Discount Model Search for Quality Diversity Optimization in High-Dimensional Measure Spaces

Bryon Tjanaka, Henry Chen, Matthew C. Fontaine, Stefanos Nikolaidis
ICAROS Lab, University of Southern California

Project Page
https://discount-models.github.io

Conference Page
https://iclr.cc/virtual/2026/poster/10007611

Abstract

Quality diversity (QD) optimization searches for a collection of solutions that optimize an objective while attaining diverse outputs of a user-specified, vector-valued measure function. Contemporary QD algorithms are typically limited to low-dimensional measures because high-dimensional measures are prone to distortion, where many solutions found by the QD algorithm map to similar measures. For example, the state-of-the-art CMA-MAE algorithm guides measure space exploration with a histogram in measure space that records so-called discount values. However, CMA-MAE stagnates in domains with high-dimensional measure spaces because solutions with similar measures fall into the same histogram cell and hence receive the same discount value. To address these limitations, we propose Discount Model Search (DMS), which guides exploration with a model that provides a smooth, continuous representation of discount values. In high-dimensional measure spaces, this model enables DMS to distinguish between solutions with similar measures and thus continue exploration. We show that DMS facilitates new capabilities for QD algorithms by introducing two new domains where the measure space is the high-dimensional space of images, which enables users to specify their desired measures by providing a dataset of images rather than hand-designing the measure function. Results in these domains and on high-dimensional benchmarks show that DMS outperforms CMA-MAE and other existing black-box QD algorithms.

Quality Diversity Optimization

Quality diversity (QD) optimization fills an archive with solutions $\boldsymbol{\theta}$ based on measures (measurable characteristics) $\boldsymbol{m}$. The goal is to maximize the objective $\boldsymbol{f}$ in each cell of the archive.

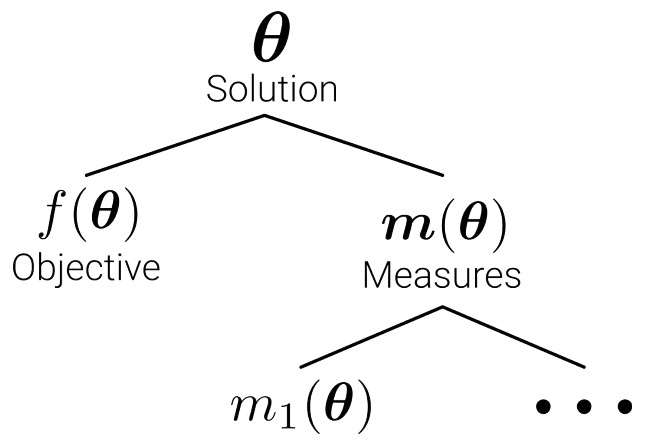

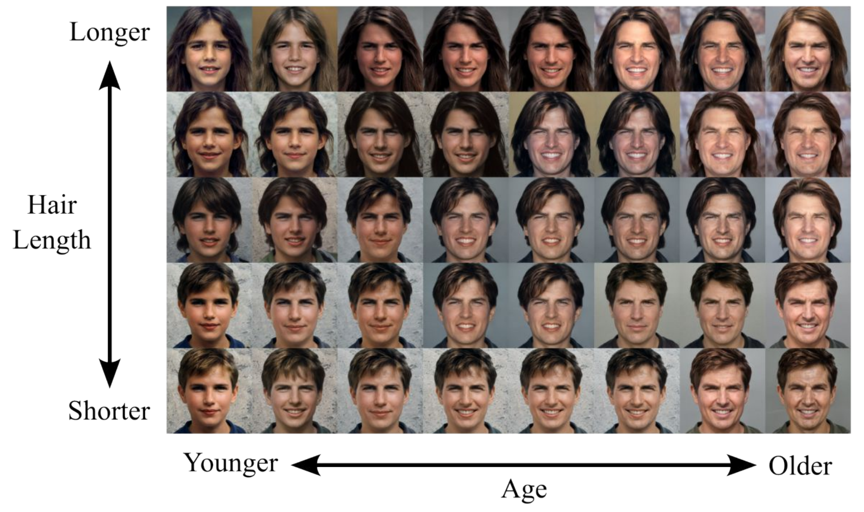

Credits: [1]

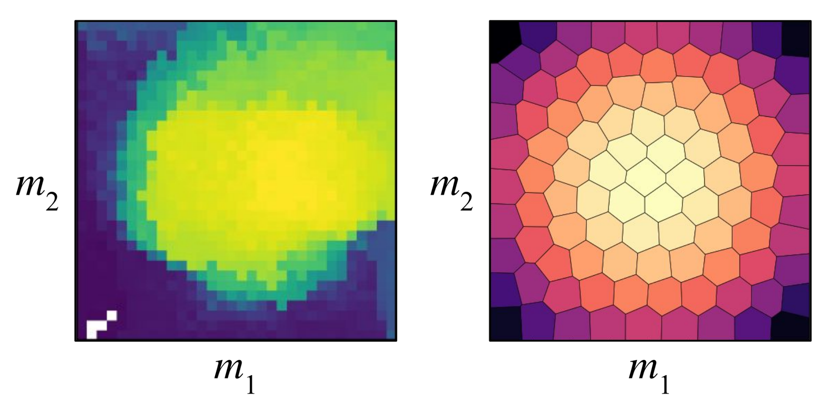

Distortion and High-Dimensional Measure Spaces

Distortion in QD refers to when large areas of solution space map to a small region of measure space [1]. Although distortion exists in low-dimensional measure spaces [2], it is more prominent in high-dimensional measure spaces because there are exponentially larger volumes to which solutions can map.

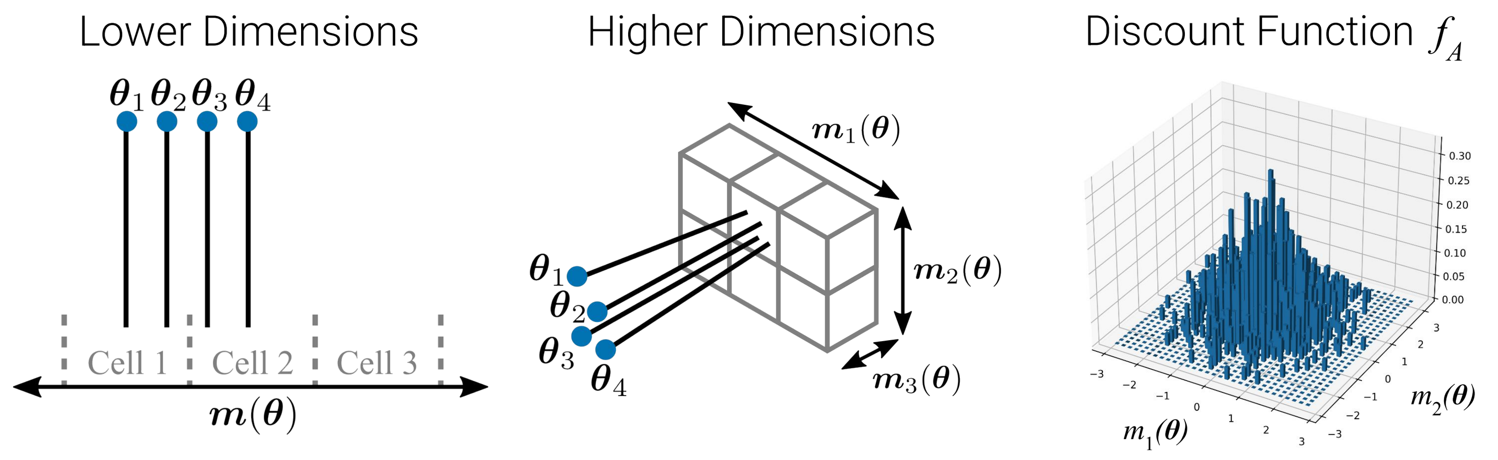

One algorithm affected by distortion is CMA-MAE [1]. CMA-MAE guides search by recording a discount value in each cell of the archive, forming a histogram over the measure space that we term the *discount function*. In high-dimensional measure spaces, distortion makes solutions fall into the same histogram cell and receive the same discount value, making it difficult for CMA-MAE to tell them apart.

Results

We primarily record QD Score, which sums the objectives of all solutions in the archive. DMS outperforms other black-box QD algorithms across domains.

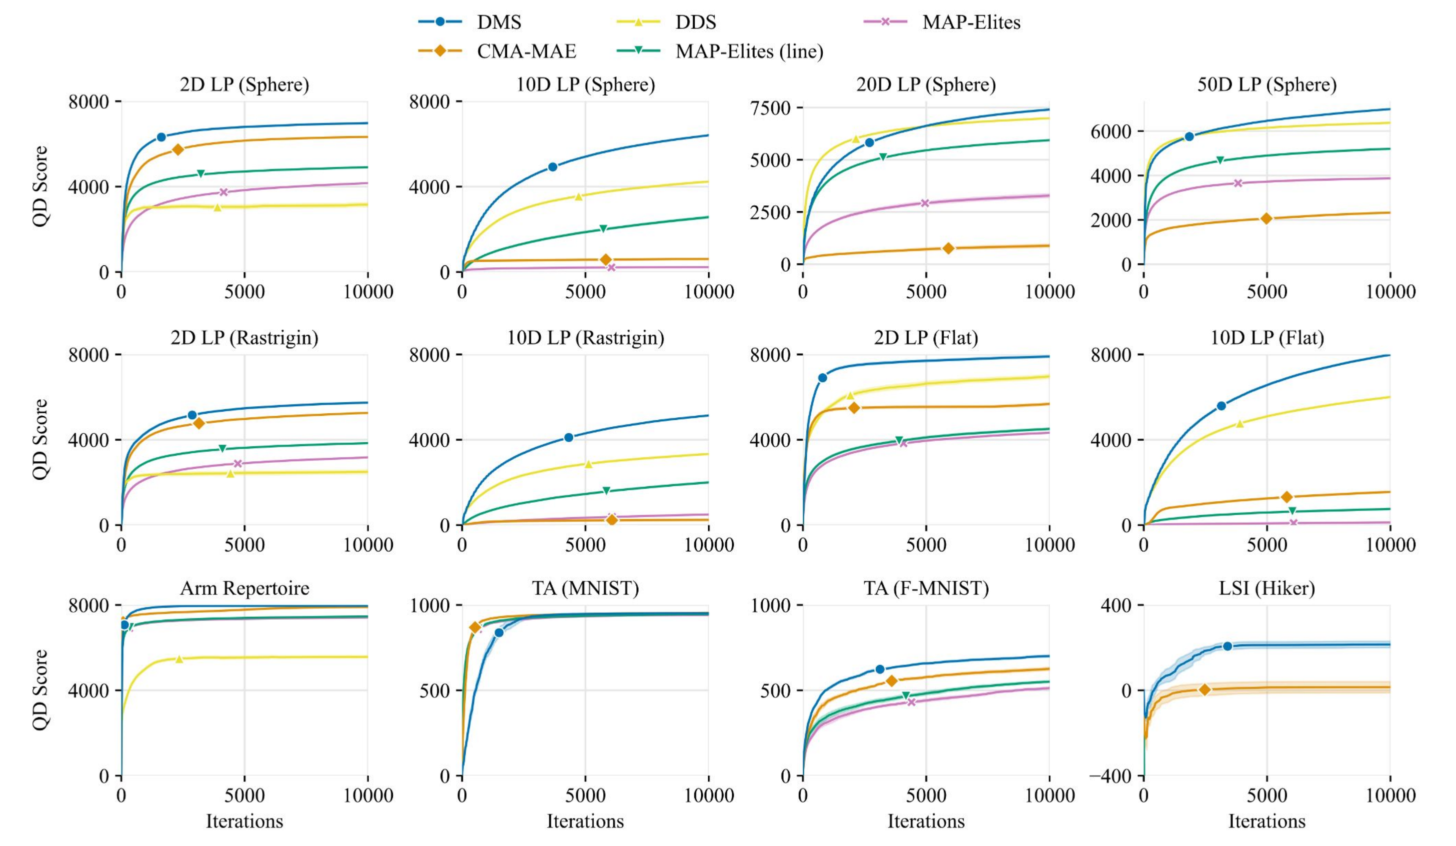

Discount Model Search (DMS)

To improve exploration in domains with distorted, high-dimensional measure spaces, we propose DMS. DMS trains a *discount model* to provide a smooth, continuous representation of the discount function.

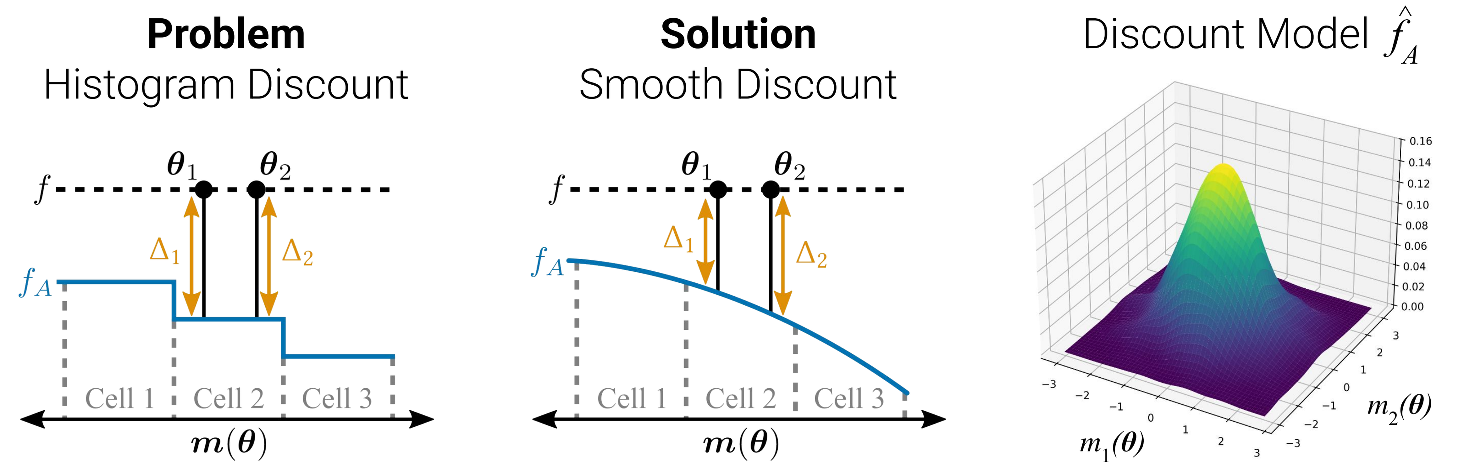

The key insight of DMS is that such a representation provides distinct discount values and improvement values, even when solutions have similar measures, making it easier to guide search towards solutions that improve the archive.

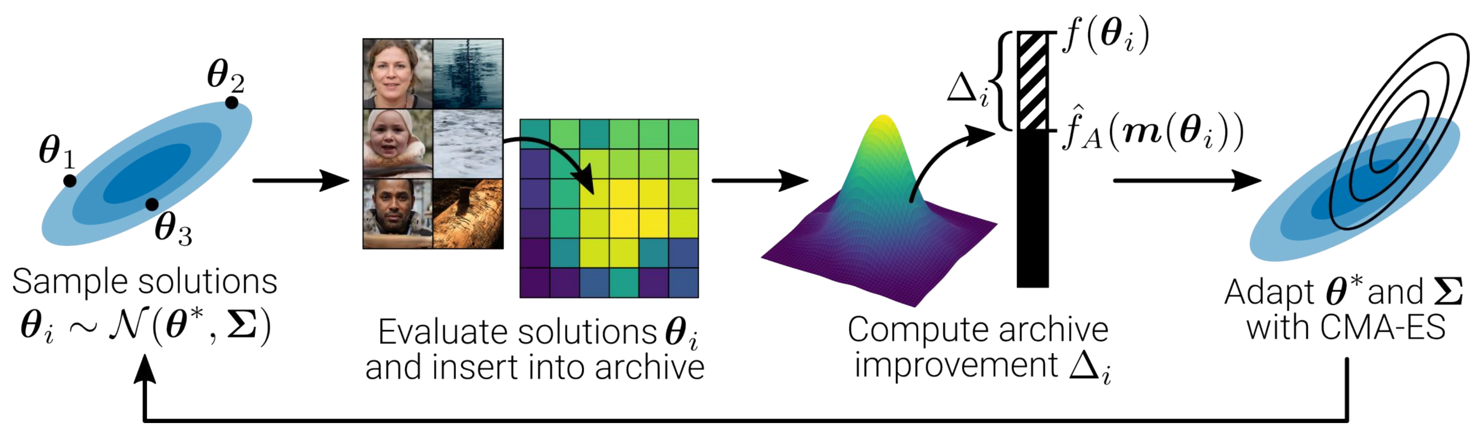

DMS regresses the discount model to match a dataset D_A of input measure values s and their corresponding discount value targets t_A. The dataset entries (s, t_A) come from two sources:

1. Solutions $\boldsymbol{\theta}$ sampled during search, where $s = \boldsymbol{m}(\boldsymbol{\theta})$ and

$$t_A = \begin{cases} \hat{f}_A(\boldsymbol{s}) & \text{if } f(\boldsymbol{\theta}) \leq \hat{f}_A(\boldsymbol{s}) \\ (1-\alpha)\hat{f}_A(\boldsymbol{s}) + \alpha f(\boldsymbol{\theta}) & \text{if } f(\boldsymbol{\theta}) > \hat{f}_A(\boldsymbol{s}) \end{cases}$$

2. *Empty points*, which are sampled among the centers of unoccupied cells in the result archive. Here, t_A is set to the min objective. These points "clamp down" the discount model in unexplored areas.

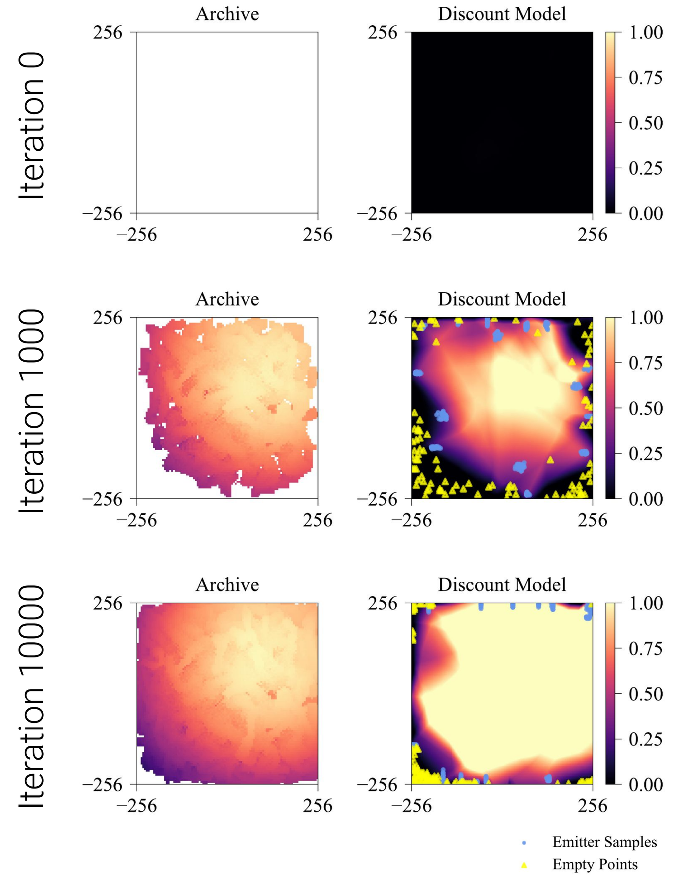

Quality Diversity with Datasets of Measures (QDDM)

DMS further enables *QDDM*, where the measure space is a high-dimensional space, e.g., of images, and desired measures are specified with a dataset. Given the ubiquity of datasets in ML, we are excited for problems that can be framed as QDDM and tackled with DMS.

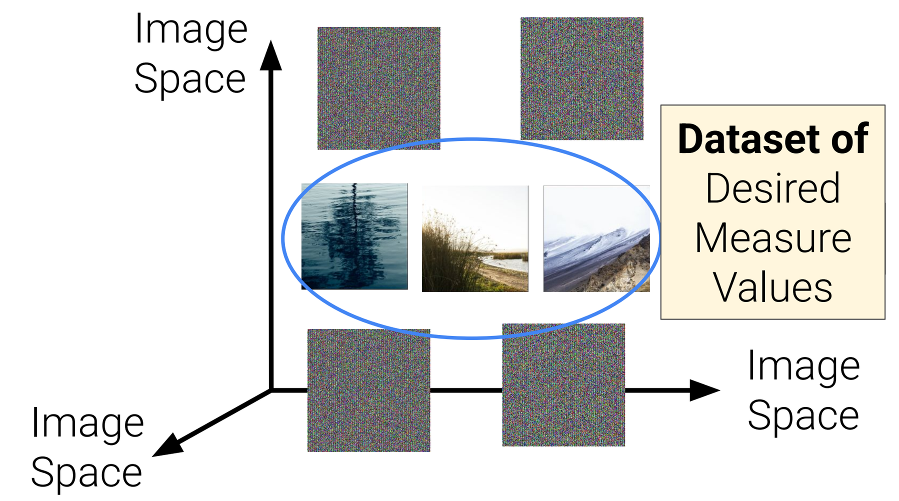

Results from the LSI (Hiker) domain, where the objective is "A photo of the face of a hiker," and the measures are images showing the landscape where the hiker belongs. The desired measures are specified with images from the LHQ dataset.

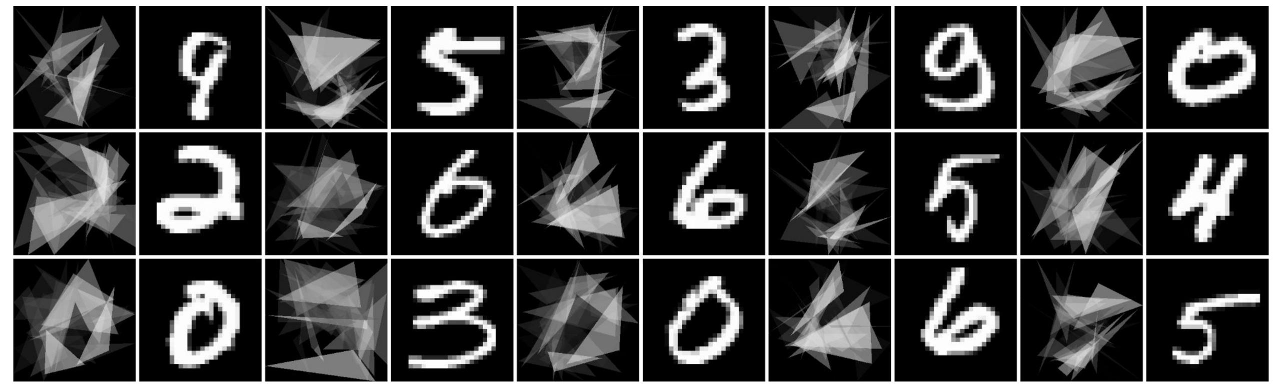

Results from the TA (MNIST) domain, where the objective is to arrange triangles such that they resemble MNIST digits.

References

1. M. C. Fontaine and S. Nikolaidis. "Covariance Matrix Adaptation MAP-Annealing." GECCO 2023.
2. M. C. Fontaine and S. Nikolaidis. "Differentiable Quality Diversity." NeurIPS 2021.
3. B. Tjanaka, M. C. Fontaine, D. H. Lee, Y. Zhang, N. R. Balam, N. Dennler, S. S. Garlanka, N. D. Klapsis, S. Nikolaidis. "pyribs: A Bare-Bones Python Library for Quality Diversity Optimization." GECCO 2023.