

DETC2018-85310

**EFFICIENT LOAD SAMPLING FOR WORST-CASE STRUCTURAL ANALYSIS
UNDER FORCE LOCATION UNCERTAINTY**

Yining Wang
Carnegie Mellon University

Erva Ulu
Carnegie Mellon University

Aarti Singh
Carnegie Mellon University

Levent Burak Kara*
Carnegie Mellon University

ABSTRACT

An important task in structural design is to quantify the structural performance of an object under the external forces it may experience during its use. The problem proves to be computationally very challenging as the external forces' contact locations and magnitudes may exhibit significant variations. We present an efficient analysis approach to determine the most critical force contact location in such problems with force location uncertainty. Given an input 3D model and regions on its boundary where arbitrary normal forces may make contact, our algorithm predicts the worst-case force configuration responsible for creating the highest stress within the object. Our approach uses a computationally tractable experimental design method to select number of sample force locations based on geometry only, without inspecting the stress response that requires computationally expensive finite-element analysis. Then, we construct a simple regression model on these samples and corresponding maximum stresses. Combined with a simple ranking based post-processing step, our method provides a practical solution to worst-case structural analysis problem. The results indicate that our approach achieves significant improvements over the existing work and brute force approaches. We demonstrate that further speed-up can be obtained when small amount of an error tolerance in maximum stress is allowed.

INTRODUCTION

As additive manufacturing technologies become increasingly popular in recent years, optimization in structure design has received much research attention. A common approach in formalizing such an optimization problem is to model the external forces as known and fixed quantities. However, in many real world applications, the external force's contact locations and magnitudes may change significantly. In order to guarantee that the resulting structure is robust under all possible force configurations, the maximum stress experienced within the current shape hypothesis under the worst-case force configuration is to be computed at each optimization step. The maximum stress is then compared against a pre-specified tolerance threshold (i.e., material yield strength) and the structure design is then updated accordingly.

Finite-element analysis (FEA) is the standard technique that computes the stress distribution for a given external force configuration and the maximum stress suffered can then be subsequently obtained. However, FEA is generally computationally expensive, and performing the FEA on every external force configuration possible is out of the question for most structures. Approximation methods that reduce the total number of FEA runs are mandatory to make the stress prediction and structure optimization problem feasible in practice.

Ulu *et al.* [1] initiated the research of applying machine

*Address all correspondence to this author: lkara@cmu.edu




FIGURE 1. Example test structures with complex geometries. Fixed boundary conditions and contact regions are indicated in blue and red, respectively.

learning techniques to the maximum stress prediction problem when there exists uncertainty in the external force locations. In particular, a small subset of force locations were sub-sampled and the stress responses for forces applied on these locations are computed by FEAs. Afterwards, a quadratic regression model was built on the sub-sampled data points, which was then used to predict the stress distribution for the other force locations not sampled. Empirical results show that, with additional post-processing steps, the machine learning based approach gives accurate predictions of the maximum stress using small number of FEAs.

Inspired by [1], we propose a new method for predicting the worst-case force configuration for problems in which there is uncertainty in the force locations. Our approach takes as input a 3D shape represented by its boundary surface mesh and a user-specified contact region where external forces may make contact and computes the force configuration that is most likely to result in structural failure. The proposed algorithm pipeline is based on a linear regression model built on the top eigenvectors of the graph Laplacian of the full boundary mesh and a simple ranking post-processing step. We apply a computationally tractable experimental design method to select the training set to reduce the number of FEAs needed without much sacrifice in the prediction accuracy of the worst-case load and the corresponding maximum stress. Our experimental results suggest that the newly proposed algorithm significantly improves over existing work, with approximately $5\times$ fewer FEAs required compared to [1] and up to $100\times$ compared against the brute-force approach.

RELATED WORK

Our review focuses on studies that highlight structural analysis and experimental design, with an emphasis on approaches involving structural mechanics.

Structural Analysis

In predicting structural soundness of an object, stress and deformation analysis using FEA often introduce expensive computational bottlenecks. Commonly, researchers alleviate this issue

by using simple elemental structures such as trusses [2–4] and beams [5]. For cases where the structure cannot be represented by these simple elements, Umetani and Schmidt [6] extend the well-known Euler-Bernoulli model to free-from 3D objects and simplify the problem into 2D cross-sectional analysis.

Specific to problems with force location uncertainty, a common approach in engineering is to use the concept of equivalent uniformly distributed static load to perform simple approximate analysis [7]. This concept is commonly encountered in bridge (for traffic load) and building (for wind load) design. However, it is limited to simple geometries, making it unsuitable for our purposes. For complex geometries, Langlois *et al.* [8] use contact force samples generated by rigid body simulations to predict failure modes of objects in real world use. However, their method is applicable to scenarios where loading is stochastic in nature (such as dropping and collisions) and it is not practical for deterministic scenarios where possible force configurations are known and no failure is tolerated for any of them.

In the context of worst-case structural analysis, Zhou *et al.* [9] present a modal analysis based heuristic to static problems in determining the structurally problematic regions that is likely to fail under an arbitrary loading. Building upon this approach, Ulu *et al.* [1] presents a machine learning technique to find the most critical force configuration efficiently for problems in which there is uncertainty in the force locations. The approach is based on a naive sub-sampling algorithm to select the training set for the machine learning model that maximizes the pairwise geodesic distance between the selected force configurations. Our approach improves over this method by incorporating a computationally tractable experimental design method, resulting in a significant reduction in the number of FEAs needed.

Experimental Design

Experimental design, also known as optimal design, is a classical problem in statistics research [10, 11]. Given a large pool of potential candidates, a small subset of design points are selected such that the statistical efficiency by regressing on the selected designs are maximized. The problem is usually formulated as discrete optimization that is computationally challenging (NP-hard) to solve, and in practice heuristics such as greedy exchange [12, 13] and sampling [14] are usually deployed.

Recently, there has been a surge of research in computationally efficient experimental design approaches that enjoy rigorous theoretical guarantees [14–18]. In this paper, we adopt the methodology developed in [17], which involves solving a continuous convex optimization problem [14, 19] followed by a greedy rounding algorithm based on graph sparsification techniques [20, 21]. Compared to other methods, the algorithm proposed in [17] has the advantages of being applicable to a wide range of optimality criteria and computationally very efficient in practice.




FIGURE 2. Overview of our algorithm. Given a structure (a) with corresponding fixed boundary conditions (blue) and contact regions (red), we use a computationally tractable experimental design method to optimally sample a number of force instants on \mathcal{F} (b) and perform FEAs to obtain corresponding maximum stress values (c). We then construct a simple regression model to estimate the largest stresses for the remaining force instants and perform a simple ranking based post-processing step to predict the worst-case force configuration (d). Corresponding maximum stress value constitutes the most critical stress that the object could experience.

PROBLEM FORMULATION

Suppose the input model is represented by a boundary surface mesh \mathcal{S} and its interior is discretized by a volumetric mesh \mathcal{W} . Our aim is to find the maximum stress generated on \mathcal{W} for any external force that might make contact within a user-specified contact regions $\mathcal{F} \subseteq \mathcal{S}$ (Figure 1). We assume that the object is anchored in space by fixing three or more non-collinear boundary nodes and forces are applied compressively along the surface normal direction.

Under these assumptions, for linear elastic structures, it has been shown in [1] that the stress is maximized at some point in \mathcal{W} when the normal forces are concentrated at a single point rather than being distributed. Therefore, the problem becomes finding the most critical contact point in \mathcal{F} .

Given \mathcal{S} , \mathcal{F} and \mathcal{W} , it is possible to calculate stress distribution over \mathcal{W} for a force applied at a particular contact point in \mathcal{F} using FEA, mathematically formulated as

$$\sigma : \mathcal{F} \times \mathcal{W} \rightarrow \mathbb{R}_+. \quad (1)$$

More specifically, running a single FEA for a force applied on a node $f \in \mathcal{F}$, one obtains $\sigma(f, \cdot)$, which contains the von Mises stress data $\sigma(f, w)$ for all nodes $w \in \mathcal{W}$. Then, the problem of calculating the maximum stress on \mathcal{W} incurred by the worst-case external force can be formulated as

$$\sigma^* := \max_{f \in \mathcal{F}} \sigma^*(f) := \max_{f \in \mathcal{F}} \max_{w \in \mathcal{W}} \sigma(f, w). \quad (2)$$

ALGORITHM

A naive method to obtain the maximum stress σ^* is to compute the stress distribution $\sigma(f, \cdot)$ for every force node f in the contact region \mathcal{F} . However, as FEAs are computationally heavy and the number of force configurations n_F can be very large, such a brute-force method is infeasible in real-world applications and approximate computations of the maximum stress σ^* is mandatory.

Overview

In this paper, we introduce a computational approach that efficiently computes σ^* using a small subset of force nodes in \mathcal{F} . Our algorithmic pipeline is a great simplification of that developed in [1] but yields much more competitive results.

Figure 2 illustrates our approach. From an input 3D shape and prescribed contact regions (Figure 2(a)), our algorithm predicts the force configuration creating the largest stress on the object. We start by sampling a small subset of the contact region $\mathcal{F}_L \subseteq \mathcal{F}$ such that \mathcal{F}_L contains $n_{FL} \ll n_F$ force nodes (Figure 2(b)). After \mathcal{F}_L is obtained, FEAs are performed on the n_{FL} subsampled force nodes to obtain the stress distribution $\sigma(f, \cdot)$ as well as the maximum stress $\sigma^*(f)$ for all $f \in \mathcal{F}_L$ (Figure 2(c)). Then, we build a simple linear regression model on these force samples and corresponding maximum stresses and predict the largest stress values for the remaining force nodes. Finally, we perform a simple ranking based post-processing step to make the worst-case force prediction and compute the corresponding maximum stress σ^* (Figure 2(d)).

Regression Model

Let $\mathbf{F} \in \mathbb{R}^{n_F \times n_F}$ be the force matrix where each row represents a force configuration. Similar to [1], we distribute the force

to a small circular area around the contact point to avoid stress singularities. Therefore, $\mathbf{F}(f, f')$ is non-zero if f' is inside the circular region centered at f and $\mathbf{F}(f, f') = 0$ otherwise. Please refer to [1] for complete description of the construction of \mathbf{F} .

In its original form, dimensionality of the force configuration is very large compared to the number of training samples ($n_{FL} \ll n_F$). Moreover, this sparse representation does not carry spatial information, i.e. forces on two spatially very close nodes that create very similar stress distributions can be equally distinct in this representation as two spatially distant forces. To solve these problems, we project the force representation onto a lower dimensional Laplacian basis in which the spatial information is implicitly carried in smaller dimensionality. Let $\mathbf{L} \in \mathbb{R}^{n_F \times p}$ be the top- p eigenvectors of the graph Laplacian of \mathcal{F} , where p is a small integer that is set as $p = 15$ in this paper. We refer the readers to [22] for an introduction to the graph Laplacian and its properties. Denote also $\sigma^* = [\sigma^*(f)]_{f \in \mathcal{F}} \in \mathbb{R}_+^{n_F}$ as the maximum stress responses for all force configurations in \mathcal{F} . We build the following linear regression model

$$\sigma^* = \overline{\mathbf{F}}\mathbf{L}\beta_0 + \varepsilon, \quad (3)$$

where $\overline{\mathbf{F}}$ is the mean-centered \mathbf{F} so that each column of $\overline{\mathbf{F}}$ sums to 0, β_0 is a p -dimensional unknown vector that models the linear map, and $\varepsilon \in \mathbb{R}^{n_F}$ represents modeling noise for each force node.

Let $\sigma_L^* = [\sigma^*(f)]_{f \in \mathcal{F}_L}$ be the maximum stress responses on the subsampled force nodes in \mathcal{F}_L , which can be calculated from the results of FEAs carried out on all $f \in \mathcal{F}_L$. Let also $[\overline{\mathbf{F}}\mathbf{L}]_L$ be the corresponding n_{FL} rows in the $n_F \times p$ matrix $\overline{\mathbf{F}}\mathbf{L}$. An ordinary least squares (OLS) estimator $\hat{\beta}$ can be calculated using σ_L^* and $[\overline{\mathbf{F}}\mathbf{L}]_L$ as follows:

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \|\sigma_L^* - [\overline{\mathbf{F}}\mathbf{L}]_L\beta\|_2^2. \quad (4)$$

Afterwards, a preliminary prediction of maximum stress $\hat{\sigma}^*$ can be obtained by applying the OLS estimate $\hat{\beta}$:

$$\hat{\sigma}^* = \overline{\mathbf{F}}\mathbf{L}\hat{\beta}. \quad (5)$$

Force Node Ranking and Final Predictions

Given the maximum stress prediction $\hat{\sigma}^*$ in Eq. (5), it is tempting to directly report $\max_{f \in \mathcal{F}} \hat{\sigma}^*(f)$ as the final prediction of the maximum stress σ^* corresponding to the worst-case force. However, our empirical results suggest that the absolute estimation error $\|\hat{\sigma}^* - \sigma^*\|_\infty$ is in fact quite large, which makes the direct predicting approach less desirable. This phenomenon was also observed in [1], which showed that a quadratic regression

model on reduced-dimensional data has a similar large error in predicting the absolute value of the maximum stress response.

On the other hand, while the absolute estimation error $\|\hat{\sigma}^* - \sigma^*\|_\infty$ can be large, we observe that the relative positions of the force nodes which result in maximum stress as predicted by $\hat{\sigma}^*$ are quite accurate. In other words, if a force at node $f \in \mathcal{F}$ is predicted to result in large maximum stress as indicated by $\hat{\sigma}^*$, then it is likely to cause large stress in the ground-truth response σ^* as well. This observation motivates us to consider a *ranking* approach for making the final stress predictions.

In particular, for a parameter $k \ll n_F$, we define $\widehat{\mathcal{F}}_K$ as the top k force nodes in \mathcal{F} that have the largest stress response according to $\hat{\sigma}^*$. We then perform FEA on every force node in $\widehat{\mathcal{F}}_K$, with the maximum stress response denoted as

$$\hat{\sigma}_K^* = \left[\max_{w \in \mathcal{W}} \sigma(f, w) \right]_{f \in \widehat{\mathcal{F}}_K}. \quad (6)$$

The predicted maximum stress is then calculated as

$$\tilde{\sigma}^* := \max_{f \in \widehat{\mathcal{F}}_K} \hat{\sigma}_K^*(f). \quad (7)$$

Compared to the fully $\hat{\sigma}^*$ based approach, the new predictor $\hat{\sigma}_K^*$ is much more accurate as fresh FEAs are run on the predicted ‘‘top’’ force nodes; therefore, as long as the relative force positions resulting in maximum stress as indicated by $\hat{\sigma}^*$ are sufficiently accurate, the predicted maximum stress $\tilde{\sigma}^*$ is very accurate. For a small k , very few number of new FEAs are needed for the final prediction, which is only a small computational overhead incurred on the entire algorithm pipeline.

Computationally Tractable Experimental Design

In order to select the contact region subset $\mathcal{F}_L \subseteq \mathcal{F}$ that we train the regression model on, we benefit from the experimental design algorithm described in [17]. This approach is based on the (optimal) experimental design question in the statistics literature, and draws tools from the areas of convex optimization and spectral graph sparsification to achieve computational tractability.

Optimal experimental design For the linear regression model $\sigma^* = \mathbf{X}\beta_0 + \varepsilon$ specified in Eq. (3), where $\mathbf{X} = \overline{\mathbf{F}}\mathbf{L} = (x_1, \dots, x_{n_F}) \subseteq \mathbb{R}^p$, the optimal design problem can be formulated as a combinatorial optimization problem as follows:

$$\min_{s_1, \dots, s_n} \Phi \left(\sum_{i=1}^n s_i x_i x_i^\top \right) \quad \text{s.t.} \quad s_i \in \{0, 1\}, \quad \sum_{i=1}^n s_i \leq n_{FL}. \quad (8)$$

TABLE 1. Statistics of the testing structures

	n_W	n_S	n_F	σ^* [MPa]
FERTILITY	11221	4494	3914	6.37
ROCKINGCHAIR	15191	5918	5348	37.0
SHARK	14152	5757	4281	26.2

Here, the binary variables $\{s_i\}_{i=1}^n$ represent the selected design points (force nodes) in \mathcal{F}_L , and the objective function Φ is the optimality criterion that reflects certain aspects of the desired statistical efficiency. As in our problem where the prediction error $\mathbf{X}(\hat{\beta} - \beta_0)$ is of primary concern, the most relevant criteria are the V- and G-optimality:

$$\text{V-optimality: } \Phi_V(A) = \frac{1}{n} \text{tr}(XA^{-1}X^\top);$$

$$\text{G-optimality: } \Phi_G(A) = \max_{1 \leq i \leq n} x_i^\top A^{-1} x_i.$$

The V-optimality Φ_V measures the variance of the prediction averaged over all design points x_i , and the G-optimality Φ_G measures the maximum prediction variance. In this paper, we opt for the V-optimality Φ_V for computational reasons, as Φ_V is differentiable and hence easier to optimize by first-order optimization methods.

A continuous relaxation As direct optimization of the combinatorial problem in Eq. (8) is difficult, we instead consider a continuous relaxation of it as described in [14, 17, 19]:

$$\min_{\pi_1, \dots, \pi_n} \Phi \left(\sum_{i=1}^n \pi_i x_i x_i^\top \right) \quad \text{s.t. } 0 \leq \pi_i \leq 1, \quad \sum_{i=1}^n \pi_i \leq n_{FL}. \quad (9)$$

The only difference here is that the optimization variables are no longer constrained to take integer values. In addition, because the Φ_V objective is convex, the optimization problem becomes a standard convex continuous optimization problem and can be solved using conventional convex optimization methods. In particular, we use the projected gradient descent method [23, 24], which is observed to converge fast in practice [14]. Details of the projected gradient descent algorithm are given in Appendix A.

Sparsifying the continuous solution The optimal solution π^* to the continuously relaxed problem in Eq. (9) is in general a dense vector, and cannot be used directly to obtain a subset $\mathcal{F}_L \subseteq \mathcal{F}$ with only n_{FL} force nodes. To address that, we turn the optimal continuous solution π^* into a sparse design set as described in [17]. The algorithm starts with an empty set \emptyset and greedily add elements $f \in \mathcal{F}$ into the design set \mathcal{F}_L , according to a carefully designed potential function. This greedy algorithm has the advantage of being completely deterministic, thus reducing the uncertainty in sampling methods caused by statistical fluctuation. It has also been shown that the resulting design

subset \mathcal{F}_L enjoys theoretical approximation guarantees. This method will be referred as GREEDY in the remainder of the paper and details of the algorithm is given in Appendix B.

EXPERIMENTS

We evaluate the performance of our algorithm on three test structures (FERTILITY, ROCKINGCHAIR and SHARK) illustrated in Fig. 1. Table 1 gives a description of statistics of these test structures, including number of nodes in \mathcal{W} , \mathcal{S} and \mathcal{F} as n_W , n_S and n_F and the maximum stress σ^* .

In our experiments, we consider 5 different methods to sample the force nodes \mathcal{F}_L , which is arguably the most important step in our algorithmic framework. We compare the GREEDY approach described previously in this paper with relatively simple baseline methods, UNIFORM and LEVSCORE, as well as the previous work, K-MEANS [1] and SAMPLING [14]:

1. UNIFORM: The sample set \mathcal{F}_L is obtained by sampling without replacement each force node $f \in \mathcal{F}$ uniformly at random, until n_{FL} samples are obtained.
2. LEVSCORE: The sample set \mathcal{F}_L is obtained by sampling without replacement each force node $f \in \mathcal{F}$ with probability proportional to its leverage score, defined as $x(f)^\top (\mathbf{X}^\top \mathbf{X})^{-1} x(f)$, until n_{FL} samples are obtained.
3. K-MEANS: The sample set \mathcal{F}_L consists of n_{FL} force nodes in \mathcal{F} such that the geodesic distance between the closest force nodes in \mathcal{F}_L is maximized. As the problem itself is NP-hard, the K-means (Lloyd's) algorithm is employed to get an approximate solution.
4. SAMPLING: The sample set \mathcal{F}_L is obtained by sampling without replacement each force node $f \in \mathcal{F}$ with probability $\pi(f)^*$ until n_{FL} samples are obtained, where π^* is the optimal solution to Eq. (9).

Evaluation Measures

The predicted maximum stress $\tilde{\sigma}^*$ is by definition less than or equal to the ground truth σ^* . The performance of an algorithm is evaluated by the smallest integer k required such that $\sigma^* \leq (1 + \delta)\tilde{\sigma}^*$, where $\delta \geq 0$ is a pre-specified error tolerance parameter. The setting of $\delta = 0$ asks for exact identification of the maximum stress caused by the worst-case external force application, while $\delta > 0$ allows for a small error margin in the prediction of $\tilde{\sigma}^*$.

RESULTS AND DISCUSSION

We report the smallest k recovered for $\tilde{\sigma}^*$ to be lower bounded by $\sigma^*/(1 + \delta)$ in Tables 2, 3 and 4 for the three test structures. We report the performance for $\delta \in \{0, 0.01, 0.05, 0.1\}$ settings, and the sizes of the sub-sampled training set n_{FL} ranging from 25 to 300. The total number of FEAs needed by an algorithm is computed as the sum of n_{FL} and k . For non-deterministic

TABLE 2. Results for the FERTILITY model. Randomized algorithms (UNIFORM, LEVSCORE and SAMPLING) are run for 10 independent trials and the median performance is reported. Best performing settings are indicated in bold.

	$n_{FL} =$	25	50	100	150	200	250	300	Total FEAs
$\delta = 0$	UNIFORM	316.5	149	78.5	37.5	98.5	42.5	39	178.5 ($n_{FL} = 100$)
	LEVSCORE	252.5	54.5	73.5	68.5	42.5	31	13.5	104.5 ($n_{FL} = 50$)
	K-MEANS	237	25	61	82	57	17	16	75 ($n_{FL} = 50$)
	SAMPLING	210.5	148.5	51	30	35.5	34	26.5	151 ($n_{FL} = 100$)
	GREEDY	12	26	13	7	11	25	33	37 ($n_{FL} = 25$)
$\delta = 0.05$	UNIFORM	285	80.5	52	10	63	10	10	130.5 ($n_{FL} = 50$)
	LEVSCORE	175	26.5	55.5	59	17	10	7	76.5 ($n_{FL} = 50$)
	K-MEANS	144	2	19	22	14	2	2	52 ($n_{FL} = 50$)
	SAMPLING	202	113	10	7	11	8	6	110 ($n_{FL} = 100$)
	GREEDY	4	3	4	7	5	2	6	29 ($n_{FL} = 25$)
$\delta = 0.1$	UNIFORM	87.5	37.5	13	7	15.5	7	6.5	87.5 ($n_{FL} = 50$)
	LEVSCORE	59	13	15	14	13	8	6	63 ($n_{FL} = 50$)
	K-MEANS	46	1	7	20	8	1	1	51 ($n_{FL} = 50$)
	SAMPLING	88	25	7.5	6	7.5	6.5	4	75 ($n_{FL} = 50$)
	GREEDY	4	1	1	4	1	1	4	29 ($n_{FL} = 25$)

TABLE 3. Results for the ROCKINGCHAIR model. Randomized algorithms (UNIFORM, LEVSCORE and SAMPLING) are run for 10 independent trials and the median performance is reported. Best performing settings are indicated in bold.

	$n_{FL} =$	25	50	100	150	200	250	300	Total FEAs
$\delta = 0$	UNIFORM	716	857	385.5	42	135.5	269.5	36	192 ($n_{FL} = 150$)
	LEVSCORE	764.5	208.5	36	36	36	36	36	136 ($n_{FL} = 100$)
	K-MEANS	4013	4400	4573	4301	4320	4620	4757	4038 ($n_{FL} = 25$)
	SAMPLING	672.5	282	38.5	38	38	36	36	138.5 ($n_{FL} = 100$)
	GREEDY	36	35	208	35	36	36	36	61 ($n_{FL} = 25$)
$\delta = 0.05$	UNIFORM	404	466	201.5	20	88	93.5	18	170 ($n_{FL} = 150$)
	LEVSCORE	444	192.5	20	18.5	18	18	18	120 ($n_{FL} = 100$)
	K-MEANS	285	466	14	24	26	161	195	114 ($n_{FL} = 100$)
	SAMPLING	540	268	21.5	20.5	20.5	20	20	121.5 ($n_{FL} = 100$)
	GREEDY	20	19	200	20	20	20	20	45 ($n_{FL} = 25$)
$\delta = 0.1$	UNIFORM	399	384.5	135.5	15.5	60.5	75.5	14	165.5 ($n_{FL} = 150$)
	LEVSCORE	437	145	14	14	14	14	14	114 ($n_{FL} = 100$)
	K-MEANS	258	395	5	13	16	140	184	105 ($n_{FL} = 100$)
	SAMPLING	535	237	16	16	16	16	16	116 ($n_{FL} = 100$)
	GREEDY	14	14	175	16	16	14	14	39 ($n_{FL} = 25$)

algorithms (UNIFORM, LEVSCORE and SAMPLING), we perform 10 repetitions and report the median performance.

Tables 2, 3 and 4 suggest that both the K-MEANS and the GREEDY algorithms outperform their competitors for most pa-

rameter settings. One important reason for the overall good performance of K-MEANS and GREEDY is their deterministic nature, which avoids poor designs due to statistical perturbations in the other randomized algorithms. Furthermore, the GREEDY al-

TABLE 4. Results for the SHARK model. Randomized algorithms (UNIFORM, LEVSCORE and SAMPLING) are run for 10 independent trials and the median performance is reported. Best performing settings are indicated in bold.

	$n_{FL} =$	25	50	100	150	200	250	300	Total FEAs
$\delta = 0$	UNIFORM	585	384	141.5	208.5	20	9	9.5	220 ($n_{FL} = 200$)
	LEVSCORE	478.5	9	9	9	9	9	9	59 ($n_{FL} = 50$)
	K-MEANS	133	102	9	9	9	9	9	109 ($n_{FL} = 100$)
	SAMPLING	963.5	87	9	9	9	9	9	109 ($n_{FL} = 100$)
	GREEDY	9	171	9	9	9	9	9	34 ($n_{FL} = 25$)
$\delta = 0.01$	UNIFORM	568.5	341	131.5	156	15	4	4.5	215 ($n_{FL} = 200$)
	LEVSCORE	416	4	4	4	4	4	4	54 ($n_{FL} = 50$)
	K-MEANS	129	84	4	4	4	4	4	104 ($n_{FL} = 100$)
	SAMPLING	872.5	69	4	4	4	4	4	104 ($n_{FL} = 100$)
	GREEDY	4	115	4	4	4	4	4	29 ($n_{FL} = 25$)
$\delta = 0.05$	UNIFORM	323	172.5	52	75	10	4	4.5	151 ($n_{FL} = 100$)
	LEVSCORE	225	4	4	4	4	4	4	54 ($n_{FL} = 50$)
	K-MEANS	129	80	4	4	4	4	4	104 ($n_{FL} = 100$)
	SAMPLING	391.5	52.5	4	4	4	4	4	102.5 ($n_{FL} = 50$)
	GREEDY	4	115	4	4	4	4	4	29 ($n_{FL} = 25$)




FIGURE 3. Sampled force nodes (\mathcal{F}_L) using the K-MEANS algorithm (top row) versus the GREEDY algorithm (bottom row) for $n_{FL} = 100$.

algorithm remains accurate and robust even when n_{FL} is very small (e.g., $n_{FL} = 25$). For such a small n_{FL} setting, the other methods require large k values to compensate for the prediction error. Therefore, the GREEDY algorithm can produce an accurate pre-

diction of the overall maximum stress using much fewer number of total FEAs in both the first and the last stages of our algorithm pipeline, as shown by the rightmost columns in the tables. Our approach uses no more than 65 FEAs to successfully recover the




FIGURE 4. Sampled force nodes (\mathcal{F}_L) using the K-MEANS algorithm (top row) versus the GREEDY algorithm (bottom row) for $n_{FL} = 200$.

maximum stress caused by worst-case external forces. In addition, when a 5% to 10% error tolerance is used, the number of FEAs can be further reduced to less than 40. This is close to a $100\times$ speed-up compared to the brute-force approach that performs FEA on the entire surface mesh. It also achieves a $5\times$ speed-up over the existing work [1] and is simpler to implement.

In Fig. 3 and Fig. 4, we plot the sub-sampled force nodes (i.e., \mathcal{F}_L) of the GREEDY algorithm for $n_{FL} = 100$ and $n_{FL} = 200$, respectively. We provide the samples obtained by the K-MEANS algorithm in [1] for comparison. The difference in the sampling patterns between GREEDY and K-MEANS are quite obvious from the figures. We explain the differences for the three structures separately:

FERTILITY: The K-MEANS algorithm emphasizes the pairwise geodesic distance between sample points and thus places samples in a uniform fashion on the bodies, necks and heads of the structure. On the other hand, the GREEDY algorithm places more samples on the arms connecting the mother and the child, which are the most fragile parts of the structure. By placing more samples on these parts the learned linear model is more accurate in predicting the maximum stress, and therefore fewer FEAs are required to achieve a certain error tolerance level in $\tilde{\sigma}^*$.

ROCKINGCHAIR: The GREEDY algorithm produces more samples on the top end of the body, the surface area of the smaller back support and the edges of the larger seat com-

pared to the equally distanced K-MEANS design. External forces applied onto these parts of this structure are likely to cause increased stress, and therefore more samples placed around this region can greatly improve the regression model built for the maximum stress.

SHARK: As reported in Table 4, most of the sampling methods can predict the maximum stress using very small number of FEAs. However, if we focus on the sample points on the fins of the shark there are some noticeable differences between K-MEANS and GREEDY. While the GREEDY algorithm places more points around the tips, samples produced by the K-MEANS algorithm are relatively uniformly distributed on the surfaces. This subtle difference makes GREEDY algorithm more robust in prediction accuracy for small n_{FL} values.

Despite the significant reduction in FEA time, one important limitation of the proposed algorithm is the lack of stopping criteria. In particular, the performance parameter k can only be evaluated once the ground-truth maximum stress σ^* is known. On the other hand, performance control in problems involving structural mechanics is of vital importance because designs with insufficient stress tolerance may actually fail in reality, leading to devastating consequences. In our examples, we empirically determined that $k = 40$ is sufficient for most of the n_{FL} and δ settings. Determining the optimum value in a more principled way remains as an open problem and an immediate future work.

CONCLUSION

We present an efficient analysis approach for 3D objects under force location uncertainty. The proposed worst-case analysis approach efficiently determines the force contact point creating the highest stress in the structure. Driven by a computationally tractable experimental design method, we show that the relationship between the force configurations and resulting largest stress can be captured using only small number of FEA evaluations. We evaluate the performance of our algorithm on a set of arbitrarily complex 3D models. The results indicate that significant improvements over a brute force approach and an existing work can be achieved.

ACKNOWLEDGMENT

We are grateful to the designers whose 3D models we used: Aim@Shape for the fertility, Qingnan Zhou for the rocking chair and Lu et al. for the shark. This work is partly funded by NCDMM America Makes Project #4058.

REFERENCES

- [1] Ulu, E., Mccann, J., and Kara, L. B., 2017. “Lightweight structure design under force location uncertainty”. *ACM Transactions on Graphics*, **36**(4), p. 158.
- [2] Smith, J., Hodgins, J., Oppenheim, I., and Witkin, A., 2002. “Creating models of truss structures with optimization”. *ACM Transactions on Graphics*, **21**(3), July, pp. 295–301.
- [3] Rosen, D. W., 2007. “Design for additive manufacturing: a method to explore unexplored regions of the design space”. In Eighteenth Annual Solid Freeform Fabrication Symposium, pp. 402–415.
- [4] Wang, W., Wang, T. Y., Yang, Z., Liu, L., Tong, X., Tong, W., Deng, J., Chen, F., and Liu, X., 2013. “Cost-effective printing of 3d objects with skin-frame structures”. *ACM Transactions on Graphics*, **32**(6), Nov., pp. 177:1–177:10.
- [5] Zhang, X., Xia, Y., Wang, J., Yang, Z., Tu, C., and Wang, W., 2015. “Medial axis tree-an internal supporting structure for 3d printing”. *Comput. Aided Geom. Des.*, **35**(C), May, pp. 149–162.
- [6] Umetani, N., and Schmidt, R., 2013. “Cross-sectional structural analysis for 3d printing optimization”. In SIGGRAPH Asia 2013 Technical Briefs, SA '13, ACM, pp. 5:1–5:4.
- [7] Choi, W.-S., and Park, G.-J., 2002. “Structural optimization using equivalent static loads at all time intervals”. *Computer Methods in Applied Mechanics and Engineering*, **191**(1920), pp. 2105 – 2122.
- [8] Langlois, T., Shamir, A., Dror, D., Matusik, W., and Levin, D. I. W., 2016. “Stochastic structural analysis for context-aware design and fabrication”. *ACM Transactions on Graphics*, **35**(6), Nov., pp. 226:1–226:13.
- [9] Zhou, Q., Panetta, J., and Zorin, D., 2013. “Worst-case structural analysis”. *ACM Transactions on Graphics*, **32**(4), July, pp. 137:1–137:12.
- [10] Pukelsheim, F., 2006. *Optimal design of experiments*. SIAM.
- [11] Chaloner, K., and Verdinelli, I., 1995. “Bayesian experimental design: A review”. *Statistical Science*, **10**(3), pp. 273–304.
- [12] Fedorov, V. V., 1972. *Theory of optimal experiments*. Elsevier.
- [13] Miller, A., and Nguyen, N.-K., 1994. “A Fedorov exchange algorithm for d-optimal design”. *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, **43**(4), pp. 669–677.
- [14] Wang, Y., Yu, W. A., and Singh, A., 2017. “On computationally tractable selection of experiments in regression models”. *Journal of Machine Learning Research*, **18**(143), pp. 1–41.
- [15] Avron, H., and Boutsidis, C., 2013. “Faster subset selection for matrices and applications”. *SIAM Journal on Matrix Analysis and Applications*, **34**(4), pp. 1464–1499.
- [16] Horel, T., Ioannidis, S., and Muthukrishnan, S., 2014. “Budget feasible mechanisms for experimental design”. In Proceedings of Latin American Symposium on Theoretical Informatics (LATIN).
- [17] Allen-Zhu, Z., Li, Y., Singh, A., and Wang, Y., 2017. “Near-optimal design of experiments via regret minimization”. In Proceedings of the International Conference on Machine Learning (ICML).
- [18] Singh, M., and Xie, W., 2017. “Approximate positive correlated distributions and approximation algorithms for d-optimal design”. In Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA).
- [19] Joshi, S., and Boyd, S., 2009. “Sensor selection via convex optimization”. *IEEE Transactions on Signal Processing*, **57**(2), pp. 451–462.
- [20] Allen-Zhu, Z., Liao, Z., and Orecchia, L., 2015. “Spectral sparsification and regret minimization beyond matrix multiplicative updates”. In Proceedings of Annual Symposium on the Theory of Computing (STOC).
- [21] Silva, M. K., Harvey, N. J., and Sato, C. M., 2016. “Sparse sums of positive semidefinite matrices”. *ACM Transactions on Algorithms*, **12**(1), p. 9.
- [22] Chung, F. R., 1997. *Spectral graph theory*. No. 92. American Mathematical Soc.
- [23] Bubeck, S., et al., 2015. “Convex optimization: Algorithms and complexity”. *Foundations and Trends® in Machine Learning*, **8**(3-4), pp. 231–357.
- [24] Nesterov, Y., 2013. *Introductory lectures on convex optimization: A basic course*, Vol. 87. Springer Science & Business Media.
- [25] Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T.,

2008. “Efficient projections onto the L1-ball for learning in high dimensions”. In Proceedings of International Conference on Machine learning (ICML).

- [26] Condat, L., 2015. “Fast projection onto the simplex and the L1-ball”. *Mathematical Programming (Series A)*, **158**, pp. 575–585.
- [27] Su, H., Yu, A. W., and Li, F.-F., 2012. “Efficient euclidean projections onto the intersection of norm balls”. In Proceedings of International Conference on Machine Learning (ICML).

Appendix A: The Projected Gradient Descent Algorithm

Using the V -optimality $\Phi_V(A) = \text{tr}(XA^{-1}X^\top)$, the partial derivative of Φ_V with respect to π_i can be calculated as

$$\frac{\partial \Phi_V}{\partial \pi_i} = -\frac{1}{n} x_i^\top \Sigma^{-1} \mathbf{X}^\top \mathbf{X} \Sigma^{-1} x_i, \quad (10)$$

where $\Sigma = \sum_{j=1}^n \pi_j x_j x_j^\top$. Let also $\Pi := \{\pi \in \mathbb{R}^n : 0 \leq \pi_i \leq 1, \sum_{i=1}^n \pi_i \leq n_{FL}\}$ be the feasible set. The projected gradient descent algorithm can then be formulated as following:

1. **Input:** $\mathbf{X} \in \mathbb{R}^{n_F \times p}$, feasibility set Π , algorithm parameters $\alpha \in (0, 1/2]$, $\beta \in (0, 1)$;
2. **Initialization:** $\pi^{(0)} = (n_{FL}/n_F, \dots, n_{FL}/n_F)$, $t = 0$;
3. **While** stopping criteria are not met **do the following:**
 - (a) Compute the gradient $g_t = \nabla_\pi \Phi(\pi^{(t)})$ using Eq. (10);
 - (b) Find the smallest integer $s \geq 0$ such that $\Phi(\pi) - \Phi(\pi^{(t)}) \leq \alpha g_t^\top (\pi - \pi^{(t)})$, where $\pi = \mathcal{P}_\Pi(\pi^{(t)} - \beta^s g_t)$;
 - (c) Update: $\pi^{(t+1)} = \mathcal{P}_\Pi(\pi^{(t)} - \beta^s g_t)$, $t \leftarrow t + 1$.

Here in Steps 3(b) and 3(c), the $\mathcal{P}_\Pi(\cdot)$ is the projection operator onto the (convex) constrain set Π in Euclidean norm. More specifically, $\mathcal{P}_\Pi(\pi) := \arg \min_{z \in \Pi} \|\pi - z\|_2$. Such projection can be efficiently computed in almost linear time up to high accuracy [14, 25–27]. Also, the step 3(b) corresponds to the Amijo’s rule (also known as backtracking line search) that automatically selects step sizes, a popular and efficient method for step size selection in full gradient descent methods.

Appendix B: The Greedy Algorithm

The GREEDY algorithm was proposed in [17] as a principled method to sparsify the continuous optimization solution π^* . The algorithm makes uses of a carefully designed potential function

for $i \in [n_F]$ and $\Lambda \subseteq [n_F]$:

$$\psi(i; \Lambda) := \frac{x_i^\top \mathbf{B}(\Lambda) x_i}{1 + \alpha x_i^\top \mathbf{B}(\Lambda)^{1/2} x_i} \quad (11)$$

$$\text{where } \mathbf{B}(\Lambda) = \left(c \mathbf{I} + \sum_{j \in \Lambda} x_j x_j^\top \right)^{-2}, \quad \text{tr}(\mathbf{B}(\Lambda)) = 1.$$

Here, $\alpha > 0$ and $c \in \mathbb{R}$ is the unique real number such that $\text{tr}(\mathbf{B}(\Lambda)) = 1$. The exact values of c can be computed efficiently using a binary search procedure, as shown in [17]. The potential function is inspired by a regret minimization interpretation of the least singular values of sum of rank-1 matrices. Interested readers should refer to [17, 20, 21] for more details.

Based on the potential function in Eq. (11), the GREEDY algorithm starts with an empty set and add force nodes one by one in a greedy manner, until there are n_{FL} elements in \mathcal{F}_L .

1. **Input:** $\mathbf{X} \in \mathbb{R}^{n_F \times p}$, budget n_{FL} , optimal solution π^* , algorithm parameter $\alpha > 0$;
2. **Whitening:** $\mathbf{X} \leftarrow \mathbf{X}(\mathbf{X} \Sigma_* \mathbf{X}^\top)^{-1/2}$, where $\Sigma_* = \sum_{i=1}^n \pi_i^* x_i x_i^\top$;
3. **Initialization:** $\Lambda_0 = \emptyset$, $\mathcal{F}_L = \emptyset$;
4. **For** $t = 1$ to n_{FL} **do the following:**
 - (a) Compute $\psi(i; \Lambda_{t-1})$ for all $i \notin \Lambda_{t-1}$ and select $i_t := \arg \max_{i \notin \Lambda_{t-1}} \psi(i; \Lambda_{t-1})$;
 - (b) Update: $\Lambda_t = \Lambda_{t-1} \cup \{i_t\}$, $\mathcal{F}_L \leftarrow \mathcal{F}_L \cup \{i_t\}$.