Differentiable Quality Diversity

Authors: Matthew C. Fontaine and Stefanos Nikolaidis

USC Viterbi School of Engineering

University of Southern California

Goal: Create an algorithm that …

- Discovers a **diverse** collection of solutions
  - According to measure functions

Goal: Create an algorithm that …

- Discovers a **diverse** collection of solutions
  - According to measure functions
- Where each solution should be high **quality**
  - According to an objective function

Goal: Create an algorithm that …

- Discovers a **diverse** collection of solutions
    - According to measure functions
- Where each solution should be high **quality**
    - According to an objective function
- While efficiently leveraging **gradients** to solve both aspects

StyleGAN+CLIP

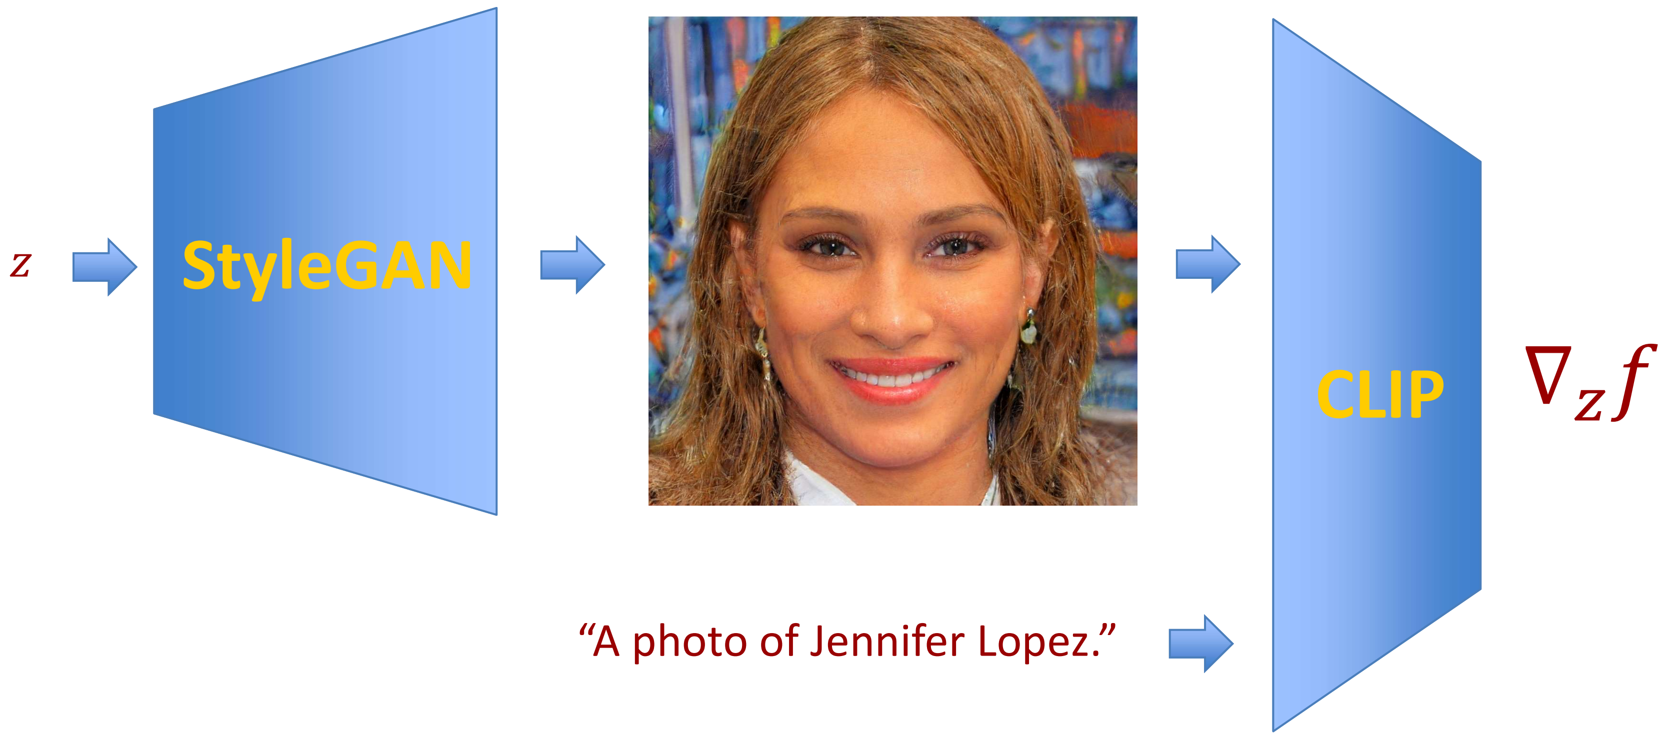

"Generating Images from Prompts using CLIP and StyleGAN" Perez (2021)
"Learning Transferable Visual Models From Natural Language Supervision" Radford et. al. (2021)
"A Style-Based Generator Architecture for Generative Adversarial Networks" Karras et. al. (CVPR 2019)

Our Prompt has Many Solutions

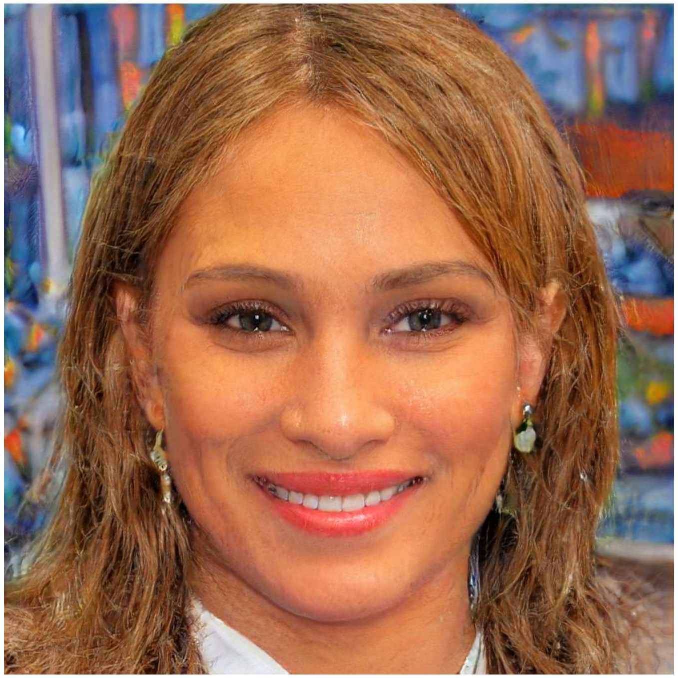

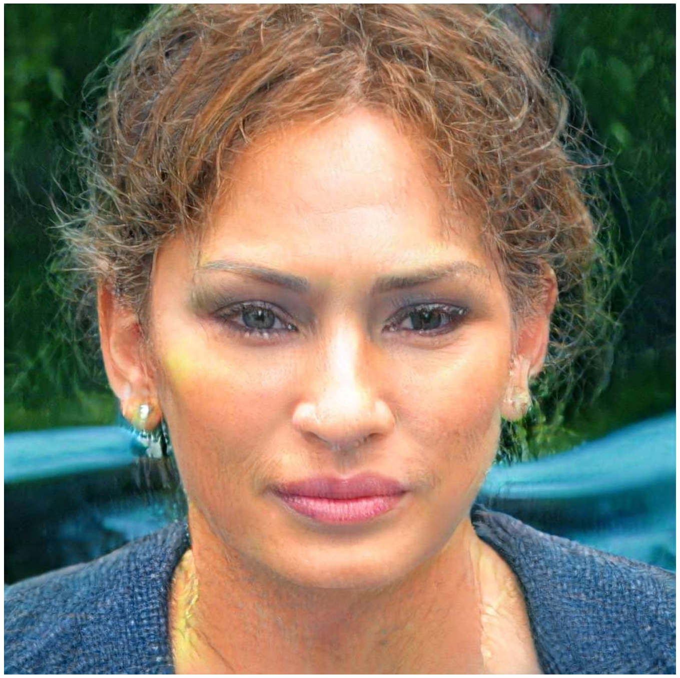

Age as a Spectrum

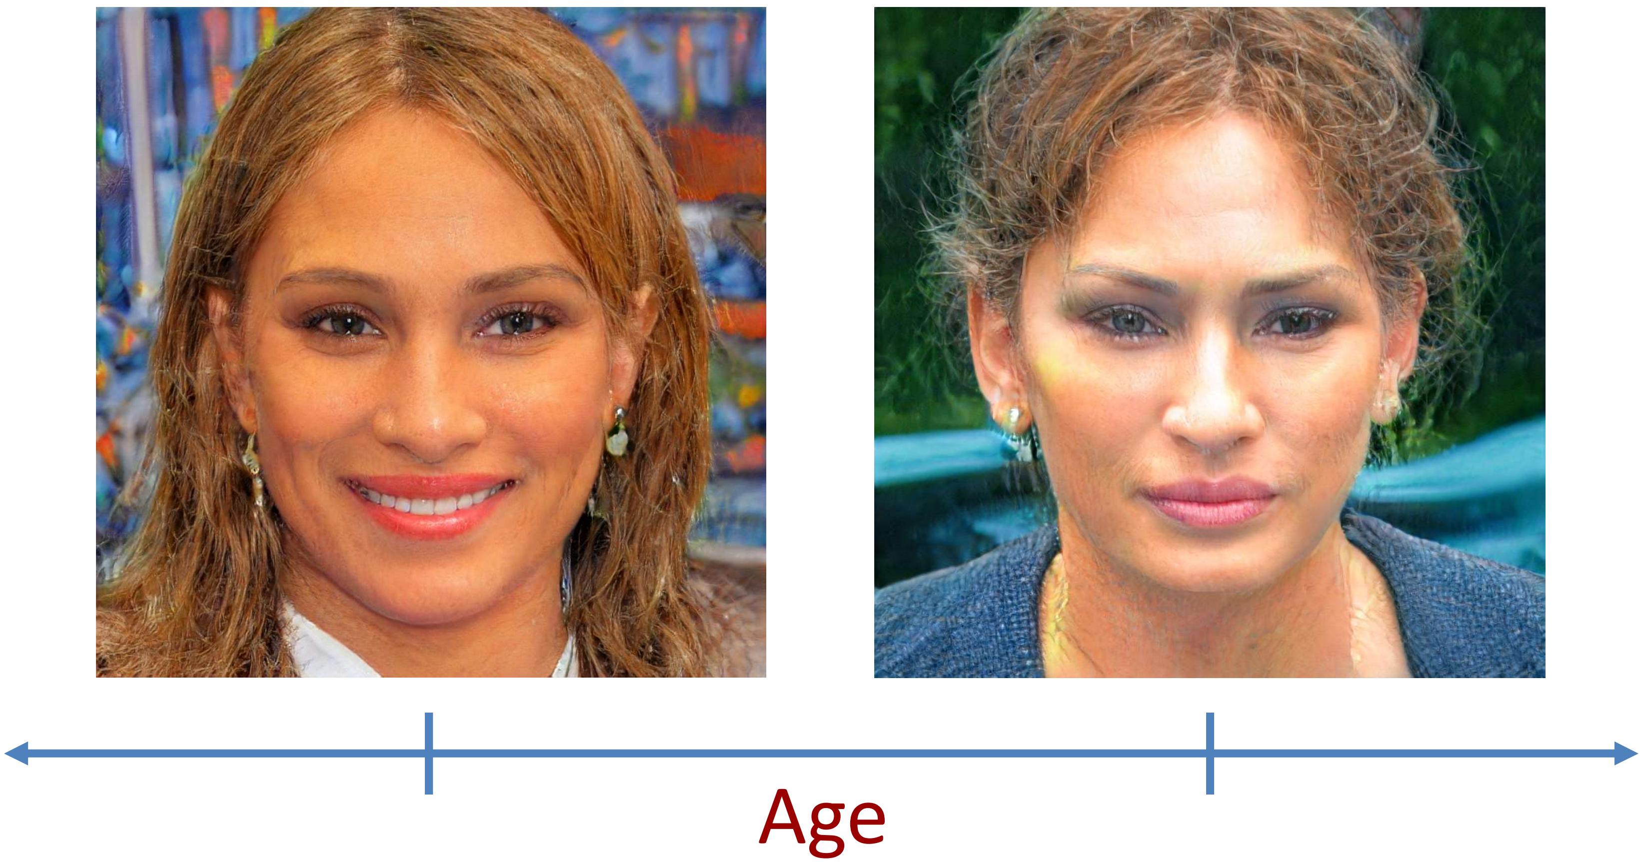

The Quality Diversity Problem

Assume we are optimizing a space $\mathbb{R}^n$.

Given functions $f: \mathbb{R}^n \to \mathbb{R}$ and k measures $m_i: \mathbb{R}^n \to \mathbb{R}$ or as a vector function $m: \mathbb{R}^n \to \mathbb{R}^k$.

“Quality Diversity: A New Frontier for Evolutionary Computation” Pugh et. al. (2016)
“Quality-Diversity Optimization: A Novel Branch of Stochastic Optimization” Chatzilygeroudis et. al. (2020)

The Quality Diversity Problem

Assume we are optimizing a space $\mathbb{R}^n$.

Given functions $f: \mathbb{R}^n \to \mathbb{R}$ and k measures $m_i: \mathbb{R}^n \to \mathbb{R}$ or as a vector function $m: \mathbb{R}^n \to \mathbb{R}^k$.

Measure Space: $S = m(\mathbb{R}^n)$
(approximated via tessellation)

The Quality Diversity Problem

Assume we are optimizing a space $\mathbb{R}^n$.

Given functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$ and k measures $m_i: \mathbb{R}^n \rightarrow \mathbb{R}$ or as a vector function $m: \mathbb{R}^n \rightarrow \mathbb{R}^k$.

Measure Space: $S = m(\mathbb{R}^n)$
(approximated via tessellation)

Objective: For each $s \in S$ find $\theta \in \mathbb{R}^n$ such that $m(\theta) = s$ and $f(\theta)$ is maximized.

The Quality Diversity Problem

Assume we are optimizing a space $\mathbb{R}^n$.

Given functions $f: \mathbb{R}^n \to \mathbb{R}$ and k measures $m_i: \mathbb{R}^n \to \mathbb{R}$ or as a vector function $m: \mathbb{R}^n \to \mathbb{R}^k$.

Measure Space: $S = m(\mathbb{R}^n)$ (approximated via tessellation)

Assumes f and m are black-boxes (derivative-free)

Objective: For each $s \in S$ find $\theta \in \mathbb{R}^n$ such that $m(\theta) = s$ and $f(\theta)$ is maximized.

QD Application: Diverse Agent Behavior

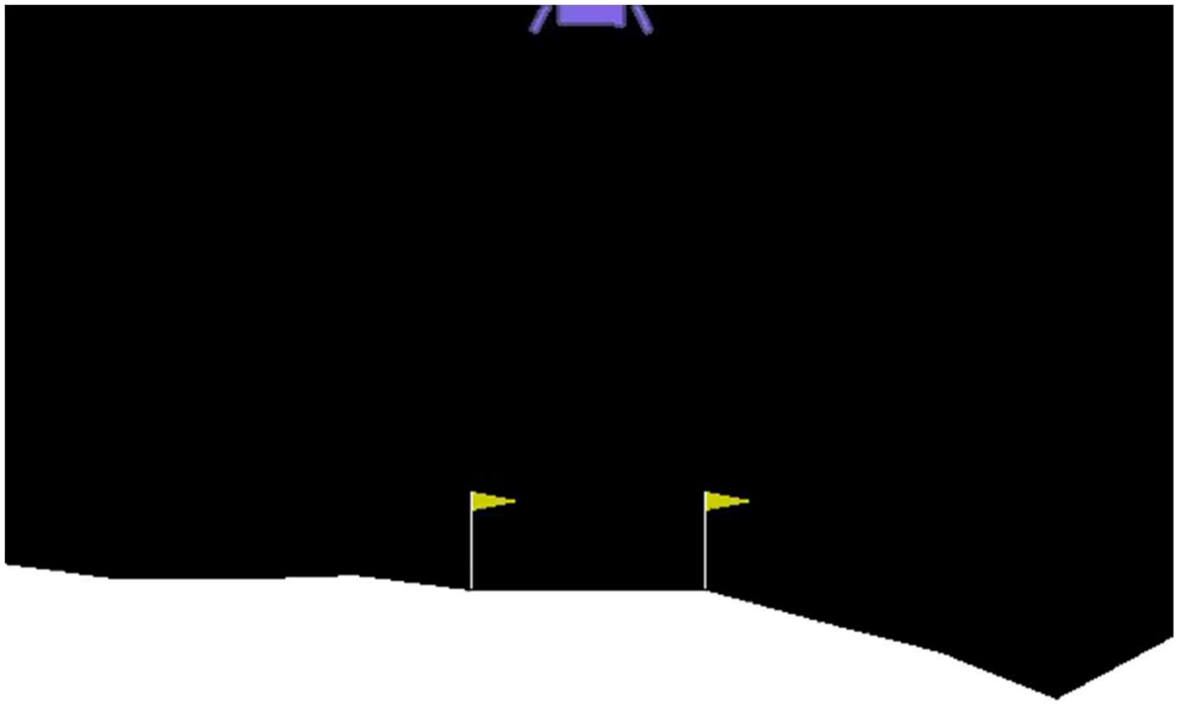

“Using CMA-ME to Land a Lunar Lander Like a Space Shuttle” Tjanaka et. al. (Pyribs.org 2021)

StyleGAN+CLIP

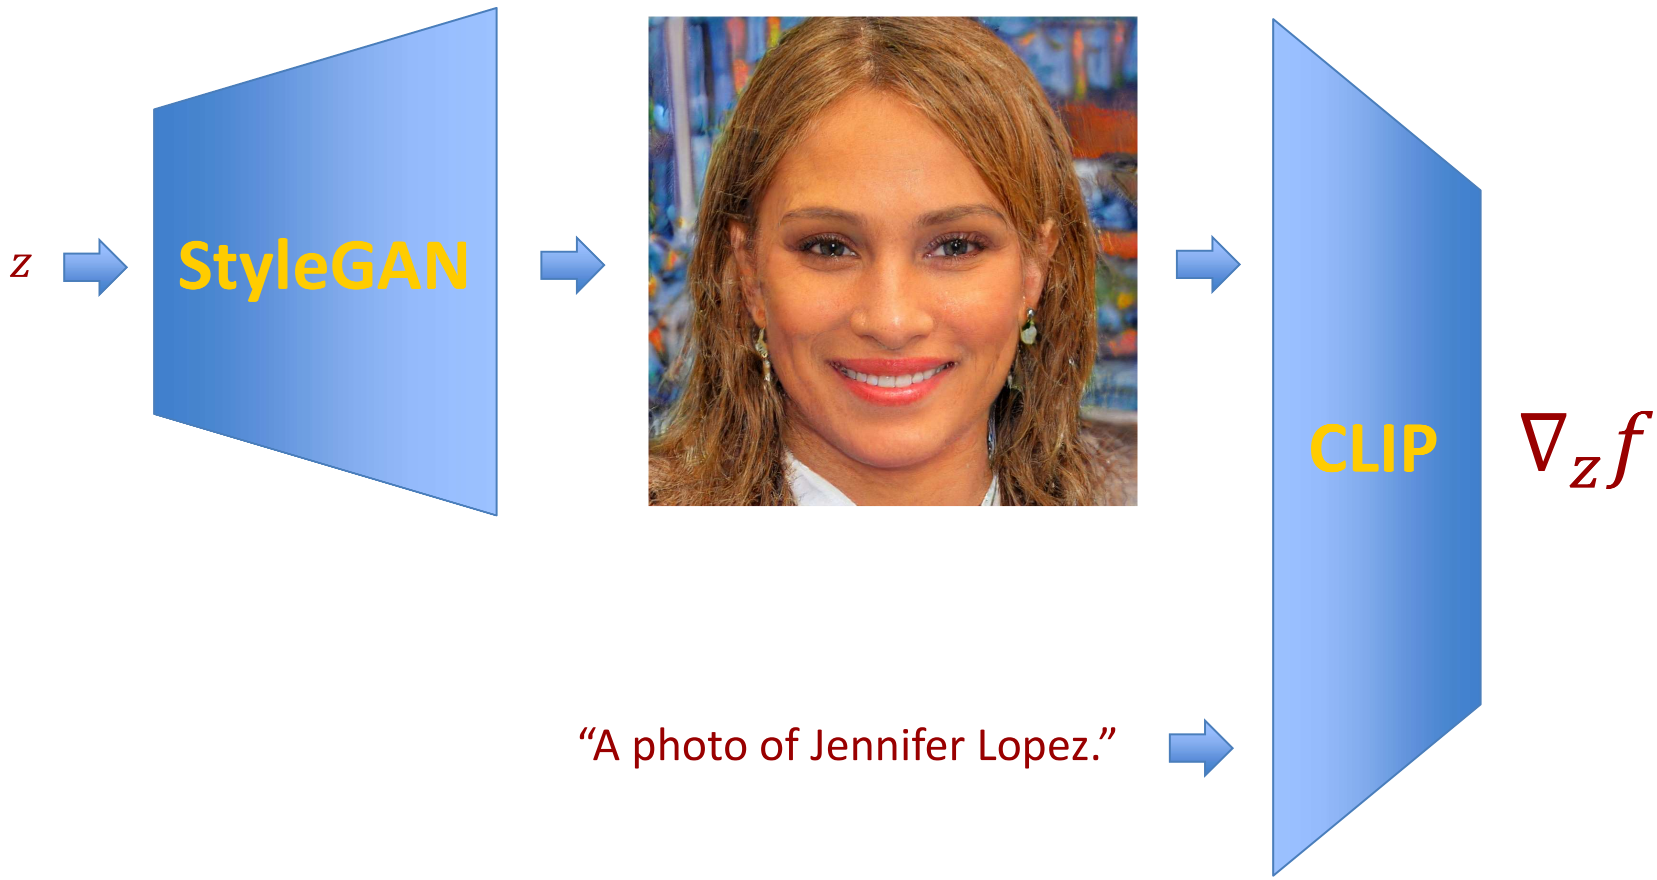

Research Question

How can we leverage gradient information in a quality diversity algorithm?

The Quality Diversity Problem

Assume we are optimizing a space $\mathbb{R}^n$.

Given functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$ and k measures $\mathrm{m_i}: \mathbb{R}^n \rightarrow \mathbb{R}$ or as a vector function $m: \mathbb{R}^n \rightarrow \mathbb{R}^k$.

Measure Space: $S = f(\mathbb{R}^n)$

Objective: For each $s \in S$ find $\theta \in \mathbb{R}^n$ such that $m(\theta) = s$ and $f(\theta)$ is maximized.

The **Differentiable** Quality Diversity Problem

Assume we are optimizing a space $\mathbb{R}^n$.

Given **differentiable** functions $f: \mathbb{R}^n \to \mathbb{R}$ and k measures $\mathrm{m_i}: \mathbb{R}^n \to \mathbb{R}$ or as a vector function $m: \mathbb{R}^n \to \mathbb{R}^k$.

Measure Space: $S = f(\mathbb{R}^n)$

Objective: For each $s \in S$ find $\theta \in \mathbb{R}^n$ such that $m(\theta) = s$ and $f(\theta)$ is maximized.

How do we derive DQD algorithms?

- Adapt ideas from **derivative-free** quality diversity algorithms
  - MAP-Elites
  - CMA-ME

MAP-Elites

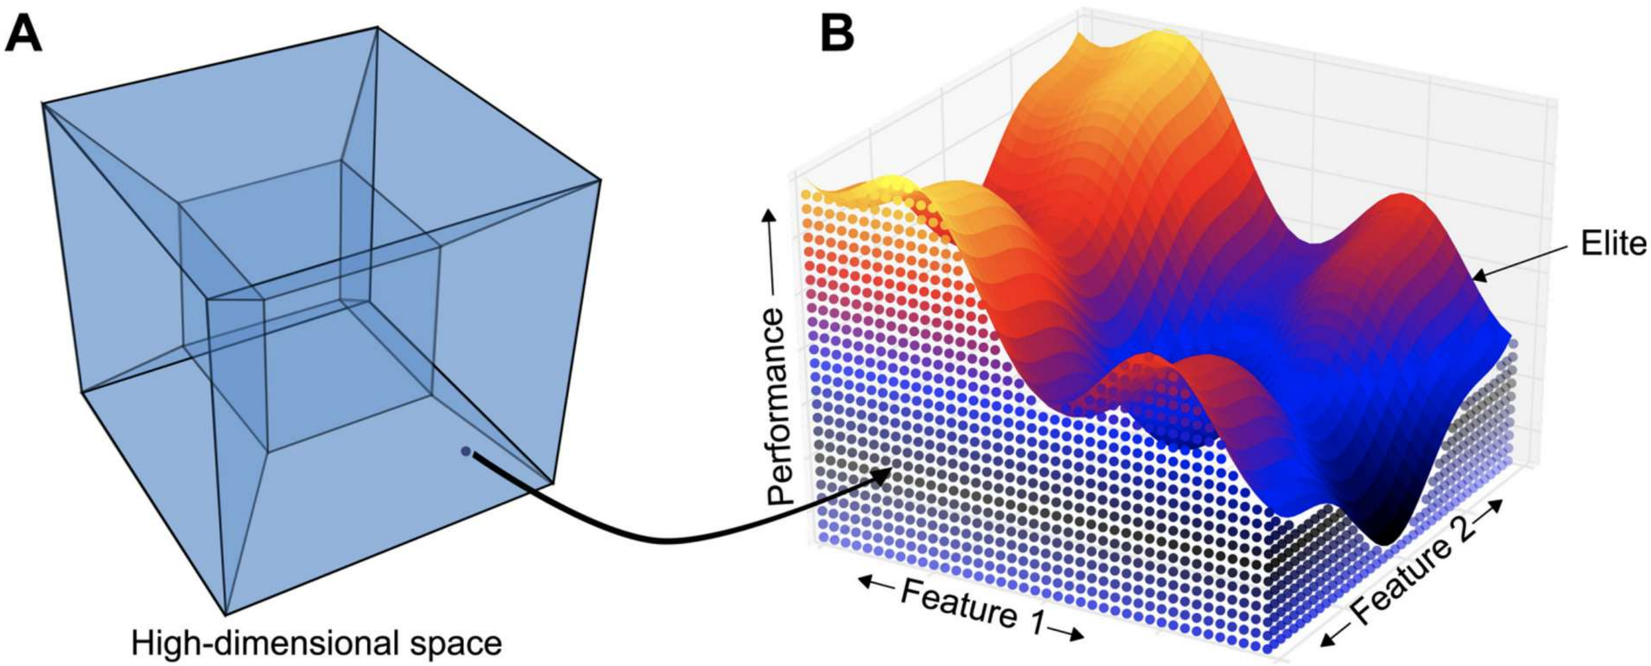

"Illuminating Search Spaces by Mapping Elites" Mouret et. al. (2015)

The MAP-Elites Algorithm

- Sample initial solutions via a fixed distribution
- Uniform random selection of elites
- Perturb solutions with isotropic Gaussian noise

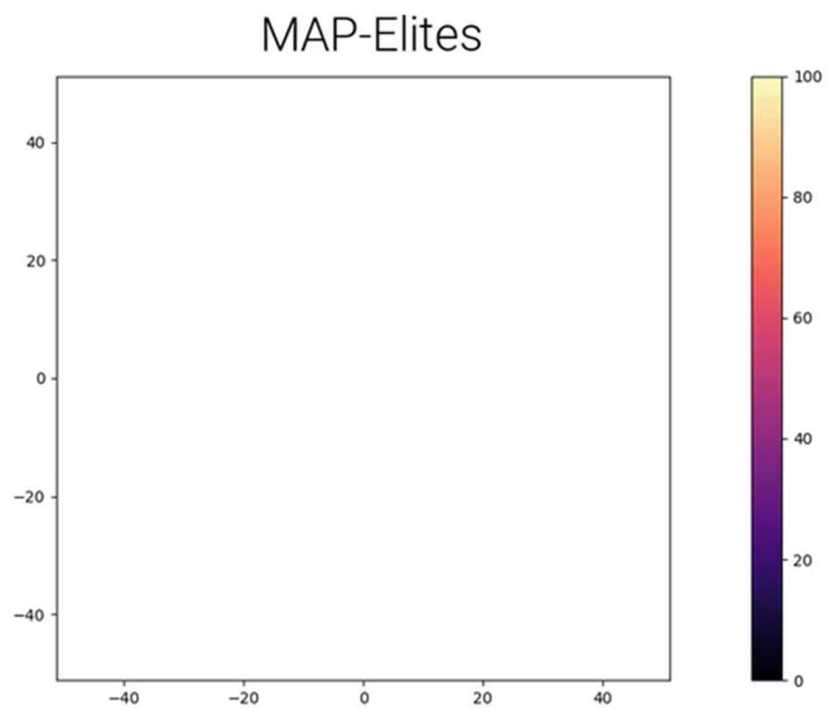

Benchmark Example ($n = 20$)

The MAP-Elites Algorithm

- Sample initial solutions via a fixed distribution
- Uniform random selection of elites
- Perturb solutions with isotropic Gaussian noise

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \sigma N(\mathbf{0}, I)$$

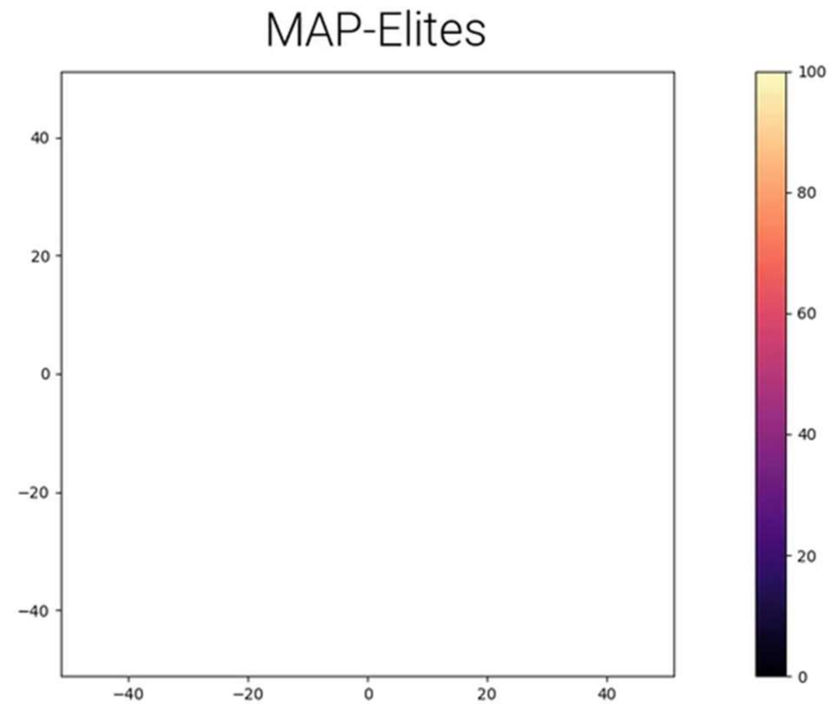

Benchmark Example ($n = 20$)

The MAP-Elites Algorithm

- Sample initial solutions via a fixed distribution
- Uniform random selection of elites
- **Perturb solutions with isotropic Gaussian noise**

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \sigma N(\mathbf{0}, I)$$

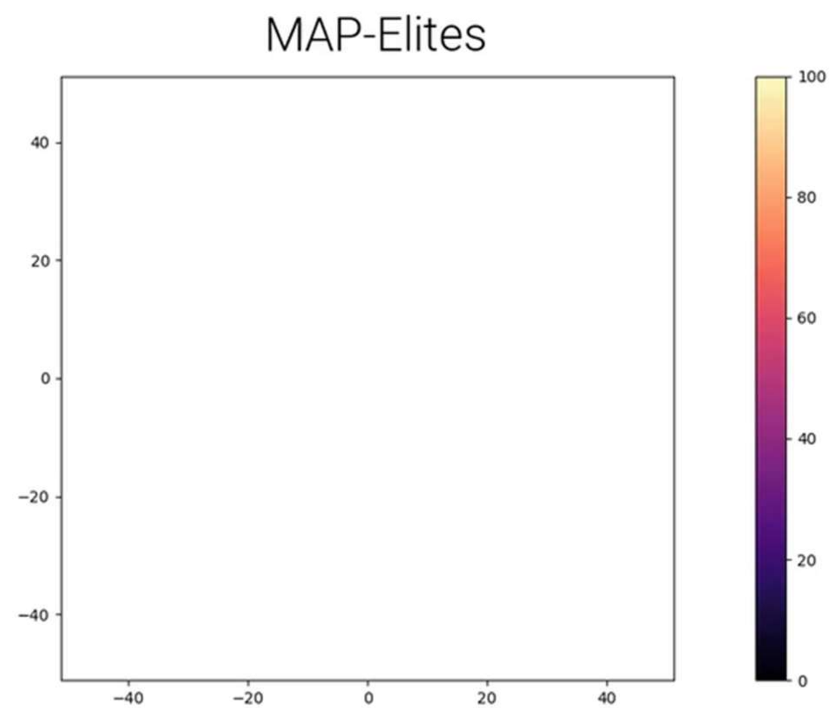

Benchmark Example ($n = 20$)

How do gradients help?

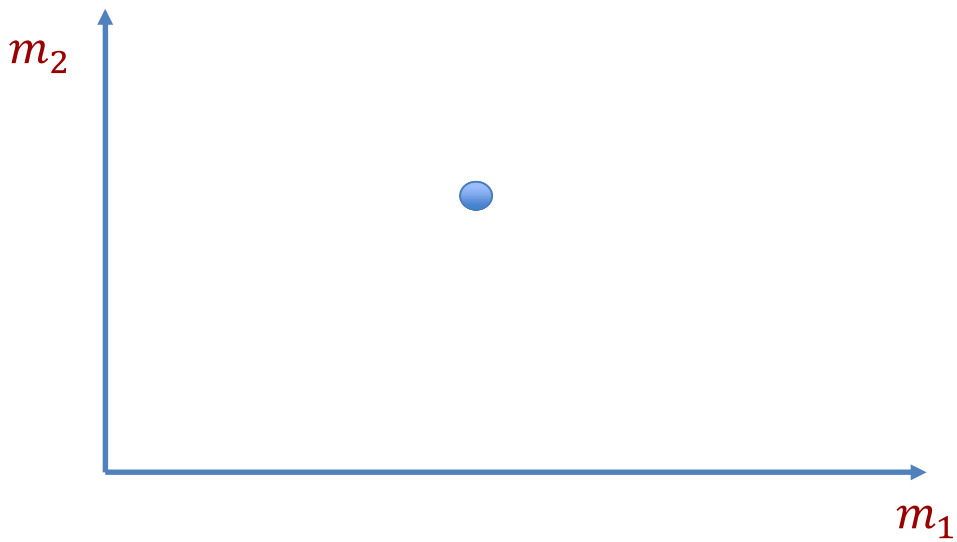

How do gradients help?

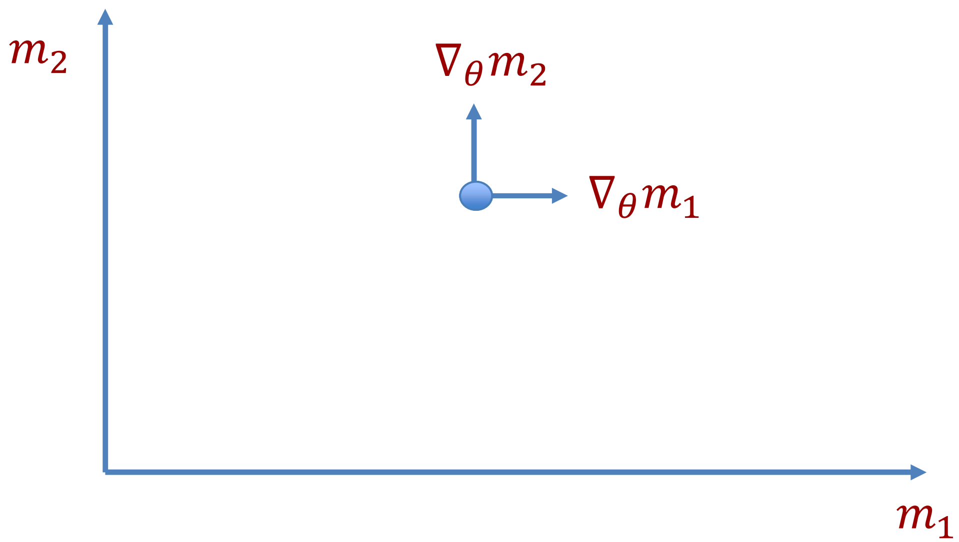

How do gradients help?

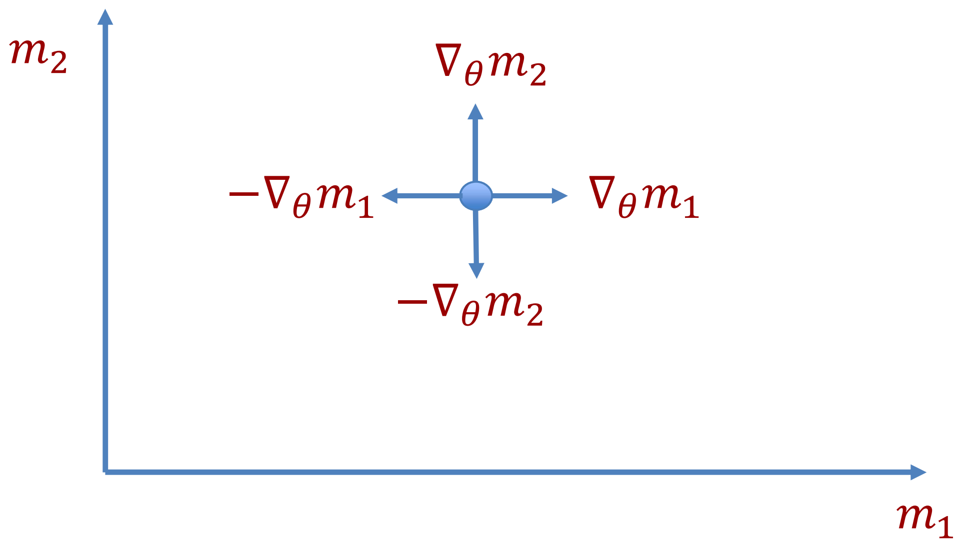

How do gradients help?

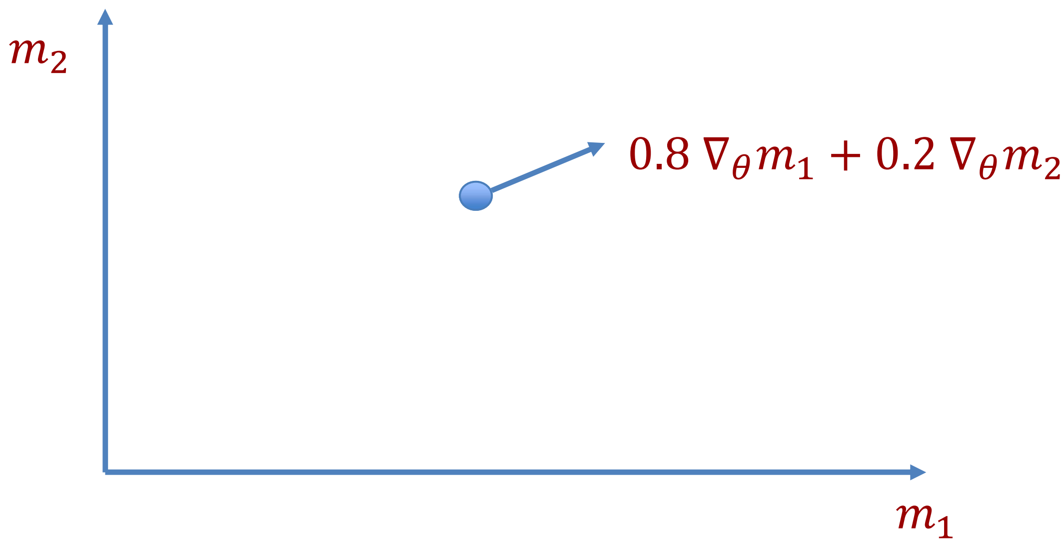

Gradient Arborescence

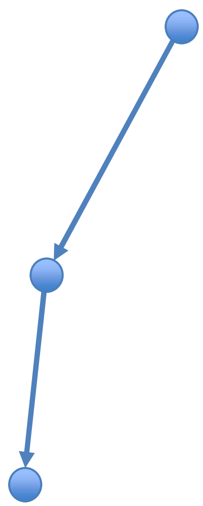

Gradient Ascent

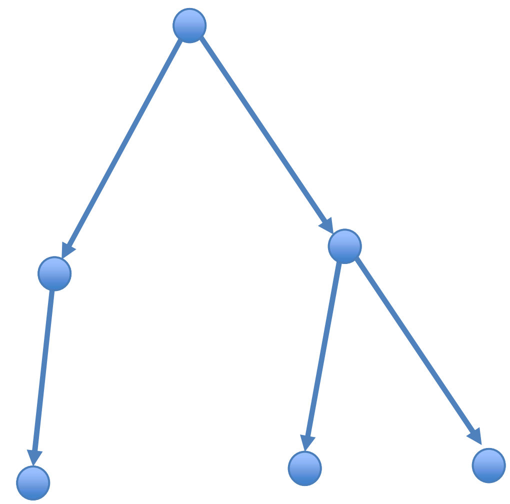

Ascending Gradient Arborescence

"On the Shortest Arborescence of a Directed Graph" Chu et. al. (Science Sinica 1965)

MAP-Elites Operator

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \sigma N(\mathbf{0}, I)$$

Gradient-based Variation Operator

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \sigma N(\mathbf{0}, I)$$

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + |c_0|\nabla f(\theta) + \sum_{i=1}^{k} c_i \nabla m_i(\theta)$$

Objective and Measure Gradient MAP-Elites via a Gradient Arborescence (OMG-MEGA)

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \sigma N(\mathbf{0}, I)$$

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + |c_0| \nabla f(\theta) + \sum_{i=1}^{k} c_i \nabla m_i(\theta)$$

$$\boldsymbol{c} \sim N(\mathbf{0}, I)$$

OMG-MEGA fails to cover the measure space

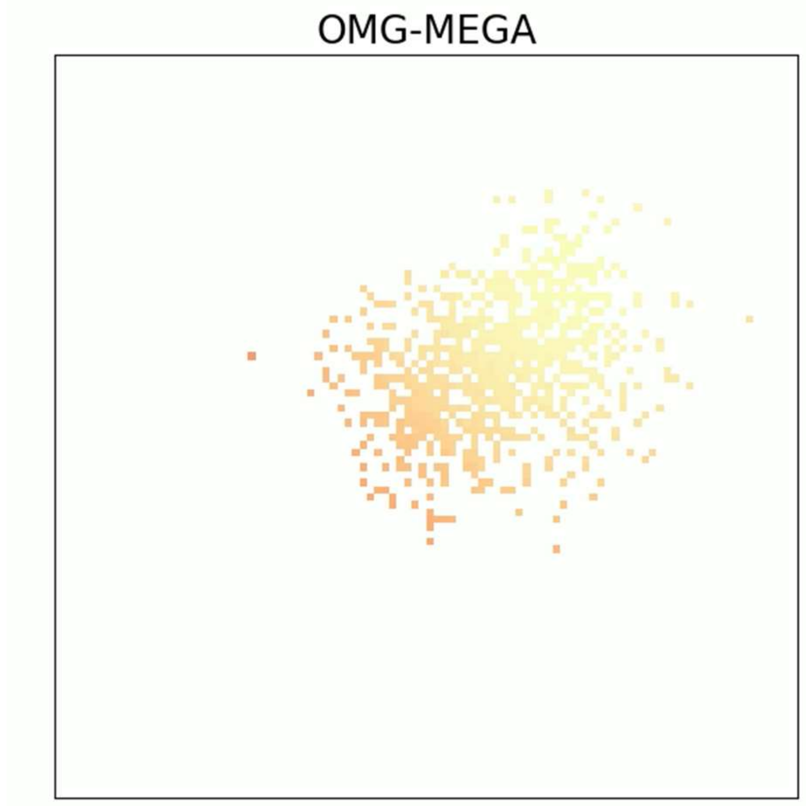

Benchmark Example ($n = 1000$)

Can we frame QD as an optimization problem?

Can we frame QD as an optimization problem?

$$f_{QD}(A) = \sum_{i=1}^{M} f(\theta_i)$$

Can we frame QD as an optimization problem?

$$f_{QD}(A) = \sum_{i=1}^{M} f(\theta_i)$$

$$\max f_{QD}$$

Can we compute the gradient of f_{QD}?

$$f_{QD}(A) = \sum_{i=1}^{M} f(\theta_i)$$

$$\max f_{QD}$$

Can we compute the gradient of f_{QD}?

$$f_{QD}(A) = \sum_{i=1}^{M} f(\theta_i)$$

$$\max f_{QD}$$

$$\nabla f_{QD}$$

Gradients are steepest ascent

$$f_{QD}(A + \boldsymbol{\theta}') = \boldsymbol{f}(\boldsymbol{\theta}') - \boldsymbol{f}(\boldsymbol{\theta}_p) + \sum_{i=1}^{M} f(\theta_i)$$

$$\max f_{QD}$$

$$\nabla_\theta f_{QD} = \boldsymbol{\theta}' \textbf{ that maximizes archive change}$$

Solve for a change in θ' to maximize f_{QD}

$$f_{QD}(A + \theta') = f(\theta') - f(\theta_p) + \sum_{i=1}^{M} f(\theta_i)$$

$$\max f_{QD}$$

$$\nabla_\theta f_{QD} = \theta' \text{ which maximizes archive change}$$

Approximate $\nabla_\theta f_{QD}$ via CMA-ES (derivative free)!

CMA-ES

- Models search directions as a multivariate Gaussian
- Updates Gaussian based on **selecting** and **ranking** best solutions
- Approximates a **natural gradient descent** of the objective modelled with uncertainty

“The CMA Evolution Strategy: A Tutorial” Nikolaus Hansen (2016)
“Bidirectional Relation between CMA Evolution Strategies and Natural Evolution Strategies” Akimoto et. al. (2010)

CMA-ME (Covariance Matrix Adaptation MAP-Elites)

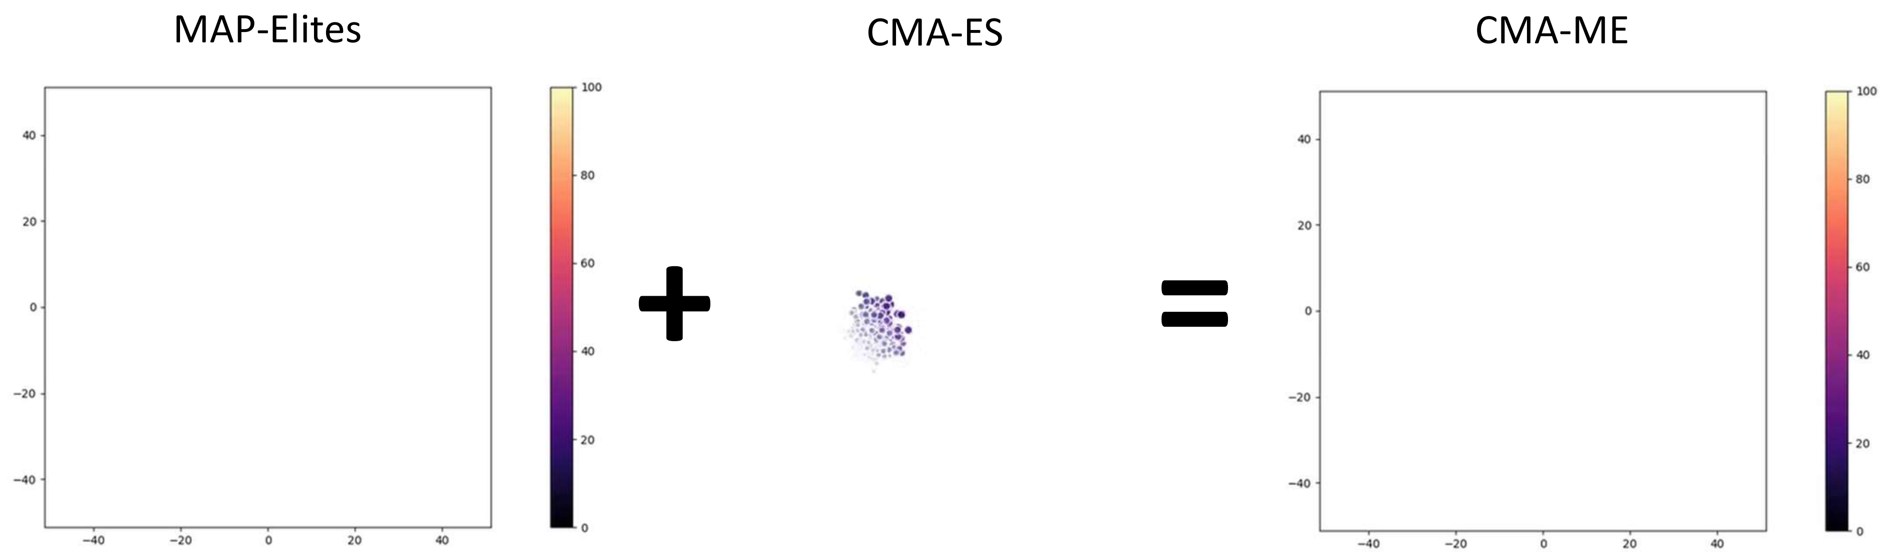

"Covariance Matrix Adaptation for the Rapid Illumination of Behavior Space"
Fontaine et. al. (GECCO 2020)

How do we leverage gradients?

How do we leverage gradients?

OMG-MEGA operator

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \mathrm{c}_0 \nabla f(\theta) + \sum_{i=1}^{k} c_i \nabla m_i(\theta)$$

How do we leverage gradients?

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \mathrm{c}_0 \nabla f(\theta) + \sum_{i=1}^{k} c_i \nabla m_i(\theta)$$

Approximates $\nabla_{\boldsymbol{c}} f_{QD}$ instead of $\nabla_{\boldsymbol{\theta}} f_{QD}$

CMA-MEGA Insight

We **branch** by sampling coefficients from a distribution $c \sim N(\mu, \Sigma)$ to **approximate** $\nabla_c f_{QD}$ and fill the archive with quality solutions at the same time.

Covariance Matrix Adaptation MAP-Elites via a Gradient Arborescence (CMA-MEGA)

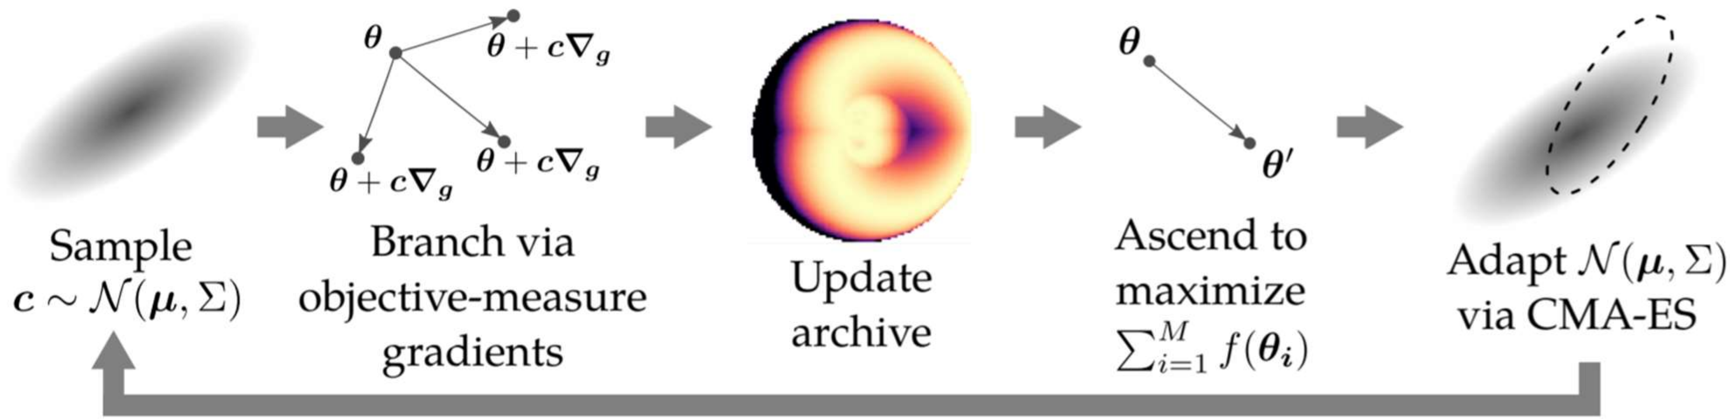

Covariance Matrix Adaptation MAP-Elites via a Gradient Arborescence (CMA-MEGA)

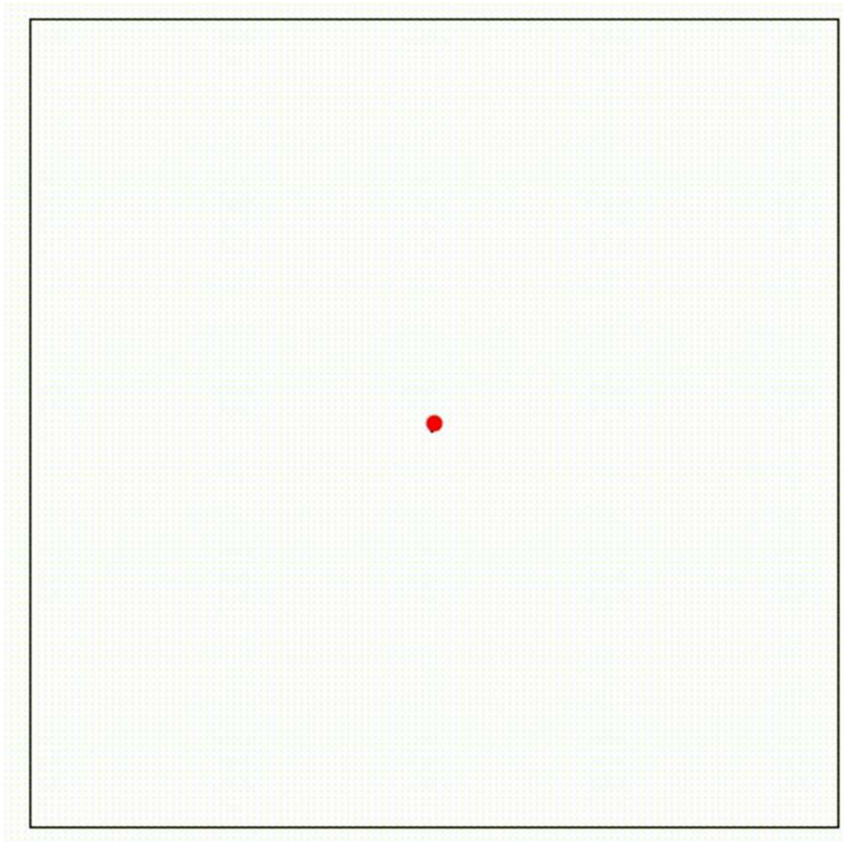

CMA-MEGA versus state-of-the-art QD algorithms

StyleGAN+CLIP

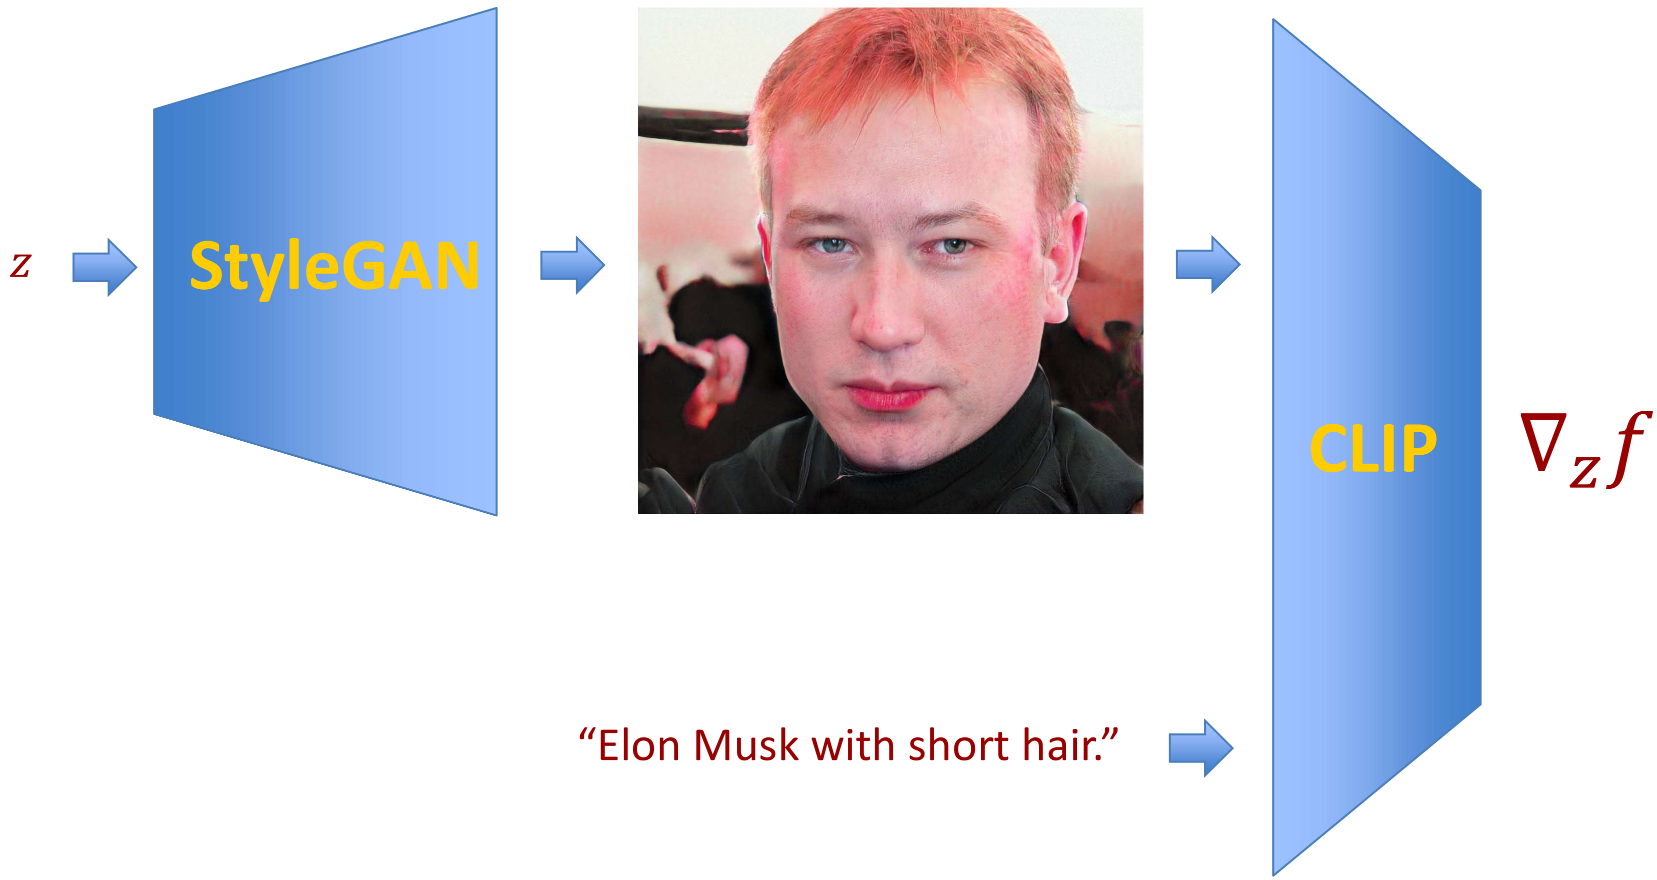

StyleGAN+CLIP

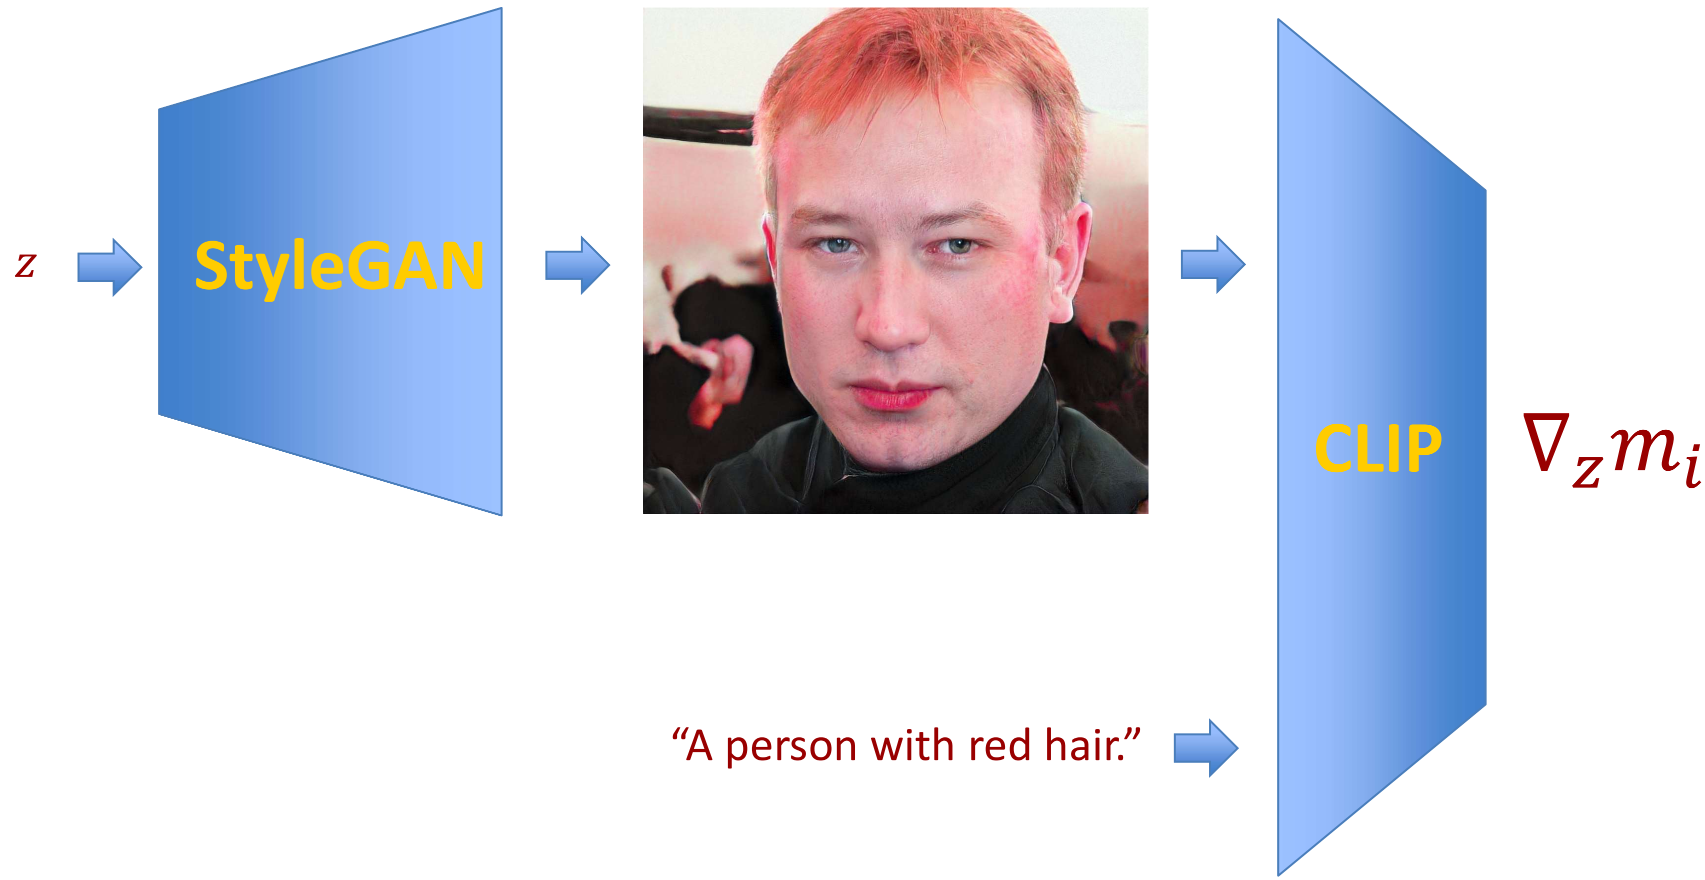

A collage of high quality and diverse solutions

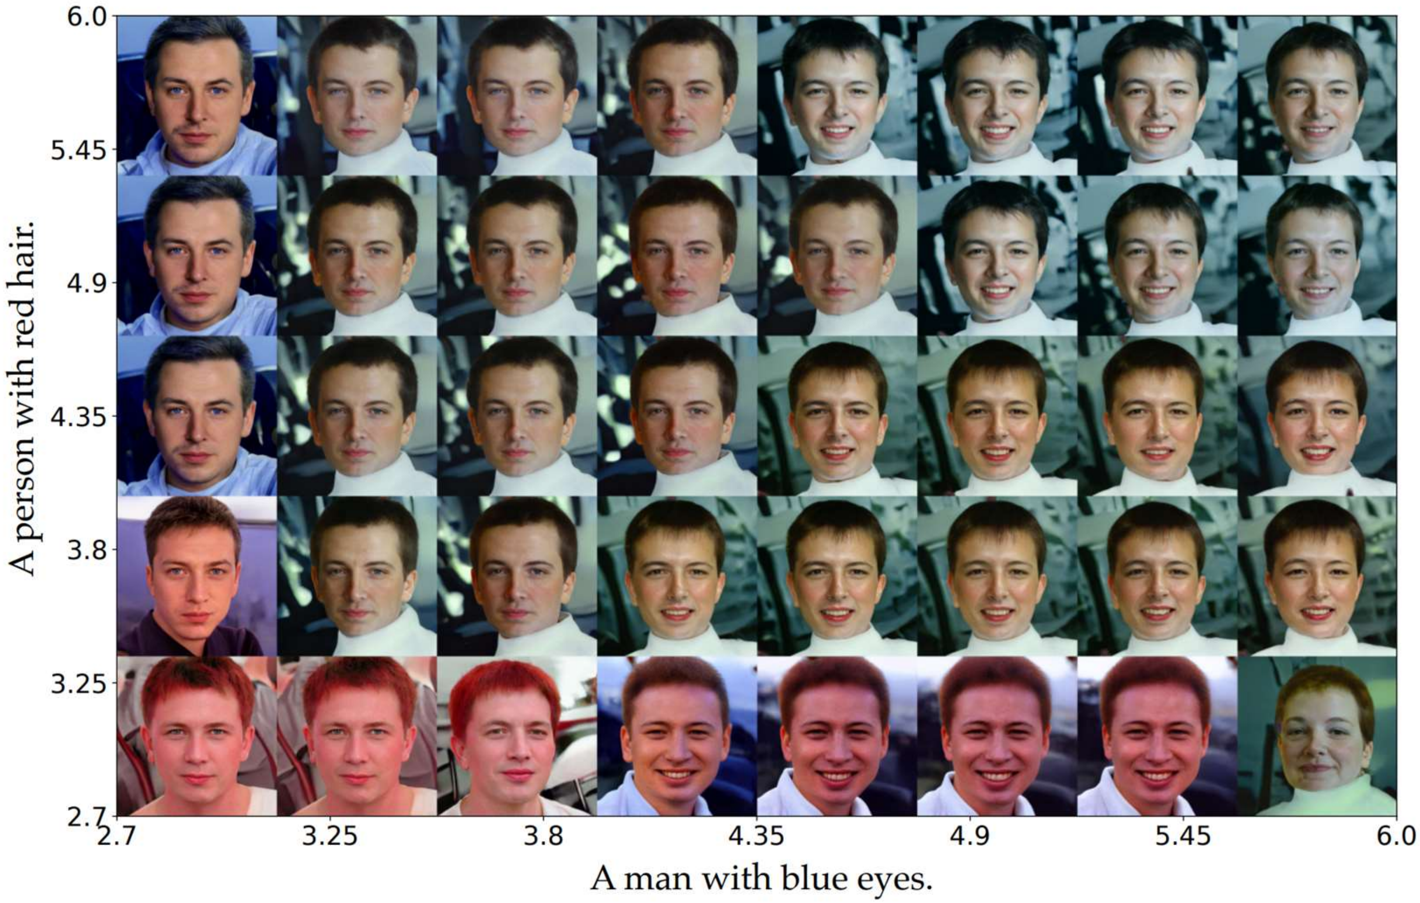

Intermediate Solutions (Cherry-picked)

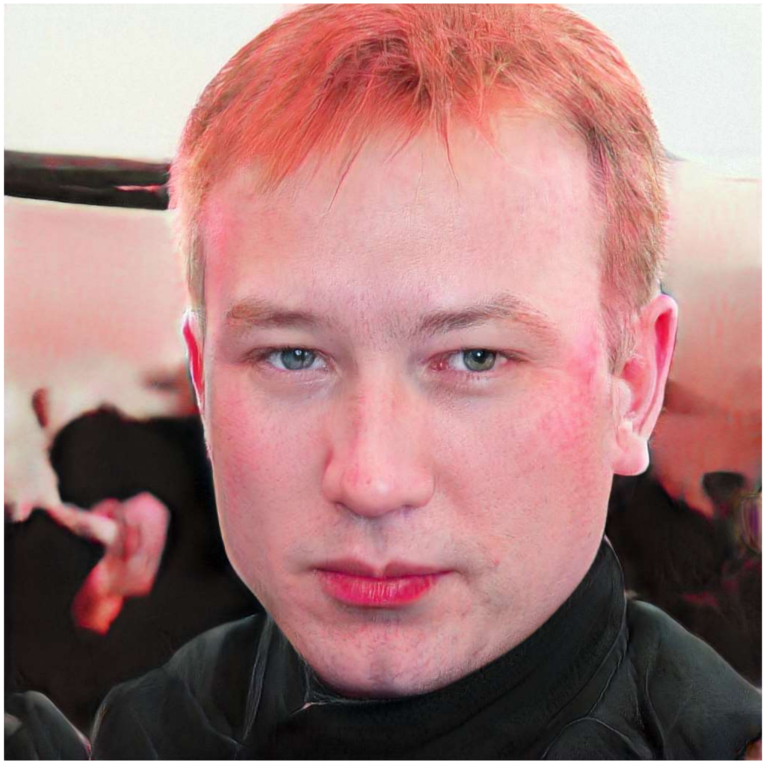

Beyoncé Example

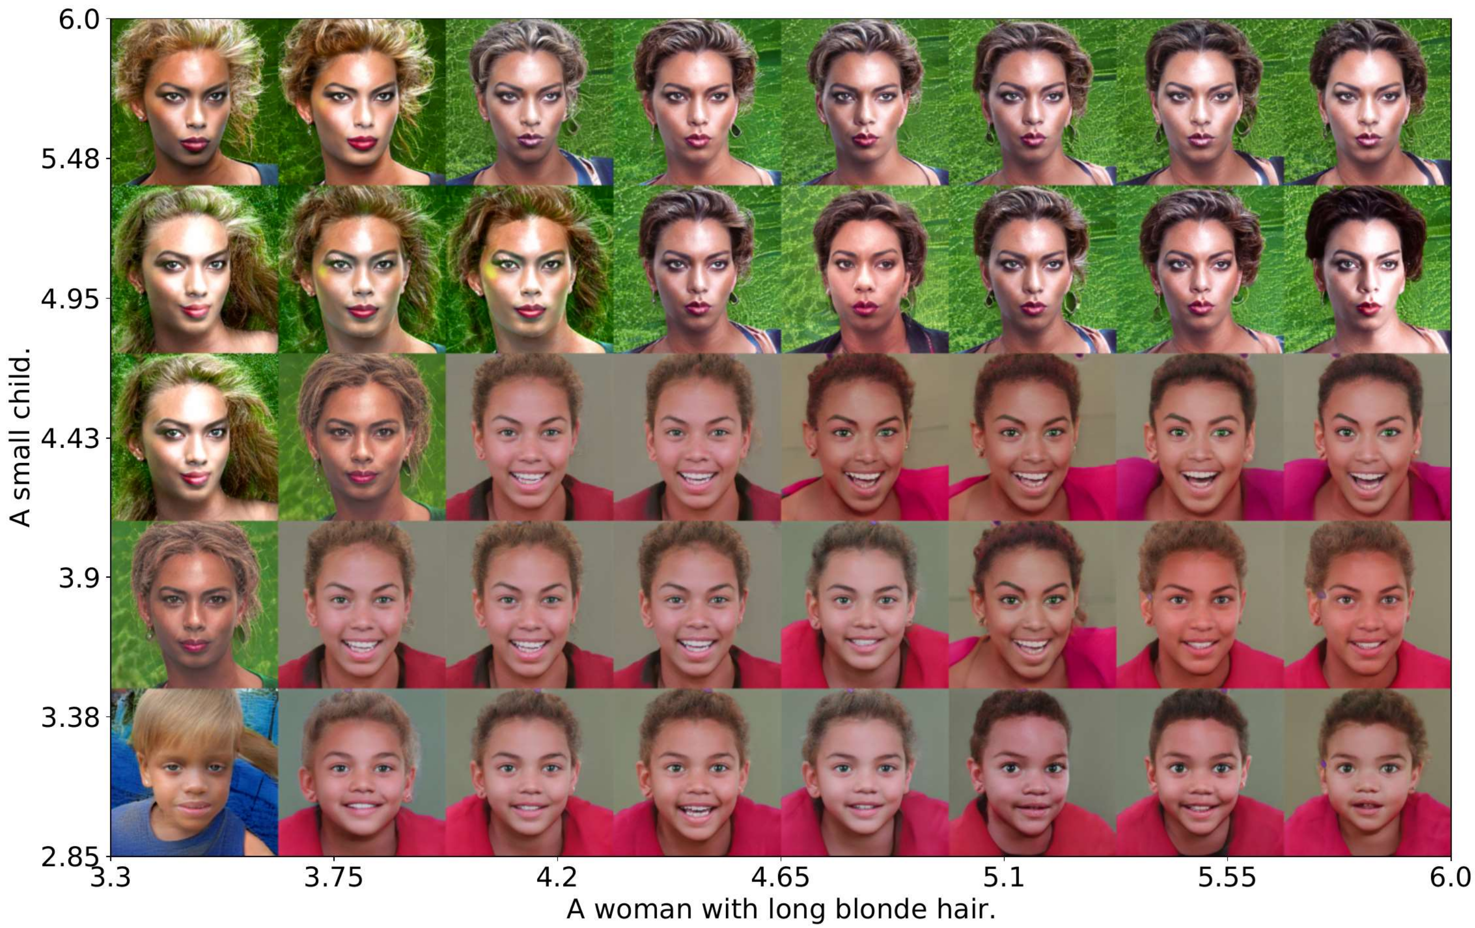

Beyoncé Example

Jennifer Lopez Example

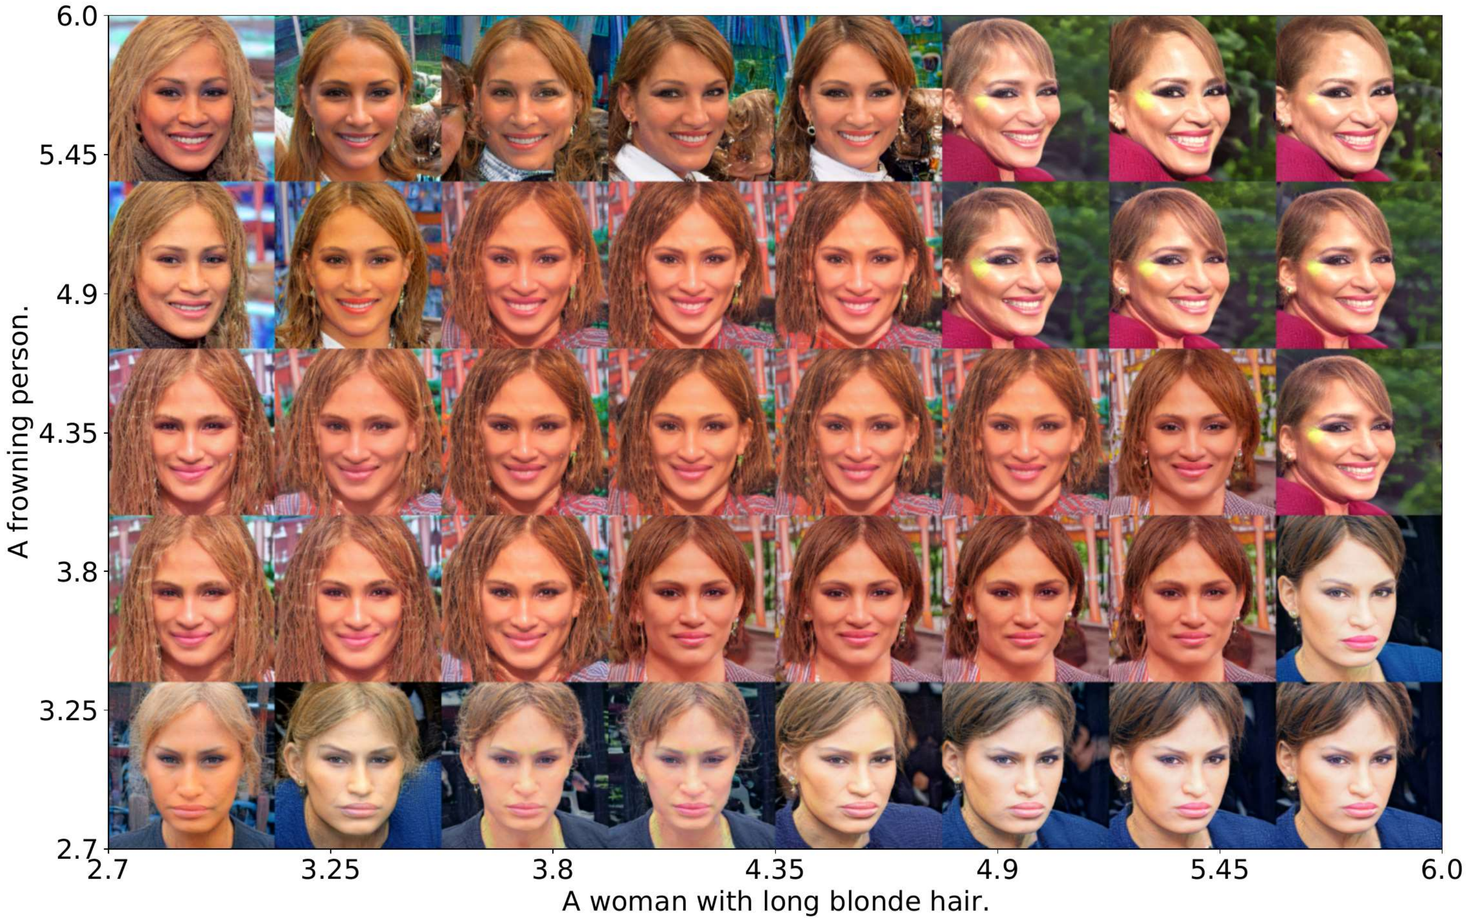

Takeaways

- Many papers at NeurIPS leverage **gradients**.
- Differentiable quality diversity (DQD) algorithms make efficient use of **gradients**.
- If a collection of high **quality** and **diverse** solutions is required, DQD may be a better tool than gradient descent.

Questions?

QD library: https://pyribs.org/
Code: https://github.com/icaros-usc/dqd