**Proceedings of the ASME 2017 International Design Engineering Technical Conferences &
Computers and Information in Engineering Conference
IDETC/CIE 2017
August 6-9, 2017, Cleveland, USA**

DETC2017-68366

DESIGN PREFERENCE PREDICTION WITH DATA PRIVACY SAFEGUARDS: A PRELIMINARY STUDY

Alexander Burnap[*]
Mechanical Engineering
University of Michigan
Ann Arbor, MI 48109
Email: aburnap@umich.edu

Panos Y. Papalambros
Mechanical Engineering
University of Michigan
Ann Arbor, MI 48109
Email: pyp@umich.edu

ABSTRACT

Design preference models are used widely in product planning and design development. Their prediction accuracy requires large amounts of personal user data including purchase and other personal choice records. With increased Internet and smart device use, sources of personal data are becoming more varied and their capture more ubiquitous. This situation leads to questioning whether there is a trade off between improving products and compromising individual user privacy. To advance this conversation, we analyze how privacy safeguards may affect design preference modeling. We conduct an experiment using real user data to study the performance of design preference models under different levels of privacy. Results indicate there is a tradeoff between accuracy and privacy. However, with enough data, models with privacy safeguards can still be sufficiently accurate to answer population-level design questions.

1 INTRODUCTION

User data is commonly employed to help designers and marketers understand user needs and corresponding design requirements quantitatively. Common approaches include individual-level design preference modeling [1, 2] and discrete choice modeling [3], as well as market-level segmentation for product planning [4]. These approaches historically were associated with single-source supermarket scanner data in the early 1980's [5]; however, modern sources of user and design data are much more varied. The advent of the Internet has brought a shift toward online information seeking and product purchase, in which online browsing behavior is recorded across distinct websites using third-party tracking networks [6, 7]. Similarly, mobile applications on smart devices persistently record the spatial and temporal details of daily habits of individuals [8–10].

The engineering design community has recognized these trends as new opportunities to improve the design process; for example, product usage data for design knowledge and analytics [11], social media data for product demand and market adoption modeling [12], large-scale purchase data for design preference modeling [13], crowdsourcing and gamification to obtain user data [14–16], and design idea filtering among online collaborating designers [17]. There is a recognition that these trends are only increasing, as the Internet of Things, fueled by pervasive networking and sensor miniaturization, promises to connect and capture data from the estimated 50 billion smart devices, 85% of which are currently unconnected [18]. With this trend, user and design data collected from smart and connected hardware designs (e.g., biometric wearables [19], humanitarian designs [20], household appliances [21], networked mobility, and smart city and structures [22]), will enable tracking of users in our physical world at the same depth that is being recorded in the Internet's virtual world, further improving our ability to understand user needs and behaviors, and develop subsequent designs.

There are an estimated 4000 companies engaged in user data aggregation and brokering [23]. Moving into this era of ubiquitous data recording raises ethical, and possibly legal, questions regarding tradeoffs between potential design improvement and

[*] Address all correspondence to this author.

potential consequences of compromised user data privacy. Previous works have detailed consequences of compromised data privacy, including highly-publicized cases such as sensitive medical data leaks [24], recommendation data [25], and connections between participant names and human genome data [26]. Even when a dataset is not explicitly compromised, privacy leaks may still occur; for example, normal operating behavior of advertisers can be used to pinpoint ∼75% accuracy of gender, as well as parental status, age group, income, political affiliation, and religious affiliation [9]. Regardless of the mechanism of compromised privacy, these numbers may be troubling given that 87% of U.S. citizens are identifiable solely with zip code, gender, and date of birth. [27].

Academic research has begun to address ethical ramifications of compromised user data privacy [28–31]. Research in the computer science community on algorithmic privacy safeguards with theoretical guarantees includes approaches such as k-anonymity [32], fully homomorphic encryption, and differential privacy [33, 34].

In this paper, we follow a similar line of research focusing on data-driven design. We conducted an experiment to assess the effect of safeguarding privacy while constructing design preference models for product planning. In particular, we adopt findings from research in local differential privacy [34–36] and inject Laplacian noise into the statistical estimation of a common family of design preference models used by engineering design researchers and practitioners.

Results indicate there is indeed a tradeoff between protecting individual privacy and population-level design preference prediction quality, leading to decreased performance in defining regions of the design space where there may be product opportunities. We quantify this tradeoff for subsequent informed decision making. The contributions offered in this paper are as follows:

> (1) Differential privacy to safeguard individual user data in the domain of data-driven engineering design is a viable concept, as shown via design preference visualization for product planning.
>
> (2) There exists a tradeoff between the accuracy of design preference models and the degree to which individual privacy is safeguarded; moreover, we show that data size affects this tradeoff.

The intent here is not to optimize or otherwise suggest a desired tradeoff between privacy and population-level design preference prediction (whether it even be zero privacy), nor to contribute to the privacy-preserving methods themselves.

The rest of the paper is structured as follows: Section 2 discusses related work in design preference modeling and theoretical approaches to data privacy. Section 3 formulates a user's design preference function satisfying differential privacy. Section 4 conducts an experiment using real design preference data. Section 5 applies these experimental results to a common usage of design data for product positioning and planning. Section 6 discusses implications and opportunities for future work. Section 7 offers conclusions.

2 RELATED WORK

Data privacy has been a serious consideration for many research domains long before engineering and product design began using user data. Data anonymization has been standard for many governmental census measurements, and social science research includes protocols in handling "sensitive" survey questions that may result in evasive biases [37]. In this work, we focus on differential privacy as applied to design preference modeling.

2.1 Design Preference Modeling

Use of quantitative models that aim to capture user preferences, purchases, or more generally, decision-making behavior over a set of designs, is well established in the design and marketing communities [1,4,5]. These models, such as conjoint analysis and discrete choice analysis, are often built off expected utility theory, with implicit assumptions on rational decision-making behavior.

Common usage cases of these models include demand modeling and product pricing [38], product positioning [4], industrial and aesthetic design [39, 40], product brand recognition [16], user needs finding [15], and engineering design optimization. Recent work in these areas has improved on earlier modeling limitations in three major ways:

First, there is recognition that assumptions on rationality and on comparative decision-making are often not valid for many types of products or design scenarios. Major departures from these initial assumptions include relaxation of rationality, including bounded rationality [41] from psychological and behavioral findings [42], as well as consideration models of design preferences [43, 44]. Preference inconsistencies also are a major direction of study, with many findings showing the same user giving inconsistent preference responses with temporal or other experimental changes [45–47].

Second, there is recognition that the linear formulations of classical models do not accurately capture human decision-making behavior. Nonlinear extensions that account for dependencies among many variables can significantly improve the fidelity of design preference models [13], often with the tradeoff of interpretability [48]. There is, however, work aiming to reconcile these two areas, including local linearization of nonlinear models [49, 50].

Third, there is recognition that previous models do not adequately capture the diversity of user preferences in a population over the set of designs due to preference inhomogeneity, and that models that account for heterogeneous users often make overly rigid parametric assumptions on a population-level preference distribution [51]. Models that allow more flexible parametric models [13, 52], or nonparametric approaches [16] have shown promise in better capture of diverse design preferences.

2.2 Differential Privacy

Differential privacy is a formal algorithmic definition capturing the idea that adding or removing a single datum in a dataset results in statistically indistinguishable changes in the the output distribution of a random algorithm [34]. The level of how statistically close this output change is bounded by a user-defined privacy parameter [33]. An important note is that differential privacy does not protect users from the results of an analysis at a population level, just that it protects users from personal identification. In fact, in the context of engineering design, the general goal is to understand population-level design preferences to create "better" design alternatives for individual users to choose from.

Differential privacy is not an algorithm itself, but a definition that a random algorithm must satisfy. There are, accordingly, a number of "mechanisms" that may be used to satisfy this definition. A common example is the Laplace mechanism, used in this work, in which a random variable taking a Laplacian density is used to inject randomness at some stage between the data source and a result by an algorithm on that data source. In the domain of engineering design, examples of such algorithms include design preference models and social media sentiment analysis algorithms [12].

The stage in which this randomness is introduced determines whether the algorithm satisfies local or global differential privacy [35, 53]. Global differential privacy ensures that the outcomes of randomized algorithm are private. Local differential privacy takes that a step further, and introduces randomness at the level of the individual user, such that their data remains private even from the analysts of the data. In this work, we adopt the local model of differential privacy, as this approach protects users and may be implementable at the hardware layer of a mechanical design.

3 PROBLEM FORMULATION

We conceptually connect users to their preferred designs using a predictive design preference model. This design preference model is next upgraded to satisfy differential privacy, thus safeguarding individual user data.

3.1 Design Preference Modeling

Our goal is to understand the diverse preferences of users $\mathbf{x}_c \in \mathscr{X}_c \subseteq \mathbb{R}^{M_c}$ over D existing product designs $\mathbf{x}_d \in \mathscr{X}_d \subseteq \mathbb{R}^{M_d}$, such that these preferences may be encoded in a mathematical function that can inform design decisions of future products or systems. We capture this user-design preference relation using a heterogeneous linear utility model $U : \mathscr{X}_c \times \mathscr{X}_d \to \mathbb{R}$ in which we assume each user c has their own design preference $\mathbf{w}_c$.

$$U(\mathbf{x}_c, \mathbf{x}_d) = \mathbf{w}_c^T(\mathbf{x}_c)\mathbf{x}_d \tag{1}$$

Note that $\mathbf{w}_c$ is a function of the user, and exists in the space of designs. The design preference for user c is their $\mathbf{w}_c$, and the inner product from this vector in the design space corresponds to how close existing designs are to the design preferences of the user. We next assume a user's design preference is a linear function of their user variables,

$$\mathbf{w}_c^T(\mathbf{x}_c) = [\mathbf{x}_c, 1]^T\Omega \tag{2}$$

such that $\Omega \in \mathbb{R}^{[M_c+1]\times[M_d+1]}$ and the utility $U(\mathbf{x}_c, \mathbf{x}_d) = [\mathbf{x}_c, 1]^T\Omega[\mathbf{x}_d, 1]$ has a bilinear form. Note that we concatenate '1' onto each user and design datum to account for main effects independent of interactions between users and designs.

We further assume a multinomial logit link function that transforms a user's utility values over designs to probabilistic design preference relations,

$$f_{pref}\left(\mathbf{x}_c^{(i)}, \mathbf{x}_d^{(j)}, \Omega\right) = \frac{e^{U(\mathbf{x}_c^{(i)}, \mathbf{x}_d^{(j)})}}{\sum_{j=1}^{K} e^{U(\mathbf{x}_c^{(i)}, \mathbf{x}_d^{(j)})}} \tag{3}$$

This model is symmetric across users and designs, and may be used to find regions of the design space that contain product design opportunities.

3.1.1 Preference Estimation

To estimate values of this user-design preference relation, we use previous data on users and designs and adopt the framework of empirical risk minimization. In particular, we collect a dataset $\mathscr{D} = \{(\mathbf{x}_c^{(i)}, \mathbf{x}_d^{(i)}, y^{(i)}\}_{i=1}^N$ made up of N data on users, designs, and an indicator variable $y \in \{1, \dots, K\}$ connecting a user with their preferred design. We then choose a loss function $l : \mathscr{Y} \times \mathscr{Y} \to \mathbb{R}$, and minimize this loss under the collected data distribution,

$$\hat{\Omega} = \arg\min_{\Omega} \frac{1}{N}\left[l(f_{pref}\left(\mathbf{x}_c^{(i)}, \mathbf{x}_d^{(i)}, \Omega\right), y^{(i)})\right] + \frac{\lambda}{2}||\Omega||_2^2 \tag{4}$$

where $||\cdot||_2$ is the L^2 norm and used to prevent overfitting as controlled by parameter λ.

We assume l as the softmax/log-loss $\sum_i \sum_j y^{(i)(j)} log(f_{pref})$, and note this this strictly convex optimization problem is necessary for later results on privacy safeguards.

3.2 Differential Privacy

We work with the definition of Dwork et. al [33]: Let A be a randomized algorithm for obtaining a result Ω using a dataset $\mathscr{D}$. We say A satisfies ε-differential privacy if for all possible Ω, and all possible datasets $(\mathscr{D}, \mathscr{D}')$ that differ by just one datum, the following relation holds:

$$\frac{P(A(\mathscr{D}) = \Omega)}{P(A(\mathscr{D}') = \Omega)} \leq e^{\varepsilon} \tag{5}$$

where P is defined over the probability space of possible outputs made by A, and ε is a privacy parameter, a positive value bounding the worst case difference between all possible values in the output distribution of A. In other words, a smaller ε results in more user privacy.

We note that this point, simply solving Equation (4) is deterministic and does not satisfy the definition of differential privacy, while Equation (5) treats Ω as a random variable. This leads to the necessity of introducing stochasticity to the algorithm that estimates values of the user-design preference relation. As described in Section 2.2, a number of mechanisms may be used to add randomness.

3.2.1 Local Model using Laplacian Noise

Adding noise at the user's end (e.g., directly to their preference survey, purchase records, or smartphone data), results in a "local model" of differential privacy [34]. We simulate this by adding noise to an online stochastic gradient descent over each user datum as given in [36,54,55].

Specifically, we can take obtained values of users $\mathbf{x}_c$ and designs $\mathbf{x}_d$, as well as their bilinear interactions by rewriting $U(\mathbf{x}_c, \mathbf{x}_d) = [\mathbf{x}_c, 1]^T \Omega [\mathbf{x}_d, 1]$ as $U = \left[\text{vec}\left(\mathbf{x}_c \otimes \mathbf{x}_d\right)^T, \mathbf{x}_c^T, \mathbf{x}_d^T\right] \omega$, where $\omega = \text{vec}(\Omega)$, $\text{vec}(\cdot)$ vectorizes matrices, and $\otimes$ is the outer product operator. Then, the online optimization routine updates as,

$$\omega^{(t+1)} = \omega^{(t)} - \eta^{(t)}[\nabla_{\omega} l(\omega^{(t)}, \mathbf{x}_c^{(i)}, \mathbf{x}_d^{(i)}, y^{(i)}) + \lambda \omega^{(t)} + Z(\varepsilon, \lambda)] \tag{6}$$

for the t iteration, where $\eta^{(t)}$ is the learning rate, and Z is a Laplacian random vector with elements sampled from,

$$z_j \sim p(z) = e^{-\frac{\varepsilon}{2}|z|} \tag{7}$$

in which ε controls the scale of the noise variable, and thus the level of individual user privacy.

4 EXPERIMENT

We conduct an experiment to test the effect of incorporating user privacy safeguards through local differential privacy on the prediction accuracy of the design preference model described in Equation (3). In particular, we look at two relationships: (1) the effect of the individual-level privacy parameter ε on population-level design preference prediction accuracy, and (2) the effect of data size N on a design preference prediction algorithm satisfying differential privacy. The data used is real user and design data from purchases of Model Year 2014 automotive vehicles in the United States, and the design preference task is to understand which users would prefer which brand.

4.1 Data

Our dataset consists of 50,000 real purchase data of Model Year 2014 (MY2014) vehicle designs, in which each datum consists of a user $\mathbf{x}_c$ and their corresponding purchased design $\mathbf{x}_c$. Each user datum exists in a 282-dimensional space, and each design datum exists in a 285-dimensional space. Variables were split into user and design data according to the author's domain expertise in the vehicle dataset. For both users and designs, real, binary, and categorical variables were partitioned and appropriately one-hot encoded. Similarly, the corresponding preference for the automobile brand was encoded in y for each pair $(\mathbf{x}_c, \mathbf{x}_d)$.

User data includes variables such as demographics (e.g., age, gender, income) as well as personal viewpoints and values important to the user. Design data includes variables such as vehicle characteristics (e.g., rounded, flashy), and technical characteristics (e.g., engine, transmission, fuel type). More detailed description is restricted due to the proprietary nature of this data. A total of 35 brands (e.g., Audi, Lexus, Cadillac) were included in the data.

4.2 Procedure

The experimental procedure includes preprocessing, parameter estimation of the user-design preference matrix Ω, and assessment of preference prediction accuracy on a held out portion of the data.

1. **Preprocessing** - The full data was split into training and testing sets at a ratio of 80% to 20%, and we normalized the training data (i.e., $\frac{x-\mu}{\sigma}$), where μ and σ are the empirical mean and standard deviation of the training data. The testing data is then normalized using μ and σ.
2. **Parameter Estimation** - The L^2 regularization hyperparameter is set at $\lambda = 0.1$ and fixed at this value for all design preference models. Similarly, the learning rate hyperparameter is set at $\eta^{(t)} = \frac{1}{\lambda t}$ as defined in [56]. We note that cross-validation of these hyperparameters are at risk of privacy leakage. Stochastic gradient descent with Laplacian noise, as described in Equation (6), is used to optimize Equation (4) with 100 epochs (passes) through the training data.
3. **Design Preference Prediction Accuracy** - The prediction accuracy is assessed on the held-out testing data using the realization of Ω from the parameter estimation step. We obtain both the "Top-1" and "Top-5" accuracy, corresponding to whether the design preference model is able to correctly predict the true preferred brand y.

4.3 Results

We display results for two analyses: (1) Effect of individual-level privacy parameter ε on population-level design preference prediction accuracy, and (2) the effect of data size N on a design preference prediction algorithm satisfying differential privacy.

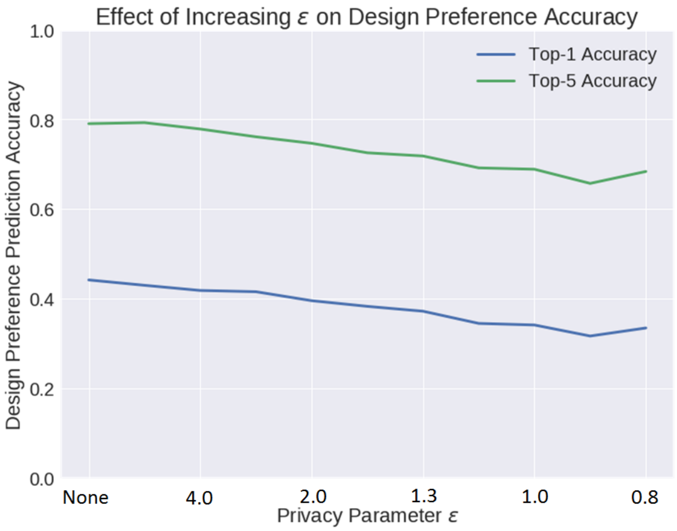

FIGURE 1. **EFFECT OF INDIVIDUAL-LEVEL PRIVACY PARAMETER ε ON POPULATION-LEVEL DESIGN PREFERENCE PREDICTION ACCURACY. SMALLER VALUES OF ε RESULT IN MORE PRIVACY.**

The effect of privacy parameter ε on the accuracy of the design preference model is given in Figure 1. The prediction accuracy of the baseline case, a design preference model with no privacy safeguards, is shown on the left-hand side as denoted by “None.” As the privacy parameter ε decreases, and thus the level of individual user privacy increases, the prediction accuracy of the design preference model decreases. This trend is given for both the “Top-1” and “Top-5” accuracies, corresponding to whether the design preference model correctly predicting a user’s preferred brand in its top guess or within the top five. Note that there are 35 possible brands to predict, leaving a $\sim 3\%$ chance of randomly guessing the correct brand.

The effect of data size N on a design preference prediction algorithm satisfying differential privacy is shown in Figure 2. In the low data regime, both the design privacy model without privacy safeguard and the design preference model with privacy safeguards ($\varepsilon = 0.8$) have low prediction accuracy. As the datasize increases, we see both design preference models improve in accuracy; however, the baseline model tends to improve faster as shown by the gap between the prediction accuracy of the two design preference models.

5 APPLICATION TO DESIGN

While quantitative design preference models have been used for a number of design applications, we focus in this work on design gap visualization for product opportunity planning. We use estimated design preference model parameters Ω from the model described in Equation (3), obtained for both a baseline case without privacy safeguards, and a privacy safeguarded case with a design preference model satisfying differential privacy.

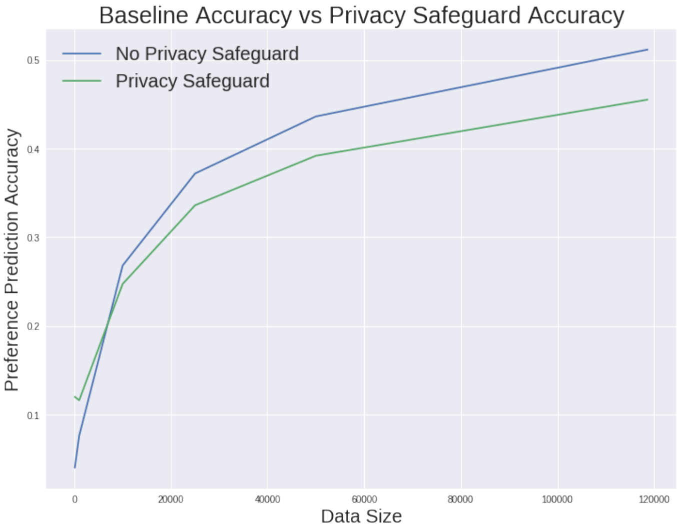

FIGURE 2. **EFFECT OF DATA SIZE N ON A DESIGN PREFERENCE PREDICTION ALGORITHM SATISFYING DIFFERENTIAL PRIVACY AS WELL AS A BASELINE PREFERENCE PREDICTION ALGORITHM. THE DESIGN PREFERENCE MODEL WITH PRIVACY SAFEGUARD SATISFIES DIFFERENTIAL PRIVACY WITH PARAMETER $\varepsilon = 0.8$.**

5.1 Visualizing Design Preferences

We visualize the design preferences of users in a 2-dimensional (2D) representation for design preference models with and without a privacy safeguard. As design preferences exist in a high-dimensional space, we aim to preserve distances between points in the high-dimensional space as best as possible in the 2-dimensional (2D) space, acknowledging that information loss is necessary.

Specifically, we aim to preserve distances of the user’s design preference $\mathbf{w}_c$. This vector is split into the preference interaction coefficients $\hat{\omega} = [\omega_b, \omega_c, \omega_d]$, corresponding to the interaction effects between users and designs, main effects for users, and main effects for designs, respectively.

$$\mathbf{w}_c = [\mathbf{x}_c, 1]^T \, \hat{\Omega} = \left[\mathbf{x}_c^T \text{vec}^{-1}(\omega_b), \mathbf{x}_c^T \omega_c, \right] \qquad (8)$$

in which $\text{vec}^{-1}(\cdot)$ is the inverse of the earlier defined vectorization operator, and $[\cdot, \cdot]$ concatenates two elements.

To obtain a 2D visual representation of the distribution of design preferences $\mathbf{w}_c^{(i)}$, or $\mathbf{W}_{c,2D}$ over all users, we minimize an objective function using t-distributed stochastic embedding [57]. This method minimizes the Kullback-Leibler divergence of Gaussian-distributed pairwise differences of design preferences in the space of $\mathbf{w}_c$, denoted p_{ij}, with Student-t-distributed pairwise differences of design preferences in the projected 2D space of $\mathbf{w}_{c,2D}$, denoted q_{ij}.

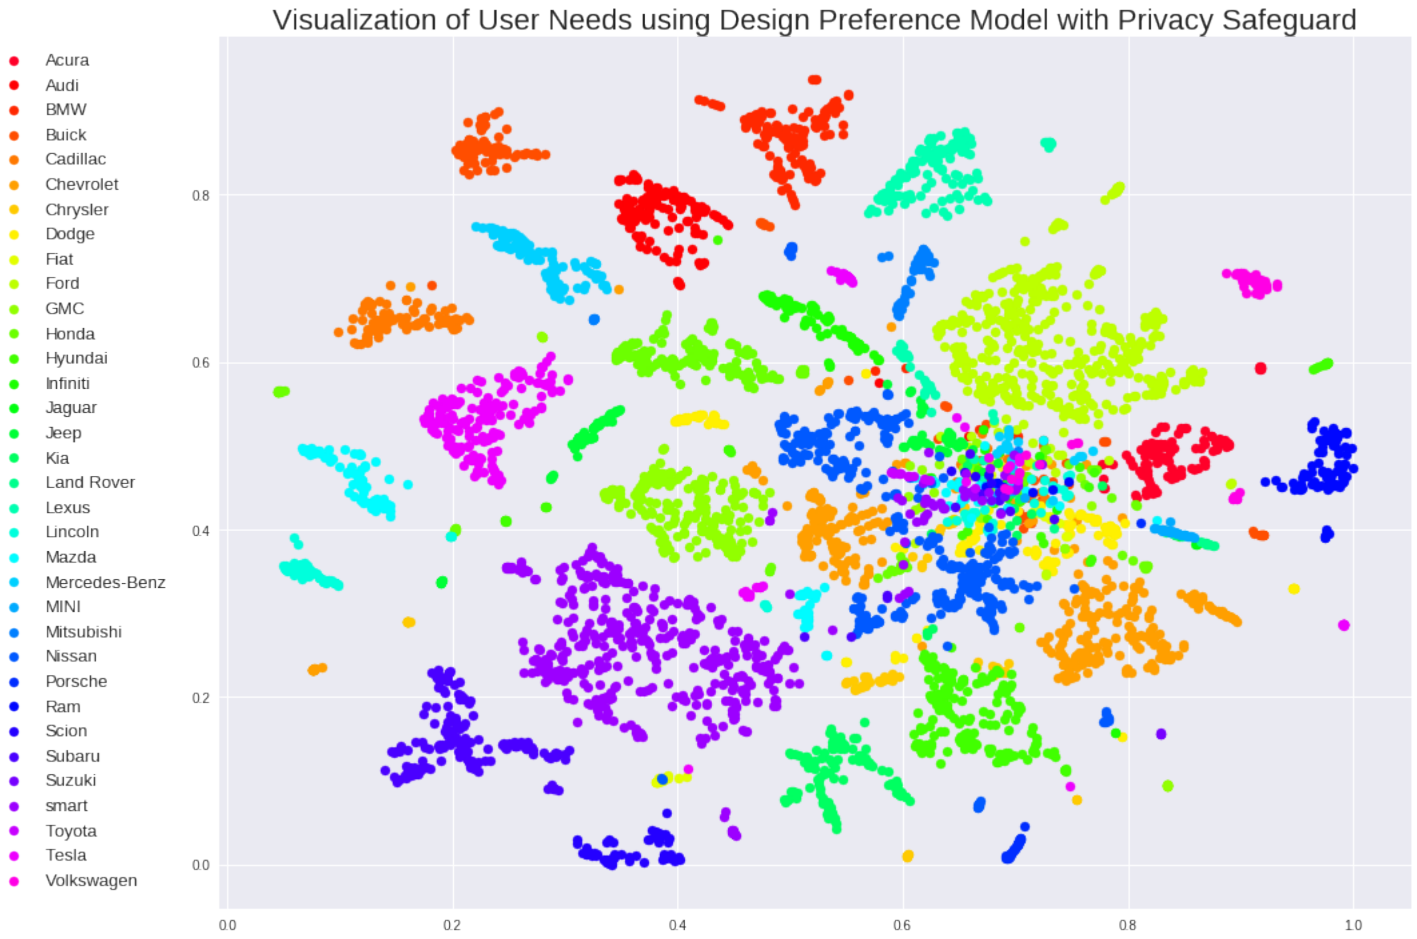

FIGURE 3. 2D VISUALIZATION OF USERS' DESIGN PREFERENCES USING DESIGN PREFERENCE MODEL SATISFYING DIFFERENTIAL PRIVACY.

$$\begin{aligned}
\mathbf{W}_{c,2D} &= \arg\min_{\mathbf{W}} \sum_{i,j=1,i\neq j}^{N} p_{ij} \log \frac{p_{ij}}{q_{ij}} \\
&\text{where,} \\
p_{ij} &= \frac{\exp(-\frac{1}{2\sigma^2}\left\|\mathbf{w}_c^{(i)}-\mathbf{w}_c^{(j)}\right\|_2^2)}{2N\sum_{i,k=1,i\neq k}^{N}\exp(-\frac{1}{2\sigma^2}\left\|\mathbf{w}_c^{(i)}-\mathbf{w}_c^{(k)}\right\|_2^2)} \\
&+ \frac{\exp(-\frac{1}{2\sigma^2}\left\|\mathbf{w}_c^{(j)}-\mathbf{w}_c^{(i)}\right\|_2^2)}{2N\sum_{j,k=1,j\neq k}^{N}\exp(-\frac{1}{2\sigma^2}\left\|\mathbf{w}_c^{(j)}-\mathbf{w}_c^{(k)}\right\|_2^2)} \qquad (9) \\
q_{ij} &= \frac{(1+\left\|\mathbf{w}_{c,2D}^{(i)}-\mathbf{w}_{c,2D}^{(j)}\right\|_2^2)^{-1}}{\sum_{k,l=1,l\neq k}^{N}(1+\left\|\mathbf{w}_{c,2D}^{(k)}-\mathbf{w}_{c,2D}^{(l)}\right\|_2^2)^{-1}}
\end{aligned}$$

Figure 3 shows the 2D visualization of design preferences over brands, allowing designers to understand user needs and discover product opportunities. Users, represented by points in this figure, are color coded according to the brand of vehicle they purchased. We note that these design preferences are defined over the space of designs (not brands) as given in Equation (8), thus this visualization gives strong evidence that we are clustering design preferences in general. We further note that this analysis is qualitative, and its purpose is to show that design preference models that safeguard user data privacy may still result in meaningful clustering of design preferences at the population level.

5.1.1 Intracluster Design Preference Dispersion

While the previous analysis was qualitative, we obtain a quantitative measure of the dispersion among various clusters of similar design preferences. Specifically, we normalize the obtained coordinates $\mathbf{w}_{c,2D}$ of the 2D visual representation to be between 0 and 1, such that design preference models may be meaningfully compared. Then, we find the mean and standard deviation of all pairwise distances between users in a given brand segment.

This dispersion measure d_k is calculated for both design preference clusters obtained using the baseline design preference

Design Preference Model	Average Dispersion	Standard Deviation
Baseline (No Privacy Safeguard)	0.144	0.0795
With Privacy Safeguard ($\varepsilon = 0.8$)	0.163	0.0759

TABLE 1. **INTRACLUSTER DESIGN PREFERENCE DISPERSION MEAN AND STANDARD DEVIATION FOR THE BASELINE DESIGN PREFERENCE MODEL AND THE DESIGN PREFERENCE MODEL WITH PRIVACY SAFEGUARDS.**

model and the design preference model that satisfies differential privacy.

$$d_k = \frac{1}{N_k} \sum_{i,j=1,i \neq j}^{N_k} \left\| \mathbf{w}_{c,2D}^{(i)} - \mathbf{w}_{c,2D}^{(j)} \right\|_2 \qquad (10)$$

where N_k is the number of users in the k^{th} design preference cluster.

Table 1 gives the values of the intracluster dispersion of design preferences. Smaller values indicate a design preference representation better captures clustering of design preferences assuming a correct clustering. Note that we have ground truth labels in this case for correct clustering according to purchased brand. This result indicates the baseline design preference model better captures the clustering of design preferences than the design preference model with privacy safeguards.

6 DISCUSSION

Since our goal is to determine the effect of including user privacy safeguards in building design preference models, we looked at three analyses. First, we analyzed how the preference parameter ε affects the accuracy of a differentially private version of the design preference model. Second, we analyzed the effect of increasing the data size on both a baseline design preference model, not satisfying differential privacy, and a version satisfying differential privacy. Third, we applied a differentially private design preference model to an application of product design, specifically, visualization of product brand positioning.

In general, these analyses show that including privacy safeguards decreases the accuracy of design preference models. The magnitude of this decrease depends on the level of individual user privacy, namely, there is a tradeoff between prediction quality and privacy safeguards.

The analysis also shows that population-level design prediction is still viable, see Figure 3. Given enough user and design data, and the correct family of design preference models, we can reach meaningful design decisions under privacy safeguards. Figure 2 shows that model accuracy under differential privacy increases with more data.

6.1 Limitations and Future Work

As discussed in Section 2, there are three general approaches to randomizing an algorithm to satisfy differential privacy–: at the objective, output, or data source itself. These approaches translate to whether a user uploads "raw data" to another party (in which case it may be safeguarded) or uploads data using a procedure that already has privacy safeguards.

In this work, we looked at the latter case. We simulated an individual user uploading data that already has privacy safeguards. Such implementation can be accomplished through software or firmware algorithms, or potentially though hardware-layer privacy safeguard algorithms. This latter direction may prove a valuable contribution to the overall user data privacy conversation from the standpoint of mechanical design.

The mechanism of differential privacy invoked here is a rather elementary choice. A significant body of work has advanced these methods, often focusing on specific classes of randomized algorithms. Recent differential privacy findings may be used in conjunction with more advanced design preference models than the model used in this work, e.g., combining differentially private singular value decomposition [58] with low-rank and sparse matrix design preference models [13]. Moreover, various clustering methods may prove useful for design tasks such as obtaining similar clusters of designs, users, and their design preferences [59].

7 CONCLUSION

Design preference models are used in a variety of product and engineering design tasks, ranging from applications of product planning to new product design and development. These models are statistically estimated on large databases of personal user data, which historically contained previous product purchases and customer survey data. This has raised the question on how to safeguard individual user privacy when making design decisions.

Using real user data we studied how design preference models satisfying differential privacy perform on a common design task—visualization of design preferences for product planning. Our results showed there is a tradeoff between the accuracy of design preference models and individual user privacy. At the same time, increasing the amount of data used to estimate the design preference model satisfying differential privacy is shown to still bring model accuracy to useful levels. This gives evidence that design preferences models with individual privacy safeguards can serve to answer population-level design questions.

Understanding the pros and cons of various user privacy safeguards in data-driven engineering design is important, as these approaches direct affect the efficiency of the design organization, particularly in a future with near-guaranteed ubiquity of data capture. The need for such understanding has been recognized at the highest levels of the U.S. government in the legal sphere [60] and in research and development [61].

ACKNOWLEDGMENTS

This work has been partially supported by the National Science Foundation under Grant No. CMMI-1266184, General Motors Corporation. This support is gratefully acknowledged.

REFERENCES

[1] Chen, W., Hoyle, C., and Wassenaar, H. J., 2013. *Decision-Based Design*. Springer London, London.

[2] Wassenaar, H. J., and Chen, W., 2001. "An approach to decision-based design". In ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers.

[3] Rossi, P. E., McCulloch, R. E., and Allenby, G. M., 1996. "The value of purchase history data in target marketing". *Marketing Science,* ***15***(4), pp. 321–340.

[4] Wedel, M., and Kamakura, W. A., 2012. *Market segmentation: Conceptual and methodological foundations*, Vol. 8. Springer Science & Business Media.

[5] Guadagni, P. M., and Little, J. D., 1983. "A logit model of brand choice calibrated on scanner data". *Marketing science,* ***2***(3), pp. 203–238.

[6] Englehardt, S., Reisman, D., Eubank, C., Zimmerman, P., Mayer, J., Narayanan, A., and Felten, E. W., 2015. "Cookies that give you away: The surveillance implications of web tracking". In Proceedings of the 24th International Conference on World Wide Web, ACM, pp. 289–299.

[7] Acar, G., Eubank, C., Englehardt, S., Juarez, M., Narayanan, A., and Diaz, C., 2014. "The web never forgets: Persistent tracking mechanisms in the wild". In Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, ACM, pp. 674–689.

[8] Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L. F., and Agarwal, Y., 2015. "Your Location has been Shared 5,398 Times!: A Field Study on Mobile App Privacy Nudging". In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, ACM, pp. 787–796.

[9] Meng, W., Ding, R., Chung, S. P., Han, S., and Lee, W., 2016. "The price of free: Privacy leakage in personalized mobile in-app ads". *Proc. of NDSS'16*.

[10] Enck, W., Gilbert, P., Han, S., Tendulkar, V., Chun, B.-G., Cox, L. P., Jung, J., McDaniel, P., and Sheth, A. N., 2014. "TaintDroid: an information-flow tracking system for realtime privacy monitoring on smartphones". *ACM Transactions on Computer Systems (TOCS),* ***32***(2), p. 5.

[11] Van Horn, D., Olewnik, A., and Lewis, K., 2012. "Design analytics: capturing, understanding, and meeting customer needs using big data". In ASME 2012 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. 863–875.

[12] Tuarob, S., and Tucker, C. S., 2013. "Fad or here to stay: Predicting product market adoption and longevity using large scale, social media data". In ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. V02BT02A012–V02BT02A012.

[13] Burnap, A., Pan, Y., Liu, Y., Ren, Y., Lee, H., Gonzalez, R., and Papalambros, P. Y., 2016. "Improving Design Preference Prediction Accuracy Using Feature Learning". *Journal of Mechanical Design,* ***138***(7), p. 071404.

[14] Ren, Y., Bayrak, A. E., and Papalambros, P. Y., 2016. "ecoracer: Game-based optimal electric vehicle design and driver control using human players". *Journal of Mechanical Design,* ***138***(6), p. 061407.

[15] Häggman, A., Tsai, G., Elsen, C., Honda, T., and Yang, M. C., 2015. "Connections between the design tool, design attributes, and user preferences in early stage design". *Journal of Mechanical Design,* ***137***(7), p. 071408.

[16] Burnap, A., Hartley, J., Pan, Y., Gonzalez, R., and Papalambros, P. Y., 2016. "Balancing design freedom and brand recognition in the evolution of automotive brand styling". *Design Science Journal,* ***2***(9).

[17] Ahmed, F., and Fuge, M., 2017. "Capturing Winning Ideas in Online Design Communities".

[18] Evans, D., 2011. "The internet of things: How the next evolution of the internet is changing everything". *CISCO white paper,* ***1***(2011), pp. 1–11.

[19] Ghosh, D. D., Kim, J., Olewnik, A., Lakshmanan, A., and Lewis, K. E., 2016. "Cyber-Empathic Design: A Data Driven Framework for Product Design". In ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers.

[20] Garrity, J., 2015. "Harnessing the Internet of Things for Global Development". *UN International Telecommunication Union*.

[21] Van Horn, D., and Lewis, K., 2015. "The use of analytics in the design of sociotechnical products". *Artificial Intelligence for Engineering Design, Analysis and Manufacturing,* ***29***(01), pp. 65–81.

[22] Whitefoot, K., and Donofrio, N., 2015. *Making Value for America: Embracing the Future of Manufacturing, Tech-*

nology, and Work. National Academies Press.

[23] Dixon, P., 2014. *What information do data brokers have on consumers, and how do they use it*. Senate Committee on Commerce, Science, and Transportation.

[24] Barth-Jones, D. C., 2012. "The're-identification'of Governor William Weld's medical information: a critical re-examination of health data identification risks and privacy protections, then and now".

[25] Narayanan, A., and Shmatikov, V., 2006. "How to break anonymity of the netflix prize dataset". *arXiv preprint cs/0610105*.

[26] Sweeney, L., Abu, A., and Winn, J., 2013. "Identifying participants in the personal genome project by name".

[27] Sweeney, L., 2000. Uniqueness of simple demographics in the US population. Tech. rep., Technical report, Carnegie Mellon University.

[28] Sollins, K., 2016. "Decomposing Data Privacy for Evaluation".

[29] Sadeh, N., Acquisti, A., Breaux, T. D., Cranor, L. F., McDonalda, A. M., Reidenbergb, J. R., Smith, N. A., Liu, F., Russellb, N. C., Schaub, F., and others, 2013. The usable privacy policy project. Tech. rep., Tech. report CMU-ISR-13-119, Carnegie Mellon University.

[30] Rainie, L., Kiesler, S., Kang, R., Madden, M., Duggan, M., Brown, S., and Dabbish, L., 2013. "Anonymity, privacy, and security online". *Pew Research Center,* ***5***.

[31] Martin, K. D., and Murphy, P. E., 2017. "The role of data privacy in marketing". *Journal of the Academy of Marketing Science*, pp. 1–21.

[32] Sweeney, L., 2002. "k-anonymity: A model for protecting privacy". *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems,* ***10***(05), pp. 557–570.

[33] Dwork, C., McSherry, F., Nissim, K., and Smith, A., 2006. "Calibrating noise to sensitivity in private data analysis". In Theory of Cryptography Conference, Springer, pp. 265–284.

[34] Dwork, C., and Roth, A., 2013. "The Algorithmic Foundations of Differential Privacy". *Foundations and Trends® in Theoretical Computer Science,* ***9***(3-4), pp. 211–407.

[35] Duchi, J. C., Jordan, M. I., and Wainwright, M. J., 2013. "Local privacy and statistical minimax rates". In Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on, IEEE, pp. 429–438.

[36] Song, S., Chaudhuri, K., and Sarwate, A. D., 2013. "Stochastic gradient descent with differentially private updates". In Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE, IEEE, pp. 245–248.

[37] Warner, S. L., 1965. "Randomized response: A survey technique for eliminating evasive answer bias". *Journal of the American Statistical Association,* ***60***(309), pp. 63–69.

[38] 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.

[39] Orsborn, S., Cagan, J., and Boatwright, P., 2009. "Quantifying Aesthetic Form Preference in a Utility Function". *Journal of Mechanical Design,* ***131***(6), p. 061001.

[40] Reid, T. N., Gonzalez, R. D., and Papalambros, P. Y., 2010. "Quantification of Perceived Environmental Friendliness for Vehicle Silhouette Design". *Journal of Mechanical Design,* ***132***(10), p. 101010.

[41] Gurnani, A., and Lewis, K., 2008. "Collaborative, decentralized engineering design at the edge of rationality". *Journal of Mechanical Design,* ***130***(12), p. 121101.

[42] Kahneman, D., 2003. "Maps of bounded rationality: Psychology for behavioral economics". *The American economic review,* ***93***(5), pp. 1449–1475. 02160.

[43] Morrow, W. R., Long, M., and MacDonald, E. F., 2012. "Consider-Then-Choose Models in Decision-Based Design Optimization". In ASME 2012 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. 405–422.

[44] Dzyabura, D., and Hauser, J. R., 2011. "Active Machine Learning for Consideration Heuristics". *Marketing Science,* ***30***(5), Sept., pp. 801–819.

[45] Thurston, D. L., 2001. "Real and Misconceived Limitations to Decision Based Design With Utility Analysis". *Journal of Mechanical Design,* ***123***(2), p. 176.

[46] MacDonald, E. F., Gonzalez, R., and Papalambros, P. Y., 2009. "Preference Inconsistency in Multidisciplinary Design Decision Making". *Journal of Mechanical Design,* ***131***(3), p. 031009.

[47] Bao, Q., El Ferik, S., Shaukat, M. M., and Yang, M. C., 2014. "An Investigation on the Inconsistency of Consumer Preferences: A Case Study of Residential Solar Panels". In Proceedings of the 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference.

[48] Lipton, Z. C., Kale, D. C., Elkan, C., Wetzell, R., Vikram, S., McAuley, J., Wetzell, R. C., Ji, Z., Narayaswamy, B., Wang, C.-I., and others, 2016. "The Mythos of Model Interpretability". *IEEE Spectrum*.

[49] Ribeiro, M. T., Singh, S., and Guestrin, C., 2016. "" Why Should I Trust You?": Explaining the Predictions of Any Classifier". *arXiv preprint arXiv:1602.04938*.

[50] Pan, Y., Burnap, A., Liu, Y., Lee, H., Gonzalez, R., and Papalambros, P., 2016. "A Quantitative Model for Identifying Regions of Design Visual Attraction and Application to Automobile Styling". In Proceedings of the 2016 Internation Design Conference.

[51] Lenk, P. J., DeSarbo, W. S., Green, P. E., and Young, M. R., 1996. "Hierarchical Bayes Conjoint Analysis: Recovery of Partworth Heterogeneity from Reduced Experimental De-

signs". *Marketing Science,* ***15***(2), pp. 173–191.

[52] Liu, L., and Dzyabura, D., 2016. "Capturing Multi-taste Preferences: A Machine Learning Approach".

[53] Sarwate, A. D., and Chaudhuri, K., 2013. "Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data". *IEEE signal processing magazine,* ***30***(5), pp. 86–94.

[54] Shokri, R., and Shmatikov, V., 2015. "Privacy-preserving deep learning". In Proceedings of the 22nd ACM SIGSAC conference on computer and communications security, ACM, pp. 1310–1321.

[55] Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L., 2016. "Deep learning with differential privacy". In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, ACM, pp. 308–318.

[56] Bottou, L., 2010. "Large-scale machine learning with stochastic gradient descent". In *Proceedings of COMPSTAT'2010*. Springer, pp. 177–186.

[57] Maaten, L. v. d., and Hinton, G., 2008. "Visualizing data using t-SNE". *Journal of Machine Learning Research,* ***9***(Nov), pp. 2579–2605.

[58] Chaudhuri, K., Sarwate, A. D., and Sinha, K., 2013. "A near-optimal algorithm for differentially-private principal components.". *Journal of Machine Learning Research,* ***14***(1), pp. 2905–2943.

[59] Park, M., Foulds, J., Chaudhuri, K., and Welling, M., 2016. "Private Topic Modeling". *arXiv preprint arXiv:1609.04120*.

[60] House, W., 2012. "Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy". *White House, Washington, DC*.

[61] House, W., 2016. National Privacy Research Strategy. Tech. rep., Networking and Information Technology Research and Development Program, National Science and Technology Council, June.