



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Automatic synthesis of experiment designs from probabilistic environment specifications

Citation for published version:

Innes, C, Hristov, Y, Kamaras, G & Ramamoorthy, S 2021, 'Automatic synthesis of experiment designs from probabilistic environment specifications', Paper presented at 10th Workshop on Synthesis, Los Angeles, United States, 19/07/21 - 19/07/21 pp. 1-4. <https://doi.org/10.48550/arXiv.2107.00093>

Digital Object Identifier (DOI):

[10.48550/arXiv.2107.00093](https://doi.org/10.48550/arXiv.2107.00093)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Automatic Synthesis of Experiment Designs from Probabilistic Environment Specifications

Craig Innes
University of Edinburgh
craig.innes@ed.ac.uk

Georgios Kamaras
g.kamaras-1@sms.ed.ac.uk
University of Edinburgh

Yordan Hristov
yordan.hristov@ed.ac.uk
University of Edinburgh

Subramanian Ramamoorthy
s.ramamoorthy@ed.ac.uk
University of Edinburgh

Keywords: computational geometry, design of experiments, probabilistic specification, verification, robotics

1 Introduction

Imagine you have designed an algorithm for a tabletop robotic arm. You had a *task specification* (e.g., move an object from the in-tray to an out-tray). You likely also had an *environment specification*—a description of the range of environments the robot should perform successfully in (e.g., “robot at table’s centre, two trays at any position on the table, object weighs 500-1000g and always starts in the in-tray”). Across all potential variations in position, size, and weight, you now wish to assess your algorithm’s performance (e.g., with an STL robustness metric based on your task specification [8, 11, 15]).

Checking all possible environments is usually impossible: one robotic experiment, even simulated, takes non-trivial time; We cannot check all values for continuous variables like object position without first discretizing them; assessing all interaction effects between parameters requires an exponential combination of values. Realistically, we have time to try only a finite subset of the possible environments.

Given a budget of N experiments, how can we pick the best subset of possible values which will give us best coverage of the possible range of environments? We could draw N environment configurations at random; Frameworks like Scenic [6] or ProbRobScene [9] provide a way to automatically generate samplers for declarative descriptions of an environment specification. However, there is no guarantee a random sample of configurations will best represent the range of parameter values, nor interactions between parameters—random samples tend to have *high discrepancy* [7].

To generate *low-discrepancy (uniform)* sample sets, there exist established methods in *experiment design*. Such methods generate sample sets (*designs*) on the d -dimensional unit-hypercube. However, in robot experiments like the one above, the parameter space is not a hypercube, but a constrained geometric region. Generating low-discrepancy samplings




Figure 1. Tabletop scene generated from PRS spec (4b).

over such regions often requires knowledge of special properties of the space [14]. Further, large environment specifications may require a design over not just one, but multiple (possibly dependent) regions. Typical users lack the expertise to construct such complex designs, despite being able to express their environment specification declaratively.

This paper augments the probabilistic programming language ProbRobScene (PRS) [9] with a module for automatically synthesizing uniform designs. It builds on methods for constructing uniform designs over irregular domains [18], dependency resolution [6], and algorithms for convex geometry [9]. We also include a case-study which uses snippets of an environment specification from a tabletop scenario. Our experiments show that our designs achieve lower discrepancy than random-sampling for a range of N . We also show how the computation time varies with N and d .

2 Uniform Designs and Environment Specs

We first review uniform designs on a unit hypercube, and the PRS language. We explain our method for synthesizing designs from arbitrary PRS specifications in the next section.




Figure 2. Visualization of $N = 50$ samples in $[0, 1]^2$ for Random vs GLP

```
t = Table on V3D(0,0,0)
r1 = Robot on (top back t)
Cube completely on t, with mass (500, 1000)

(a) "Cube-Table"

t = Table on V3D(0,0,0)
r1 = Robot on (top back t)
tr_1 = Tray completely on t, ahead of r1, left of t
Cube completely on tr_1

(b) "Cube-Tray-Table"
```

Figure 3. Examples of PRS environment specifications

2.1 Uniform Designs on the Unit HyperCube

Say we have a budget for N experiments, $\mathcal{X}_d^N = \{x^{(i)} \in D \mid i = 0 \dots N\}$ where our domain D is defined by d parameters in the unit range (i.e., $D = [0, 1]^d$). We wish to pick an \mathcal{X}_d^N with *low discrepancy* over D . *Discrepancy* measures how much \mathcal{X}_d^N differs from the ideal uniform design—one where, in every subspace $\Delta D \subset [0, 1]^d$, the proportion of points in ΔD is proportional to the size of ΔD [7].

There exist standard methods to generate low discrepancy unit hyper-cube designs. For example the *Good Lattice Point Method* (GLP) [17] leverages the relative primes of N to produce reliably low-discrepancy designs. Figure (2) shows random sampling versus GLP on the unit-square.

2.2 PRS Specs as Dependent Convex Regions

Probabilistic programming languages can be used declaratively specify environment variations in domains such as autonomous vehicles [6], image generation [10], and procedural modelling [13]. Here, we focus on PRS, which is well suited to setting up tabletop manipulation environments. Figure (3) shows two examples of PRS environment specifications. PRS leverages the *specifier syntax* (from Scenic [6]) to let the user describe the value ranges for object properties in three ways. First, they can provide a scalar (e.g., set the z -coordinate of t to 0). Second, they can provide a uniform range (e.g., setting mass of Cube to between 500 and 1000).

Third, they can define range of values for a property via a *geometric relation* (e.g., tr_1 is *ahead of* $r1$).

PRS generates samples from this specification in several steps. First, PRS converts each geometric relation into a convex set. Since some properties are defined as the combination of multiple relations, this may also involve using algorithms for intersecting convex polytopes. For example, in (3b), the tray position is intersection of the table surface, and the two half-spaces in front of the robot and to the left of the table centre. Next, PRS determines which order to evaluate the specifiers, as some depend on each other. Again in figure (3b): Cube’s position depends on the size and position of tr_1 , which depends on the size and position of t . Finally, once the dependency order and sampling regions are defined, PRS draws samples from the constrained space using hit-and-run—a uniform sampling method for convex polytopes [16].

Using this framework, users can declaratively specify the range of environments they expect their robot to operate in, and automatically generate samples which conform to this specification. However, as discussed in the previous section, generating arbitrary samples from a space tends to result in a high-discrepancy designs, and is therefore typically a poor representation of the performance of the system. We now show how we extend this system such to automatically synthesize uniform designs.

3 Design Synthesis from PRS Specifications

Recall our setup—we have N experiments to generate a design \mathcal{X}_d^N with low discrepancy over D . However, now D is not a unit hypercube $[0, 1]^d$, but a combination of dependent convex regions. Can we generate a uniform design over this new domain by building on our ability to generate a unit-hypercube design $\mathcal{U}_d^N = \{u^{(i)} \in [0, 1]^d \mid i = 0 \dots N\}$? Recent work in transforming uniform designs gives us an insight: Let $x \in \mathcal{D} \subseteq \mathbb{R}^d$ be a random vector with density:

$$f(x_1, \dots, x_d) = f_1(x_1)f_{2|1}(x_2|x_1) \dots f_{d|1\dots}(x_d|x_1, \dots) \quad (1)$$

If we can write the cumulative distribution functions (CDFs) for each factor of (1) (denoted $F_1, F_{2|1} \dots F_{d,1\dots}$), then we can map \mathcal{D} to $[0, 1]^d$:

$$\begin{cases} u_1 = F_1(x_1) \\ u_j = F_{j|1\dots j-1}(x_j|x_1, \dots, x_{j-1}), j = 2, \dots, d \end{cases} \quad (2)$$

Given (2), we can convert \mathcal{U}_d^N to a uniform design \mathcal{X}_d^N over D using the *Inverse Rosenblatt Transformation* [18]:

$$\begin{cases} x_1^{(i)} = F_1^{-1}(u_1^{(i)}) \\ x_j^{(i)} = F_{j|1\dots j-1}^{-1}(u_j^{(i)} \mid x_1^{(i)} \dots x_{j-1}^{(i)}), j = 2, \dots, d \end{cases} \quad (3)$$

To construct the CDFs for our specification, recall that each geometric relation in PRS can be represented by a convex-set. Each region can therefore be written compactly as:

$$R_x = \{x \in \mathbb{R}^k \mid Ax + b \leq 0\} \quad (4)$$

$A \in \mathbb{R}^{m \times k}$ and $b \in \mathbb{R}^m$ are the coefficients of m convex inequalities. Using (4) we can construct a CDF for a uniform distribution over a convex region by integrating over R_x . For a 3-dimensional convex polyhedron, F_1 would be given by:

$$F_1(x_1^{(i)}) = \iiint_{x_1^{(i)}} \frac{1}{V(R_x)} dx_1 dx_2 dx_3 \quad (5)$$

Where $V(R)$ is the volume of region R . The conditional CDFs such as $F_{2|1}(x_2^{(i)} | x_1^{(i)})$ can be computed similarly by first projecting R_x onto the relevant hyper-plane. While we lack a closed form representation for the inverse CDFs required to calculate (3), we can approximate each transformation from $u^{(i)}$ to $x^{(i)}$ numerically using Brent's method [1].

We can now generate uniform designs over convex regions. To synthesize a uniform design over a full PRS specification, we have two further difficulties.

First, PRS environment specifications are actually comprised of several convex regions, whose boundaries depend on one another (e.g., recall from Figure (3b)— the Cube position depends on tr_1 , which depends Table). We must therefore construct a tree of these dependencies, similar to the process followed when sampling. This dependency tree limits on the possible factorings of (1) and, by extension, the orders we can condition the CDFs in (2, 3).

The second difficulty is that our design is dependent on the order we condition each CDF dimension — $x_1^{(i)} = F_1^{-1}(u_1^{(i)})$, $x_2^{(i)} = F_{2|1}^{-1}(u_2^{(i)} | x_1^{(i)})$ produces a different design from $x_2^{(i)} = F_2^{-1}(u_2^{(i)})$, $x_1^{(i)} = F_{1|2}^{-1}(u_1^{(i)} | x_2^{(i)})$. To decide the best, we pick the \mathcal{X}_d^N with minimal *central composite discrepancy* (CCD) [3]:

$$\text{CCD}(\mathcal{X}_d^N) = \left\{ \frac{2^{-d}}{V(\mathcal{D})} \int_{\mathcal{D}} \sum_{k=1}^{2^d} \left| \frac{|\mathcal{X}_d^N \cap \mathcal{D}_k(z)|}{N} - \frac{V(\mathcal{D}_k(z))}{V(\mathcal{D})} \right|^2 dz \right\}^{\frac{1}{2}} \quad (6)$$

Here, $\mathcal{D}_k(z)$ is k -th section if we partition \mathcal{D} into 2^d sections, splitting across every axis from centre z . Intuitively, (6) measures the difference between the proportion of points in each sub-region, and the sub-region's relative volume.

Algorithm (1) summarizes the full process. First, the PRS specification is converted into individual convex regions, and an object property dependency graph is extracted. Next, we loop through viable dimension orderings (respecting restrictions imposed by the dependency graph), generate candidate design via (3), then choose the design with minimal CCD.

Algorithm 1 Uniform Design from PRS Specification

```

1: function SYNTHESIZE-DESIGN(prs-spec, N)
2:   viable-orders ← EXTRACT-DEPENDENCIES(prs-spec)
3:    $\mathcal{D}$  ← CONVERT-TO-CONVEX-REGIONS(prs-spec)
4:    $\mathcal{U}_d^N$  ← GOOD-LATTICE-POINT(N, d)
5:   for d-order ∈ viable-orders do
6:     designsi ← INVERSE-ROSENBLATT( $\mathcal{U}_d^N$ , d-order)
7:     CCDi ← CCD(cand-designsi)
8:   return  $\mathcal{X}_d^N$  in designs with lowest CCD

```

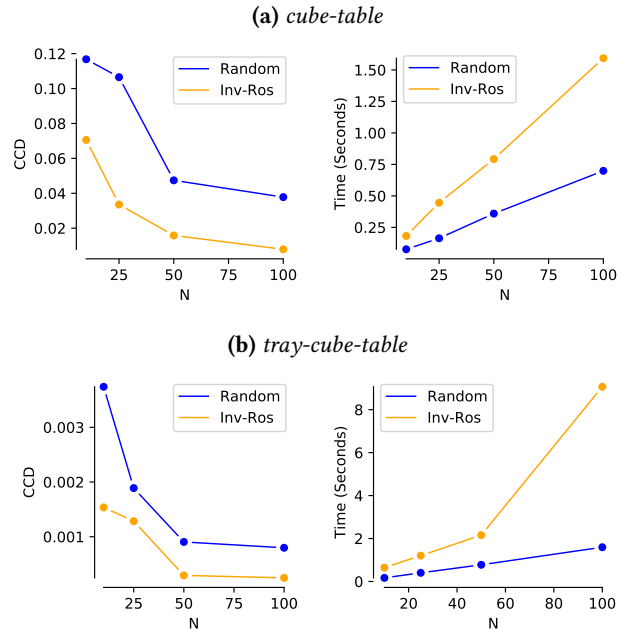


Figure 4. CCDs and times for $N = \{10, 25, 50, 100\}$

4 Case Study: Tabletop Setups

This section demonstrates our augmented PRS module on two environment specification snippets from tabletop manipulation examples (Figures (3a) and (3b)). These cover the single convex region case and the dependent region case.

Figure (4) shows the CCD and computation time for our method for $N = \{10, 25, 50, 100\}$. Our system achieves reliably low-discrepancy designs at all values of N , with lower discrepancies as N increases. Contrast this with random sampling, which has consistently higher discrepancies.

One thing to note in the change from (4a) to (4b) is the change in relative computation time. While random sampling remains quick and scales linearly with N , our method appears to increase both in absolute time and relative growth. The reason for this is that the bulk of the computation resides in the inverse-rosenblatt transformation, and the number of transformations to consider grows factorially with d . In

practice, the computation time taken to generate the experimental designs is usually completely dominated by the time taken to actually run the experiments. However, for complex specifications with large d , heuristics for searching the space of possible transformations may have to be used to preserve scalability.

5 Related Work and Discussion

Most existing probabilistic programming languages offer only the ability to draw Monte-Carlo samples [6, 10, 13]. VerifAI [4] allows users to manually configure external quasi-random samplers (e.g., Halton sampling [12]) which tend to achieve lower discrepancies. None offer the ability to synthesize uniform designs directly from specifications.

Our module builds off a generic technique for synthesizing designs using Inverse Rosenblatt transforms [18], but there are also specialized techniques based on particle swarms [2], and identifying specific sub-classes of convex polyhedrons [14]. Such specialized methods tend not to generalize to arbitrary specifications, but could be used as pruning heuristics to improve computational efficiency in larger specifications.

Our synthesis approach can be classified as *system-free*, as it favours low discrepancy designs without reference to a task-specific metric. However, in the field of robot manipulation, there exist numerous metrics for evaluating task success with respect to a formal task-specification [8, 11, 15]. One direction for future work is to use such metrics to drive *adaptive experiment design*. Such methods choose samples based on uncertainty over predicting the task metric [5].

Acknowledgments

This work is supported in part by funding from the Alan Turing Institute, as part of the Safe AI for Surgical Assistance project, and from EPSRC for the UKRI Research Node on Trustworthy Autonomous Systems Governance and Regulation (EP/V026607/1).

References

- [1] Richard P Brent. 2013. *Algorithms for minimization without derivatives*. Courier Corporation.
- [2] Ray-Bing Chen, YH Shu, Ying Hung, and Weichung Wang. 2014. Central composite discrepancy-based uniform designs for irregular experimental regions. *Computational Statistics and Data Analysis* 72 (2014), 282–297.
- [3] SC Chuang and YC Hung. 2010. Uniform design over general input domains with applications to target region estimation in computer experiments. *Computational Statistics & Data Analysis* 54, 1 (2010), 219–232.
- [4] Tommaso Dreossi, Daniel J Fremont, Shromona Ghosh, Edward Kim, Hadi Ravanbakhsh, Marcell Vazquez-Chanlatte, and Sanjit A Seshia. 2019. Verifai: A toolkit for the formal design and analysis of artificial intelligence-based systems. In *International Conference on Computer Aided Verification*. Springer, 432–442.
- [5] John Eason and Selen Cremaschi. 2014. Adaptive sequential sampling for surrogate model generation with artificial neural networks. *Computers & Chemical Engineering* 68 (2014), 220–232.
- [6] Daniel Fremont, Xiangyu Yue, Tommaso Dreossi, Shromona Ghosh, Alberto L Sangiovanni-Vincentelli, and Sanjit A Seshia. 2018. Scenic: Language-based scene generation. *arXiv preprint arXiv:1809.09310* (2018).
- [7] Sushant S Garud, Iftekhar A Karimi, and Markus Kraft. 2017. Design of computer experiments: A review. *Computers & Chemical Engineering* 106 (2017), 71–95.
- [8] Iman Haghighi, Noushin Mehdipour, Ezio Bartocci, and Calin Belta. 2019. Control from signal temporal logic specifications with smooth cumulative quantitative semantics. In *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 4361–4366.
- [9] Craig Innes and Subramanian Ramamoorthy. 2020. ProbRobScene: A Probabilistic Specification Language for 3D Robotic Manipulation Environments. *arXiv preprint arXiv:2011.01126* (2020).
- [10] Tejas D Kulkarni, Pushmeet Kohli, Joshua B Tenenbaum, and Vikash Mansinghka. 2015. Picture: A probabilistic programming language for scene perception. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4390–4399.
- [11] Noushin Mehdipour, Cristian-Ioan Vasile, and Calin Belta. 2019. Arithmetic-geometric mean robustness for control from signal temporal logic specifications. In *2019 American Control Conference (ACC)*. IEEE, 1690–1695.
- [12] Harald Niederreiter. 1992. *Random number generation and quasi-Monte Carlo methods*. SIAM.
- [13] Daniel Ritchie. 2014. Quicksand: A Lightweight Embedding of Probabilistic Programming for Procedural Modeling and Design. In *3rd NIPS Workshop on Probabilistic Programming*. 164.
- [14] Guo-Liang Tian, Hong-Bin Fang, Ming Tan, Hong Qin, and Man-Lai Tang. 2009. Uniform distributions in a class of convex polyhedrons with applications to drug combination studies. *Journal of multivariate analysis* 100, 8 (2009), 1854–1865.
- [15] Peter Varnai and Dimos V Dimarogonas. 2020. On robustness metrics for learning STL tasks. In *2020 American Control Conference (ACC)*. IEEE, 5394–5399.
- [16] Zeld B. Zabinsky and Robert L. Smith. 2013. *Hit-and-Run Methods*. Springer US, Boston, MA, 721–729. https://doi.org/10.1007/978-1-4419-1153-7_1145
- [17] Stanisław K Zaremba. 1966. Good lattice points, discrepancy, and numerical integration. *Annali di matematica pura ed applicata* 73, 1 (1966), 293–317.
- [18] Mei Zhang, Aijun Zhang, and Yongdao Zhou. 2020. Construction of Uniform Designs on Arbitrary Domains by Inverse Rosenblatt Transformation. In *Contemporary Experimental Design, Multivariate Analysis and Data Mining*. Springer, 111–126.