Two-year post-doc offer: Active few-shot learning for fault detection in technical equipment

Contacts:
Jean Ponce (INRIA, jean.ponce@inria.fr)
Philippe Calvez (ENGIE, philippe.calvez1@engie.com)

Background for the research
ENGIE is one of the world leaders in the fields of energy and the environment. Strongly involved in the energy transition and in the development of renewable energies, ENGIE has numerous infrastructures that must be managed and maintained in good working order. This maintenance notably involves the detection of malfunctions, but also the detection of faults on technical equipment. For these tasks, artificial intelligence can prove to be very interesting by providing automatic solutions to detect these visual defects. However, this requires capturing images as well as annotating them. However, the acquisition of images is not always easy, requiring specific installations adapted to the different cases. Likewise, labeling defects is takes time and requires experts. Acquiring data, particularly labeled data, can quickly become quite costly, but the progress that artificial intelligence has made in exploiting this data is very promising.

Objectives
The proposed post-doc project aims a harnessing the power of artificial intelligence to take into account material and financial constraints as well as time and personal constraints. The objective is to develop an active learning method coupled with a few-shot learning approach with applications notably to the visual detection of faults in technical equipment. The principle of active learning methods is to identify and select the most relevant images to provide to annotators with only non-redundant high-quality data in order to minimize their labeling time. This allows the model to learn with less data while maintaining similar performance. The process is iterative: as new image selections are made by active learning and their labeling by annotators, the model is updated, then the active learning algorithm selects a new batch of images which is then annotated and injected into the model training set and so on. Few-shot learning approaches learn from a small sample of data for each of the classes. In addition, they must allow not only the learning of the classes initially planned, but also learning for new classes. Few-shot learning therefore not only makes it possible to get rid of the ever increasing amount of data necessary for learning artificial intelligence models, but also to get rid of the limitation of learning for the fixed number of classes initially considered. The idea of the post-doc is therefore to develop a method which selects in a relevant way a small number of relevant images to annotate and which will form the small learning set of a few-shot learning approach to develop. The detection of visual defects (images from cameras, faults on a wind turbine blade, solar panel, etc.) on equipment will be one of the applications of the group on which the approach will be validated.

More on active learning and few-shot learning
Active learning

With the development of artificial intelligence and the need for data and in particular for labeled data, the interest in active learning is growing. In particular, with the explosion of the use of deep learning techniques, the demand for data for these approaches is always greater, but these are rarely available and their quality is not always in line with the needs. Since annotation is expensive, selecting the relevant data to annotate is important. Active learning methods applicable to the vast majority of machine learning methods exist. Indeed, some methods rely only on the probabilities returned at the output of a model. We can cite for example the "least confident" approach which consists in selecting the data not yet labeled for which the model predictions are not certain or the "query by committee" approach which selects the data for which the predictions of several models differ the most. Burr Settles' paper [1] outlines a number of active learning methods and associated issues. More recently, researchers have looked at task-specific active learning methods. Take the example of object detection, which can be seen as localization coupled with object classification. The paper [2] proposes an active learning method which takes into account not only the probabilities associated with each class for a given bounding box, but also considers the model's certainty as to the location of the box. Similarly, in [3], the architecture of the SSD network [4] is exploited using different convolutional layers as being as many different outputs as if they were different models which must vote on the location of the bounding boxes, comparable to a "query by committee" method. Some researchers [5] suggest relying instead on annotations of the verification type with questions such as: is the box found by the model well predicted? is it badly predicted? rather than asking the annotator to actually locate the bounding boxes. The paper [6] highlights the advantage of their active learning method for learning new classes that were not initially learned.

Few-shot learning

The success of neural networks and deep learning is indisputable, especially for computer vision. However, these methods require large amounts of data. Furthermore, most of the approaches proposed today only consider a well-defined number of classes. So, for object detection, these methods recognize only the objects present in the learning set, and it is not trivial to adapt them to recognize new objects. One-shot learning and few-shot learning aim to learn with little data. In the case of one-shot learning, this is the extreme case of learning with only one element per class. For the few-shot learning, there are a few instances per class. In both cases (one-shot and few-shot), the method must make it possible to learn new classes when they appear. The first occurrence of one-shot learning dates from 2006 with the paper by Fei Fei Li [7] where a Bayesian approach is proposed. In 2015, the Omniglot dataset was made available to encourage research on the few-shot learning. It is made up of images of letters from 50 alphabets. 20 annotators each draw a single letter each time. Each class is therefore poorly represented. An approach with a Siamese network [8] coupled with a loss function favoring bringing together positive pairs and separating negative ones makes it possible to deal with the one-shot question. Indeed, the model evaluates whether two elements are of the same class or not, which consequently makes it possible to identify new classes. The paper [9] proposes an approach where the neural network has a more flexible external memory which would allow to better apprehend and recognize the properties of known classes and those of new classes.

Post-doc setting

The post-doc will be carried out within the Willow team at the INRIA Center in Paris, under the supervision of Jean Ponce (INRIA) and Yuxiong Wang (University of Illinois at Urbana-Champaign), in collaboration with Philippe Calvez of the CSAI laboratory (ENGIE Computer Science and Artificial Intelligence). Willow (http://www.di.ens.fr/willow) is an internationally recognized research team that addresses a wide range of problems in artificial vision. CSAI is a research and development team that produces solutions for many ENGIE entities, particularly around the problems of computer vision and automatic natural language processing.

The post-doc's missions will be as follows:

- Conduct a bibliographic study on active learning methods as well as on few-shot learning methods for computer vision.
- Make a benchmark of existing methods on one or more public datasets.
- Analyze and compare these methods in order to synthesize their advantages and shortcomings.
- Suggest avenues for work and improvement in relation to the state of the art.
- Propose a combined method of active learning and few-shot learning.
- Implement the improvements identified in the environments and architectures of the CSAI lab. Specific use cases (associated datasets) in the ENGIE business areas in connection with the CSAI Lab research work may be considered for this implementation.
- Validate the contributions by appropriate experiments.
- Write and submit scientific publications.
- Communication and oral presentation of the work.
- Document the code.

The post-doc will be done in 2 years with a flexible start date.

The application must be sent as soon as possible to:

- Jean Ponce (jean.ponce@inria.fr)
- Philippe Calvez (philippe.calvez1@engie.com)

It must contain the following documents:

- a detailed CV (all experiences and technologies mastered)
- a cover letter
- letters of recommendation (optional)