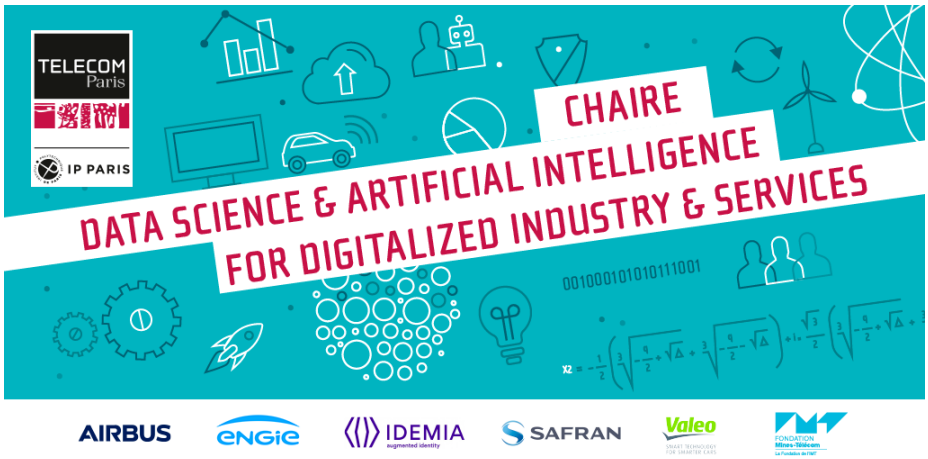

Paris, le xx/xx/xx

Offre de stage

Sujet : Active learning at the lowest cost

Possibilité de poursuivre sur une thèse

La Chaire Data Science and Artificial Intelligence for Digitized Industry and Services

Portée par Florence d'Alché-Buc, enseignante-chercheur dans le département Image, Données, Signal de Télécom ParisTech, la chaire DSAI réunit cinq partenaires industriels : Airbus Defence & Space, Engie, Idemia, Safran et Valeo Finance. Son objectif général est de développer, en liaison étroite entre les Parties, une formation et une recherche de niveau international.

Ses quatre principaux axes de recherche sont :

1. Analyse et prévision de séries temporelles (Predictive Analytics on Time Series) ;

2. Exploitation de données hétérogènes, massives et partiellement étiquetées (Exploiting Large Scale and Heterogeneous, Partially Labelled Data) ;

3. Apprentissage pour une prise de décision robuste et fiable (Learning for Trusted and Robust Decision) ;

4. Apprentissage dans un environnement dynamique (Learning through Interactions with a Changing Environment).

Description du stage

Encadrement

Florence d'Alche, Ekhine Irurozki

Lieu et dates du stage

Telecom Paris, 19 Place Marguerite Perey, 91120 Palaiseau

Date de début du stage : Mai 2022

Équipe(s) d'accueil de la thèse

département DSI, équipe S2A

Mots clés

Active learning, multi-class learning, structured output, partially labelled data,

Sujet détaillé

1 Goal and problem setting

Object recognition tasks are crucial in several applications such as self driving cars. Frequently, we have access to a small labeled dataset while a very large number of unlabeled data are available. Object recognition from partially labeled data can be addressed either by semi-supervised learning or active learning. We consider here active learning [2] which consists in choosing sequentially datapoints to label and re-training the classifier with the augmented labeled dataset. We further challenge this setting by assuming a structured-multiclass scenario [3]. In the self-driving car example this means that each image can be recognized as one out of many object, for example it is either a tree, or another car, a traffic signal, . . . By structured we denote the fact that the classes can be grouped hierarchically, for example, all the traffic signals form a group. Another specificity of our challenge is that the pool of annotators is hetereogeneous: some of the them are experts annotators and more expensive and some are not. The final goal is to choose the sequence of pairs (annotator, question) in such a way that the total price for all the questions is minimized while the accuracy is larger than a given threshold.

2 Challenges

Active learning can be described in the context of Reinforcement Learning [4]. The inherent challenges are related to the action space representation, the definition of the reward, the scaling issue and theoretical guarantees about the approach.

In this internship we consider a multi-armed contextual bandit approach [1] which enjoys theoretical guarantees and provides a very general framework to learn how to select data points and query annotators. Each arm (which corresponds to a possible action) is associated to a context vector describing the encoded action. A Thomson sampling approach can be applied to sample the bandit arms and learn the weight vector applied to the context vector through reward maximization. A first challenge consists in modeling the problem by choosing an appropriate definition of the reward and the context vector. The reward is multi-objective, relying on accuracy of the learned classifier and the annotation budget devoted to training. The context vector which must encode information about the pair (datapoint, annotation question) including the relevance of the datapoint and the interest of a question concerning the datapoint. An on-going work has shown that a two-step approach, picking a data point and then selecting an annotator gives promising results, leveraging two distinct bandits, one devoted to data selection [2] and the other to annotation question. An interesting question will be to improve these first results, either by improving modeling and data/question representation or by designing a single multi-armed bandits with more complex context vectors.

The overall approach is parametrized by the choice of the cost of annotation questions and the nature of the questions to ask (binary, one superclass against the other superclasses, one class against another one). A second challenge is therefore to provide tools to measure the impact of these choices on the algorithm performance.

3 Practical information

This internship will take place in Telecom Paris following an ongoing research line on the group. Moreover, there are several open research directions suitable for a PhD.

Profil du candidat

Etudiant titulaire d'un master 2 rechercheM

- Apprentissage statistique / reconnaissance des formes
- Traitement de la parole, traitement du langage naturel
- Bon niveau en programmation (Python)
- Bon niveau d'anglais

Candidatures

A envoyer à florence.dalche@telecom-paris.fr, irurozki@telecom-paris.fr

- Curriculum Vitae
- Lettre de motivation personnalisée expliquant l'intérêt du candidat sur le sujet (directement dans le corps du mail)
- Relevés de notes des années précédentes
- Contact d'une personne de référence

Les candidatures incomplètes ne seront pas examinées.

Références

[1] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In International conference on machine learning, pages 127–135. PMLR, 2013.

[2] Djallel Bouneffouf, Romain Laroche, Tanguy Urvoy, Raphael Feraud, and Robin Allesiardo. Contextual bandit for active learning: Active Thompson sampling. In International Conference on Neural Information Processing, pages 405–412. Springer, 2014.

[3] Peiyun Hu, Zachary C Lipton, Anima Anandkumar, and Deva Ramanan. Active learning with partial feedback. arXiv preprint arXiv:1802.07427,2018.

[4] Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction. MIT press, 2018.