From MetaGAN to MetaGAIL

Feng Xu

LAMDA5, Nanjing University

December 8, 2020

Table of Contents

Table of Contents

Introduction

MetaGAN: An Adversarial Approach to Few-Shot Learning. [1]
Ruixiang Zhang et al. (Yoshua Bengio)
Motivation:How to form generalizable decision boundaries from a small number of traning samples in few shot learning?

Contribution

- Propose MetaGAN that unify supervised/semi-supervised few-shot learning naturally

Formalization

Distribution of Tasks P(T)
Training tasks $(T_i)_{i=1}^N$, where $T = (S_T, Q_T)$
$S_T = S_T^s \cup S_T^u$, $Q_T = Q_T^s \cup Q_T^u$, S/Q denote support/query set, s/u denote superversied/unsupervised
$S_T^u = \{x_1, x_2, ...X_M\}$, $Q_T^s = \{(x_1, y_1), ...(X_O, y_O)\}$
Objective: Minimize prediction loss on a query set given support set

Trouble of GANs

- GAN have trouble generating realistic samples in complex datasets
- Easy to run into mode collapsing

Algorithm

Assume a few-shot classifier $p_D(x;T) = (p_1(x), p_2(x), ...p_n(x))$, augmented with an additional output $p_{n+1}(x)$ to model the probability that input data is fake. The objectives are

Discriminator:

$$L_D^T = L_{superversied} + L_{unsupervised}$$

$$L_{supervised} = \mathbb{E}_{x,y\sim Q_T^s} log p_D(y|x, y \leq N)$$

$$L_{unsupervised} = \mathbb{E}_{x\sim Q_T^u} log p_d(y \leq N|x) + \mathbb{E}_{x\sim p_G^T}(N+1|x)$$

Generator:

$$L_G^T(D) = -\mathbb{E}_{x\sim P_G^T}[log(p_D(y \leq N|x))]$$

Overall:

$$L_D = max_D \mathbb{E}_{T\sim P(T)} L_D^T$$

$$L_G = min_G \mathbb{E}_{T\sim P(T)} L_G^T$$

Practical implementation of meta Discriminator

- MAML:$\theta'_d = \theta_d - \alpha \nabla_{\theta_d} l_D^T$
  $l_D^T = -\mathbb{E}_{x,y \sim S_T^s} logp_D(y|x, y \leq N) - \mathbb{E}_{x \sim S_T^u} logp_D(y \leq N|x) - \mathbb{E}_{x \sim p_G^T} logp_D(N+1|x)$
- Relation Network [2]: do classification via a deep distance metric
  Let $r_{i,j} = g_\psi(C(f_\phi(x_i), f_\phi(x_j))), x_i \in S_T^s, x_j \in Q_T^s$ be the relevance score between query x_j and support x_i, where g_ϕ is the relation module, f_ϕ is the embedding network, C is the concatenation operator. $r_{i,j}$ is computed via softmax classification

$$p_D(y = k|x_j) = \frac{exp(r_{k,j})}{1 + \sum_{i=1}^{N} exp(r_{i,j})}$$

Practical implementation of Generator

A conditional generative mode: compress support dataset to a vector h_T and concatenate it with random noise z as input

1. Instance-Encoder Module:
   $x_i \sim S_T^s \rightarrow e_i = Instance - Encoder(x_i)$
2. Feature-Aggregation Module: element wise operators such as average pooling, max pooling.

To make it harder for the generator to simply reconstricut its inputs

An insight

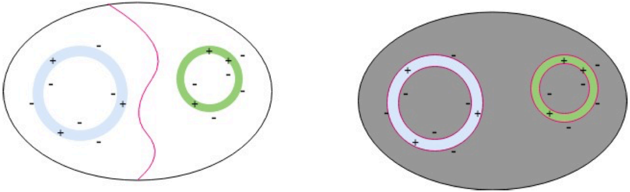

Figure 1: Left: decision boundary without metaGAN. Right: decision boundary with metaGAN. We use red curves to denote the decision boundary. Blue area in figure represents class A, green area represents class B, and gray area represents fake class. We use + to denote real samples and – to denote fake samples generated.

(a) MetaGAN

Performance

Model	5-way Acc. 1-shot	5-way Acc. 5-shot
Prototypical Nets	49.42 ± 0.78	68.20 ± 0.66
MAML(5 gradient steps)	48.70 ± 1.84	63.11 ± 0.92
MAML(5 gradient steps, first order)	48.07 ± 1.75	63.15 ± 0.91
MAML(1 gradient step, first order)	43.64 ± 1.91	58.72 ± 1.20
Ours: MetaGAN + MAML(1 step, first order)	46.13 ± 1.78	60.71 ± 0.89
Relation Net	50.44 ± 0.82	65.32 ± 0.7
Ours: MetaGAN + RN	$\mathbf{52.71 \pm 0.64}$	$\mathbf{68.63 \pm 0.67}$

Table 2: Few-shot classification results on Mini-Imagenet.

(b) Supervised Case

Model	Omniglot 1-shot 5-way	Mini-Imagenet 1-shot 5-way
Prototypical Net(Supervised)	93.66 ± 0.09	42.28 ± 0.32
Relation Net(Supervised)	93.82 ± 0.07	43.87 ± 0.20
Ours: MetaGAN + RN	$\mathbf{97.12 \pm 0.08}$	$\mathbf{47.43 \pm 0.27}$

Table 4: Task-level Semi-Supervised 1-shot classification results on Omniglot and Mini-Imagenet.

(c) Semi-Supervised Case

Table of Contents

NANJING UNIVERSITY

Learning And Mining from DatA

GAN/GAIL Framework

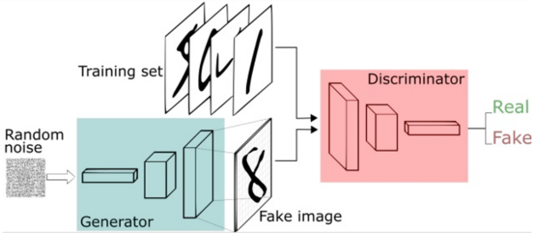

(d) GAN framework

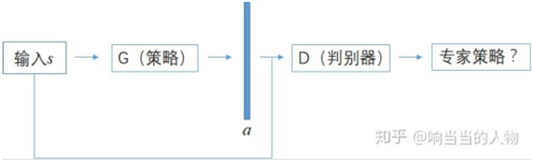

(e) GAIL framework

each GAN's corresponding GAIL

- GAN [3] & GAIL [4]
- CGAN [5] & CGAIL (additional label input)
- InfoGAN [6] & InfoGAIL (learnt latent variable)
- ACGAN & ACGAIL (augment D with an auxiliary classifier)
- f-GAN & f-GAIL (minimize f-divergence of data distribution)
- GoalGAN & GOALGAIL (Generate goals)
- TripleGAN& TripleGAIL

Table of Contents

Table of Contents

References

- Ruixiang Zhang et al. "MetaGAN: An Adversarial Approach to Few-Shot Learning". In: *NIPS*. 2018, pp. 2371–2380.
- Flood Sung et al. "Learning to Compare: Relation Network for Few-Shot Learning". In: *CVPR*. IEEE Computer Society, 2018, pp. 1199–1208. URL: http://openaccess.thecvf.com/content_cvpr_2018/html/Sung_Learning_to_Compare_CVPR_2018_paper.html.
- Ian J. Goodfellow et al. "Generative Adversarial Nets". In: *NIPS*. Ed. by Zoubin Ghahramani et al. 2014, pp. 2672–2680.
- Jonathan Ho and Stefano Ermon. "Generative Adversarial Imitation Learning". In: *NIPS*. 2016, pp. 4565–4573.
- Mehdi Mirza and Simon Osindero. "Conditional Generative Adversarial Nets". In: *CoRR* abs/1411.1784 (2014).
- Xi Chen et al. "InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets". In: *NIPS*. 2016, pp. 2172–2180.