Journal of Mechanical Design. Received August 02, 2019;
Accepted manuscript posted August 11, 2019. doi:10.1115/1.4044690

Special Issue: Machine Learning for Engineering Design

Jitesh H Panchal
Associate Professor
School of Mechanical Engineering
Purdue University
West Lafayette, Indiana 47907
E-mail: panchal@purdue.edu

Mark Fuge
Assistant Professor
Mechanical Engineering
University of Maryland
College Park, Maryland 20740
E-mail: fuge@umd.edu

Ying Liu
Reader
School of Engineering
Cardiff University
Cardiff CF10 3AT, UK
E-mail: LiuY81@cardiff.ac.uk

Samy Missoum
Professor
Aerospace and Mechanical Engineering
University of Arizona
Tucson, Arizona 85721
E-mail: smissoum@email.arizona.edu

Conrad Tucker
Associate Professor
Engineering Design and Industrial and Manufacturing Engineering
The Pennsylvania State University
University Park, Pennsylvania 16802
E-mail: ctucker4@psu.edu

Modern Machine Learning (ML) techniques are transforming many disciplines ranging from transportation to healthcare by uncovering patterns in data, developing autonomous systems that mimic human abilities, and supporting human decision-making. Modern ML techniques, such as deep neural networks, are fueling the rapid developments in artificial intelligence. Engineering design researchers have increasingly used and developed ML techniques to support a wide range of activities from preference modeling to uncertainty quantification in high-dimensional design optimization problems. This special issue brings together fundamental scientific contributions across these areas.

The special issue consists of 24 papers spread over two issues of the *Journal of Mechanical Design*. The papers use various ML techniques, including artificial neural networks, Gaussian processes, reinforcement learning, clustering techniques, and natural language processing. Based on their research objective, the papers can be broadly classified into four groups: (i) ML to support surrogate modeling, design exploration and optimization, (ii) ML for design synthesis, (iii) ML for extracting human preferences and design strategies, and (iv) comparative studies of ML techniques and research platforms to help design researchers. The papers are summarized in Sections 1 through 4. An analysis of the themes covered in the special issue, and the potential opportunities for future research in ML for Engineering Design are presented in Section 5.

1 ML to support surrogate modeling, design exploration and optimization

In their paper titled *Multi-Fidelity Physics-Constrained Neural Network and Its Application in Materials Modeling*, Liu and Yang address how to incorporate multi-fidelity, physics-based constraints into Neural Network predictions. The paper contributes two key insights. First, the paper extends existing Physics-Constraints Neural Network architectures by imposing a multi-fidelity constraint scheme wherein an auxiliary network minimizes discrepancies between low and high fidelity models—essentially learning how to correct the low-fidelity one. Second, it proposes an adaptive weighting scheme to control the convergence of indi-

Journal of Mechanical Design. Received August 02, 2019;
Accepted manuscript posted August 11, 2019. doi:10.1115/1.4044690

vidual losses among the different fidelities. They demonstrate the impact of these improvements on several fundamental multi-scale material modeling challenges including two-dimensional heat transfer, phase transition, and dendritic growth problems. On these problems, the proposed multi-fidelity, physics-based constraints decrease the prediction error up to order of magnitude lower compared to networks without such constraints. This achieves comparable accuracy to that of direct numerical solutions of the underlying equations.

Sarkar *et al.* present a multi-fidelity modeling and information-theoretic sequential sampling strategy for optimization in their paper titled *Multi-fidelity and Multi-scale Bayesian Framework for High-dimensional Engineering Design and Calibration*. The approach is based on modeling of the varied fidelity information sources via Gaussian processes, augmented with efficient active learning strategies which involve sequential selection of optimal points in a multi-scale architecture. The strategy is demonstrated using the design optimization of a compressor rotor and calibration of a microstructure prediction model.

In the paper titled *A Case Study of Deep Reinforcement Learning for Engineering Design: Application to Microfluidic Devices for Flow Sculpting*, Lee *et al.* address how to design micro-fluidic flow sculpting devices by overcoming some of the key weaknesses of evolutionary optimization-based methods; namely poor sample efficiency and slow optimization convergence. The paper adapts Deep Reinforcement Learning (DRL) techniques to the flow sculpting task, and also studies the effectiveness of transfer learning on accelerating the design of target flow shapes. The paper demonstrates that DRL is able to match 90% of the target flow shapes using significantly fewer sculpting pillars than comparable GA models, as well as provides a means to interpret the learned model (using Principal Components) that existing approaches to fluidic sculpting do not provide.

Lynch *et al.*, in their paper *Machine Learning to Aid Tuning of Numerical Parameters in Topology Optimization*, present a ML-based meta-learning framework to determine tuning parameters in topology optimization. The parameters are learned from similar optimization problems carried out in the past and adjusted for the problem at hand. This helps in avoiding costly trial-and-error involved in manual parameter tuning.

In the paper *Data-Driven Design Space Exploration and Exploitation for Design for Additive Manufacturing*, Xiong *et al.* present a data-driven approach for design search and optimization at successive stages in the design process. They use Bayesian network classifier in the embodiment design stage, and Gaussian process regression in the detailed design phase. The approach is illustrated through the design of a customized ankle brace design.

Odonkor and Lewis apply data-driven design to the design of operational strategies of complex systems, specifically distributed energy resources. The paper is titled *Data-Driven Design of Control Strategies for Distributed Energy Systems*. The problem of maximizing arbitrage value is formulated as an optimization problem, and solved using reinforcement learning. The approach is demonstrated for shared distributed energy resources in multi-building residential clusters.

In *Globally Approximate Gaussian Processes for Big Data with Application to Data-Driven Metamaterials Design* by Bostanabad *et al.*, a globally approximate Gaussian process (GAGP) is introduced for the purpose of handling large data sets. A GAGP is constructed by pooling several Gaussian processes using identical hyperparameters, but built from different subsets of the training data. The predictive capability of GAGPs is shown to be at least as good as state-of-the-art supervised learning methods. It is demonstrated on the unit cell design of metamaterials though inverse optimization.

Liu *et al.* present a method for the design for crashworthiness involving categorical multimaterial structures in their paper titled *Design for Crashworthiness of Categorical Multimaterial Structures using Cluster Analysis and Bayesian Optimization*. Following a topology optimization, the dimensionality of the problem is reduced through clustering followed by a Bayesian optimization to assign a given material to a specific cluster. The approach is applied to the maximization of absorbed energy of an S-rail.

Garriga *et al.* propose a framework to assist the optimization of aircraft systems at the early design stages. The approach in their paper titled *A machine learning enabled multi-fidelity platform for the integrated design of aircraft systems* is based on the screening of designs using clustering followed by an identification of the best candidate on a Pareto front. The framework enables the use of models of various fidelities and is demonstrated on a primary flight control system and a landing gear.

2 ML for Design Synthesis

In *Synthesizing Designs with Inter-Part Dependencies using Hierarchical Generative Adversarial Networks*, Chen and Fuge present a method for synthesizing hierarchical designs with inter-part dependencies using generative models learned from examples. The method constructs multiple generative models using generative adversarial networks (GANs) while satisfying the dependencies through part dependency graphs. The paper lays the foundation for extending the use of generative models from creative individual parts to more realistic engineering systems.

The objective in *Evolving a Psycho-physical Distance Metric for Generative Design Exploration of Diverse Shapes* by Khan *et al.* is to incorporate humans' psychological perceptions about design into the design exploration process. A psycho-physical distance metric is proposed that enables the augmentation of CAD designs, based on feedback from users. Results reveal that the proposed method generates more distinct variations of CAD designs, compared to a baseline Euclidean distance method.

Oh *et al.*, in their paper titled *Deep Generative Design: Integration of Topology Optimization and Generative Models*, present a deep generative design framework for creating diverse aesthetic designs that are optimized for engineering

Journal of Mechanical Design. Received August 02, 2019;
Accepted manuscript posted August 11, 2019. doi:10.1115/1.4044690

performance. The framework integrates topology optimization and generative adversarial networks (GANs) to generate large numbers of design options from limited previous design data. The approach is validated using a 2D wheel design problem.

Deshpande and Purwar, in their paper *Computational Creativity via Assisted Variational Synthesis of Mechanisms using Deep Generative Models*, present an approach for variational synthesis of mechanisms and an End-to-End synthesis pipeline that accepts raw, high-level input from users and provides them with distinct concept solutions. The approach is based on learning the probability distribution of linkage parameters and their interdependence to perform tasks such as input conditioning, imputation, and variational synthesis. The approach is a step in the direction of enhancing users' computational creativity for engineering design.

Stump *et al.* in their paper, *Spatial Grammar-based Recurrent Neural Network for Design Form and Behavior Optimization*, present a method for simultaneous optimization of form and behavior through a combination of physics-based models and ML techniques. Specifically, they use character-Recurrent Neural Networks to embody spatial grammars and reinforcement learning to optimize the behavior. The design of a modular multi-hull sailing craft is used as a demonstration problem.

3 Extracting human preferences and design strategies

Suryadi and Kim utilize machine-learning algorithms for customer choice modeling in their paper titled *A Data-driven Methodology to Construct Customer Choice Sets Using Online Data and Customer Reviews*. They present an approach that utilizes publicly available online data and customer reviews from e-commerce websites to construct customer choice sets in the absence of both an actual choice set and customer socio-demographic data. The approach consists of clustering (i) products based on their attributes, and (ii) customers based on their reviews, and constructing the choice-sets based on a sampling probability scenario that relies on product and customer clusters. The approach generates choice models with higher predictive ability than randomly constructed choice sets.

In their paper *Extracting Customer Perceptions of Product Sustainability from Online Reviews*, El Dehaibi *et al.* seek to extract perceived sustainable design features from online reviews. Annotators from Amazon's Mechanical Turk are used to annotate product reviews and develop a natural language processing model that predicts the positive/negative sentiment of sustainable phrases. The results reveal that the model is more efficient at predicting positive sentiment pertaining to sustainable product features, compared to negative sentiments.

Raina *et al.* take a step towards transfer learning from human designers to computational agents in their paper titled *Transferring Design Strategies From Human To Computer And Across Design Problems*. They present an approach where design strategies are represented using a probabilistic model that provides a general mechanism to transfer strategies from human designers to computational design agents, and to generate new designs. The approach is illustrated using a configuration design problem.

The goal in *Learning to Design From Humans: Imitating Human Designers Through Deep Learning* by Raina *et al.* is to teach computational agents to generate designs without the need for explicit information about objective or performance metrics. A deep learning model is proposed that learns from historical human data and identifies the important regions of a design space. The results reveal that the machine learning agent learns to create designs that are comparable to human-generated ones, despite not having the same explicit feedback that humans do to guide them through the design exploration process.

He *et al.* address the challenge of mining large numbers of design ideas generated from the crowd in their paper titled *Mining and Representing the Concept Space of Existing Ideas for Directed Ideation*. The authors use natural language processing to extract keywords as elementary concepts, and represent the concepts in a way that they can be recombined to generate new ideas.

In their paper *A Data-Driven Approach to Product Usage Context Identification from Online Customer Reviews*, Suryadi and Kim use machine learning and natural language processing to identify and cluster usage contexts from a large volume of customer reviews. The methodology also captures sentiments towards a particular usage context in a sentence. The methodology enables designers to effectively use online product reviews by focusing on several specific reviews regarding particular usage contexts, and potentially to identify market opportunities for new products that excel in specific usage contexts.

4 Comparative studies of ML techniques and platforms

Sharpe *et al.* illuminate differences between Supervised Learning algorithms in terms of how and where different algorithms may apply to different Engineering Design applications. Their paper titled *A Comparative Evaluation of Supervised Machine Learning Classification Techniques for Engineering Design Applications* does this by comparing four common supervised learning approaches—Support Vector Machines, Random Forests, Gaussian Näive Bayes, and shallow depth Neural Networks—across six example problems which demonstrate different facets or challenges classifiers may face within Engineering design. The results from the work are multi-faceted with different algorithms performing better or worse under different conditions and performance measures. However, this leads to the general notion of strong problem dependence for the classifier choice and highlight the importance of understanding appropriate benchmark problems within Engineering Design that can shed light on such issues in the future.

The availability of data enables not just designers, but also design researchers. Rahman *et al.*, in their paper *A CAD-Based Research Platform for Design Thinking Studies*, present a research platform to support data-driven design-thinking and decision-making research. Through the use of

Journal of Mechanical Design. Received August 02, 2019;
Accepted manuscript posted August 11, 2019. doi:10.1115/1.4044690

fine-grained design action data and unsupervised clustering methods in conjunction with design process models, the authors show how the platform enables data-driven research studies on designers sequential decision-making behaviors.

In *Design Repository Effectiveness for 3D Convolutional Neural Networks: Application to Additive Manufacturing*, Williams *et al.* address the question of whether or not a data repository is useful for training effective ML tools. The authors experimentally test the effects of changes in CAD datasets on the precision and generalizability of trained convolutional neural networks (CNNs) for additive manufacturing applications. The study sheds light on how standardization of design repositories can influence the performance of ML tools.

Cunningham *et al.* study the construction of a performance surrogate based on 3D point cloud representations. In their paper titled *An Investigation of Surrogate Models for Efficient Performance-Based Decoding of 3D Point Clouds*, a Radial Basis Function (RBF) surrogate is used to link performance and cloud representation mapped onto a latent vector. The proposed RBF-based approach was found to be more efficient and accurate than traditional neural network-based approaches.

5 A look ahead

In the call for proposals for this special issue, the guest editors posed three primary questions:

1. How to effectively use ML for new design applications that are not well-supported by existing ML practice or tools?
2. How to leverage the unique aspects of Engineering Design in creating new ML approaches?
3. How to share benchmark problems or datasets that can measure ML progress in design?

Looking back at the papers collectively within this *Special Issue* helps shed light on the areas that are receiving significant attention within the design research community, and the areas where there are opportunities for future research.

5.1 Using ML for design applications

Engineering design research community has a strong tradition of using machine learning. Clustering techniques have been used for data-driven design; techniques such as Gaussian Process Regression and neural networks have seen applications in learning complex mappings between design and performance spaces; natural language processing has been used for mining customer reviews; and reinforcement learning is an integral part of control systems design. Many of these techniques are reflected in the papers received for the special issue (e.g., Bostanabad *et al.*; Sarkar *et al.*; Suryadi and Kim; Xiong *et al.*). Some of the recent deep learning techniques such as convolutional neural networks (CNNs) and generative adversarial networks (GANs) are finding their way into engineering design research (Chen and Fuge; Deshpande and Purwar; and Oh *et al.*).

The papers in the special issue address diverse applications including materials modeling, additive manufacturing, distributed energy systems, parametric and topology optimization, shape synthesis, mechanism design, ideation, preference modeling, and learning from human designers. Some of the areas where there is potential for leveraging ML techniques, but are not well represented in this special issue, include modeling human decision making, design of market systems, interactions of products with humans or the surrounding environment/economy, design for reliability, use of real-time data from product usage to inform design, design-for-X (manufacturing, environment, assembly, *etc.*), using data from one aspect of the lifecycle (*e.g.*, conceptual design/synthesis, manufacturing, in-use product or consumer data, re-use, *etc.*) or across different aspects of the lifecycle. There are opportunities for addressing non-traditional design challenges in engineering design, such as security, privacy, and cyber-resilience and reliability that arise from the emerging smart products and systems. Additionally, ML can also be used to support engineering design research, particularly, for supporting validation studies, and testing the generalizability of research outcomes.

5.2 Leveraging unique aspects of Engineering Design to advance ML

In terms of driving our understanding and improvement of machine learning methods, the papers in the special issue concentrate on the following areas: (1) Addressing the integration of multi-fidelity or multiple source data or simulations, with the most common application being surrogate modeling and optimization; (2) Incorporating Physics-based constraints into ML models; and (3) Building ML-based models of human preferences and behavior.

Many of the papers in the special issue present advanced surrogate modeling strategies. While this has been an active area of research for the past two decades, the approaches are being extended using advanced ML techniques. A number of papers tackled issues that arise when placing ML models in multi-fidelity situations (Liu and Wang; Sarkar *et al.*). The availability of multi-fidelity models is pervasive in Engineering Design. In some sense, designers' roles are to construct multiple types and fidelities of models or approximations of a system so as to deal with it at the appropriate level. This is not something that typical ML systems are frequently asked to do. In general, this idea of "How do I combine multiple abstractions or fidelities?" seemed important to those who submitted to this special issue and is still an active area of research.

Many of the approaches we saw here focused on adapting existing architectures/method with changes to, *e.g.*, regularization terms, that factored in, say, physical constraints. One of the gaps here is that many ML systems are not used to handling many of the multi-task problems that we need for Engineering Design. For example, in Stump *et al.*we saw how the interaction between form and behavior governs the generated designs, or in Odonkor and Lewis the interaction between Control and Design – these kinds of multi-task interactions between Structure, Behavior, and Function arise

Journal of Mechanical Design. Received August 02, 2019;
Accepted manuscript posted August 11, 2019. doi:10.1115/1.4044690

frequently within Engineering Design, and existing ML approaches often do not need to account for this.

Many papers tackled the role of design or physics constraints in a system. For example, constraints described by PDEs or ODEs (Liu and Wang), hierarchical or geometric constraints (Stump *et al.*; Chen and Fuge), or Kinematic constraints (Deshpande and Purwar). This is important since the strength of many ML systems relies on their abilities to encode strong, but specific kinds of inductive bias into the model. This is one place where Engineering Design researchers are well equipped to contribute. There are opportunities for further research in mathematically unifying prior engineering knowledge with ML models, such as physics-based models (*e.g.*, from first-principles or via more abstract system models), ontological models, formal logic/constraints, and human-gathered data/preferences. Most studies still see a problem and its physics as a black-box. In the future, can ML be used with a stronger coupling between design and physics? A related under-addressed area is learning from multiple representations of a design, *e.g.*, mathematically defining transfer or multi-task learning among design representations with multiple data structures and physics.

Many of the papers submitted to this special issue used either Supervised- or Reinforcement-Learning where label data were either provided or available via simulation. Fewer submissions used Unsupervised Learning approaches—for example, clustering and dimensionality reduction. Many recent advances in other fields like Computer Vision, Natural Language, and Speech have benefited from unsupervised learning of representations (whether using Deep Learning or other methods) and this could be a fruitful avenue for future work in Design. Further, there are opportunities for principled approaches to leveraging or managing uncertainty: for example, in ML-guided uncertainty quantification, dealing with limited, small, incomplete, or expensive-to-collect data. Likewise, verification, validation, and calibration of ML models remains an open challenge.

Many papers focused on the role of ML for assisting optimization, and specifically how to think about generative models of optima. There were two broad classes of approach to this, depending on the paper's view of the problem, which ultimately drove the approach used: (1) Optimization as low-dimension manifolds, as typified by approaches that used Auto Encoders or other latent variable methods (e.g., Oh *et al.*and Lynch *et al.*), or (2) Optimization as interactive search, such as those that formulated the problem using Reinforcement Learning or inverse problem formulations (Lee *et al.*). These avenues are still active and promising areas of research, in addition to other metaphors for ML-Assisted optimization that did not occur in this special issue, such as direct Inverse Design methods.

Some papers tackled how to best encode either designer or consumer behavior into an ML model. This is a tough challenge. There are many opportunities for blending human information and strategies into ML models, integrating cognitive and behavioral models of designers, teams, and organizations, computational creativity, or ML-models for predicting/testing designer behavior. This draws parallels to approaches people have currently used to embed known physics models into ML. Could the same be done for other types of models, such as those involving human or designer behavior?

When attempting to build ML models of human behavior (*e.g.*, designers or consumers) the design research community still heavily depends on advances in both Natural Language Processing and Sequence Learning, and this is reflected within the special issue. However, there are additional constraints that Engineering Design has on top of this, *e.g.*, a notion that text could be referring to function and to separate that out as a separate object of study. Traditional approaches to Natural Language do not emphasize such aspects and this represents a gap that the Engineering Design community can eventually fill. This also goes for designer behavior—what are the fundamental limits of viewing designers as (possibly context-dependent) Markov chains? Under what conditions and situations does this break down?

5.3 Benchmark problems and common data-sets for measuring progress at the intersection of ML and Design

Common datasets have fundamentally transformed ML research and provide a common standard for assessing performance. For example, collectively MNIST, ImageNet, ShapeNet, CIFAR, the Penn TreeBank, the KITTI Vision Benchmark Suite, among others, have been collectively cited tens of thoursdands of times. They have enabled modern advances in Vision, Natural Language, Robotic perception and others. This special issue set out to try to understand what gaps new datasets for Engineering Design might fill.

The papers in the special issue propose data-sets or platforms around Optimization problems (Sharpe *et al.*), 2D and 3D shapes (Cunningham *et al.*; Chen and Fuge), and manufacturing (Williams *et al.*). Other papers, while not providing datasets directly, do use platforms such as *Thingiverse* to gather data or conduct experiments.

Beyond these papers within the Special Issue, there is still a wide gap to be filled by future work that can contribute useful datasets or benchmarks for Machine Learning within Engineering Design. For example, here are some of the application areas that papers in the Special Issue and elsewhere have used ML for, but for which few good benchmark datasets exist:

1. Ground Truth Datasets for human or designer behavior, including how designers make decisions or evaluate alternatives.
2. Easily controllable surrogates for more complex phenomena than continuous optimization models. For example, we do not have the equivalent of, say, the Erdos-Reyni model of small world graphs for Engineering style problems such as collaboration phenomena or Mutli-Fidelity simulation models.
3. Datasets that link or use multiple types of design representations. For example, paired datasets of, say, a sketch, its corresponding CAD model, and a text

Journal of Mechanical Design. Received August 02, 2019;
Accepted manuscript posted August 11, 2019. doi:10.1115/1.4044690

description by the designer do not exist. Paired datasets have transformed ML approaches in other domains—*e.g.*, LIDAR + Camera + semantic labels in the KITTI Vision Benchmark have driven advances in Robotic perception for self-driving vehicles and paired Video and Caption datasets for Computer Vision approaches to automated closed captioning.

Beyond specific applications, in general, we do not have a good sense of what "bounds" engineering problems and when a given dataset is sufficient to cover what we need. This is not unique to Engineering Design. How do we bound or quantify the scope of, say, ImageNet, and say that it "represents" a realistic benchmark for Computer Vision problems? In some sense, to be exhaustive is perhaps impossible. So instead, the design community might focus on finding ways of quantifying or otherwise crisply stating the bounds of where and when design datasets apply and how their results translate to other problems.

Such datasets and rigorous understanding of their bounds or limitations will be critical for any future discussions of how to practically deploy such models within Engineering Design. For example, how should we handle Verification and Validation of ML models within design? What affects should this have, if any, on the development of standards? What liability or risk concerns manifest themselves once ML models influence the design of a system? What privacy concerns arise? How does this affect Design Education?

These are some of the open questions, and the papers in this special issue highlight the breadth of research opportunities in this rapidly advancing frontier. The papers in the special issue provides many launching points from which future researchers and practitioners may set sail.

The guest editorial team thanks the reviewers for all their effort in reviewing the papers, and providing constructive feedback to the authors. Special thanks to Associate Editors Dr. Xiaoping Qian, Dr. Yaoyao Fiona Zhao, and Dr. Xiaoping Du for their help with review coordination. Finally, we thank the Editor, Dr. Wei Chen, for her guidance throughout the process of realization of this special issue, and Ms. Amy Suski for all the administrative support.