Quality-Diversity Generative Sampling for Learning with Synthetic Data

Allen Chang[1], Matthew C. Fontaine[1], Serena Booth[2], Maja J. Matarić[1], Stefanos Nikolaidis[1]

[1] University of Southern California [2] Massachusetts Institute of Technology

Introduction

Synthetic datasets can supplement real datasets to improve training.

However, **synthetic training datasets should be balanced to prevent a transfer of bias to a downstream task.**

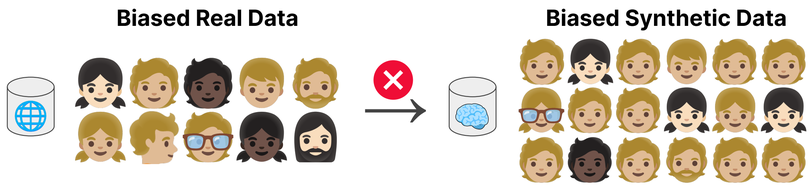

Contributions:

→ **Use quality-diversity algorithms to simultaneously balance data features while creating synthetic training datasets**

→ **Improve fairness without harming accuracy when the data are used to pretrain downstream facial recognition classifiers**

Intuition

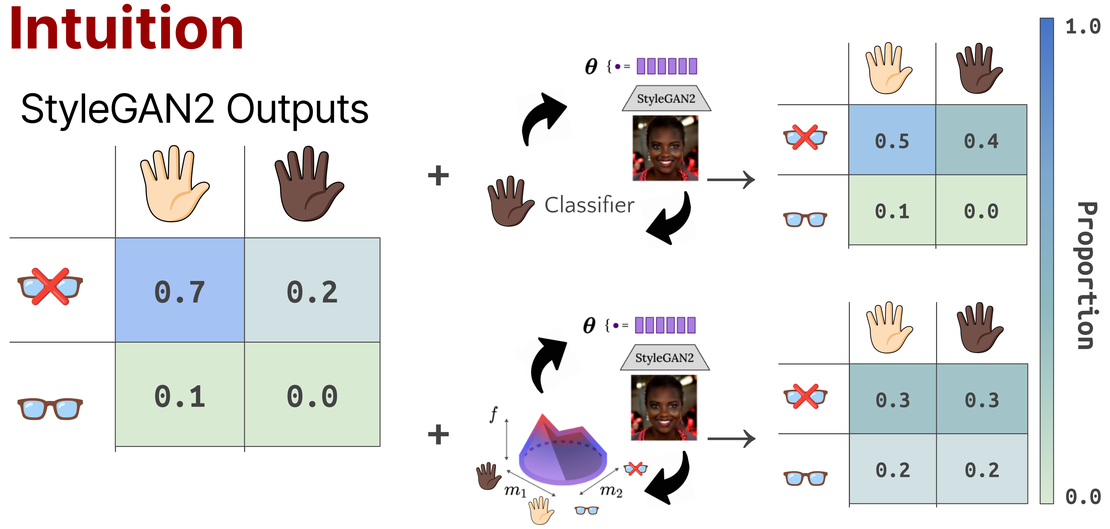

Prior Work:
Single-factor diversity optimization

QDGS:
Quality-diversity (QD) optimization for multiple measures of diversity optimization + text prompts to describe desired attributes

Quality-Diversity Generative Sampling (QDGS)

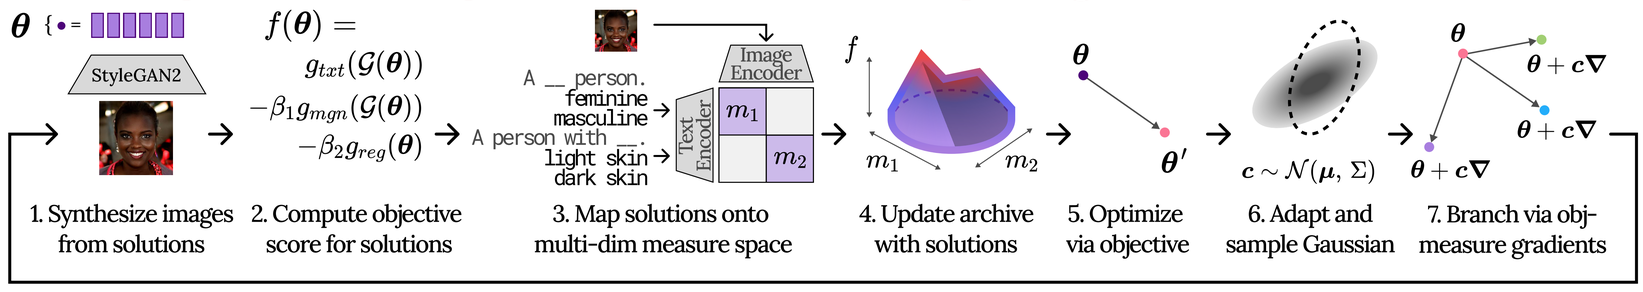

Results *(1) More Uniformity in QDGS Distributions*

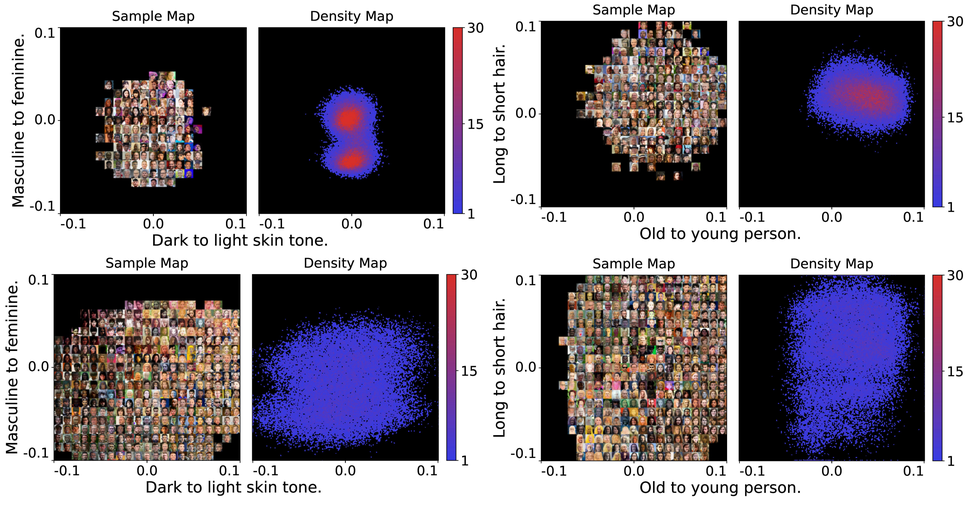

QDGS datasets ...achieve more uniformity than random sampling.

(2) > Fairness Increase, and (3) ≍ Accuracy Increase

Pretraining	Dark-skinned (RFW)	Light-skinned (RFW)	ACC (LFW ... AgeDB)*
Train: CASIA (Imbalanced)			
None	88.08 ± 0.07	94.05 ± 0.06	94.52 ± 0.03
Rand15	88.17 ± 0.08	94.41 ± 0.07	94.83 ± 0.03
Rand50	88.69 ± 0.10	$\mathbf{94.67 \pm 0.07}$	94.89 ± 0.03
QD15 (Ours)	88.25 ± 0.09	94.42 ± 0.06	94.83 ± 0.03
QD50 (Ours)	$\mathbf{88.94 \pm 0.07}$	94.62 ± 0.10	$\mathbf{94.92 \pm 0.02}$
Train: BUPT (Balanced)			
None	96.22 ± 0.03	97.58 ± 0.05	96.27 ± 0.02
Rand15	96.22 ± 0.05	97.72 ± 0.05	$\mathbf{96.35 \pm 0.02}$
Rand50	96.17 ± 0.09	$\mathbf{97.82 \pm 0.04}$	96.31 ± 0.03
QD15 (Ours)	96.29 ± 0.06	97.75 ± 0.03	96.32 ± 0.03
QD50 (Ours)	$\mathbf{96.32 \pm 0.06}$	97.80 ± 0.04	96.30 ± 0.02

* ACC is the averaged accuracy of LFW, CFPFP, CPLFW, CALFW, and AgeDB

...are helpful for balanced and imbalanced training sets.

Conclusions

Synthetic datasets created from generative models can be helpful to train classifiers, especially when they are balanced.

QDGS uses QD optimization and text prompts to control for which attributes should be made more balanced in these datasets.

QDGS + facial recognition classifiers ⟶ more fairness without decrease in accuracy!