

Integrating Sequence Learning and Game Theory to Predict Design Decisions under Competition

Alparslan Emrah Bayrak

School of Systems and Enterprises
Stevens Institute of Technology
Hoboken, NJ 07030
Email: ebayrak@stevens.edu

Zhenghui Sha*

Department of Mechanical Engineering
University of Arkansas
Fayetteville, AR 72701
Email: zsha@uark.edu

Design can be viewed as a sequential and iterative search process. Fundamental understanding and computational modeling of human sequential design decisions are essential for developing new methods in design automation and human-AI collaboration. This paper presents an approach for predicting designers' future search behaviors in a sequential design process under an unknown objective function by combining sequence learning with game theory. While the majority of existing studies focus on analyzing sequential design decisions from the descriptive and prescriptive point of view, this study is motivated to develop a predictive framework. We use data containing designers' actual sequential search decisions under competition collected from a black-box function optimization game developed previously. We integrate the long short-term memory networks with the Delta method to predict the next sampling point with a distribution, and combine this model with a non-cooperative game to predict whether a designer will stop searching the design space or not based on their belief of the opponent's best design. In the function optimization game, the proposed model accurately predicts 82% of the next design variable values and 92% of the next function values in the test data with an upper and lower bound, suggesting that a long short-term memory network can effectively predict the next design decisions based on their past decisions. Further, the game-theoretic model predicts that 60.8% of the participants stop searching for designs sooner than they actually do while accurately predicting when the remaining 39.2% of the participants stop. These results suggest that a majority of the designers show a strong tendency to overestimate their opponents' performance, leading them to spend more on searching for better designs than they would have, had they known their opponents' actual performance.

1 Introduction

Design is an iterative process with multiple stages that can broadly be categorized into identifying a need, generating design concepts, detailed design and implementation as exemplified by multiple formal processes [1, 2]. During this process, designers usually do not have a complete understanding of the design or the evaluation space but rather acquire knowledge in a sequential manner through multiple design evaluations. Particularly in detailed design where the design domain is more or less bounded by the concept selected for further development, by using either simulation models or experiments, designers evaluate the candidates they “sample” from the design domain to develop an understanding of where to “search” for promising designs. This search is usually not random if any learning occurs throughout the sampling process.

In reality, such a design process is not simply an individual decision-making process but can be influenced by how other designers make decisions, too. *Design under competition* is one of the most common interactive decision-making scenarios where design decisions are heavily influenced by competitors' decisions [3]. Using car design as an example, in order to identify the design features that are most preferred by customers and determine the values of design attributes (e.g., engine size, cargo space, etc.), not only do the designers need to consider the information from their own company, they must have information from their competitors so as to produce competitive products in the market for profit. Because of the existence of competitiveness, designers' extrinsic motivation could be biased, thereby their decisions may become irrational. Therefore, the research question of how designers make sequential design decisions under competition does not only have practical implications in many engineering applications, but also attracts scientific in-

*Address all correspondence to this author.

quires. The answer to this question can directly advance the understanding of human cognition in complex design scenarios and inform the development of forward-looking decision support systems for human-artificial intelligence (AI) collaboration. While human cognition has been known to suffer from several biases under uncertainty such as availability heuristic (i.e., using the information that easily comes to mind) or anchoring bias (i.e., relying on the earlier information to make future decisions) [4–6], a collaborative AI system developed considering human decision-making processes can potentially help designers (or enterprises) make rational decisions from design to business.

The study of human decision-making process can be approached from three perspectives, namely, descriptive, prescriptive and predictive, depending on the research interests and the questions to be answered. Descriptive analysis uses data aggregation and data mining to provide insight into the past design behaviors and answers: “What has happened and how were the design decisions made?”. Prescriptive analysis uses modeling and optimization to advise on possible outcomes and answers: “How should the design decisions be made?” Finally, predictive analysis uses statistical models and forecasting techniques to understand the future and answers: “What could the future design decisions be?”.

Prior literature has studied sequential design processes from different perspectives. For example, in support of product development and project management, design structure matrices [7–9] have been used for task sequencing to identify the sequence that minimizes expected project completion time. Other studies have primarily focused on the design search process based on optimization algorithms. Multi-objective formulations have been introduced in the literature to study the design process sequentially advancing through smaller sets of alternatives using models of increasing fidelity [10,11]. Additionally, the expected value of perfect information (EVPI) [12], Bayesian optimization [13], generic algorithm [14], and optimal learning [15], have been adopted in the literature to study optimal design sequences. These studies, however, are very different from the work presented in this paper in that these studies investigate the optimal sequential decision process using normative models which study how the design decisions should be made.

In the present paper, we investigate *humans*’ actual sequential design decisions. Chaudhari et al. [16] recently performed study to identify models that provide the best description of a designer’s sequential decisions when multiple information sources are present and the total budget is limited. So, this study tries to find the appropriate *descriptive* model for sequential design decisions. Computational models have also been used to learn and “replay” designers’ sequential decision-making using Markov chains [17–20], simulated annealing [21], Gaussian process [22–24], and more recently the deep learning-based methods [25, 26]. Yet, no computational models have been developed to date for predicting designers’ sequential decisions with the consideration of competition.

In [24], the authors have used a function optimization game for design research and studied designers’ information

acquisition decisions under competition. That study has developed normative models of design decisions under competition to answer questions more related to the descriptive and prescriptive aspects of sequential design decisions. In the present paper, we adopt the experimental settings (detailed in Section 5) of [24] and propose a predictive approach to model what future decisions individuals could make in a sequential design process under competition given their past decisions. The uniqueness of this approach is that it first integrates sequence learning (using a long short-term (LSTM) unit) with game theory (using a non-cooperative game) in studying engineering design under competition.

Specifically, we use prior data from human subject experiments with the function optimization game introduced in [27] to learn search behaviors of the participants (i.e., designers) under one-to-one competition and to predict their future decisions based on their past behaviors. The *research question* we aim to answer in this study is: “to what extent can the game theory and sequence learning-based models predict and explain designers’ sequential decisions under competition?” The answer to this question can provide insights as to how effective collaborative AI systems could be developed to pro-actively guide human design decision-making.

The rest of the paper is organized as follows. In Section 2, the technical background on the non-cooperative game and the LSTM model is provided. These two models are the core components in the proposed predictive framework. Section 3 presents how the research problem is formulated and the proposed research approach. In Section 4, we introduce the function optimization game, and describe the experimental settings based on such a game for the collection of the sequential design data. In Section 5, the proposed model which integrates LSTM and the non-cooperative game is introduced in detail. In sections 6 and 7, the results from both analysis and validation are presented and discussed. Insights from the analysis are also summarized in Section 7. At the end, we conclude this paper with a further discussion on how our approach can be transferred to other engineering design scenarios and the future work.

2 Technical Background

2.1 Non-Cooperative Games

A design competition is typically modeled as a non-cooperative game. Non-cooperative games are often inherently zero-sum where what one wins the other loses. Designers’ behaviors are studied at the equilibrium of such games. Various game-theoretic models have been developed in design literature to provide insights into designers’ (or enterprises’) behaviors, their design options and strategies. For example, market competition has been modeled with game-theory and the optimal price and design decisions under competition has been studied with long-run and short-run equilibrium solutions [3, 28, 29]. Further, game-theoretic models have been used to model rationality of the designers for collaborative, decentralized design scenarios [30].

In game theory, each player (or decision maker) is as-

sumed to act with rational behaviors. That means they are self-interested agents whose goal is to maximize their own payoff. When using non-cooperative games to model a competition, the payoff (π_i) is dependent on the prize of the competition (e.g., the total revenue of a certain product if that product wins the market) and the probability of winning that prize. The probability of winning (P_i) is a function of the quality of the submissions from every competitor, and depends on the competitors' characteristics (e.g., expertise) and inputs (e.g., effort, time investment). Hence, a general model of a design competition based on game theory has three main parts: quality function, winning probability function, and payoff function [27].

1). *Quality function* describes the quality of a design solution (q_i) as a function of the designers' characteristics, such as their expertise (K_i) and the inputs, such as effort (e_i). For a real-time competition where competitors' design decisions are known, the effort e_i is affected by the design characteristics and inputs from other designers in a rational decision-making process. Designers can adjust the amount of effort to spend on improving their design to find a design solution better than their competitors' for maximum pay-off. In this paper, we assume that competitions are not in real time, i.e., there is no real-time information exchange, and designers make decisions based on their past experience with competitors, which is included in K_i . Thus, in this paper, q_i is assumed to be independent from the characteristics and the inputs of other designers in the competition.

$$q_i = q_i(e_i, K_i) \quad (1)$$

2). *Winning probability function* defines each competitor's probability of winning as a function of the quality of all submitted design solutions [31], and is typically modeled in an additive form,

$$P_i = \begin{cases} \frac{f(q_i)}{\sum_{j=1}^N f(q_j)} & \text{if } \sum_{j=1}^N f(q_j) > 0 \\ \frac{1}{2} & \text{otherwise,} \end{cases} \quad (2)$$

where $f(q_i)$ is a non-negative increasing function. The empirical results obtained from previous studies indicate that using a power function for the winning probability and an exponential function for the quality well captures the designers' winning status in a two-player competition game [27].

3). *Payoff function* defines the expected value of the prize, e.g., in a winner-takes-all game, the payoff of an individual is

$$E(\pi_i) = \Pi P_i - C_i, \quad (3)$$

where Π is the amount of the prize and C_i is the cost incurred in developing the solution. The Nash equilibrium of

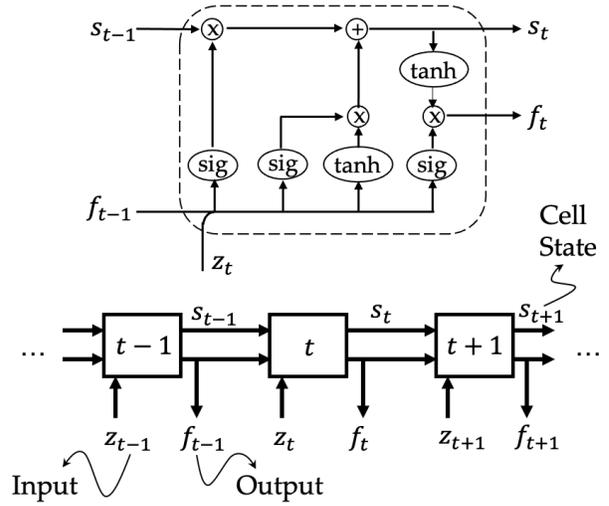


Fig. 1: Overview of an LSTM network architecture (adapted from [36])

the game is generally used as the solution of the game. At the Nash equilibrium, a rational player i chooses the input (e_i) that is the best response to other players' best responses.

2.2 Sequence Learning

Several machine learning methods have been designed to learn from independent and identically distributed data. Sequential design decisions are a particular type of data that does not fall within that category since subsequent decisions are correlated with the previous ones. Recurrent neural networks (RNNs) in sequence learning literature have been developed for such problems [32, 33]. These networks use the output of the hidden layer corresponding to the previous sequence as an input to the next sequence to retain dependencies. Drawbacks of the classical RNNs such as vanishing gradients where gradients become smaller in each layer led to the emergence of a more powerful variant of RNNs known as long short term memory (LSTM) networks [34, 35].

LSTM networks consist of a repeated chain of cells each of which represents the network for a time-step. Adapted from [36], Fig. 1 depicts the overall network architecture and the contents of a cell. These networks are capable of capturing long-term relationships in sequential data by taking inputs from both the output and the internal state of the previous cell. Three *sigmoid* functions in Fig. 1 output values in $[0, 1]$ and allow selecting what information to retain during learning. Two *tanh* functions in the figure scale the input and output data into $[-1, 1]$ and allow the gradients to sustain for long time periods. Multiple variants of the LSTM networks exist. We refer readers to [35] for a detailed discussion on the practical use of these variants.

In this paper, we use LSTM networks to predict future design decisions under a competition based on past decisions, assuming that past behavior can be related to future behaviors. Earlier, normative models have been used to explain designer behaviors under the same setting [24]. While

Function Optimization Game

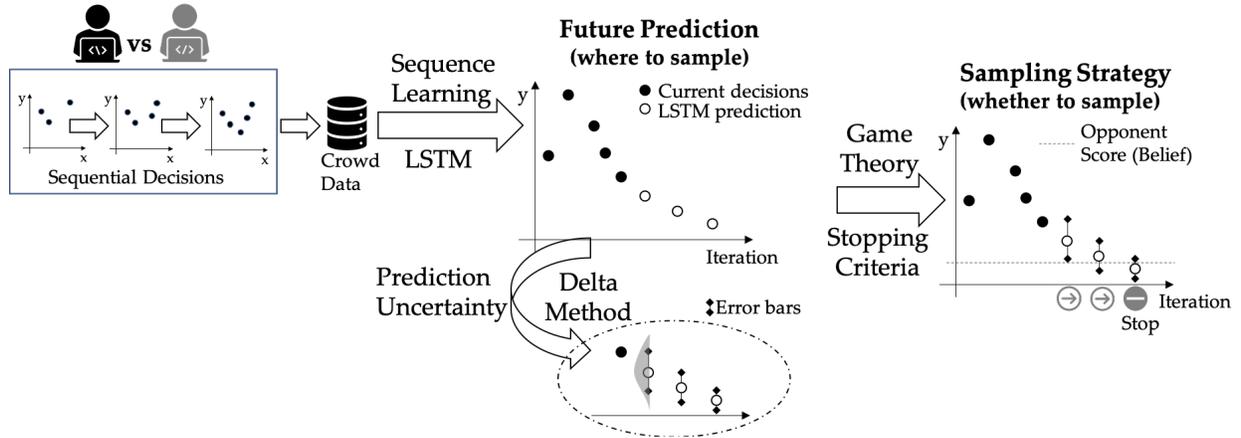


Fig. 2: Overview of the Learning and Prediction Process

these models provide interpretable outcomes, they make several assumptions regarding designer behavior such as assuming certain bounds for future design decisions. LSTM networks could extract complex behaviors without making such assumptions and use this knowledge to predict future design decisions. Also, a machine learning-based approach is more generalizable to different design problems.

Classical LSTM networks output a deterministic value for a given set of inputs, i.e., a point estimate of the expected value, while a model with uncertainty is necessary to predict future design decisions due to the natural variation in designer behaviors. We combine the classical LSTM network models with an existing method proposed for estimating prediction intervals from the literature, known as Delta method [37], to develop a probabilistic model that predicts the designer decisions with a distribution.

Fig. 2 shows an overview of the learning process in the predictive model presented in this paper. We use the sequential design data from a crowd to train an LSTM network that predicts the expected next designs to sample and the corresponding quality (i.e., function value in the game) based on past history of design decisions. Note that Fig. 2 depicts a minimization problem where smaller values of y in the vertical axis represent better designs. We use the LSTM network with the Delta method to estimate the uncertainty in the prediction. Finally, in a competition, designers form an opinion or a belief of how good the competitor's design could be and determine how much to spend on improving the existing design based on that belief. We combine the future prediction and the corresponding uncertainty using game theory to estimate the designers' belief regarding the quality of their opponent's design and predict whether the designers will make the predicted design decision. For instance, the proposed model predicts that the designer will stop searching for new designs if the predicted design quality is better than the designer's belief of the opponent's design quality as depicted in Fig. 2. We present the details of our mathematical approach in Section 5.

3 Problem Formulation and Research Approach

In this study, we are particularly interested in designers' decision-making behavior in the parametric design stage where a design problem has been already defined (i.e., the objective is set and the key design variables are known), and a designer's task is to determine the value of a specific design variable (e.g., a dimension or a material property). The goal of the parametric design, therefore, is to search the design space to find a design that satisfies the requirements and constraints in an optimal way with respect to a given objective [38]. While there could be other views on this design process, this optimization-based view is adopted in this study. As highlighted by Papalambros and Wilde [39], "philosophically, optimization formalizes what humans (and designers) have always done. Operationally, it can be used in design, in any situation where analysis is used". Based on this view, we focus on a design process with the following characteristics [24]:

- C1. A designer's goal is to find the best design quantified by certain objective values.
- C2. Designers evaluate the performance of candidate designs, either through simulations or physical experiments.
- C3. There is a cost associated with searching. The term "cost" is used more generally. Costs can either be monetary cost or effort (computational, personnel, etc.).
- C4. More costs incurred in exploring the design space result in a better understanding of the design space, and therefore, designs of equal or potentially better quality.

In this paper, we focus on a design problem solved by individuals in a design competition that possesses all the characteristics we list in C1-C4. In this competition, the expected payoff is not only determined by the quality of a designer's own decisions but also by the decisions of other competing designers, see Equation (2). For a contestant, further experimentation may yield better design quality, and hence a higher probability of winning (C1-C3), but also greater cost (C4). This situation is similar to the real-world competitive

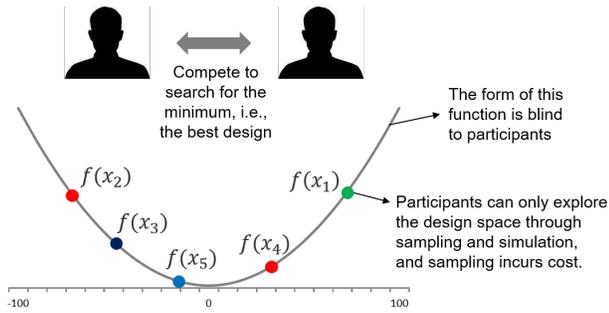


Fig. 3: The function optimization game - Minimization of an unknown function

markets where companies need to balance the added value of building a new prototype and the corresponding cost in terms of time and budget.

We present a model to predict future design decisions in such a competition. We use the function optimization game discussed in Section 4.1 that emulates such a tournament to collect data for sequential design decisions from participants. The predictive model integrates probabilistic models of sequence learning with game theory. The model predicts whether a designer will continue searching for better designs, if so, what part of the design space this designer will explore and how much improvement the designer will achieve.

4 Experiment Description

4.1 Function Optimization Game

Sha and co-authors [27] developed a function optimization game that creates a simplified scenario of design under competition yet capturing the essence of the design process characteristics specified in Section 3. So in this paper, we adopt the data collected from this game experiment in support of our investigation. Here, we summarize the key features of this game briefly and refer the reader to [27] for the details.

In the optimization game, the participants are asked to optimize a design characterized by a single variable $x \in [-100, 100]$, and its performance is evaluated by an unknown function, $f(x)$ as shown in Fig. 3. In this specific case, the participants are asked to minimize this unknown $f(x)$. This resembles many real-world design problems where the functional behavior of an artefact is not completely known. Each participant can query the value of the function for a specified x at a cost of c tokens. Each participant plays the game against one other player, randomly selected during each period. At the end of each period, the participant whose design achieves a smaller function value wins the fixed prize (Π). This game embodies the sequential information acquisition decisions, and enables the study of strategic decisions. The domain independent nature of the problem reduces the variations among designers due to the diversity in knowledge background, thus reduces noises in the data which will be beneficial to testing the models.

4.2 Experimental Setup

Based on this game setting, an experiment was carried out with 44 senior undergraduate Mechanical Engineering students at Purdue University by Sha and co-authors for an earlier study [27]. The present paper uses the same data to illustrate the application of the proposed model. The following bullets summarize the key experimental settings, and further details of the experimental data are provided in Section 6.1.

1. The control factor is the cost. Each subject participated in two *treatments*: low cost treatment ($c = 10$ tokens) and high cost treatment ($c = 20$ tokens). In both treatments, the participants start with 200 tokens. The experiments consists of 4 sessions where each session has two treatments as follows. Session 1: low treatment first and then high treatment; Session 2: high treatment first and then low treatment; Session 3: high treatment first and then low treatment; and Session 4: low treatment first and then high treatment.
2. The repetition is realized by 15 *periods*. A period refers to one full competition cycle between two players based on a randomly generated quadratic function. The coefficients a and b of the function $F(x) = (x - a)^2 + b$ are randomly drawn from a uniform distribution in each period. So, a participant plays the game 15 times in low cost treatment and 15 times in high cost treatment with 30 different functions in total.
3. The matching mechanism is random. In experiments involving partners, there may be chances for tit-for-tat strategies (learn from your partner). This could eventually affect the analysis of competitive decisions. So, in this experiment, two participants competing with each other are randomly matched; and at the beginning of each period, every pair of participants will be re-matched.
4. The awarding mechanism is winner-takes-all. For each period, the winning player receives the prize amount (200 tokens) minus the cost of sampling, whereas losing player gets nothing. At the end of the experiment, the final prize is accumulated from the prize received from each period.

With these experimental settings, a participant needs to sequentially make two essential decisions multiple times in each period of the game, a) the decision to choose next x , and b) the decision to whether stop or not. The information available for these decisions are: at the end of each period, the participants are informed whether they won or not, the best point archived by the winner, and the actual optimal value of the function in that period. Please note that this feedback does not directly reveal the opponent's searching strategy in the future periods because the function is randomly generated and the competitors are randomly matched at the beginning of a new period. However, this allows the players to form a belief (or a guess) about their opponents' best designs in the subsequent periods.

5 Sequence Learning with Game Theory

We use the dataset obtained from the experiment described in Section 4 to train an LSTM network that gives a probabilistic output, and combine it with the game theoretic model proposed in [24]. The LSTM network and the game-theoretic model serve complementary purposes. While LSTM network predicts a distribution for the next design sample and the corresponding function value, the game theoretic model calculates the belief that a designer has regarding the performance of the competitor and determines whether a designer will continue sampling based on the predicted performance improvement. When using LSTM networks and game theory, we make the following assumptions regarding designer behaviors:

- A1. Predictability. Past design decisions and the corresponding outcomes are the predictors of future decisions and outcomes.
- A2. Rational behaviors. Individuals stop sampling when the expected improvement in their payoff given by Eqn. (3) is negative.

5.1 Prediction of Next Design Samples and Outcomes

We use three types of historical information available to the players to predict the future behaviors: a given number of previous design samples, function values and cost of sampling. Instead of training different models for each treatment, we use cost as an input to combine the information regarding both the number of samples a participant has taken (i.e., the number of design iterations) and treatment. Cost linearly increases with the number of samples where the treatment determines the unit (marginal) cost of each design sample. Integration of cost into our model is also an advantage over the normative approach since we provide a single unified model that describes designer behavior under different experimental conditions.

The same LSTM network can be trained to predict the expected value of both the next design sample $\mu_{\tilde{x}_0}$ and the next function value $\mu_{\tilde{y}_0}$ based on past information as follows:

$$\begin{aligned}\mu_{\tilde{x}_0} &= f_x(\mathbf{x}_p, \mathbf{y}_p, \mathbf{c}_p, \mathbf{w}), \\ \mu_{\tilde{y}_0} &= f_y(\mathbf{x}_p, \mathbf{y}_p, \mathbf{c}_p, \mathbf{w}),\end{aligned}\quad (4)$$

where $\mathbf{x}_p, \mathbf{y}_p, \mathbf{c}_p$ are vectors containing p number of past design samples, their corresponding function and cost values, respectively, and \mathbf{w} is the vector of network weights that minimizes the training error. Also, $\tilde{\cdot}$ denotes a predicted quantity. The experiment data contains sequences of varying length since each player freely determines how many samples to take from the design space. For training, we organize the entire data set into samples of length $p + 1$. For instance if a player takes q samples from the design space in an experiment where $q > p$, that experiment gives $q - p$ samples for training. Here, the number of steps to look back to obtain a good prediction depends on the application. While a longer sequence of past history could allow learning more complex behaviors, it also reduces the number of data points to use for training.

5.2 Prediction of Uncertainty

The LSTM network provides a deterministic output for a given set of inputs whereas in reality, due to several human factors that are not included in the LSTM model, future decisions from different designers vary even if their past decisions are the same. Therefore, while the deterministic output of the LSTM network could serve as a point estimate of the expected value, a probabilistic output is necessary to account for the natural variation in human behaviors. Also, the stopping criteria presented in Section 5.3 requires a model of uncertainty in the future predictions. There are multiple approaches in the literature to estimate the distribution of the output prediction from neural network models as summarized in [40]. In the present study, we use the Delta method [37] for its computational efficiency. This method makes an assumption of normality of the overall prediction error and calculates the variance of the prediction output based on the training error and the sensitivity of the network output with respect to the prediction weights. We provide a summary of the final output and refer readers to the original references for derivations [37, 40].

The total variance of the prediction output can be calculated as follows.

$$\sigma_0^2 = \sigma_\epsilon^2(1 + g_0^T(J^T J)^{-1}g_0), \quad (5)$$

where σ_ϵ^2 is the variance of the training error, g_0 is the gradient of the output of the LSTM network with respect to the prediction weights evaluated at a test point \mathbf{z}_0 given by

$$g_0^T = \left[\frac{\partial f(\mathbf{z}_0, \mathbf{w})}{\partial w_1}, \frac{\partial f(\mathbf{z}_0, \mathbf{w})}{\partial w_2}, \dots, \frac{\partial f(\mathbf{z}_0, \mathbf{w})}{\partial w_m} \right], \quad (6)$$

and J is the Jacobian matrix of the LSTM network with respect to the prediction weights evaluated at the training samples $\mathbf{z}_1, \dots, \mathbf{z}_k$ given by

$$J = \begin{bmatrix} \frac{\partial f(\mathbf{z}_1, \mathbf{w})}{\partial w_1} & \frac{\partial f(\mathbf{z}_1, \mathbf{w})}{\partial w_2} & \dots & \frac{\partial f(\mathbf{z}_1, \mathbf{w})}{\partial w_m} \\ \frac{\partial f(\mathbf{z}_2, \mathbf{w})}{\partial w_1} & \frac{\partial f(\mathbf{z}_2, \mathbf{w})}{\partial w_2} & \dots & \frac{\partial f(\mathbf{z}_2, \mathbf{w})}{\partial w_m} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial f(\mathbf{z}_k, \mathbf{w})}{\partial w_1} & \frac{\partial f(\mathbf{z}_k, \mathbf{w})}{\partial w_2} & \dots & \frac{\partial f(\mathbf{z}_k, \mathbf{w})}{\partial w_m} \end{bmatrix}. \quad (7)$$

We model the final prediction output using a normal distribution with the expected value from Eqn. (4) and the variance from Eqn. (5).

5.3 Stopping Criteria

We combine the LSTM-based probabilistic prediction model with a game theoretic formulation proposed by Panchal et al. [24] to predict a player's belief regarding the opponent's best response and the decision to stop sampling or not. Let the superscript (i) denotes the information from

the player i and $(-i)$ denotes the information from the corresponding opponent. At a given time, let $\tilde{y}_*^{(-i)}$ be the best function value the player i believes that the opponent has achieved, and $y_*^{(i)}$ is the best function that the player i achieved up to that point. The stopping criteria from [24] is given by:

$$\text{stop if } \begin{cases} \phi \left(\frac{\tilde{y}_*^{(-i)} - \mu_{\tilde{y}_0^{(i)}}}{\sigma_{\tilde{y}_0^{(i)}}} \right) < \frac{c}{\Pi} \quad \text{and} \quad y_*^{(i)} > \tilde{y}_*^{(-i)} \\ \text{or} \\ y_*^{(i)} < \tilde{y}_*^{(-i)}, \end{cases} \quad (8)$$

where ϕ is the standard cumulative normal distribution function. This criteria can be used to predict whether a player will continue sampling or not for given $y_*^{(-i)}$, or it can be used to estimate a lower and upper bound on $\tilde{y}_*^{(-i)}$ using the data at the point where a player stopped sampling. This criteria combines the information learned from the crowd through the LSTM network and the information from the individual participants through their own beliefs regarding their opponents.

6 Results

6.1 Descriptive Analysis

In this section, we present key statistics to describe the original experiment data obtained from the function optimization game described in Section 4. We report median and interquartile range (IQR) instead of mean and standard deviation for these statistics since a few outliers in the data impact the latter significantly. Since the experiment data consists of inputs from 44 participants and each participant plays the function optimization game for 15 periods per each of the two cost treatments, this results in a data set for 1320 periods in total. We omit the first five periods in each participant (i.e., remove 440 periods) in our analysis for offsetting participants' learning curves.

For the remaining 880 periods, the median number of samples the participants have taken is 7.0 for winning players and 6.0 for losing players, and the IQR is 4.0 for both. The median of the absolute distance between participants' best design sample and true optimal design for winning and losing players are 0.48 and 2.79 with an IQR of 1.38 and 8.96, respectively. Similarly, the median distance between participants' best function value and true optimum for winning and losing players are 0.05 and 1.56 with an IQR of 0.45 and 19.65, respectively. These results are summarized in Tbl. 1. Note that these numbers refer to the absolute quantities without any normalization.

The smaller variation in the performance (quantified by the distance between the design variable values/function values and the true optimum) of the winning players as compared to that in the losing players indicates that the players attempt to be competitive regardless of their opponent's score. This result suggests that the players have developed a

belief of a strong opponent regardless of the actual ability of their opponent.

Table 1: Summary of descriptive statistics from the original dataset

Statistic	Median	Median	IQR	IQR
	(Win)	(Loss)	(Win)	(Loss)
Number of Samples	7.0	6.0	4.0	4.0
Absolute Distance of the Best x from the Optimum	0.48	2.79	1.38	8.96
Absolute Distance of the Best y from the Optimum	0.05	1.56	0.45	19.65

6.2 Predictive Analysis

In this section, we present the predictive analysis results obtained from the mathematical models described in Section 5. We provide the interpretation of these results along with a discussion in Section 7.

6.2.1 LSTM Network

We use the history of previous three design samples, function values and the cost incurred to train an LSTM network that predicts the next design sample and the function value. We remove the data points where participants take less than four samples, which is the minimum number of samples needed for LSTM model. Using the data after the fifth period in all treatments, we pre-process the data from 880 game periods to organize them into chunks of four sequences where we use the first three to predict the last. This pre-processing gives 2481 data points to use for the LSTM model. Note that the number of samples the participants take (reported in Tbl. 1) directly affects the number of data points to use for training. Had the participants taken more samples, it would have provided a larger and richer data set. We use 6-fold cross validation and present the results for one of the subsets in this section. All the results in this section are evaluated on the test data.

Using a normal distribution, we can estimate a lower and upper bound on the predicted quantity, i.e., the next design sample or function value, with a given confidence. Fig. 4 shows the prediction intervals for the next design sample and function value (from Eqn. (4)) with one standard deviation (from Eqn. (5)) away from the mean. Filled dots and hollow circles represent the upper and lower bounds, respectively, and the solid line corresponds to the perfect prediction. The x and y axes in the figure represent normalized design variable and function values. In Fig. 4a the design interval of $[-100, 100]$ is mapped to $[0, 1]$, and in Fig. 4b, 0 corresponds to the true optimum of the function that participants try to minimize and 1 corresponds to the maximum function value observed by two opposing participants. The prediction intervals shown in Fig. 4a accurately bounds 82% of the test

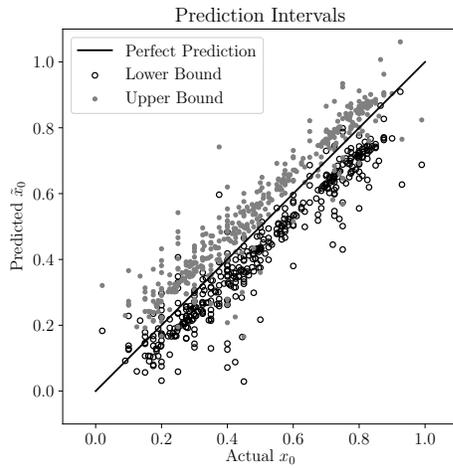
points from above and below, and in Fig. 4b, the accuracy is 92%. These prediction intervals can be narrowed down at the expense of losing accuracy.

Next, we present the prediction results for the overall behavior of the participants throughout the search process. The plot in Fig. 5a shows that the difference between two consecutive function values in the search process $|\Delta y|$ gradually decreases with the number of iterations (i.e., the samples that participant take) in the actual data. These results mean that the participants tend to make larger changes in their design early in the process compared to later. The plot in Fig. 5b depicts the same results from the prediction model. A similar trend is observed for the design samples but we omit those results for brevity. Since the LSTM model needs a past history of three samples, the prediction results start from the fourth sample. The overlap between two plots shows that the predictive model can reasonably capture the exploration and

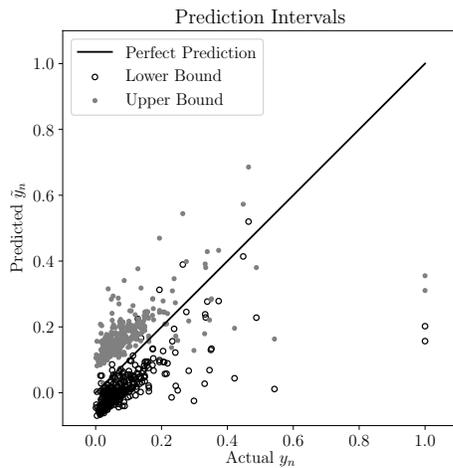
exploitation behavior as a function of the design iteration in this experiment. As opposed to the normative models from the literature such as simulated annealing [21] or Gaussian processes [24,27], the proposed model learns the exploration and exploitation from the input data without “hard-wiring” these behaviors.

6.2.2 Game Theory

We also present results for the ability of the game theoretic model in Eqn. (8) to predict and explain some of the participant behaviors under competition. First, we show the results for the participants’ belief for their opponent’s best function value $\tilde{y}_*^{(-i)}$. We apply Eqn. (8) that uses the predictions from the LSTM network to the data points where the participants stop sampling. Using the values for $\mu_{\tilde{y}_0}$ and $\sigma_{\tilde{y}_0}$ from the LSTM network and the participants best score at the time $y_*^{(i)}$ from the dataset, we obtain an upper and lower bound for $\tilde{y}_*^{(-i)}$ from the two inequalities in Eqn. (8). Fig. 6 compares these bounds for the participants’ belief with the

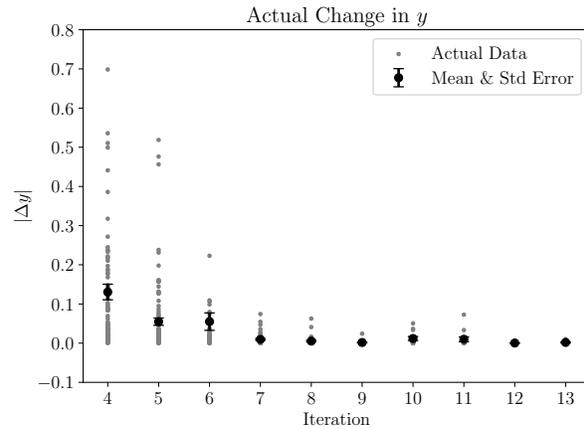


(a) Prediction of Next Design Sample

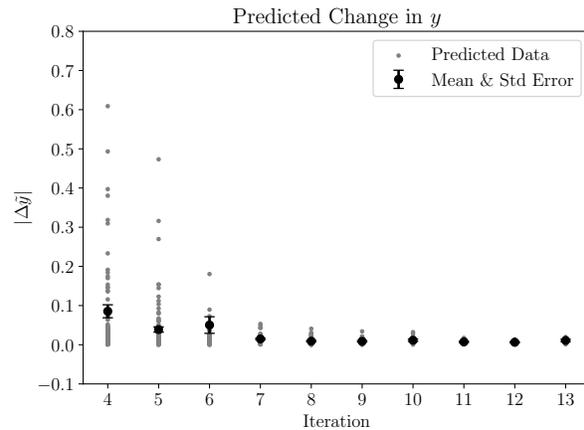


(b) Prediction of Next Function Value

Fig. 4: Prediction intervals for the next design sample and function value with one standard error away from the mean



(a) Results from the Actual Data



(b) Results from the Prediction Model

Fig. 5: Change in the function value (i.e., the difference between the current and next function values) with the number of samples the participants take

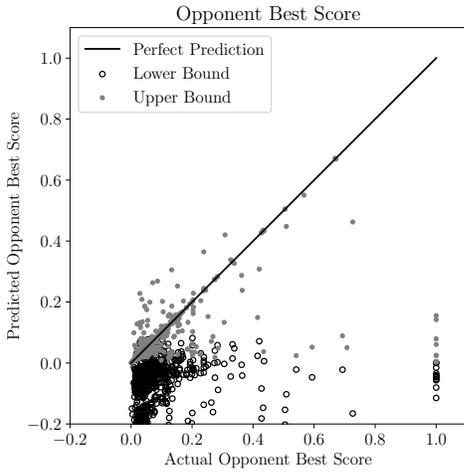


Fig. 6: Prediction results from the game theoretic model for the participants' belief of their opponent's best score

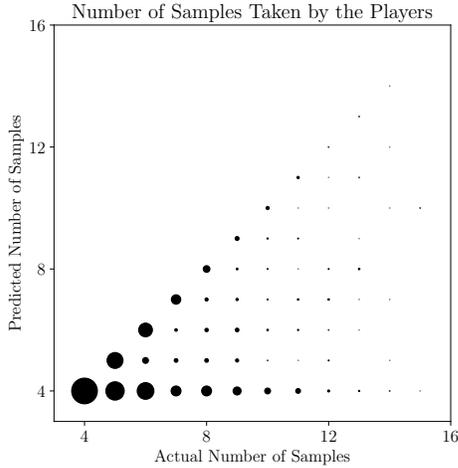


Fig. 7: Prediction results from the game theoretic model for when the participants will stop sampling. The size of the circles represents the frequency of occurrence.

actual best score of the opponent. The error in the upper and lower bounds increase as the opponent's best function value increases. 53% of the actual opponent scores fall within the bounds shown in Fig. 6.

We also use the game theoretic model to predict when the participants stop sampling given their opponent's best score. Fig. 7 shows this prediction in comparison to the actual point where the participants stopped sampling. The comparison starts from the fourth sample, which is the earliest point in the sampling process our model can predict. The figure shows that the game theoretic model generally predicts an earlier stop compared to the actual data.

7 Discussion

In this section, we interpret and discuss the results presented in Section 6 and provide some key take-aways.

The prediction of the next design sample using our LSTM network provides an accurate interval when the next design sample yields a small function value. Fig. 4b shows that if the next function value is small (e.g., $y_0 < 0.1$), the estimate for the upper bound is mostly above and the lower bound is mostly below the line (with an accuracy of 96%). The error in the prediction interval tends to increase if the next function value is large. This error can partially be attributed to the lack of sufficient data points with large y_0 value. Considering that smaller function values represent better designs in this study, learning from the crowd data in the LSTM network is biased toward behaviors that yield better design outcomes. Such a result is expected since the primary goal of the participants in the design contest is to improve their designs and the crowd data contains that behavior.

The present LSTM network also has the ability to capture some interpretable designer behaviors in the experiment. The results in Fig. 5 show the a gradual transition from exploration to exploitation as the participants take more samples from the design process. Since the participants gain a better knowledge of the design space with more samples, it is expected for participants to search for a promising region in the design space with exploration early in the process and focus on fine-tuning their results later. Recall that each sample incurs a cost. Therefore, it is not reasonable to make large changes in the design after identifying a promising region to search. This conclusion is consistent with the obtained from the analysis by the Wiener Process model (a normative model) previously studied by Panchal et al. [24].

The results from the game theoretic model show a clear tendency to assume all opponents to be well-performing designers in this experiment. In Fig. 6 the upper and lower bounds for the participants' belief regarding the opponent score obtained from the game-theoretic model are generally lower than the actual score. Since lower scores represent better designs, these bounds predict the opponents to perform better than what they actually achieve. The gap between the results from the model and actual data increases as the opponent score becomes worse.

There are several reasons for the difference between the game-theoretic prediction and the real opponent scores. First, the game-theory assumes the players to stop sampling when the expected improvement in their payoff is negative as discussed in Section 5. The actual behaviors may violate that assumption. As a result, if the participants keep sampling at the expense of unnecessary cost resulting in a prediction of lower opponent score than reality. While the overall accuracy of these bounds is 53%, the accuracy increases to 69% when we use only the opponents who achieved a relatively good score of $y_*^{(-i)} < 0.1$. Second, the prediction results are compared with the actual opponent score rather than the participants' belief, which is not collected during the experiment. Some of the inaccuracy may stem from the error in the

participants' belief. Finally, a small amount of error coming from the LSTM prediction (a mean squared error of 0.005 between predicted $\mu_{\hat{y}_0}$ and the actual y_0 in the test data) also has a contribution.

The results in Fig. 7 depict the same characteristics of our game-theoretic model from a different perspective. Using the actual score of the opponent, this figure shows that the prediction from our game-theoretic model is always less than or equal to the actual number of samples that participants take, i.e., there are no circles above the main diagonal as shown in Fig. 7. The size of the circles represents the frequency of occurrence. The circles on the main diagonal correspond to the cases where the prediction matches the reality. In this figure, 39.2% of the data points fall on the main diagonal and the remaining 60.8% of the points are below the diagonal. Circles become smaller as the actual number samples participants take increase since there are relatively fewer participants who take a large number of samples in a single period. This figure shows that if the participants knew the actual opponent score, they would stop earlier under the assumptions of the game-theoretic model. The figure also supports our claim that the participants' belief regarding their opponent's function value is smaller (or better) than what their opponents actually achieved.

Note that the results we present represent a particular subset from a 6-fold cross-validation approach. Each subset provides similar outcomes where the interpretations and conclusions are not affected by which subset is used. For instance, the prediction accuracy in Fig. 4b is 92% for the subset we present. The other subsets yield an accuracy of 93%, 88%, 89%, 88% and 95%. Therefore, we avoid the presentation of the entire 6-fold cross validation results for brevity.

8 Conclusions

We presented a mathematical model to predict sequential design decisions from human designers based on their past decision. The model uses LSTM networks to predict expected future design decisions, the Delta method to represent the prediction uncertainty and game-theory to model the condition to stop searching for better designs based on the designers' belief of their opponents' performance. We use data collected from a function optimization game developed previously in the literature to illustrate the application of the proposed method.

The results for that application indicate that a long short-term memory network can predict the next design decisions and corresponding outcomes based on their past decisions. The results further suggest that the designers show a strong tendency to overestimate their opponents' performance, leading them to spend more on searching for better design than they would have, had they known their opponents' actual performance. Our result is further supported and explained by previous empirical studies in social cognition. For example, a study in psychology performed a controlled experiment and observed that people were inclined to underestimate how good their partners were but to overesti-

mate their opponents [41].

The application we choose is abstract enough to be domain independent. The mathematical framework we present is generalizable to the design problems that possess the characteristics of C1-C4 under the assumptions of A1-A2.

8.1 Take-aways

This study further advances the prior research on understanding human sequential decision-making. The results complement the normative models in the existing literature by [24, 27]. Similar conclusions regarding exploration and exploitation behaviors of participants and the participants' belief for a strong opponent further validates the model. Learning these behaviors from the input data automatically without explicitly building them into the model is a key advantage of the proposed approach over the normative models. This paper presents implications of such behaviors on prediction outcomes.

The present results suggest some practical implications for the development of AI-based decision support systems. The difference between predicted participant behaviors from game theory (based on the assumption of rationality) and the actual behaviors highlights a need for influencing human beliefs regarding opponent performance. A forward-looking decision support agent might provide a more objective view of the opponent and guide humans to make better informed decisions. Additionally, certain mechanisms or intervention (e.g., fake opponents) can be developed to artificially influence participants' behaviors towards the direction where system designer desires.

Further, the results from our application show that the model generally learns behaviors that result in design improvements from the crowd since a significant majority of data points reflect such behaviors. In other applications, if the goal is to also learn low-performing design behaviors, a separate model trained with such decisions might be necessary.

8.2 Limitations and Future Work

There are multiple ways to study design competition. The present set of results represent a non-real-time competition as opposed to real-time competition where there is a continuous information sharing between opponents. It will be interesting to study how the participants' behavior, particularly their perception of the opponents' performance, vary when there is a real-time information sharing. The results are also obtained from a one-on-one competition whereas a team-based competition may yield different competition dynamics. Further, a team-based competition could be studied to analyze both inter-team and intra-team competition.

The present study uses LSTM networks to predict human design decisions and Delta method to estimate the uncertainty in the network models. A comparison of alternative machine learning approaches such as deep neural networks to predict sequential design decisions and alternative uncertainty prediction methods such as Bayesian approaches is left to a future study. These methods could offer varying lev-

els of success and computational efficiency in different design problem contexts. The approach presented in this paper is general enough for testing these methods without making any changes in the game-theoretic construct when investigating 1-1 competitions.

The implications of our results for human-AI collaboration need verification with further controlled experiments. It is interesting to study how human factors play a role in collaboration with a rational decision support agent. In addition, the conclusions we made based on the abstract function optimization game also need empirical validation with more concrete design problems.

Acknowledgements

The authors greatly appreciate Dr. Jitesh Panchal for sharing the data collected from the function optimization game. The authors also gratefully acknowledge the financial support from the U.S. National Science Foundation (NSF) through grants CMMI-1842588. Any opinions, findings, and conclusions or recommendations expressed in this publication, however, are those of the authors and do not necessarily reflect the views of NSF.

References

- [1] Walden, D. D., Roedler, G. J., Forsberg, K. J., Hamelin, R. D., and M., S. T., 2015. *INCOSE Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities*, 4 ed. Wiley, Hoboken, NJ.
- [2] VDI, 1993. 2221: Systematic approach to the design of technical systems and products. Standard DINGCD22, Deutsches Institut für Normung, Düsseldorf, Germany.
- [3] Shiau, C.-S. N., and Michalek, J. J., 2009. “Should designers worry about market systems?”. *Journal of Mechanical Design*, **131**(1), p. 011011.
- [4] Tversky, A., and Kahneman, D., 1974. “Judgment under uncertainty: Heuristics and biases”. *Science*, **185**(4157), pp. 1124–1131.
- [5] Jacowitz, K. E., and Kahneman, D., 1995. “Measures of anchoring in estimation tasks”. *Personality and Social Psychology Bulletin*, **21**(11), pp. 1161–1166.
- [6] Tversky, A., and Kahneman, D., 1973. “Availability: A heuristic for judging frequency and probability”. *Cognitive psychology*, **5**(2), pp. 207–232.
- [7] Yassine, A., 2004. “An introduction to modeling and analyzing complex product development processes using the design structure matrix (dsm) method”. *Quaderni di Manage. (Italian Manage. Rev.)*, **9**, pp. 71–88.
- [8] Yassine, A. A., Whitney, D. E., and Zambito, T., 2001. “Assessment of rework probabilities for simulating product development processes using the design structure matrix (dsm)”. In ASME International Design Engineering Technical Conferences, no. DETC2001/DTM-21693, American Society of Mechanical Engineers.
- [9] Browning, T. R., 2001. “Applying the design structure matrix to system decomposition and integration problems: a review and new directions”. *IEEE Transactions on Engineering Management*, **48**(3), pp. 292–306.
- [10] Unal, M., Miller, S. W., Chhabra, J. P., Warn, G. P., Yukish, M. A., and Simpson, T. W., 2017. “A sequential decision process for the system-level design of structural frames”. *Structural and Multidisciplinary Optimization*, **56**(5), pp. 991–1011.
- [11] Miller, S. W., Yukish, M. A., and Simpson, T. W., 2018. “Design as a sequential decision process: A method for reducing design set space using models to bound objectives”. *Structural and Multidisciplinary Optimization*, **57**(1), pp. 305–324.
- [12] Griffin, S., Welton, N. J., and Claxton, K., 2010. “Exploring the research decision space: the expected value of information for sequential research designs”. *Medical Decision Making*, **30**(2), pp. 155–162.
- [13] Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N., 2016. “Taking the human out of the loop: A review of bayesian optimization”. *Proceedings of the IEEE*, **104**(1), pp. 148–175.
- [14] Meier, C., Yassine, A. A., and Browning, T. R., 2007. “Design process sequencing with competent genetic algorithms”. *Journal of Mechanical Design*, **129**(6), pp. 566–585.
- [15] Duff, M. O., and Barto, A., 2002. “Optimal learning: Computational procedures for bayes-adaptive markov decision processes”. PhD thesis, University of Massachusetts at Amherst.
- [16] Chaudhari, A. M., Bilionis, I., and Panchal, J. H., 2020. “Descriptive models of sequential decisions in engineering design: An experimental study”. *Journal of Mechanical Design*, **142**, p. 081704.
- [17] McComb, C., Cagan, J., and Kotovsky, K., 2017. “Capturing human sequence-learning abilities in configuration design tasks through markov chains”. *Journal of Mechanical Design*, **139**(9).
- [18] McComb, C., Cagan, J., and Kotovsky, K., 2017. “Mining process heuristics from designer action data via hidden markov models”. *Journal of Mechanical Design*, **139**(11), p. 111412.
- [19] Gero, J. S., Kan, J. W., and Pourmohamadi, M., 2011. “Analysing design protocols: Development of methods and tools”. In *Research into Design - Supporting Sustainable Product Development.*, A. Chakrabarti, ed. Research Publishing, Bangalore, India, pp. 3–10.
- [20] Rahman, M. H., Schimpf, C., Xie, C., and Sha, Z., 2019. “A computer-aided design based research platform for design thinking studies”. *Journal of Mechanical Design*, **141**(12).
- [21] McComb, C., Cagan, J., and Kotovsky, K., 2015. “Lifting the veil: Drawing insights about design teams from a cognitively-inspired computational model”. *Design Studies*, **40**, pp. 119–142.
- [22] Chaudhari, A. M., and Panchal, J. H., 2019. “An experimental study of human decisions in sequential information acquisition in design: Impact of cost and task complexity”. In *Research into Design for a Connected*

- World*, Vol. 134. Springer, pp. 321–332.
- [23] Sexton, T., and Ren, M. Y., 2017. “Learning an optimization algorithm through human design iterations”. *Journal of Mechanical Design*, **139**(10), p. 101404.
- [24] Panchal, J. H., Sha, Z., and Kannan, K. N., 2017. “Understanding design decisions under competition using games with information acquisition and a behavioral experiment”. *Journal of Mechanical Design*, **139**(9), p. 091402.
- [25] Raina, A., McComb, C., and Cagan, J., 2019. “Learning to design from humans: Imitating human designers through deep learning”. *Journal of Mechanical Design*, **141**(11).
- [26] Rahman, M. H., Xie, C., and Sha, Z., 2019. “A deep learning based approach to predict sequential design decisions”. In International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, Vol. 59179, American Society of Mechanical Engineers, p. V001T02A029.
- [27] Sha, Z., Kannan, K. N., and Panchal, J. H., 2015. “Behavioral experimentation and game theory in engineering systems design”. *Journal of Mechanical Design*, **137**(5), p. 051405.
- [28] Shiau, C.-S. N., and Michalek, J. J., 2009. “Optimal product design under price competition”. *Journal of Mechanical Design*, **131**, p. 071003.
- [29] Frischknecht, B. D., Whitefoot, K., and Papalambros, P. Y., 2010. “On the suitability of econometric demand models in design for market systems”. *Journal of Mechanical Design*, **132**(12), p. 121007.
- [30] Gurnani, A., and Lewis, K., 2008. “Collaborative, decentralized engineering design at the edge of rationality”. *Journal of Mechanical Design*, **130**(12).
- [31] Skaperdas, S., 1996. “Contest success functions”. *Economic theory*, **7**(2), pp. 283–290.
- [32] Hopfield, J. J., 1982. “Neural networks and physical systems with emergent collective computational abilities”. *Proceedings of the National Academy of Sciences*, **79**(8), pp. 2554–2558.
- [33] Rumelhart, D. E., Hinton, G. E., and Williams, R. J., 1986. “Learning representations by back-propagating errors”. *Nature*, **323**(6088), pp. 533–536.
- [34] Hochreiter, S., and Schmidhuber, J., 1997. “Long short-term memory”. *Neural Computation*, **9**(8), pp. 1735–1780.
- [35] Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., and Schmidhuber, J., 2016. “LSTM: A search space odyssey”. *IEEE transactions on neural networks and learning systems*, **28**(10), pp. 2222–2232.
- [36] Olah, C., 2015. Understanding LSTM networks. <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>.
- [37] De Veaux, R. D., Schumi, J., Schweinsberg, J., and Ungar, L. H., 1998. “Prediction intervals for neural networks via nonlinear regression”. *Technometrics*, **40**(4), pp. 273–282.
- [38] Baldwin, C. Y., and Clark, K. B., 2000. *Design Rules*, Vol. 1: The Power of Modularity. The MIT Press, Cambridge, MA.
- [39] Papalambros, P. Y., and Wilde, D. J., 2017. *Principles of optimal design: modeling and computation*, 3 ed. Cambridge University Press, Cambridge, UK.
- [40] Khosravi, A., Nahavandi, S., Creighton, D., and Atiya, A. F., 2011. “Comprehensive review of neural network-based prediction intervals and new advances”. *IEEE Transactions on neural networks*, **22**(9), pp. 1341–1356.
- [41] Chan, E., and Ybarra, O., 2002. “Interaction goals and social information processing: Underestimating one’s partners but overestimating one’s opponents”. *Social Cognition*, **20**, pp. 409–439.

List of Figures

1	Overview of an LSTM network architecture (adapted from [36])	3
2	Overview of the Learning and Prediction Process	4
3	The function optimization game - Minimization of an unknown function	5
4	Prediction intervals for the next design sample and function value with one standard error away from the mean	8
5	Change in the function value (i.e., the difference between the current and next function values) with the number of samples the participants take	8
6	Prediction results from the game theoretic model for the participants' belief of their opponent's best score	9
7	Prediction results from the game theoretic model for when the participants will stop sampling. The size of the circles represents the frequency of occurrence.	9

List of Tables

1 Summary of descriptive statistics from the original dataset 7