LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-791997

High-Dimensional Spectral Sampling

J. J. Thiagarajan

September 28, 2019

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

LDRD FINAL REPORT
High-Dimensional Spectral Sampling
PI: Jayaraman J. Thiagarajan (17-ERD-009)

IM Release Number: LLNL-JRNL-xxxxxx

Abstract

A wide variety of applications in science and engineering depend on effectively and efficiently exploring high- dimensional parameter spaces. Common examples are uncertainty quantification, inverse problems, or various optimization approaches. In such applications, the first step is to create an initial uniform, random sampling of the space to create a baseline of knowledge. The simulations evaluated at these samples are then used for subsequent processing. In its most generic form, the goal of sampling is to produce the maximal amount of information with the minimal number of samples. However, despite being the first step and thus influencing the performance of the entire workflow, the quality of the initial sample designs is rarely analyzed or optimized. This project develops a new spectral sampling theory for analyzing and creating experiment designs in high-dimensional spaces. Further, new algorithms for estimating sample properties and sample synthesis are designed. Finally, applications of the proposed spectral sampling tools in surrogate modeling, scientific machine learning and image analysis are presented.

1 Background and Research Objectives

Exploratory analysis and inference in high dimensional parameter spaces is a ubiquitous problem in science and engineering. As a result, a wide-variety of machine learning tools and optimization techniques have been proposed to address this challenge. In its most generic formulation, one is interested in analyzing a high-dimensional function $f : \mathcal{D} \rightarrow \mathbb{R}$ defined on the d-dimensional domain $\mathcal{D}$. A typical approach for such an analysis is to first create an initial sampling $\mathcal{X} = \{\mathbf{x_i} \in \mathcal{D}\}_{i=1}^{N}$ of $\mathcal{D}$, evaluate f at all $\mathbf{x_i}$, and perform subsequent analysis and learning using only the resulting tuples $\{(\mathbf{x_i}, f(\mathbf{x_i}))\}_{i=1}^{N}$. Despite the widespread use of this approach, a critical question that still persists is: how should one obtain a high quality initial sampling $\mathcal{X}$ for which the data $f(\mathcal{X})$ is collected? This challenge is typically refered to as Design of Experiments (DoE) and solutions have been proposed as early as (Fisher 1935) to optimize agricultural experiments. Subsequently, DoE has received significant attention from researchers in different fields (Garud, Karimi, and Kraft 2017) as it is an important *building block* for a wide variety of applications, such as surrogate modeling, image reconstruction, reinforcement learning, or data analysis. Currently, a plethora of sampling solutions exist in the literature with a wide-range of assumptions and statistical guarantees; see (Garud, Karimi, and Kraft 2017; Art B Owen 2009) for a detailed review. Conceptually, most approaches aim to cover the sampling domain as uniformly as possible, in order to generate the so called *space-filling* designs. However, it is well know that uniformity alone does not necessarily lead to high performance. For example, optimal sphere packings lead to highly uniform designs, yet are well known to cause strong aliasing artifacts most easily perceived in many computer graphics applications. Instead, a common assumption is that a good design should balance uniformity and randomness. Unfortunately, an exact definition for what should be considered a good space-filling design has remained elusive.

Most common approaches use various scalar metrics to encapsulate different notions of ideal sampling properties. One popular metric is the discrepancy of an experimental design, defined as an appropriate ℓ_p norm of the ratio of points within all (hyper-rectangular)

sub-volumes of $\mathcal{D}$ and the corresponding volume ratio. In other words, discrepancy quantifies the non-uniformity of a sample design. The most prominent examples of so called *discrepancy sequences* are Quasi-Monte Carlo (QMC) methods and their variants (Caflisch 1998). In their classical form, discrepancy sequences are deterministic though extensions to incorporate randomess have been proposed, for example, using digital scrambling (Art B. Owen 1995). Nevertheless, by optimizing for discrepancy these techniques focus almost exclusively on uniformity, and consequently even optimized QMC patterns can be quite structured and create aliasing artifacts. Furthermore, even the fastest known strategies for evaluating popular discrepancy measures require $O(N^2d)$ operations making evaluation, let alone optimization, for discrepancy difficult even for moderate dimensions. Finally, for most discrepancy measures, the optimal achievable values are not known. This makes it difficult to determine whether a poorly performing sample design is due to the insufficiency of the chosen discrepancy measure or due to ineffective optimization.

Another class of metrics to describe sample designs are based on geometric distances. These can be used directly by, for example, optimizing the maximin or minimax distance of a sample design (Schlömer, Heck, and Deussen 2011) or indirectly by enforcing empty disk conditions. The latter is the basis for the so-called Poisson disk samples (Lagae and DutrÃ© 2008), which aim to generate random points such that no two samples can be closer than a given minimal distance r_{min}, i.e. enforcing an empty disc of radius r_{min} around each sample. Typically, Poisson-type samples are characterized by the *relative radius*, ρ, defined as the ratio of the minimum disk radius r_{min} and the maximum possible disk radius r_{max} for N samples to cover the sampling domain. Similar to the discrepancy sequences, maximin and minimax designs exclusively consider uniformity, are difficult to optimize for especially in higher dimensions, and often lead to very regular patterns.

1.1 Objectives

In general, scalar metrics used to evaluate the quality of a sample design tend not to be very descriptive. Especially in moderately high dimensions different designs converge to random samples (Morokoff and Caflisch 1994; Wang and Sloan 2008). Furthermore, one rarely knows the best achievable value of the metric, e.g., lowest possible discrepancy, for a given problem which makes evaluating and comparing sampling designs difficult. Finally, most metrics are expensive to compute and not easily optimized. This makes it challenging in practice to create good designs in high dimensions and with large sample sizes. Hence, the overall research objectives of this project can be summarized as follows:

- Develop a new technique to quantify the space-filling property, which enables us to systematically trade-off uniformity and randomness, consequently producing better quality sampling designs.
- Design a theoretical framework for sample designs in high-dimensions that connects the qualitative performance of a sampling pattern to its spatial characteristics.
- Obtain theoretical bounds on achievable performance for a given sample size and provide guidelines for experiment design.
- Develop computationally efficient, yet accurate, estimators for characterizing properties of samples.
- Design families of sampling patterns for different applications and devise a principle optimization framework for sample synthesis.
- Demonstrate the effectiveness of the new space-filling designs in surrogate modeling, hyperparameter search in deep neural networks, image reconstruction, sampling non-linear manifolds and studying generalization of ML algorithms.

2 Scientific Approach and Accomplishments

2.1 Methods

2.1.1 A New Spectral Sampling Framework

Without any prior knowledge on the task being solved, for example recovering an unknown regression function f or selecting hyper-parameters for training a neural network, a reasonable objective for exploratory sampling is to ensure that the samples are random, thus providing an equal chance of finding meaningful configurations anywhere in the domain of interest. However, to avoid sampling only parts of the parameter space, a second objective is often considered, which is to cover the space uniformly. Formally,

Definition A space-filling design is a set of samples that are distributed according to a uniform probability distribution (*Randomness*) but no two samples are closer than a given minimum distance r_{min} (*Coverage*).

In order to use this abstract definition to construct experimental designs, it is imperative to quantify the randomness and coverage properties of a given sampling pattern. To this end, we propose a novel metric based on the spatial statistic, *pair correlation function*(PCF) (B. Kailkhura et al. 2016) denoted as $G(r)$, which characterizes the distribution of sample distances r, thus providing a comprehensive description of sample designs. Further, we also propose to analyze the spectral properties of sample designs, using tools from Fourier analysis, namely the power spectral density. For isotropic samples, a metric of interest is the radially-averaged power spectral density, which describes how the signal power is distributed over frequencies. The spectral property determines which frequencies that the sample design can recover, in term of function approximation.

Definition For a finite set of N points, $\{\mathbf{x}_j\}_{j=1}^N$, in a region with unit volume, the power spectral density of the sampling function $\sum_{j=1}^N \delta(\mathbf{x}-\mathbf{x}_j)$ is formally defined as

$$P(\mathbf{k}) = \frac{1}{N}|S(\mathbf{k})|^2 = \frac{1}{N}\sum_{j,\ell} e^{-2\pi i \mathbf{k}.(\mathbf{x}_\ell - \mathbf{x}_j)}, \tag{1}$$

where $|.|$ is the ℓ_2-norm and $S(\mathbf{k})$ denotes the Fourier transform of the sampling function.

The connection between spectral properties of a d-dimensional isotropic sample design and its corresponding pair correlation function can be obtained via the d-dimensional Fourier transform or more efficiently using the 1-d Hankel transform.

Proposition 1 *For an isotropic sample design with N points, $\{\mathbf{x}_j\}_{j=1}^N$, in a d-dimensional region, the radially averaged power spectral density $P(k)$ and the pair correlation function $G(r)$ are related as follows:*

$$P(k) = 1 + \frac{N}{V}(2\pi)^{\frac{d}{2}} k^{1-\frac{d}{2}} H_{\frac{d}{2}-1}\left[r^{\frac{d}{2}-1}G(r) - 1\right] \tag{2}$$

where V is the volume of the sampling region and $H_d[.]$ denotes the Hankel transform

$$H_d(f(r))(k) = \int_0^\infty r J_d(kr) f(r) dr,$$

with $J_d(.)$ denoting the Bessel function of order d.

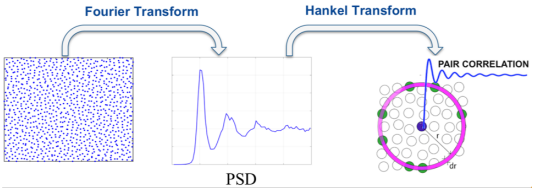

$$P(k) = 1 + \rho(2\pi)^{\frac{d}{2}} k^{1-\frac{d}{2}} H_{\frac{d}{2}-1}\left(r^{\frac{d}{2}-1}(G(r)-1)\right)$$

Figure 1: Illustration of the core idea of the proposed spectral sampling framework. By establishing a mathematical link between pair correlation function (spatial) and the 1−D power spectral density (spectral), we enable improved analysis of sampling distributions

This connection is central to using space-filling properties directly for optimization, and obtaining effective experimental designs (Figure 1). It is important to note that, not every PCF (or PSD) is physically realizable by a sample design. In fact, there are two necessary conditions [1] that a sample design must satisfy to be realizable: the PCF and PSD should be non-negative. These two conditions limit the space of realizable space-filling designs.

2.1.2 Space-Filling Spectral Designs

Definition A set of N point samples $\mathcal{X}$ in a sampling domain $\mathcal{D}$ can be characterized as a space-filling design, if $\mathcal{X} = \{\mathbf{x_i} \in \mathcal{D};\ i = 1, \cdots N\}$ satisfy the following two objectives:

- $\forall \mathbf{x_i} \in \mathcal{X},\ \forall \triangle\mathcal{D} \subseteq \mathcal{D}\ :\ P(\mathbf{x_i} \in \triangle\mathcal{D}) = \int_{\triangle\mathcal{D}} \mathbf{dx}$
- $\forall \mathbf{x_i}, \mathbf{x_j} \in \mathcal{X}\ :\ ||\mathbf{x_i} - \mathbf{x_j}|| \geq r_{\min}$

where $r_{\min}$ is referred to as the coverage radius.

In the above definition, the first objective states that the probability of a uniformly distributed random sample $\mathbf{x_i} \in \mathcal{X}$ falling inside a subset $\triangle\mathcal{D}$ of $\mathcal{D}$ is equal to the hyper-volume of $\triangle\mathcal{D}$. The second condition enforces the minimum distance constraint between point sample pairs for improving coverage. A space-filling design can be defined conveniently in both spatial and spectral domains, and we refer to this as the space-filling spectral design.

Proposition 2 *Given the desired coverage radius r_{min}, a space-filling spectral design is defined in the spatial domain as*

$$G(r - r_{min}) = \begin{cases} 0 & \text{if } r < r_{min} \\ 1 & \text{if } r \geq r_{min}. \end{cases}$$

Proposition 3 *Given the desired coverage radius r_{min}, a d-dimensional space-filling spectral design $\mathcal{X}$, with N sample points in a sampling domain $\mathcal{D}$ of volume V, can be defined in the PSD domain as*

$$P(k) = 1 - \frac{N}{V}\left(\frac{2\pi r_{min}}{k}\right)^{\frac{d}{2}} J_{\frac{d}{2}}(k r_{min})$$

where $J_{\frac{d}{2}}(.)$ is the Bessel function of order $d/2$.

This general framework allows tractable analysis of the properties for samples, for example we can derive upper bounds on achievable sample sizes while adhering to a given r_{min},

1. Whether or not these two conditions are not only necessary but also sufficient is still an open question (however, no counterexamples are known).

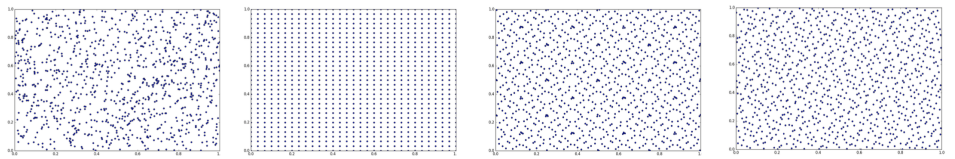

(a) Random (b) Regular (c) Sobol (d) Halton

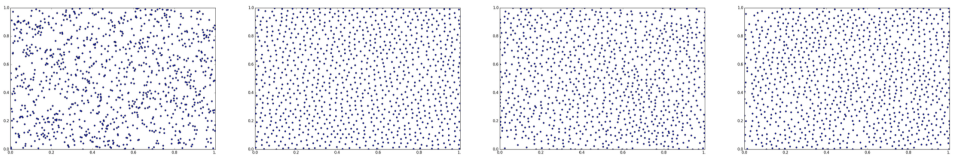

(e) LHS (f) MPDS (g) Step PCF (h) Stair PCF

Figure 2: Visualization of 2-d point distributions obtained using different sample design techniques. In all cases, the number of samples N was fixed at 1000.

Proposition 4 *For a fixed disk radius r_{min}, the maximum number of point samples possible for a realizable space-filling spectral design in the sampling region with volume V is*

$$N = \frac{V\Gamma\left(\frac{d}{2}+1\right)}{\pi^{\frac{d}{2}} r_{min}^{d}}.$$

More importantly, this framework can be utilized to design specialized sampling families targeted towards different applications.

2.1.3 A Synthesis Algorithm

We propose to iteratively transform an initial random input sample design such that its PCF matches the desired PCF characteristics. More specifically, we propose a non-linear least squares formulation. Let us denote the target PCF by $G^*(r)$. We discretize the radius r into m points $\{r_j\}_{j=1}^{m}$ and minimize the sum of the weighted squares of errors between the target PCF $G^*(r_j)$ and the curve-fit function $G(r_j)$ over m points. Sample synthesis can be posed as a non-linear least squares problem as follows:

$$\min \sum_{j=1}^{M} \left(G(r_j) - G^*(r_j)\right)^2.$$

To estimate the PCF of point samples $G(r)$, we employ a kernel density estimator similar to (Öztireli and Gross 2012). Given the PCF estimator, the PCF matching problem can be efficiently solved using gradient descent algorithm. However, the PCF matching problem is a highly non-convex and we find that vanilla gradient decent with constant learning rate (GD-CLR) (Bhavya Kailkhura et al. 2018) performs very poorly on the PCF matching task. To overcome this limitation, we developed a variant of gradient decent with adaptive learning rate, GD-ALR. Specifically, we use the following learning rate rule: $\lambda = 0.1e^{-0.1\sqrt{t}}$, for iteration t, which demonstrates good convergence behavior for all N and d (Figure 2).

2.2 Accomplishments

Optimal Sample Designs for Surrogate Modeling: Surrogate modeling is a widely-adopted first step towards analyzing uncertainties and parameter sensitivities of a complex simulation. We designed a specialized family of sampling distributions referred as *Stair* PCF

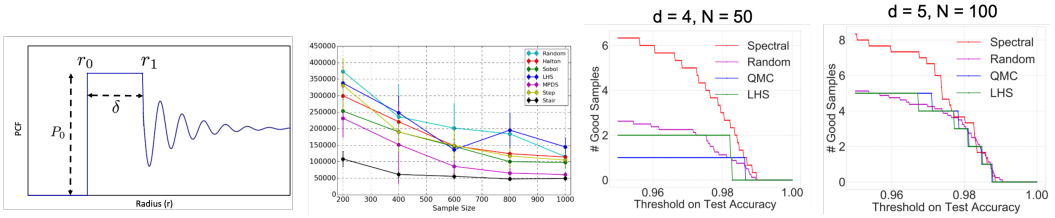

(a) Stair PCF for Surrogate Modeling (b) Stair PCF for Hyperparameter Search

Figure 3: A novel stair PCF based design, which we show to be effective for surrogate modeling and hyperparameter search in deep neural networks. We find that the proposed design requires 50% lesser samples than existing designs to achieve the same fidelity.

based space-filling design, which we prove to be effective in surrogate modeling tasks. The Stair PCF construction is defined as:

$$G(r; r_0, r_1, P_0) = f(r - r_1) + P_0 \left(f(r - r_0) - f(r - r_1) \right), \tag{3}$$
$$\text{with } f(r - r_0) = \left\{ \begin{array}{ll} 0 & \text{if } r \leq r_0 \\ 1 & \text{if } r > r_0 \end{array} \right\}, r_0 \leq r_1 \text{ and } P_0 \geq 1.$$

As illustrated in Figure 3(a), the proposed stair PCF family outperforms existing designs (in terms of mean squared error) on a synthetic test function (in 5 dimensions). We find that our design requires 50% lesser samples in order to achieve the same model fidelity. Consequently, our spectral designs have been adopted in a number of ongoing project at LLNL.

Studying Generalization in ML via Sample Design: Analyzing the generalization error of a learning algorithm is essential for estimating how well the generated hypothesis will apply to unknown test data. Traditionally, generalization error is analyzed based on the complexity of the function class, such as, the Vapnik-Chervonenkis (VC) dimension and the Rademacher complexity, or properties of the learning algorithm, such as uniform stability, and upper bounds on the error are derived. However, studying generalization from the standpoint of sample design can provide key insights into the practice of scientific machine learning and enable us to obtain performance bounds of ML algorithms. To this end, we developed an analysis framework for studying generalization error behavior of a learning algorithm through the lens of spectral properties of the sample design. We obtained key theoretical results that reveal the regimes in which one can do better than uniform random sampling in high dimensions.

Hyperparameter Optimization in DNN: With the increased adoption of modern machine learning approaches such as deep neural networks, the problem of choosing the appropriate hyper-parameters (number of layers, batch size, learning rate etc.) has become hugely critical, since they are known to impact the overall validation performance in non-trivial ways. Hence, we propose to repose the problem of choosing an optimal configuration as a high-dimensional sampling problem and utilize the proposed spectral designs for hyper-parameter search. As illustrated in Figure 3(b), the proposed stair sfsd consistently identifies optimal settings for different dimensions and sample sizes. We also successfully applied the proposed techniques for auto-tuning of system configurations in HPC systems.

Sample Design for Image Reconstruction: A common solution to reducing visible aliasing artifacts in image reconstruction is to employ sampling patterns with a blue noise power spectrum. These sampling patterns can prevent discernible artifacts by replacing them with incoherent noise. We extended the proposed spectral sampling framework to design a new family of blue noise distributions (Stair blue noise), which is mathematically tractable and enables parameter optimization to obtain the optimal sampling distribution. Furthermore, for a given sample budget, our design achieves a significantly larger alias-free

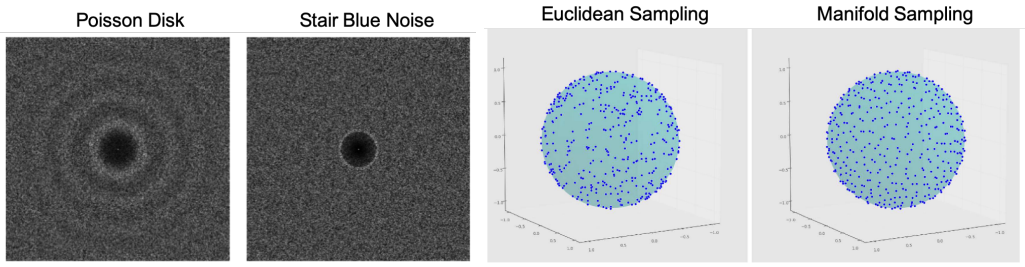

(a) Stair Blue Noise for Image Reconstruction (b) Sampling a sphere in $\mathbb{R}^3$

Figure 4: Extensions to the proposed sampling framework - (a) We develop a novel stair blue noise sampling pattern that demonstrates superior characteristics (indicated by 2D PSD) for image reconstruction tasks.; (b) We extend the proposed framework to generate Poisson disk samples on non-linear manifolds.

low-frequency region compared to existing approaches, without introducing visible artifacts in the mid- frequencies (Figure 4(a).

Sampling Non-Linear Manifolds: All techniques developed in this project focused on Euclidean domains and extending them to non-linear domains is non-trivial. Though developing a rigorous analysis framework for arbitrary manifolds is beyond the scope of this project, we designed sampling strategies for simple analytical manifolds and the Grassmannian (space of linear subspaces), a commonly used geometric space in machine learning problems. As illustrated in Figure 4(b), spectral sampling can produce uniform sampling distribution even on Riemannian manifolds (sphere in this example), when the metric is known.

3 Impact on Mission

This project falls directly in the Data Science Core Competency by developing new theories and algorithms useful for a broad range of predictive sciences. Due to its fundamental nature the results from this project have the potential to impact almost all five mission focus areas of the laboratory to a greater or smaller degree. Whether it is uncertainty quantification for stockpile stewardship or weapon design, better weather or climate ensembles, optimizing the energy grid, or supporting sensitivity analysis and inverse modeling in material science and high energy physics, better sampling strategies will lead to high quality results using significantly fewer resources. Our sampling tools have already been adopted by ongoing efforts at Livermore including LDRD SI project on cognitive simulations, CRADA project with ExxonMobil, HPC performance optimization, LDRD on collaborative autonomy and LDRD SI project in material science. We expect the software library that has been developed as part of this project will enable a wider adoption of these tools in other mission-critical applications. Furthermore, several of the research objectives of this project are well aligned with DOE's vision for scientific machine learning.

4 Conclusions

In this project, we developed a fundamental theoretical framework to study the interplay between the properties of sampling distributions and their fidelity (bounds of function approximation, generalization error etc.) in high dimensions. Based on the proposed framework we designed new families of sampling distributions that are specifically targeted to applications,

such as surrogate modeling, image reconstruction and Monte Carlo integration, and developed computational algorithms for synthesizing optimal samples. In this process, we have identified design guidelines for constructing optimal sampling patterns for a given problem. The results from this project are very promising as demonstrated by significant reduction in sample requirements in several applications, and its broad applicability. We believe that there are still interesting questions that remain to be explored in the future work such as an analysis of the generalization error for cases where data comes from non-linear manifolds. Note that the analytical methodologies developed in this project can be leveraged for studying the effect of sample design on generalization error in such cases. Other questions such as PSD/PCF parameterizations for other variants of space-filling designs, namely adaptive and importance sampling, and optimization approaches to synthesize them are also crucial directions of research in the context of DOE's missions.

References

Caflisch, Russel E. 1998. "Monte Carlo and quasi-Monte Carlo methods." *Acta Numerica* 7:1â49. doi:10.1017/S0962492900002804.

Fisher, Ronald A. 1935. "The design of experiments. 1935." *Oliver and Boyd, Edinburgh.*

Garud, Sushant S., Iftekhar A. Karimi, and Markus Kraft. 2017. "Design of computer experiments: A review." ESCAPE-26, *Computers and Chemical Engineering* 106 (Supplement C): 71–95. ISSN: 0098-1354.

Kailkhura, B., J. J. Thiagarajan, P. T. Bremer, and P. K. Varshney. 2016. "Theoretical guarantees for poisson disk sampling using pair correlation function." In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP),* 2589–2593. March. doi:10.1109/ICASSP.2016.7472145.

Kailkhura, Bhavya, Jayaraman J Thiagarajan, Charvi Rastogi, Pramod K Varshney, and Peer-Timo Bremer. 2018. "A Spectral Approach for the Design of Experiments: Design, Analysis and Algorithms." *Journal of Machine Learning Research (to appear).*

Lagae, Ares, and Philip DutrÃ©. 2008. "A Comparison of Methods for Generating Poisson Disk Distributions." *Computer Graphics Forum* 27 (1): 114–129. ISSN: 1467-8659. doi:10.1111/j.1467-8659.2007.01100.x.

Morokoff, William J., and Russel E. Caflisch. 1994. "Quasi-Random Sequences and Their Discrepancies." *SIAM J. Sci. Comput* 15:1251–1279.

Owen, Art B. 1995. "Randomly Permuted (t,m,s)-Nets and (t, s)-Sequences." In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing: Proceedings of a conference at the University of Nevada, Las Vegas, Nevada, USA, June 23–25, 1994,* edited by Harald Niederreiter and Peter Jau-Shyong Shiue, 299–317. New York, NY: Springer New York. ISBN: 978-1-4612-2552-2. doi:10.1007/978-1-4612-2552-2_19. https://doi.org/10.1007/978-1-4612-2552-2_19.

Owen, Art B. 2009. "Monte Carlo and quasi-Monte Carlo for statistics." *Monte Carlo and Quasi-Monte Carlo Methods 2008:* 3–18.

Öztireli, A. Cengiz, and Markus Gross. 2012. "Analysis and Synthesis of Point Distributions Based on Pair Correlation." *ACM Trans. Graph.* (New York, NY, USA) 31, no. 6 (November): 170:1–170:10. ISSN: 0730-0301. doi:10.1145/2366145.2366189.

Schlömer, Thomas, Daniel Heck, and Oliver Deussen. 2011. "Farthest-point Optimized Point Sets with Maximized Minimum Distance." In *Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics,* 135–142. HPG '11. Vancouver, British Columbia, Canada: ACM. ISBN: 978-1-4503-0896-0. doi:10.1145/2018323.2018345. http://doi.acm.org/10.1145/2018323.2018345.

Wang, Xiaoqun, and Ian H. Sloan. 2008. "Low discrepancy sequences in high dimensions: How well are their projections distributed?" *Journal of Computational and Applied Mathematics* 213 (2): 366–386. ISSN: 0377-0427. doi:https://doi.org/10.1016/j.cam.2007.01.005. http://www.sciencedirect.com/science/article/pii/S0377042707000374.

5 Publications and Presentations

[1] Kailkhura, Bhavya, Jayaraman J. Thiagarajan, Peer-Timo Bremer, and Pramod K. Varshney. "Stair blue noise sampling." ACM Transactions on Graphics (TOG) 35, no. 6 (2016): 248.

[2] Kailkhura, Bhavya, Jayaraman J. Thiagarajan, Peer-Timo Bremer, and Pramod K. Varshney. "Theoretical guarantees for poisson disk sampling using pair correlation function." In 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2589-2593. IEEE, 2016.

[3] Jayaraman J. Thiagarajan, Kailkhura, Bhavya and Peer-Timo Bremer. "Impact of Spectral Sampling Techniques on Surrogate Modeling." SIAM Conference on Uncertainty Quantification, 2016.

[4] Anirudh, Rushil, Bhavya Kailkhura, Jayaraman J. Thiagarajan, and Peer-Timo Bremer. "Poisson disk sampling on the grassmannnian: Applications in subspace optimization." In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 11-19. 2017.

[5] Kailkhura, Bhavya, Jayaraman J. Thiagarajan, Charvi Rastogi, Pramod K. Varshney, and Peer-Timo Bremer."A spectral approach for the design of experiments: Design, analysis and algorithms." The Journal of Machine Learning Research 19, no. 1 (2018): 1214-1259.

[6] Thiagarajan, Jayaraman J., Nikhil Jain, Rushil Anirudh, Alfredo Gimenez, Rahul Sridhar, Aniruddha Marathe, Tao Wang, Murali Emani, Abhinav Bhatele, and Todd Gamblin. "Bootstrapping parameter space exploration for fast tuning." In Proceedings of the 2018 International Conference on Supercomputing, pp. 385-395. ACM, 2018.

[7] Jayaraman J. Thiagarajan, Kailkhura, Bhavya and Peer-Timo Bremer. "A Spectral Approach for the Design of Experiments: Design, Analysis and Algorithms." SIAM Conference on Uncertainty Quantification, 2018.

[8] Muniraju, Gowtham, Bhavya Kailkhura, Jayaraman J. Thiagarajan, Peer-Timo Bremer, Cihan Tepedelenlioglu, and Andreas Spanias. "Coverage-Based Designs Improve Sample Mining and Hyper-Parameter Optimization." arXiv preprint arXiv:1809.01712 (2019).

[9] Kailkhura, Bhavya, Jayaraman J. Thiagarajan, Qunwei Li, and Peer-Timo Bremer. "A Look at the Effect of Sample Design on Generalization through the Lens of Spectral Analysis." arXiv preprint arXiv:1906.02732 (2019).