Ergodic Trajectory Optimization at Scale Using Continuous-Space Representations

Christian Sulana   Christian Hughes   Dr. Ian Abraham

Yale ENGINEERING

Coverage-Based Planning on Continuous Occupancy Maps

Motivation

We want robots to fully explore spaces by themselves.

Ergodic exploration methods have been shown to excel at achieving coverage in small to medium-sized spaces [1].

Existing ergodic exploration approaches optimize trajectories using spatial data, such as domain bounds or points that discretize the search space.

However, in large environments where domain bounds are unknown, discrete space representations (e.g., points) require massive amounts of data to maintain high resolution.

Continuous-space representations offer infinite spatial resolution in arbitrarily-sized spaces, but many existing methods lack mathematical tractability for gradient-based optimization methods [3][4].

Challenge: How do we perform coverage-based exploration in large spaces without sacrificing computational speed or resolution?

Approach

Replace cross-similarity term of MMD for proportionality between a discrete trajectory and a utility-biased Hilbert Map [1].

New Objective Function:

$$\mathcal{E}_\mu(x) = \frac{1}{T^2}\sum_{t=0}^{T-1}\sum_{t'=0}^{T-1} k(x_t, x_{t'}) - \boxed{\frac{2}{T}\bar{W}_{HM}\sum_{t=0}^{T-1}\sum_{j=0}^{N_c-1} \mathcal{J}(c_j)k(x_i, c_j)}$$

New Cross Similarity

T = Trajectory length
$g(x_t)$: feature representation of x_t
c : Inducing points of Hilbert map
$\bar{W}_{HM}$: Absolute value of Hilbert map weights
x : Trajectory point
k(,) : Kernel function
N_c : Number of clusters
$\mathcal{J}(\circ)$: Normalized utility function

Hilbert maps have smooth gradients → suitable for optimization!

Output of Hilbert maps is probability measure: easily translatable for MMD.

References

[1] R. Senanayake and F. Ramos, 'Bayesian Hilbert Maps for Dynamic Continuous Occupancy Mapping', in *Proceedings of the 1st Annual Conference on Robot Learning*, 13--15 Nov 2017, vol. 78, pp. 458–471.
[2] A. Mavrommati, E. Tzorakoleftherakis, I. Abraham, and T. D. Murphey, 'Real-Time Area Coverage and Target Localization using Receding-Horizon Ergodic Exploration', *arXiv [cs.RO]*. 2017.
[3] G. Bahati, R. M. Bena, and A. D. Ames, 'Dynamic Safety in Complex Environments: Synthesizing Safety Filters with Poisson's Equation', *arXiv [cs.RO]*. 2025.
[4] D. Verbin, P. Hedman, B. Mildenhall, T. Zickler, J. T. Barron, and P. P. Srinivasan, 'Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields', *arXiv [cs.CV]*. 2021.

Ergodic Exploration

Ergodic exploration methods seek to proportionally align a robot's visitation time to the utility of a given region.

Over infinite time, ergodic methods guarantee full exploration of all reachable space.

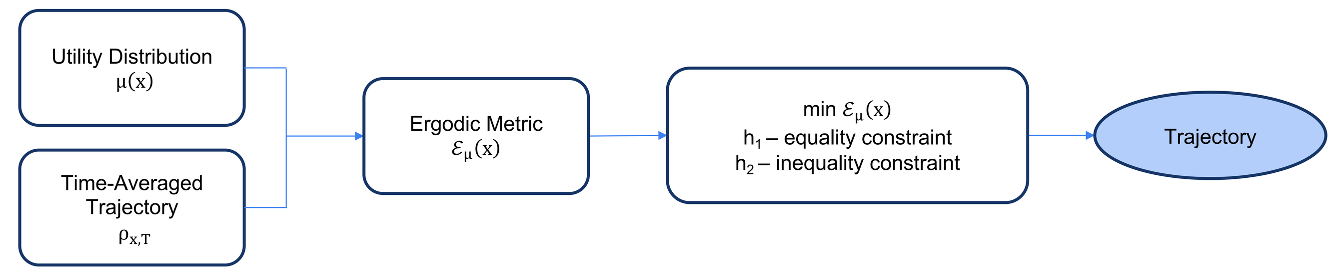

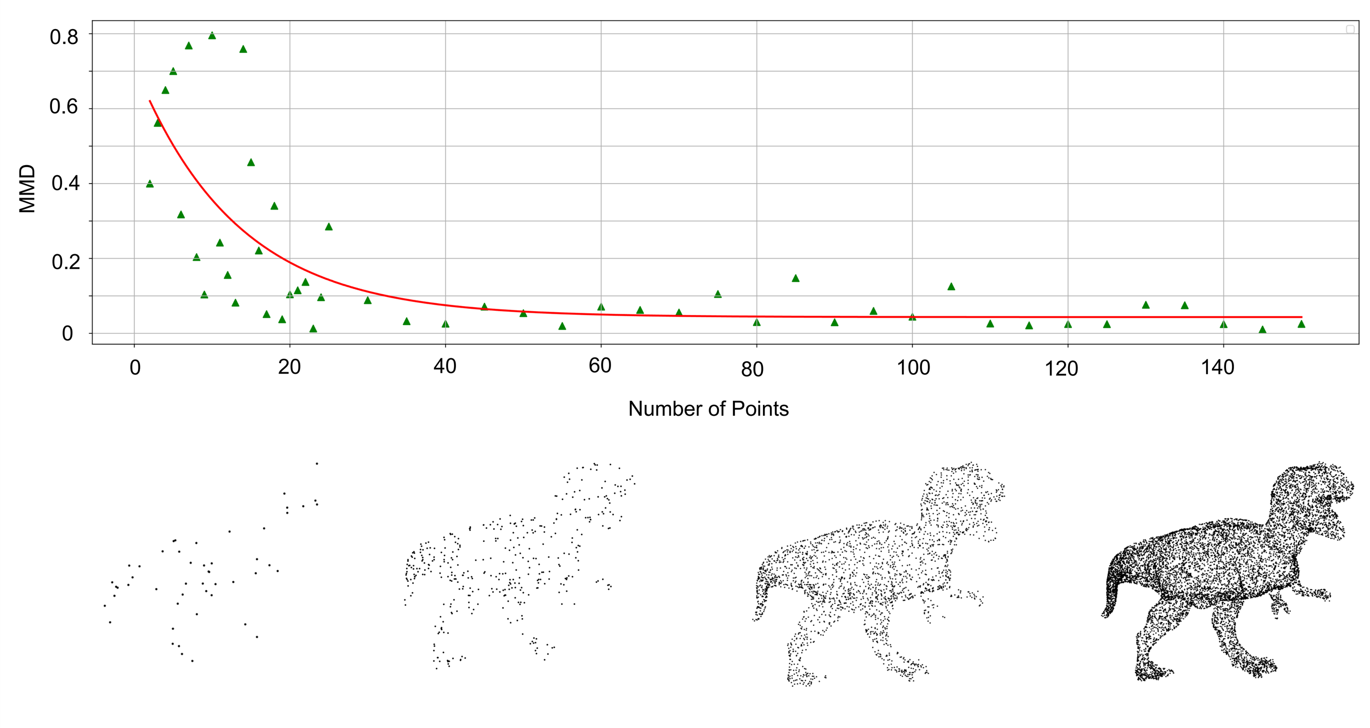

Fig. 1: Impact of Feature Representation on Exploration Quality. Effective exploration relies on accurate representations of the complex features within a space. Discrete space representations have a **minimum threshold of resolution** at which all features are adequately represented (e.g., 8 points to represent a cube), and beyond this point, exploration that respects spatial geometry is not guaranteed.

Ergodic Maximum Mean Discrepancy

Maximum mean discrepancy (**MMD**) is a kernel-based (i.e. transform-based) statistical test that measures the difference between two distributions.

Ergodic MMD generates ergodic trajectories by minimizing the MMD between our visitation and the utility of visiting different places.

$$\mathcal{E}_\mu(x) = \mathbf{MMD}_k^2\left(\rho_{x,T}, \mu(x)\right) =$$

$$\underbrace{\frac{1}{T^2}\sum_{t=0}^{T-1}\sum_{t'=0}^{T-1} k(g\circ x_t, g\circ x_{t'})}_{\text{Self-Similarity}} - \underbrace{\frac{2}{TM}\sum_{t=0}^{T-1}\sum_{j=1}^{M} k(g\circ x_t, \omega_j)}_{\text{Cross-Similarity}} + \frac{1}{M^2}\sum_{i=1}^{M}\sum_{j=1}^{M} k(\omega_j, \omega_{j'})$$

$\rho_{x,T}$: *Time-Averaged Trajectory* $\mu(x)$: *Utility Distribution* $g(\cdot): \chi \rightarrow \Omega$ Ω : *Arbitrary Compact Domain*

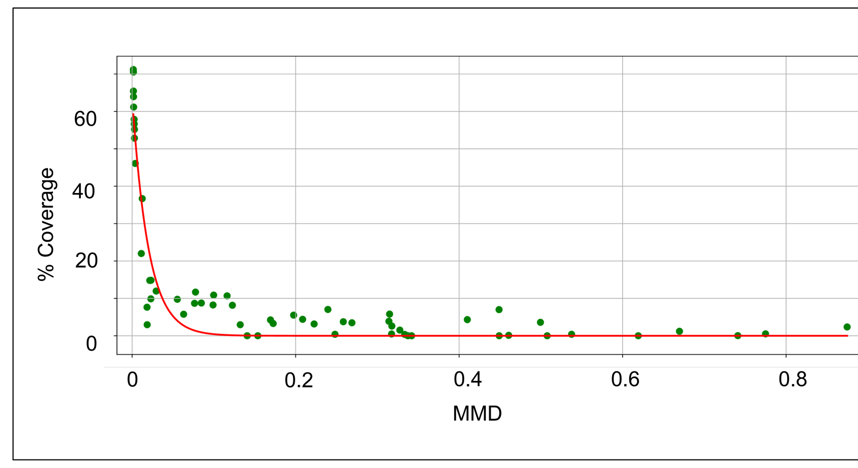

Fig. 2: Validation of MMD as Coverage Metric. Coverage is measured as all points within an arbitrary footprint during robot exploration. MMD has an inverse relationship with coverage quality **(low MMD, high coverage)**.

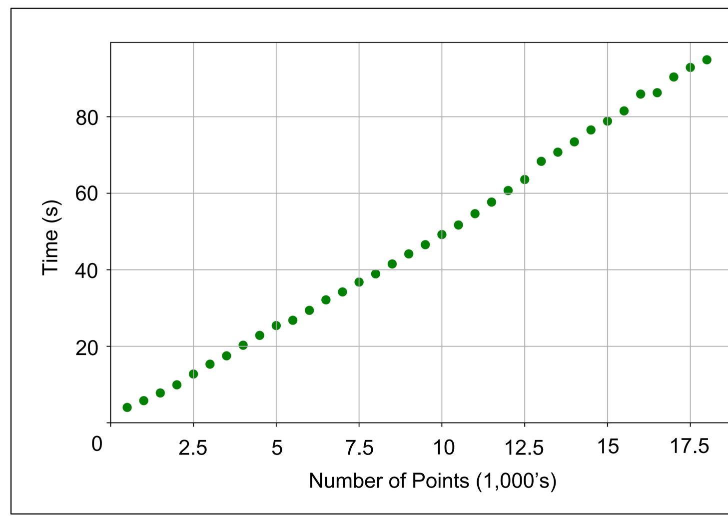

Fig. 3: Computational Scaling of Discrete-Space Representations. Computational analysis is performed using ergodic MMD trajectory optimization with static initialization settings. Optimization time has a direct relationship with the number of points that represent the search space: **more points means more computation**.

Continuous Space Representations

These methods represent spaces by fitting points to mathematical representations with infinite resolution.

Signed Distance Functions (SDF's):
Distance from a point to a surface of an object. Negative if inside. Positive if outside.
- Suffers from sharp gradients [3]

Neural Radiance Fields (NeRFs):
Volumetric maps using neural networks to create representation of space.
- Noisy, non-local gradients [4]

Hilbert Maps:
Mapping method that calculates probability of surface presence with infinite resolution. Gaussian process setting kernel function to overlays.
- Retraining requires access to all data in environment

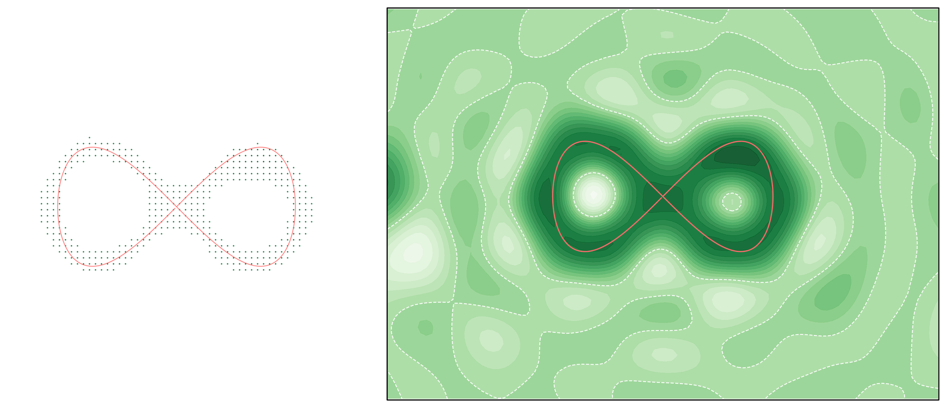

Fig. 4: Hilbert Map Occupancy Contour. Shown right, the points (green) represent where the probability of occupancy is 70% or greater around the infinity symbol. On left, 2D gradient of a Hilbert Map. **The darker green shows a higher concentration of points.**

Results

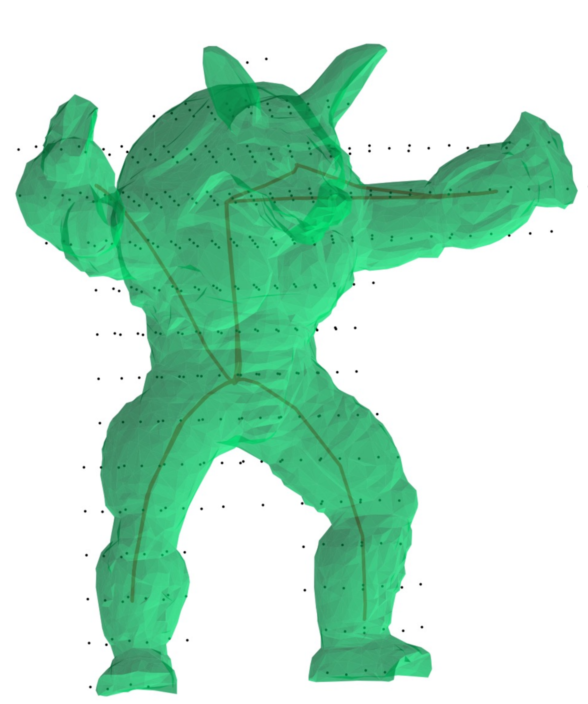

Fig. 6: Ergodic Trajectory Optimization on Reconstructed Hilbert Map. Ergodic exploration inside the surface of the Stanford Armadillo. Trajectories explore various regions of the armadillo given only the parameters of the Hilbert map reconstruction.

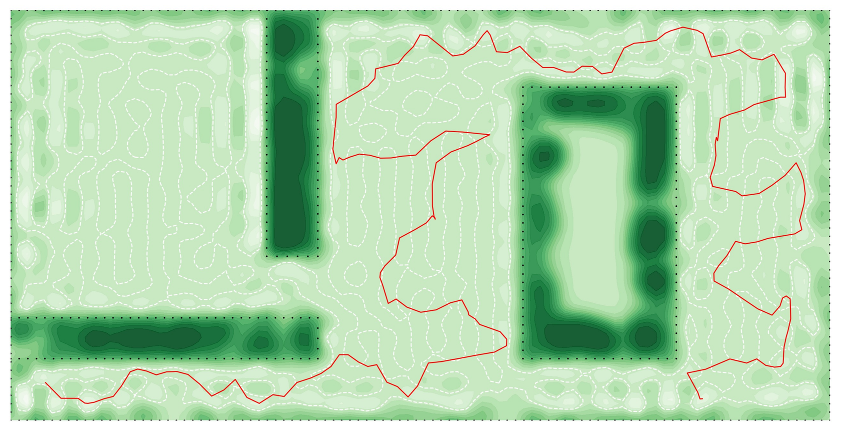

Fig. 5: Application. Given extracted surface data, we can generate trajectories that balance explorative behavior with soft **object avoidance** constraints.

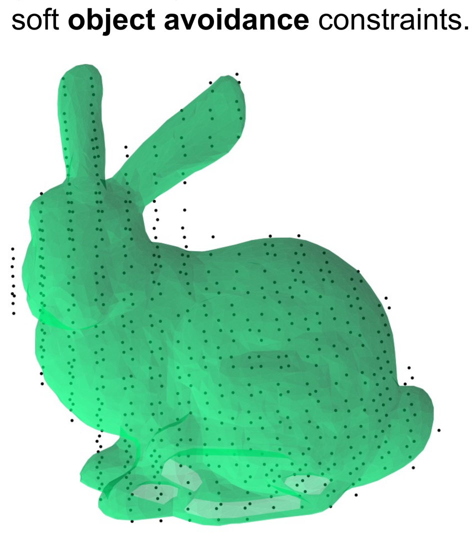

Fig. 7: Hilbert Map Reconstruction of Stanford Bunny Mesh. Hilbert maps provide a measure of the likelihood (probability) that a given location in space is occupied. Surface points (black) represent locations where the probability of occupancy is 70% or greater.

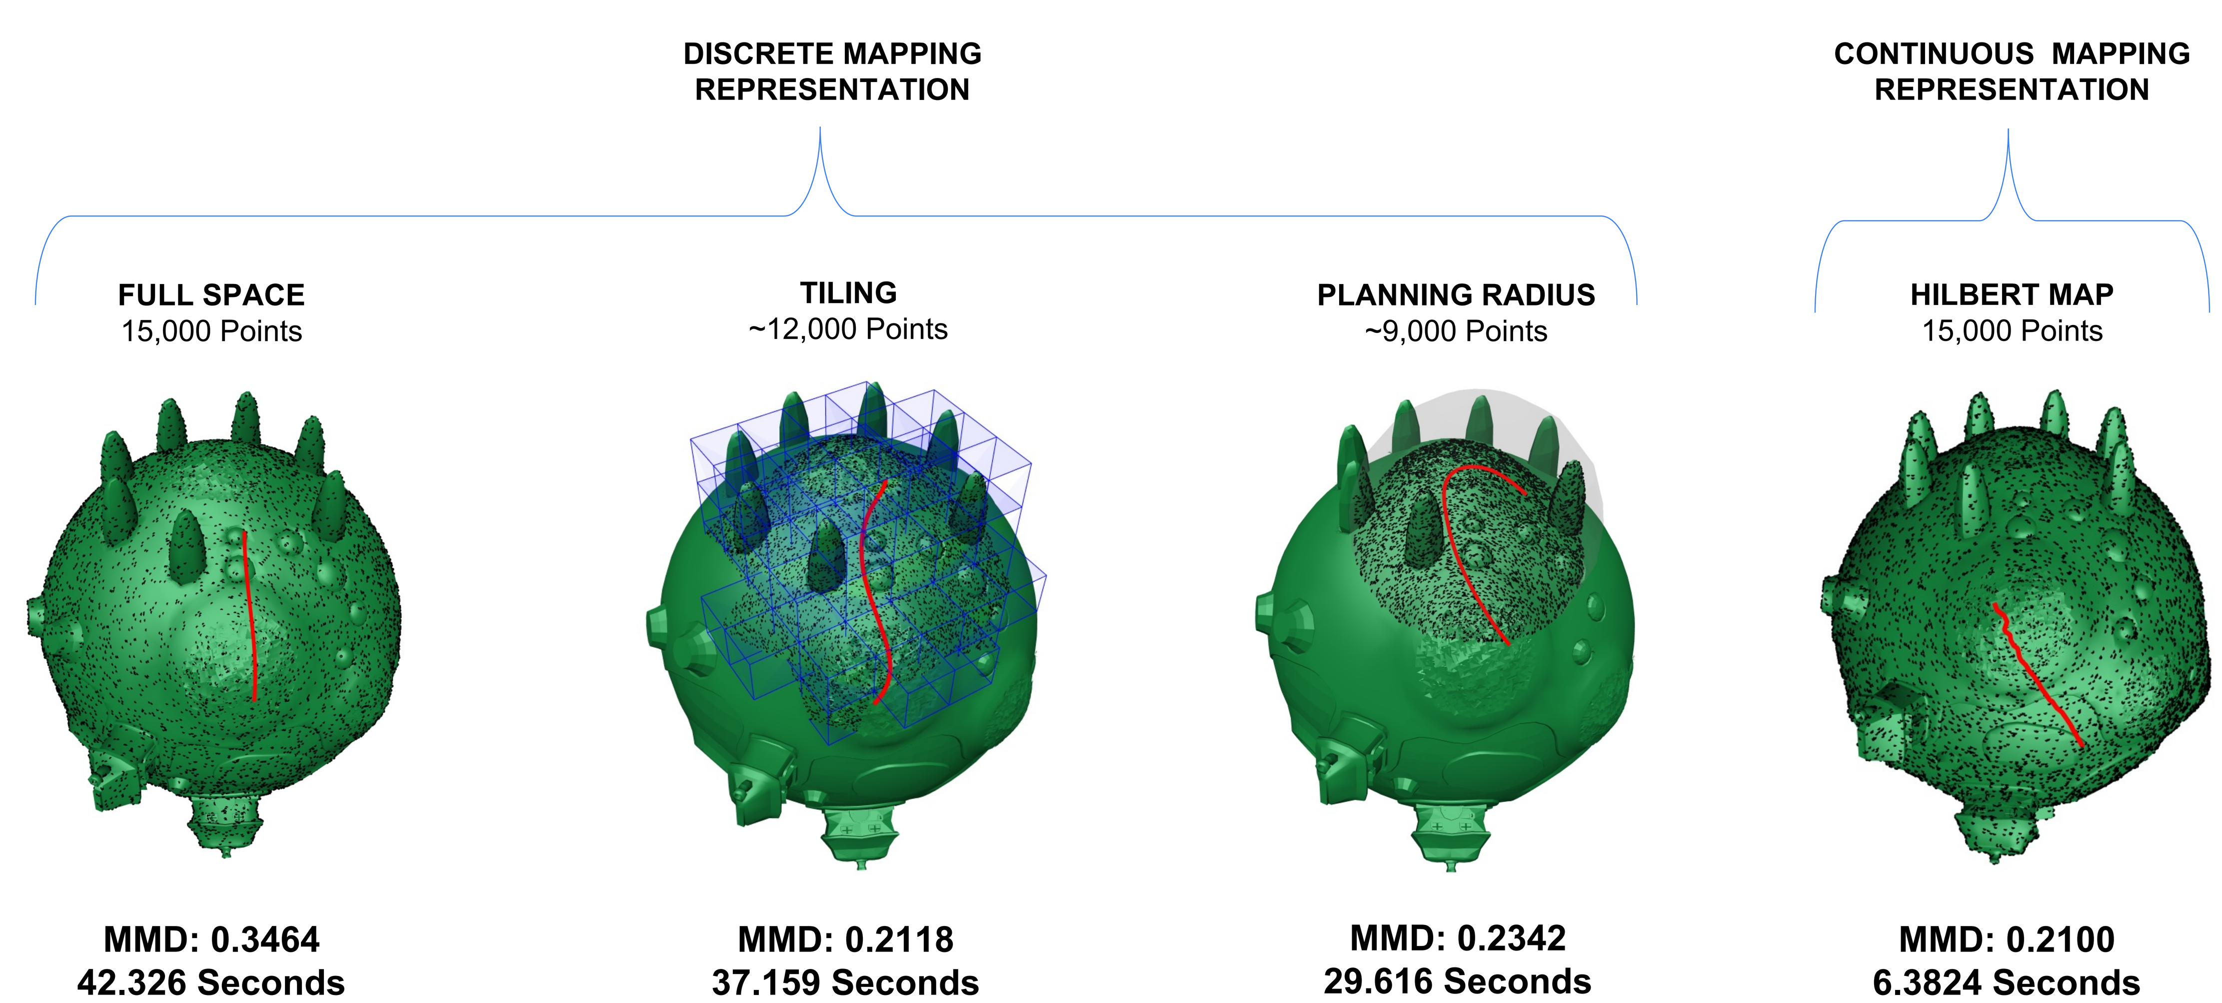

Fig. 8: Comparison of Discrete and Continuous Space Representations on Ergodic Exploration. The computational time required to generate ergodic trajectories is a function of the amount of spatial data considered (see Fig. 3). Discrete space representations that strictly consider reachable regions of the environment can decrease overall computation requirements with negligible loss in exploration quality, but highly complex spaces may still require large data stores to adequately represent spatial geometry. Optimization directly on **Hilbert maps** allows for **similar exploration quality** with significant **decreases in computational requirements**.

Acknowledgements

My mentor, Christian Hughes. My Principal Investigator, Dr. Ian Abraham.
REVU Director, Dr. Marla Geha. REVU Program Manager, Dr. Jeremy Bradford.
The Students at The Intelligent Autonomy Lab at Yale