Annonce

2 mai 2017

Intricate Learning and Storing in Deep Neural Networks

Catégorie : Doctorant

Dans cette rubrique

- Flux RSS
- Actualités
- Propositions de poste
- Evénements
- Appels à projets
- Conférences
- Ecoles thématiques
- Publications

Context

In terms of number of operations per second (and disregarding the fact the type of operations they perform are very different), modern computers are now on a par with the human brain. In this context it is interesting to note that there is still a significant difference between the problems that are considered difficult for artificial intelligence and the corresponding ones for natural intelligence. This finding motivates for a neuro-inspired approach, given the fact the human brain is the only device at our disposal capable of natural intelligence.

Nowadays, the vast majority of proposed artificial neural networks techniques relies on the denotational approach (including deep neural networks). The idea is that many challenging problems in machine learning can be seen as looking for a function that takes as input some vector x and associate it with an output vector y. Among the infinite space of such functions, one of them is ideal for a given problem. For example, there exists such a function that associates with any input photo a value 1 if the photo contains at least one pixel of a cat and 0 otherwise. Finding this function is ill-posed, as we only have a finite amount of information to train the system. Denotational Neural network approaches consist in constraining functions to a given subset, then finding within this restricted set one that performs well on given examples by using optimization techniques, while expecting it to be a fair approximation of the ideal function.

This strategy led to state-of-the-art solutions for many machine learning problems, including winning Go against human champions and classifying photos of objects in thousands of categories. However, there are still many open problems in Machine Learning. Firstly, accurate neural networks usually contain millions of parameters to be trained, requiring a large number of examples. As a consequence, there is still much progress to do in unsupervised and semi-supervised learning. Secondly, systems are trained to match input with corresponding outputs, with no control on intermediate layers. Therefore there is no explicit notion of information -- neither of semantics. Thirdly, training takes a lot of resources, leading to difficulties in dealing with novelty.

Instead of looking at neural networks as processing elements, it is also possible to consider them as devices able to store, communicate and manipulate pieces of information. This original point of view led to the proposal of other families of artificial neural networks, including self-organizing maps and associative memories. These architectures of neural networks are known to perform very well when facing indexation, memory and/or clustering problems, even nearing the limits given by Shannon's entropy.

The human brain is unarguably able to both compute and memorize, and these functionalities are not split as in Von Neumann's Model. But to our knowledge little literature has been devoted to artificial models targeting both information and computation. The object of the project is to investigate such models, with the hope of achieving architectures of neural networks able to reach state-of-the-art performance in both supervised and unsupervised challenges, while using incremental learning.

Project

With this goal in mind, a promising avenue of research consist in combining binary associative memories and deep neural networks. As a matter of fact, binary associative memories are devices that can be shown mathematically to achieve optimal performance in storing. They do not require training, and they embody pieces of information in graphical binary patterns in neural networks, which means that they use the existence of connections and not their values. On the other hand, deep neural networks are devices that are known to reach state-of-the-art in many supervised learning challenges. They capture data thanks to their millions of parameters, which are values of connection weigths. The key idea of the project is to propose hybrid neural networks in which both existence and value of connections are used, the first ones to embody pieces of information and the other one to capture processing capabilities.

In order to tackle this ambitious problem, the project proposes to first combine deep neural networks and associative memories in different layers of artificial neural networks. Namely, the idea is to use state-of-the-art techniques to binarize (or coarsely quantize) the activities of neurons in deep layers. Then, following the same ideas as in transfer learning (or fine tuning) methods, to adjoin binary associative memories. In this first step, already trained deep neural networks will be used (such as Inception), while the associative memory will store examples of a novel dataset using the deep neural network as a feature extractor. The hope is to obtain systems able to reach state-of-the-art accuracy on supervised learning benchmarks, while using an incremental binary learning method. The second step consists in designing layers able to both memorize and learn, either by removing connections or refining weights on novel examples given to the neural network.

Venue and dates

The PhD will take place in IMT-Atlantique (a french "Grande École") in Brest for a start in October 2017 and a duration of three years.

Details

As a PhD Student, the net salary is about 1450€ a month (tuition fees are very low: ~400€ a year). It is expected that the candidate has a solid background in signal processing and/or computer science and is familiar with deep learning methods and libraries. The candidate must have a Master level or equivalent.

Contact

If you are interested, please send an email **before May 23rd** to vincent.gripon@imt-atlantique.fr along with 1) a

motivation letter, 2) an up-to-date CV and 3) a recommendation letter.

Retour