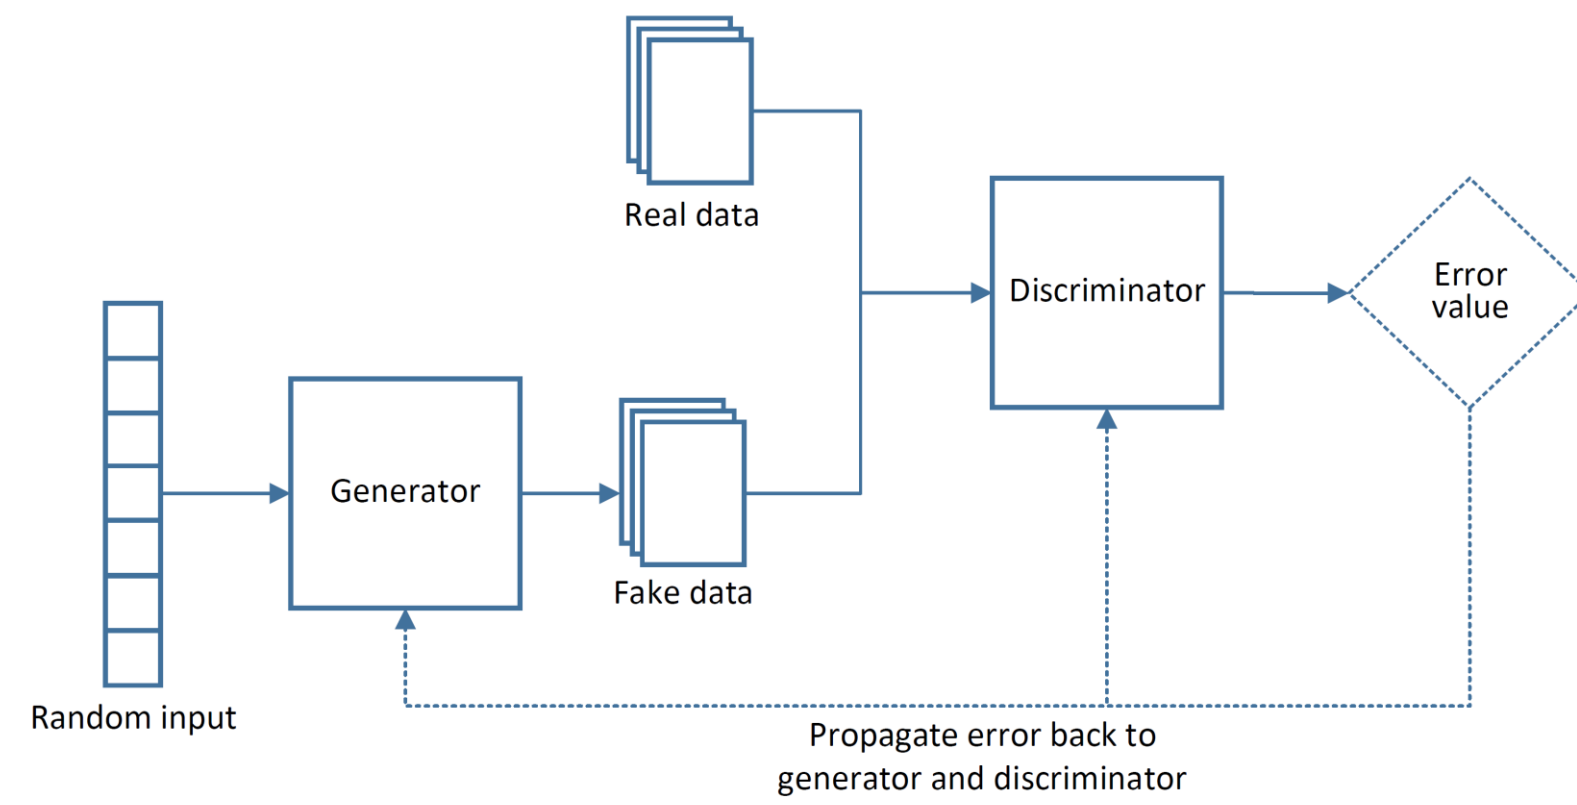




MOTIVATION

Generative Adversarial Networks



$$\min_G \max_D \mathcal{L}(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\phi(D(x))] + \mathbb{E}_{z \sim p_z(z)} [\phi(1 - D(G(z)))]$$

- Simultaneous gradient updates in GAN training lead to **unstable dynamics**
- Similar degenerate behaviors have been studied by the coevolutionary computing for minimax optimization
 - **Mode Collapse** ↔ **Focusing**
 - **Discriminator Collapse** ↔ **Relativism**
 - **Vanishing Gradients** ↔ **Loss of Gradients**



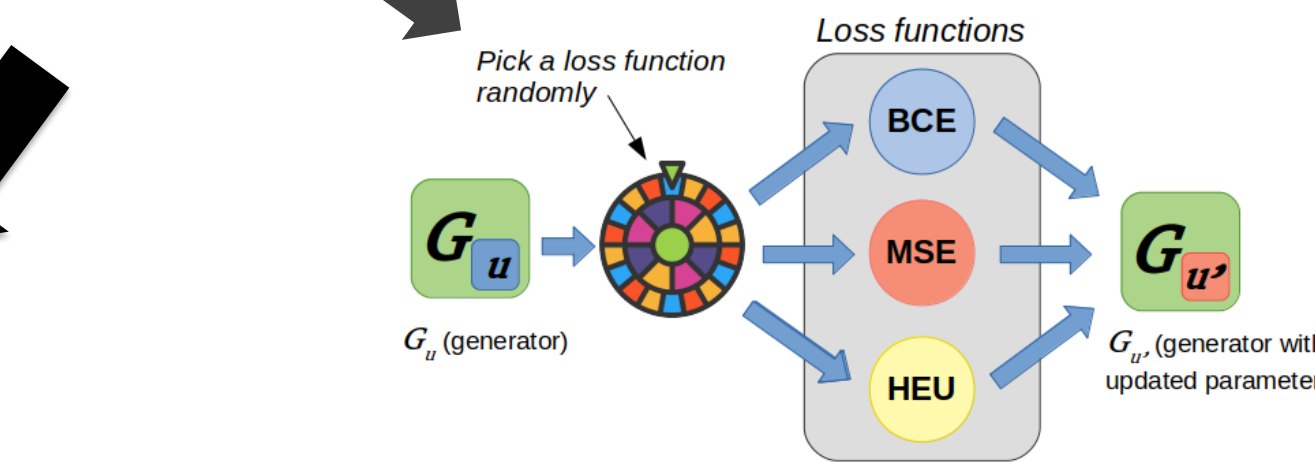
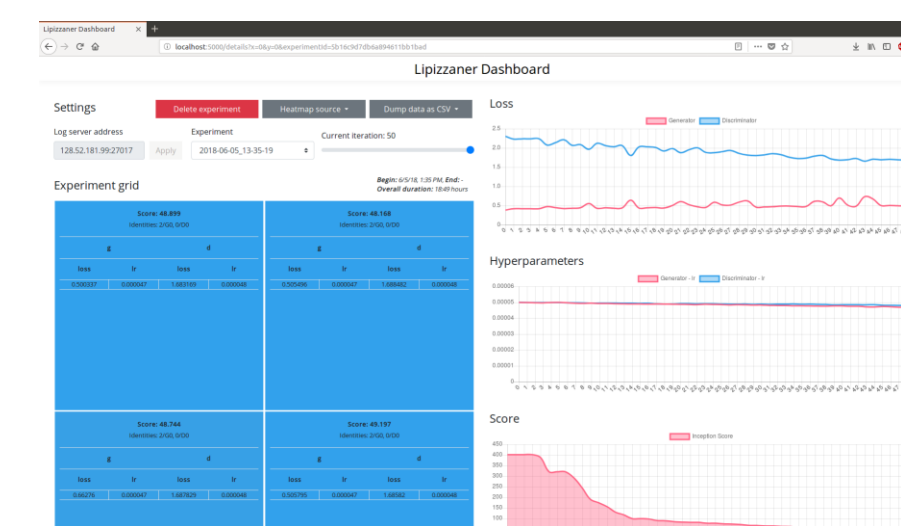
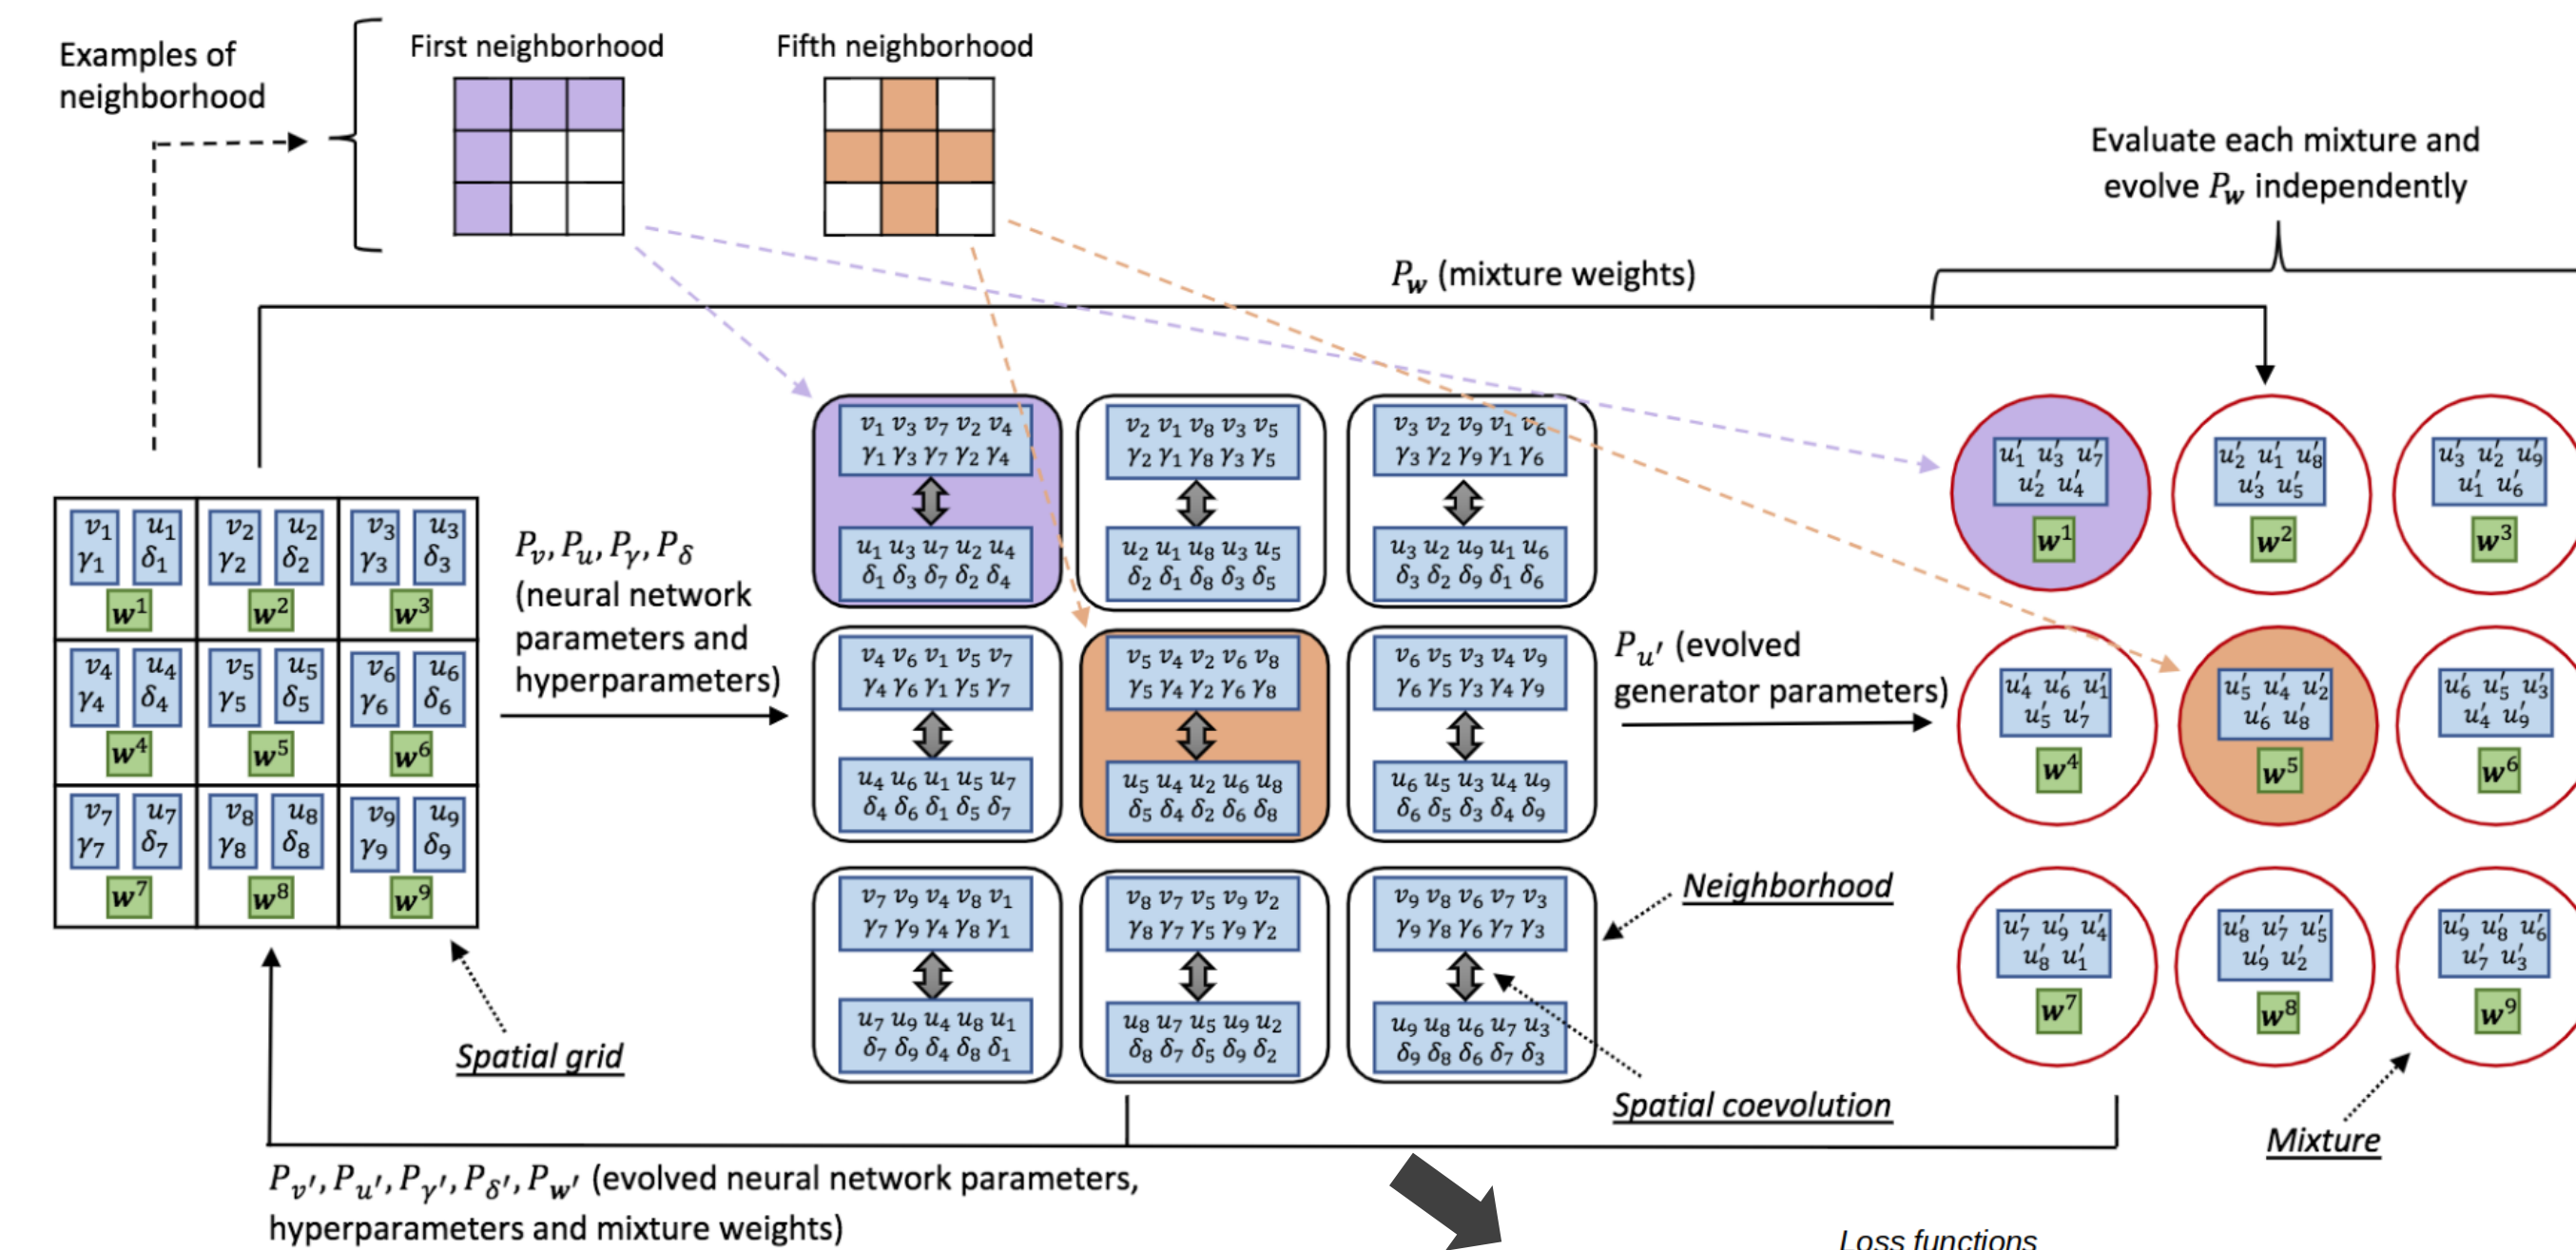
- Supplementing GAN training with established coevolutionary techniques proposed to address the unstable dynamics

RELEVANT PAPERS

1. J. Toutouh, E. Hemberg, and U. O'Reilly. Spatial Evolutionary Generative Adversarial Networks. In Genetic and Evolutionary Computation Conference (GECCO '19), July 13–17, 2019.
2. T. Schmiedlechner, I. Yong, A. Al-Dujaili, E. Hemberg, U. O'Reilly. Lipizzaner: A System That Scales Robust Generative Adversarial Network Training. NeurIPS 2018 Workshop on System for ML, 2018.
3. A. Al-Dujaili, T. Schmiedlechner, E. Hemberg, U. O'Reilly. Towards distributed coevolutionary GANs. AAAI 2018 Fall Symposium, 2018.

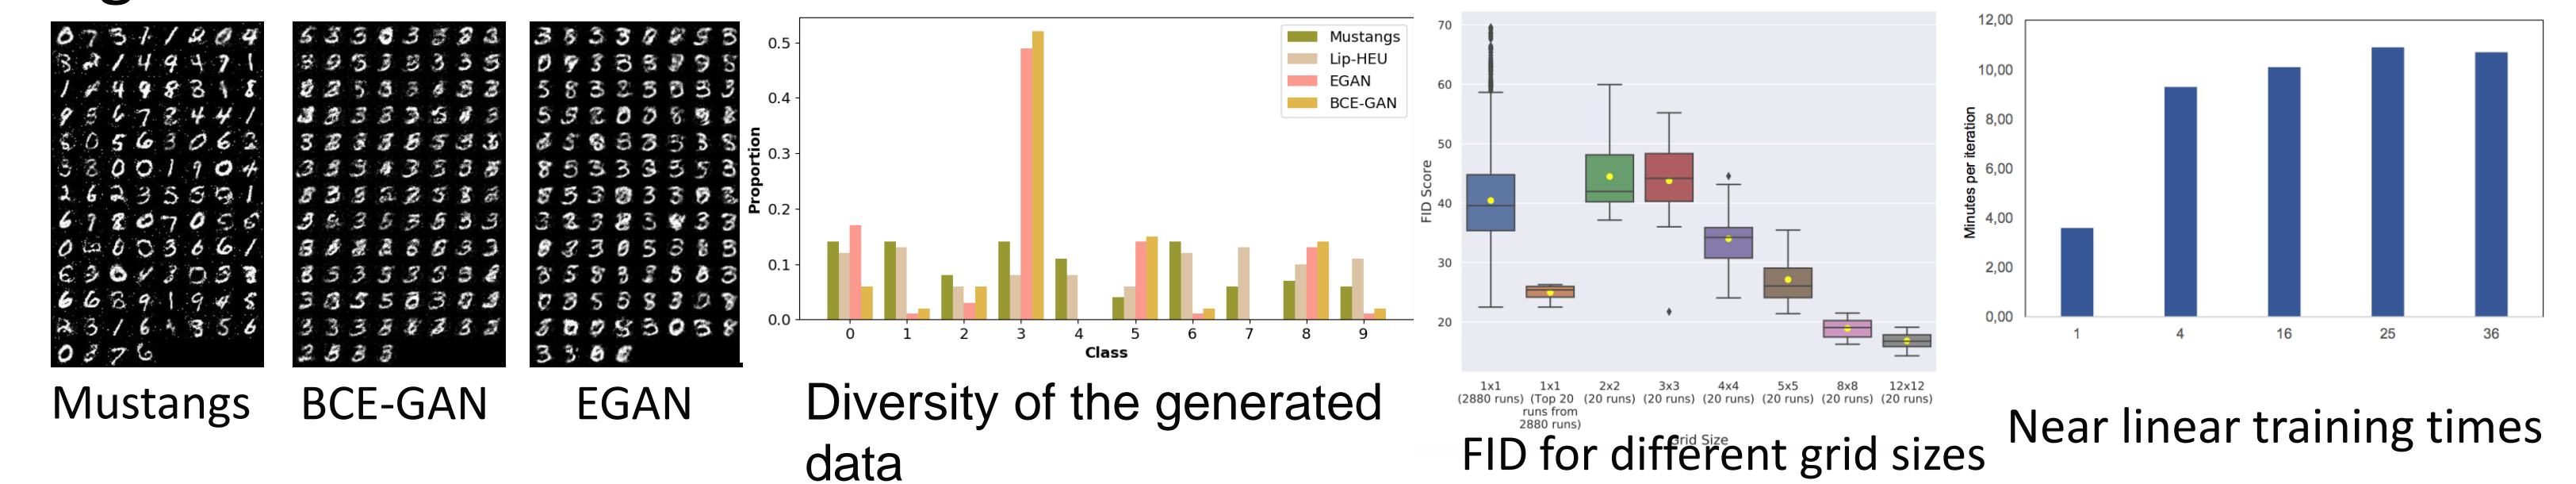
MUSTANGS: DISTRIBUTED COEVOLUTIONARY GANs TRAINING

Lipizzaner is a distributed, coevolutionary framework to simultaneously train multiple GANs with gradient-based optimizers



RESULTS

Figure: MNIST results



Mustangs distributes populations over a spatial grid

- Grid cells represent population of training instances
- Cells communicate only with neighboring cell
- Computational benefits
 - **Linear** instead of quadratic scaling
 - Large datasets can be split onto different cells

For each iteration, each network **randomly** picks a **loss function** to optimize its weights according to different objectives

Open source implementation:
<https://github.com/mustang-gan/>

Figure: CelebA results

