HyAIAI Post-doc
Declarative Constraints for privacy-friendly ML Systems

June 24, 2020

1 Context

This position is funded in the context of the HyAIAI "Hybrid Approaches for Interpretable AI" INRIA project lab (https://project.inria.fr/hyaiai/). With this subject, we would like to investigate how declarative languages (e.g. languages stating constraint satisfaction problems, or query languages) can help clarifying both what is expected of the model and what the model truly satisfies. These are aspects of the "model interpretability" not yet tackled within the HyAIAI project.

Keywords Constraint Optimization, Constraint Programming, privacy, Explainable AI, Declarative Languages.

1.1 Supervision and position location

- Jan Ramon (jan.ramon@inria.fr), INRIA lne - MAGNET Team (Lille, France)
- Elisa Fromont (elisa.fromont@irisa.fr), IRISA/INRIA rba – Lacodam team (Rennes, France)
- Siegfried Nijssen (siegfried.nijssen@uclouvain.be) - UCLouvain (Louvain-la-Neuve, Belgium)

Location INRIA LNE, team MAGNET, Lille. The recruited post-doc will be based in Lille with the MAGNET Team. (S)he will however visit the teams in Rennes and Louvain-la-Neuve and participate to the HyAIAI meetings in Paris.

1.2 Recruitment

The recruited person will benefit from the expertise of the MAGNET team on declarative languages, privacy (https://team.inria.fr/magnet/), of the LACODAM team (https://team.inria.fr/lacodam/) in declarative languages and interpretable AI and of the expertise of Siegfried Nijssen in constraint programming and its applications in machine learning and data mining (https://www.info.ucl.ac.be/ snijssen/).

We would ideally recruit a post-doc for 1 year (with possibly one additional year) with the following preferred skills:

- Knowledgeable in constraint programming
- Knowledgeable in machine learning in general
- Good programming skills
- Very good English (understanding and writing)

However, good applications by candidate PhD students will also be considered and, in this case, the position will last 3 years. The position will be funded by INRIA (https://www.inria.fr/en/). See the INRIA web site for the post-doc and PhD wages.

The candidates should send a CV, 2 names of referees and a cover letter to the three supervising researchers mentioned above. Please indicate if you are applying for the post-doc or the PhD position. The selected candidates will be interviewed before August 2020 for an expected start in October 2020.

2 Subject

Differential privacy is the most well-known measure of privacy [1], but several generalizations have been proposed to allow for more precise modeling of privacy properties (see e.g. [3]). Languages have been proposed to describe algorithms and their privacy requirements, often together with techniques for automatically proving that actions satisfy privacy requirements (see e.g., [4]), but theorem provers are often limited to a small class of rather simple problems. In this project, we want to follow another direction, compiling privacy specifications to constraint programs that can then be checked for satisfiability and/or optimized by standard solvers. Such approaches have yielded good results in the past in several applications, e.g., for pattern mining or for optimizing decision trees .

References

[1] Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

[2] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5):93:1–93:42, 2019.

[3] Daniel Kifer and Ashwin Machanavajjhala. Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems*, 39(1):3:1–3:36, 2014.

[4] Joseph P. Near, David Darais, Chike Abuah, Tim Stevens, Pranav Gaddamadugu, Lun Wang, Neel Somani, Mu Zhang, Nikhil Sharma, Alex Shan, and Dawn Song. Duet: An Expressive Higher-order Language and Linear Type System for Statically Enforcing Differential Privacy. *Proc. ACM Program. Lang.*, 3(OOPSLA):172:1–172:30, October 2019.

[5] Helene Verhaeghe, Siegfried Nijssen, Gilles Pesant, Claude-Guy Quimper, and Pierre Schaus. Learning optimal decision trees using constraint programming. In *Principles and Practice of Constraint Programming*, 2019.