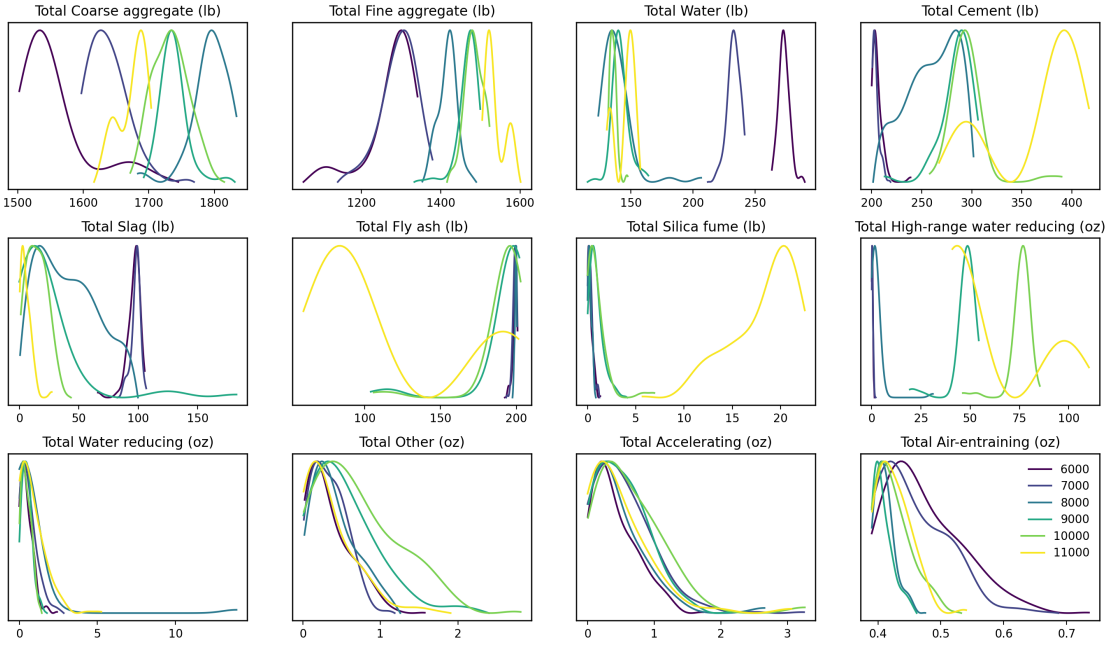

Figure 3-10: Minimum cost mixture constituent distribution comparison between various target strengths. Lines represent kernel density estimations along a shared x-axis.

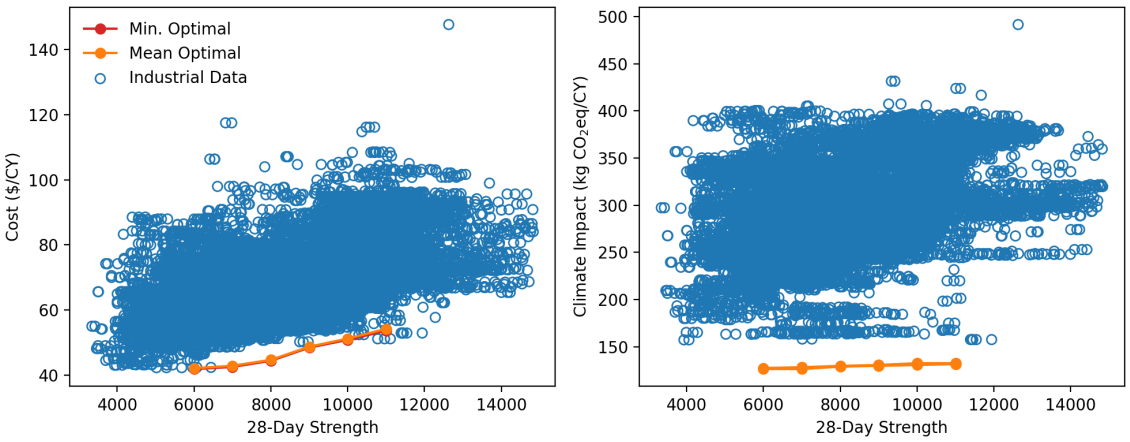

Figure 3-11: Comparison of optimal cost and climate impact and industrial data.

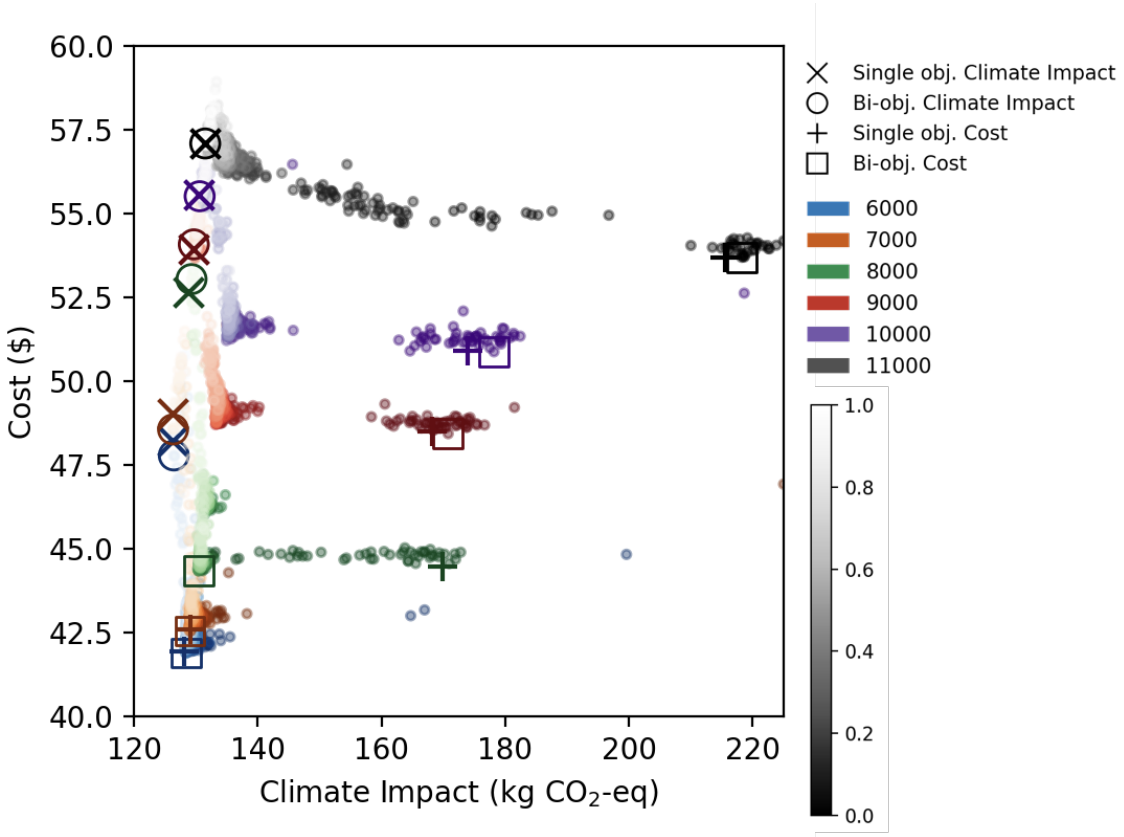

Figure 3-12: Pareto fronts for the bi-objective minimization of climate impact and cost.

Chapter 4

Opportunities for policy

Concrete is an indispensible staple of our world, and it will remain as such for the foreseeable future. As a result, it will be important to achieve low-carbon production of cement. The combustion of fossil fuels generates the heat necessary for the calcination of limestone and clay, and accounts for approximately 35% of the cement production's footprint. It is possible to use alternative fuels, such as biomass or waste, however, the replacement is limited due to the lower calorific value of these materials [31]. Hydrogen is another fuel source which may provide the heat, however, investigations are still at an early stage. Thus, incentives for further research will be necessary to drive work in this key area.

Further, nearly 75% of energy used in the production of cement is thermal; the remaining 25% is electricity [32]. Cement clinker production experiences high levels of heat loss via the kiln exhaust and cooling procedures, and approximately 40% of process heat and 26% of input heat is lost overall [32]. This makes cement plants suitable candidates for waste heat recovery (WHR) systems, which often involve retrofitting a steam turbine to capture the exhaust. Such systems can improve the energy efficiency of a cement plant, enabling it to generate from 4 to 6 MW of electric power. While most large cement producers currently utilize WHR systems in their facilities, with the majority existing in China, there are still a number of locations where it is not yet economical to retrofit the plant. Industrial electricity tariffs and sustainability initiatives can drive increased consideration of such systems [33].

Alternatively, the electrification of cement production is another potential pathway which is being explored. The replacement of fossil fuel-based energy when possible, such as in griding processes or those that have low heat requirements, with electricity can result in some emissions savings as long as the electricity grid has a reasonable proportion of clean resources in the mix. In addition, the electrochemical synthesis of cement has been demonstrated [2] as an alternative method for producing cement. This process generates a higher concentration of CO_2 in the resulting flue gas (67%, as opposed to 25% in the conventional method), lending itself to be coupled with CO_2 capture [34]. The success of such processes will nonetheless be dependent on the policy support for carbon capture and sequestration and utilization. In the United States, the federal tax credit 45Q currently offers $35 per ton of CO_2 captured for beneficial use.

Lastly, public buildings make up a large portion of our global infrastructure, giving government agencies both leverage and responsibility on the front of sustainable development. In the U.S, an increasing number of states are requiring environmental product declarations, which report the embodied carbon impacts of materials, in their procurement processes. Legislation which continues to require transparency in the building construction industry will enable the transition to cleaner infrastructure.

Chapter 5

Conclusions

Efforts to reduce the carbon intensity of the concrete industry will be essential for climate change mitigation and adaptation. Machine learning can serve as an effective screening tool for concrete designers to evaluate candidate mixtures, before investing in the labor and equipment to carry out the time-intensive experiments and reducing the trial-and-error process. In concrete design, there are necessarily structural requirements, however, additional parameters, such as the cost and climate impact of the mixture, should be considered, and their trade-offs analyzed.

In this work, we have achieved successful prediction of compressive strength of concrete from the mixture attributes using various machine learning models. Notably, we find that there is alignment between the important features of our random forest model and the existing materials science knowledge surrounding concrete mixture proportioning. Additionally, we determine that prediction accuracies on small datasets can be improved by using a hierarchical model structure which learns both dataset-specific transformations and shares population information across datasets. In our optimization procedure, we find that it is possible to achieve low climate impacts of approximately 130 kg CO_2-eq, down from an industrial baseline 300 kg CO_2-eq, at a wide variety of costs by incorporating higher amounts of supplementary cementitious materials. We also find that cost is more sensitive than climate impact to increased levels of required compressive strength strength, and that moderately less expensive mixtures can be produced at extremely low climate impacts. This suggests

that there is a great opportunity for greenhouse gas emissions abatement at little to no additional cost, on the scale of 44.7 million metric tons of CO_2-eq per year in the U.S. alone.

Appendix A

Appendix

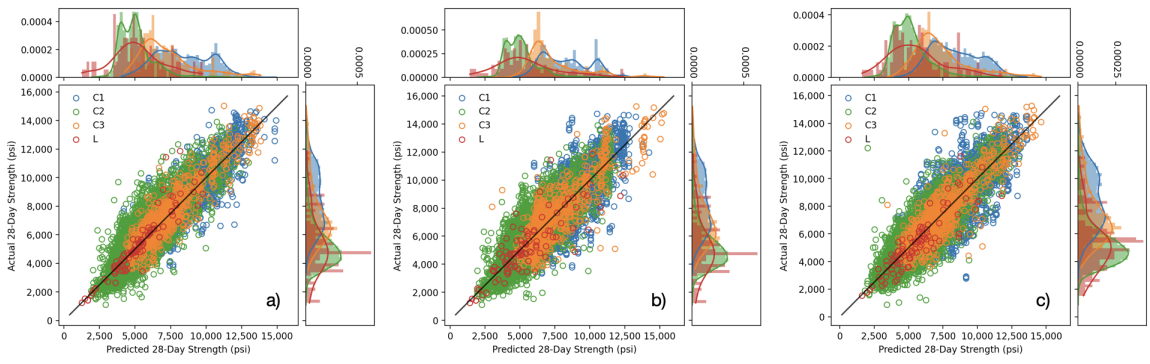

Figure A-1: Scatter plot of 28-day strength testing predictions vs. actual values for (a) multi-trunk MLP, (b) pooled MLP, (c) pooled random RF.

References

1. Monteiro, P. J. M., Miller, S. A. & Horvath, A. Towards sustainable concrete. *Nature Materials* **16,** 698–699. ISSN: 1476-1122, 1476-4660. http://www.nature.com/articles/nmat4930 (2021) (July 2017).
2. Ellis, L. D., Badel, A. F., Chiang, M. L., Park, R. J.-Y. & Chiang, Y.-M. Toward electrochemical synthesis of cement—An electrolyzer-based process for decarbonating CaCO $_3$ while producing useful gas streams. *Proceedings of the National Academy of Sciences* **117,** 12584–12591. ISSN: 0027-8424, 1091-6490. https://pnas.org/doi/full/10.1073/pnas.1821673116 (2022) (June 9, 2020).
3. Ben Chaabene, W., Flah, M. & Nehdi, M. L. Machine learning prediction of mechanical properties of concrete: Critical review. *Construction and Building Materials* **260,** 119889. ISSN: 09500618. https://linkinghub.elsevier.com/retrieve/pii/S0950061820318948 (2021) (Nov. 2020).
4. Young, B. A., Hall, A., Pilon, L., Gupta, P. & Sant, G. Can the compressive strength of concrete be estimated from knowledge of the mixture proportions?: New insights from statistical analysis and machine learning methods. *Cement and Concrete Research* **115,** 379–388. ISSN: 00088846. https://linkinghub.elsevier.com/retrieve/pii/S0008884617313807 (2021) (Jan. 2019).
5. Zhang, J., Huang, Y., Aslani, F., Ma, G. & Nener, B. A hybrid intelligent system for designing optimal proportions of recycled aggregate concrete. *Journal of Cleaner Production* **273,** 122922. ISSN: 09596526. https://linkinghub.elsevier.com/retrieve/pii/S095965262032967X (2022) (Nov. 2020).
6. Nunez, I., Marani, A. & Nehdi, M. L. Mixture Optimization of Recycled Aggregate Concrete Using Hybrid Machine Learning Model. *Materials* **13,** 4331. ISSN: 1996-1944. https://www.mdpi.com/1996-1944/13/19/4331 (2022) (Sept. 29, 2020).
7. Kalman Šipoš, T., Miličević, I. & Siddique, R. Model for mix design of brick aggregate concrete based on neural network modelling. *Construction and Building Materials* **148,** 757–769. ISSN: 09500618. https://linkinghub.elsevier.com/retrieve/pii/S0950061817309972 (2022) (Sept. 2017).

8. Omran, B. A., Chen, Q. & Jin, R. Comparison of Data Mining Techniques for Predicting Compressive Strength of Environmentally Friendly Concrete. *Journal of Computing in Civil Engineering* **30,** 04016029. ISSN: 0887-3801, 1943-5487. http://ascelibrary.org/doi/10.1061/%28ASCE%29CP.1943-5487.0000596 (2022) (Nov. 2016).

9. Ghafari, E., Costa, H. & Júlio, E. Statistical mixture design approach for eco-efficient UHPC. *Cement and Concrete Composites* **55,** 17–25. ISSN: 09589465. https://linkinghub.elsevier.com/retrieve/pii/S0958946514001346 (2022) (Jan. 2015).

10. Zhang, J., Huang, Y., Ma, G. & Nener, B. Mixture optimization for environmental, economical and mechanical objectives in silica fume concrete: A novel frame-work based on machine learning and a new meta-heuristic algorithm. *Resources, Conservation and Recycling* **167,** 105395. ISSN: 09213449. https://linkinghub.elsevier.com/retrieve/pii/S0921344921000021 (2022) (Apr. 2021).

11. Pereira Dias, P., Bhagya Jayasinghe, L. & Waldmann, D. Machine learning in mix design of Miscanthus lightweight concrete. *Construction and Building Materials* **302,** 124191. ISSN: 09500618. https://linkinghub.elsevier.com/retrieve/pii/S0950061821019516 (2022) (Oct. 2021).

12. Huang, X., Zhang, J., Sresakoolchai, J. & Kaewunruen, S. Machine Learning Aided Design and Prediction of Environmentally Friendly Rubberised Concrete. *Sustainability* **13,** 1691. ISSN: 2071-1050. https://www.mdpi.com/2071-1050/13/4/1691 (2022) (Feb. 4, 2021).

13. Yeh, I.-C. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research* **28,** 1797–1808. ISSN: 00088846. https://linkinghub.elsevier.com/retrieve/pii/S0008884698001653 (2022) (Dec. 1998).

14. Breiman, L. Random Forests. *Machine Learning* **45,** 5–32. ISSN: 08856125. http://link.springer.com/10.1023/A:1010933404324 (2022) (2001).

15. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12,** 2825–2830 (2011).

16. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. Publisher: arXiv Version Number: 9. https://arxiv.org/abs/1412.6980 (2022) (2014).

17. Paszke, A. *et al.* in *Advances in Neural Information Processing Systems 32* (eds Wallach, H. *et al.*) 8024–8035 (Curran Associates, Inc., 2019). http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

18. Lundberg, S. M. & Lee, S.-I. in *Advances in Neural Information Processing Systems 30* (eds Guyon, I. *et al.*) 4765–4774 (Curran Associates, Inc., 2017). http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf.

19. Lundberg, S. M. *et al.* From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence* **2,** 2522–5839 (2020).

20. Severson, K. *et al. Amortized inference of Gaussian process hyperparameters for improved concrete strength trajectory prediction* in *NeurIPS 2021 Workshop on Tackling Climate Change with Machine Learning* (2021). https://www.climatechange.ai/papers/neurips2021/32.

21. Rasmussen, C. E. & Williams, C. K. I. *Gaussian processes for machine learning* OCLC: ocm61285753. 248 pp. ISBN: 978-0-262-18253-9 (MIT Press, Cambridge, Mass, 2006).

22. Roberts, S. *et al.* Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **371,** 20110550. ISSN: 1364-503X, 1471-2962. https://royalsocietypublishing.org/doi/10.1098/rsta.2011.0550 (2022) (Feb. 13, 2013).

23. Wernet, G. *et al.* The ecoinvent database version 3 (part I): overview and methodology. *The International Journal of Life Cycle Assessment* **21,** 1218–1230. ISSN: 0948-3349, 1614-7502. http://link.springer.com/10.1007/s11367-016-1087-8 (2022) (Sept. 2016).

24. *A Cradle-to-Gate Life Cycle Assessment of Ready-Mixed Concrete Manufactured by NRMCA Members – Version 3* (The Athena Sustainable Materials Institute, 2020). https://www.nrmca.org/wp-content/uploads/2020/02/NRMCA_LCA_ReportV3_20200416.pdf.

25. Und Umwelt e.V. (IBU), I. B. *ENVIRONMENTAL PRODUCT DECLARATION - Concrete admixtures – Set Accelerators* EPD-EFC-20150087-IAG1-EN (Sept. 14, 2015). https://usa.sika.com/content/dam/dms/us01/i/Set%20Accelerators%20-%20SikaSet%20HE-EPD.pdf.

26. Und Umwelt e.V. (IBU), I. B. *ENVIRONMENTAL PRODUCT DECLARATION - Concrete admixtures – Air entrainers* EPD-EFC-20150086-IAG1-EN (Sept. 14, 2015). https://usa.sika.com/content/dam/dms/us01/v/Air%20Entrainers%20-%20SikaControl%20AIR-160-EPD.pdf.

27. Und Umwelt e.V. (IBU), I. B. *ENVIRONMENTAL PRODUCT DECLARATION - Concrete admixtures – Plasticisers and Superplasticisers* EPD-EFC-20150091-IAG1-EN (Sept. 14, 2015). https://usa.sika.com/content/dam/dms/us01/3/Plasticizers%20and%20Superplasticizers%20-%20Plastocrete%2010N-EPD.pdf.

28. Storn, R. & Price, K. Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *Journal of Global Optimization* **11,** 341–359. ISSN: 09255001. http://link.springer.com/10.1023/A:1008202821328 (2022) (1997).

29. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17,** 261–272 (2020).

30. Marler, R. & Arora, J. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization* **26,** 369–395. ISSN: 1615-147X, 1615-1488. http://link.springer.com/10.1007/s00158-003-0368-6 (2022) (Apr. 1, 2004).

31. Murray, A. & Price, L. *Use of Alternative Fuels in Cement Manufacture: Analysis of Fuel Characteristics and Feasibility for Use in the Chinese Cement Sector* LBNL-525E (Lawrence Berkeley National Laboratory, June 30, 2008).

32. Karellas, S., Leontaritis, A.-D., Panousis, G., Bellos, E. & Kakaras, E. Energetic and exergetic analysis of waste heat recovery systems in the cement industry. *Energy* **58,** 147–156. ISSN: 03605442. https://linkinghub.elsevier.com/retrieve/pii/S0360544213003174 (2022) (Sept. 2013).

33. Corporation, I. F. *Waste Heat Recovery for the Cement Sector: Market and Supplier analySiS* (June 2014). https://www.ifc.org/wps/wcm/connect/f0394a25-3645-4765-8291-ea33d9f09594/IFC+Waste+Heat+Recovery+Report.pdf?MOD=AJPERES&CVID=kqgTRfZ.

34. Vatopoulos, K. & Tzimas, E. Assessment of CO2 capture technologies in cement manufacturing process. *Journal of Cleaner Production* **32,** 251–261. ISSN: 09596526. https://linkinghub.elsevier.com/retrieve/pii/S0959652612001473 (2022) (Sept. 2012).