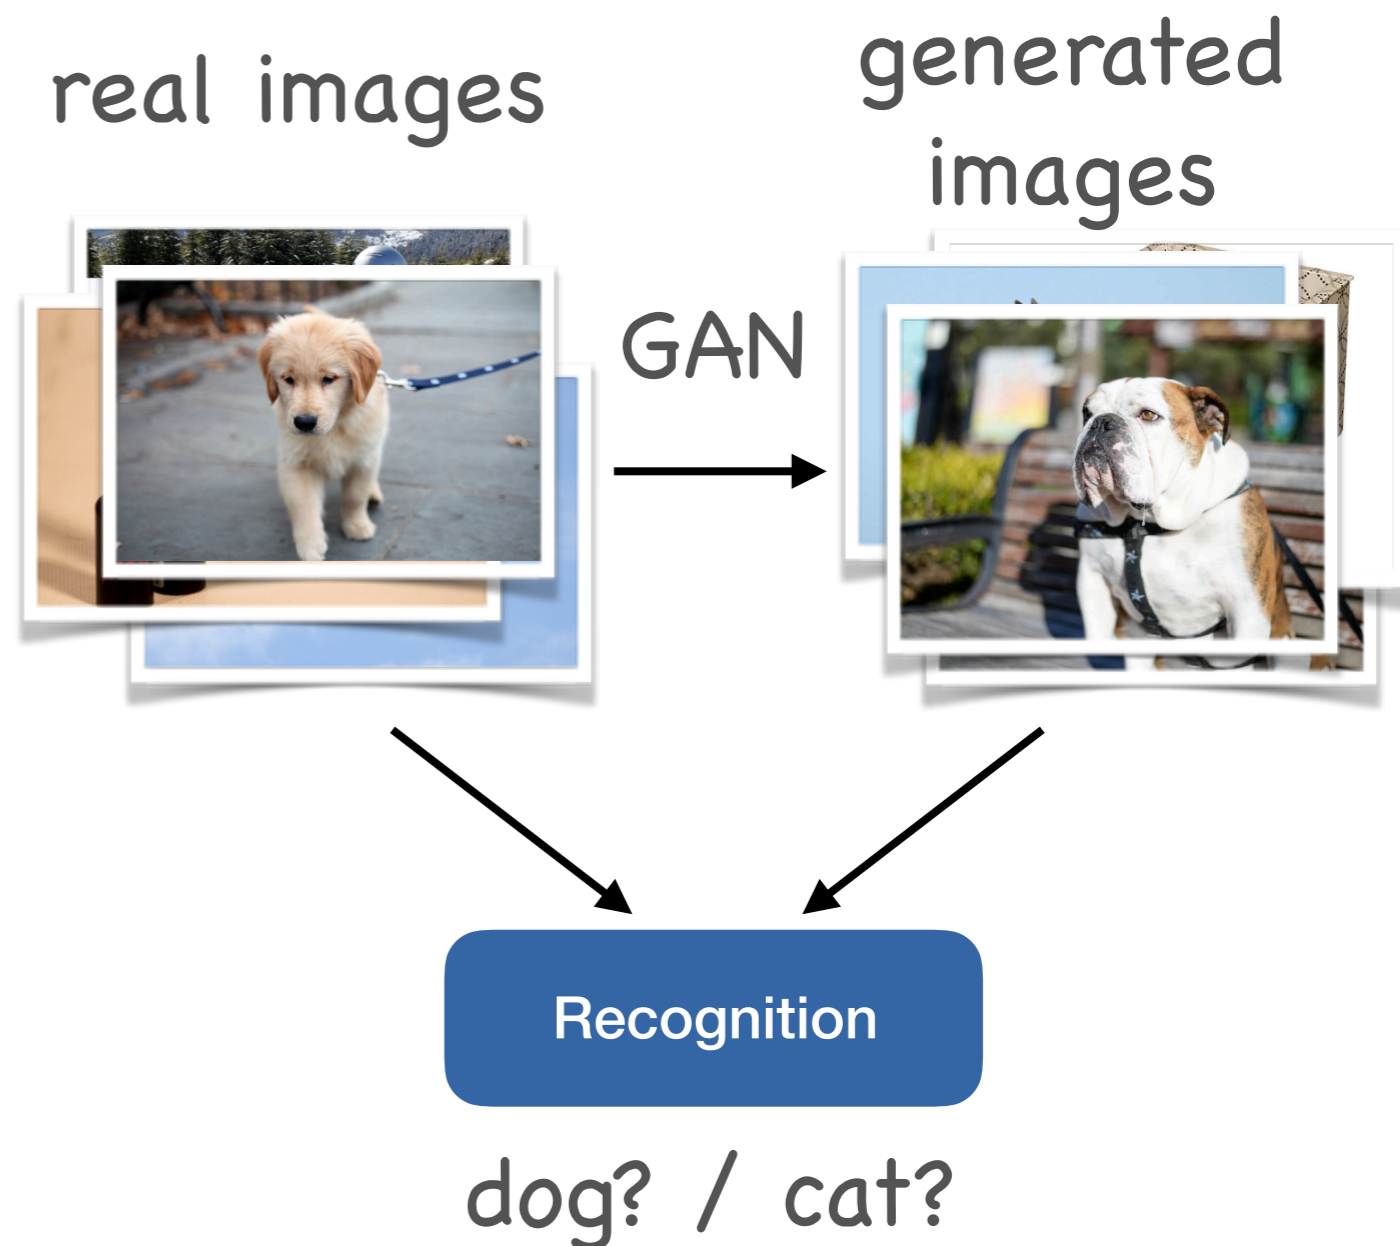


Open Problem: Understanding generative models and invariances

© 2019 Philipp Krähenbühl and Chao-Yuan Wu


Can we use generative model to create more training data?

- Step 1: Train a GAN (for each class)
- Step 2: Generate images from that GAN (with free labels)
- Step 3: Train a recognition on generated images



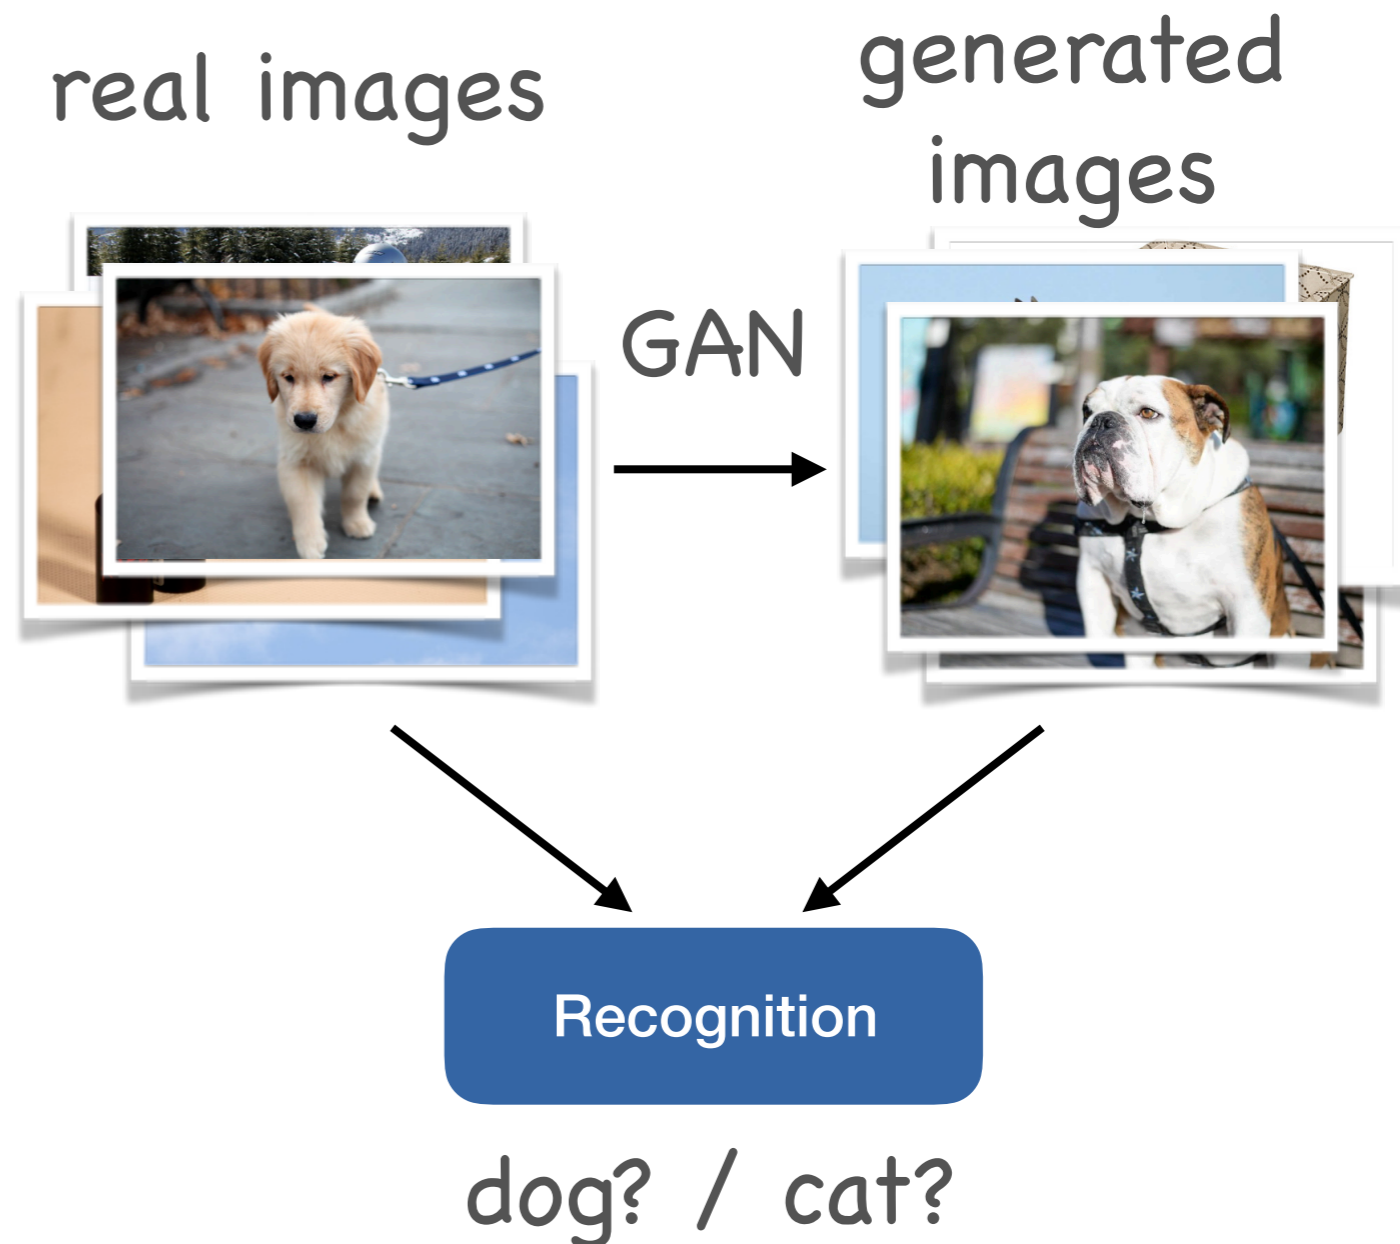
Does this give a better recognition model?

- Argument in favor
 - More data
 - Less overfitting
 - Better generalization



Does this give a better recognition model?

- Argument against
- Final training pipeline does not see more data



S

- Counter argument
 - Maybe GAN generalizes better?
- Counter counter argument
 - Why?!

