One Shot Learning: Meta-learning approaches

Armand Nicolicioiu
University Politehnica of Bucharest
armand.nicolicioiu@gmail.com

1. Introduction

Visual classification tasks can be solved with high accuracy by using large convolutional neural network models. Existing models require a large training dataset, increased computing power and often result in a relatively high inference time. Humans show the ability to acquire information fast and generalize a visual concept only from a few examples, or even one. A computational model that behaves this way is desirable and makes the object of study of One Shot Learning.

2. Visualization

Matching Networks — Prototypical Networks — Model-Agnostic Meta-Learning

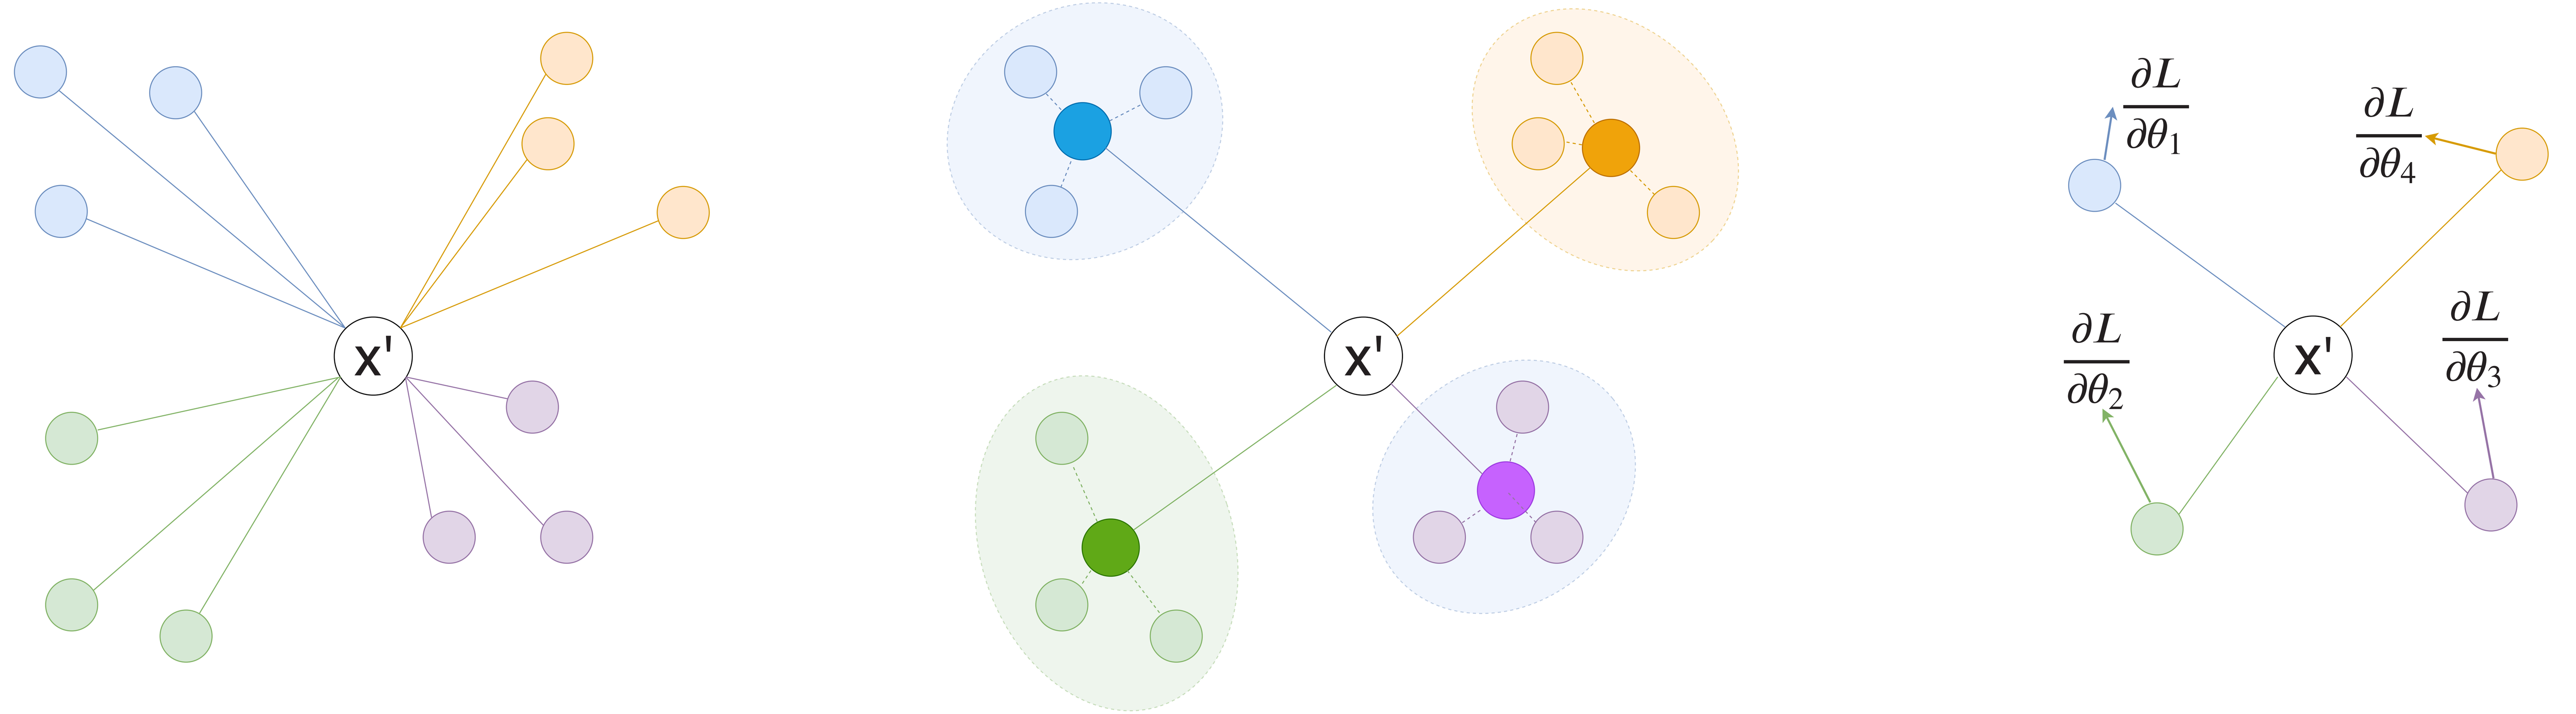

3. Matching networks [1]

- Learns an embedding function f
- For a new example x':
  1. compute distances between x' and all the examples in the training set in the embedding space
  2. compute predictions as a linear combination of the labels in the training set weighted by the similarities
- Update θ_f in order to minimize error on validation set

4. Prototypical networks [2]

- Learns an embedding function f
- For a new example x':
  1. compute the centers in the embedding space as the mean of all examples from each class
  2. compute predictions as the similarities with the centers in the embedding space
- Update θ_f in order to minimize error on validation set

5. Model-Agnostic Meta-Learning [3]

- Learns an embedding function f
- Add another fully connected layer that makes the prediction
- Inner Loop: obtain θ' that lowers error on D^{tr}.
- Outer Loop:
  1. predict on D^{val} using θ'
  2. compute $\nabla_\theta \mathcal{L}^{val}$
  3. update θ using this gradient

6. Learning trial

N-way K-shot classification task

- N possible classes
- K examples from each class
- the classes can differ between learning trials

Each task $\mathcal{T}_i \sim p(\mathcal{T})$ contains:

- a **training set** $\mathcal{D}^{tr}$
  - contains K examples from each of the N possible classes
  - used to train a N-way prediction model
- a **validation set** $\mathcal{D}^{val}$
  - contains other examples from the same N classes
  - used to estimate the generalization performance of the model trained on the $N * K$ examples

The total set of tasks is divided into:

- a **training meta-set** of tasks, $\mathcal{S}^{tr}$
- a **validation meta-set** of tasks, $\mathcal{S}^{val}$
- a **test meta-set** of tasks, $\mathcal{S}^{test}$

7. MiniImagenet Results 5-way

Method	Metric	1-shot	5-shot
Matching Nets	Dot Product	45.10	58.80
Matching Nets	Euclidean	46.96	60.40
Prototypical Nets	Dot Product	45.10	60.38
Prototypical Nets	Euclidean	46.96	62.10
MAML	Implicit Dot	44.59	-
MAML–difBN [4]	Implicit Dot	45.81	-
MAML–Matching	Dot Product	46.06	-
MAML–Matching	Euclidean	45.80	-
MAML–Proto	Dot Product	46.06	59.09
MAML–Proto	Euclidean	45.80;	62.06

8. MiniImagenet

9. References

[1] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. NeurIPS2016.

[2] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. NeurIPS 2017.

[3] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. ICML2017.

[4] Antreas Antoniou, Harrison Edwards, and Amos Storkey. "How to train your MAML."