A Self-supervised GAN for Unsupervised Few-shot Object Recognition

Khoi Nguyen and Sinisa Todorovic

Oregon State University

25th INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION
ICPR 2020
Milan, Italy 10 | 15 January 2021

Problem:
Given unlabeled data in training, classify a query image into one of the classes defined by a few support images per class. The unlabeled and support images do not share the same object classes.

Challenge:
No access to a large set of labeled images to enable the episodic training of standard few-shot classification.

Contribution 1:
A GAN architecture with reconstruction loss for the discriminator to explicitly capture the most relevant latent codes that generate “fake” images.

Contribution 2:
Enforce the discriminator to produce image encodings that respect similarity of images via self-supervised learning which uses image masking.

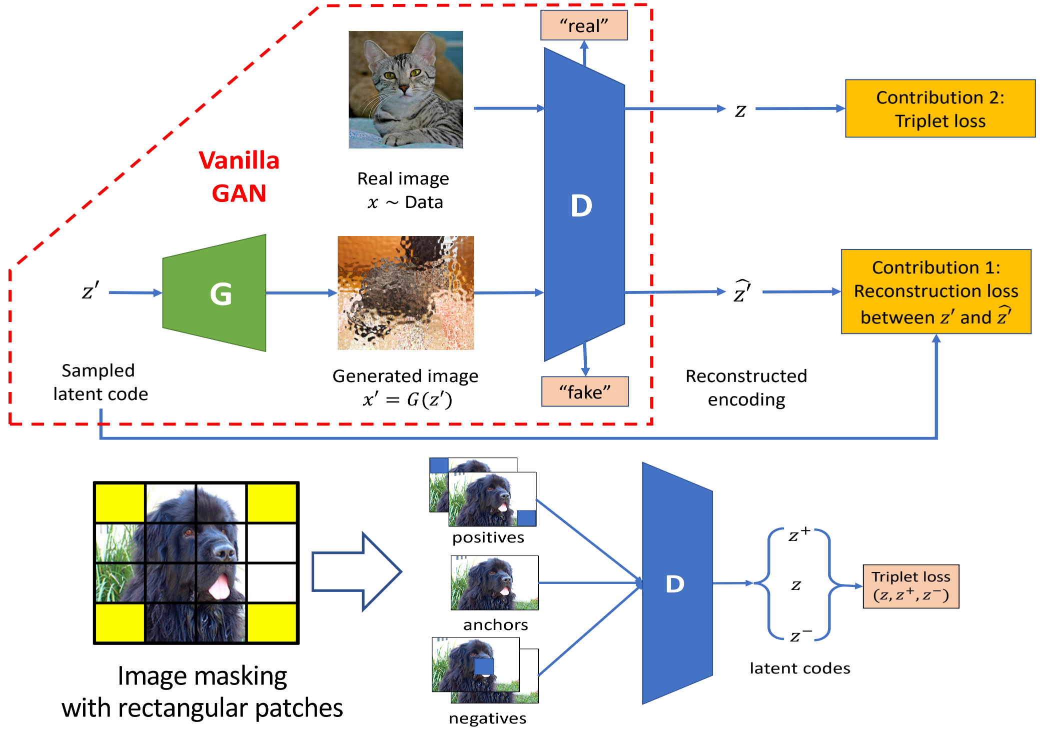

Table: Unsupervised few-shot classification on Mini-Imagenet and Tiered-Imagenet

	Mini-Imagenet, 5-way		Tiered-Imagenet, 5-way	
Unsupervised Methods	1-shot	5-shot	1-shot	5-shot
SN-GAN (Miyato et al., ICLR 2018)	34.84 ± 0.68	44.73 ± 0.67	35.57 ± 0.69	49.16 ± 0.70
AutoEncoder (Vincent et al., JMLR 2010)	28.69 ± 0.38	34.73 ± 0.63	29.57 ± 0.52	38.23 ± 0.72
Rotation (Gidaris et al., ICLR 2018)	35.54 ± 0.47	45.93 ± 0.62	36.90 ± 0.54	51.23 ± 0.72
BiGAN kNN (Donahue et al,. ICLR 2017)	25.56 ± 1.08	31.10 ± 0.63	-	-
AAL-ProtoNets (Antonious et al., Arxiv 2019)	37.67 ± 0.39	40.29 ± 0.68	-	-
CACTUs-ProtoNets (Hsu et al., ICLR 2019)	39.18 ± 0.71	53.36 ± 0.70	-	-
Our GdBT2	**48.28 ± 0.77**	**66.06 ± 0.70**	**47.86 ± 0.79**	**67.70 ± 0.75**
Fully-supervised Methods				
ProtoNets (Snell et al., NeurIPS 2017)	46.56 ± 0.76	62.29 ± 0.71	46.52 ± 0.72	66.15 ± 0.74

Figure: Every row shows images in the descending order from left to right by their estimated distance to the original (unmasked) image.

Acknowledgement: DARPA XAI Award N66001-17-2-4029 and DARPA MCS Award N66001-19-24035.