Master 2 internship in Artificial Intelligence and Operations Research

Convex optimization for solving Graphical Models

Supervisors:
Thomas Schiex, MIAT, INRA, thomas.schiex@inra.fr
Simon de Givry, MIAT, INRA, simon.de-givry@inra.fr
George Katsirelos, MIA Paris, INRA, gkatsi@gmail.com
Internship level: Master 2 in Computer Science or Applied Mathematics
Salary: ≈ 550 euros / month
Duration: 4–6 months
Location: Institut National de la Recherche en Agronomie (INRA)
Laboratoire de Mathématiques et Informatique Appliquées de Toulouse (MIAT)
24, Chemin de Borde-Rouge, Castanet-Tolosan (near Toulouse)
Requested skills: Algorithms, Theory of linear and integer programming, C++

Constraint programming (CP [Rossi et al, 2006]) is an AI *Automated Reasoning* technology with tight connections with propositional logic. It offers a problem modeling and solving framework where the set of solutions of a complex (NP-hard) problem is described by discrete variables, connected by constraints (simple Boolean functions).Together with propositional satisfiability, it is one of the automated reasoning approaches of AI, where problems are solved *exactly* to provide rigorous solutions to hardware or software testing and verification, system configuration, scheduling or planning problems.

Discrete Stochastic Graphical Models (GMs [Koller et al, 2009]) define a *Machine Learning* technology where a probability mass function is described by discrete variables, connected by potentials (simple numerical functions). GMs can be learned from data and the NP-hard problem of identifying a Maximum a Posteriori (MAP) labelling is often solved *approximately* to tackle several problems in Image [Kappes et al, 2013] and Natural Language Processing [Bilmes, 2004], among others.

The Cost Function Network framework [Cooper et al, 2010] with its associated C++ open source award-winning solver toulbar2, http://www.inra.fr/mia/T/toulbar2, developed in our team, combines the ideas of Constraint Programming and Stochastic Graphical Models. By solving the so-called Weighted Constraint Satisfaction problem, toulbar2 is capable of simultaneously reasoning on logical information described as Boolean functions and gradual, possibly Machine Learned, information described as local numerical functions.

To properly deal with the available information, the solver relies on a guaranteed Branch and Bound-based algorithm [Allouche et al, 2015] where pruning follows from efficient algorithms, processing local information, known as “local consistency filtering” in CP or “message passing” in GMs [Koller et al, 2009]. Because feasibility alone is NP-complete in CP, efficiency is crucial and depends a lot on the *strength* and *computational cost* of the pruning mechanisms used during search. Our experience is that the ideal compromise needed to solve a specific problem depends on the problem to be solved. Most existing pruning mechanism are very fast but not always sufficiently strong enough for the hardest problems.

The aim of this Master internship is therefore to explore new compromises between pruning strength and computational cost by considering new strong pruning mechanisms. The local consistency enforcing algorithms used inside toulbar2 for pruning are now well-understood as fast incremental approximate solvers of the dual of a specific linear program (LP) known as the "local polytope" [Cooper et al, 2010]. This essential pruning mechanism can be reinforced by strengthening existing bounds so that the approximate dual solution is tighter or satisfies additional constraints such as Semi-Definite-Positivity [Nesterov, 2015]. While similar directions have been partly explored in the context of approximate resolution in Image Processing [Savchynskyy et al, 2012], very little has been done in the context of Branch and Bound resolution, mixing logical and numerical information [Guibas et al 2014, Wang et al 2015, Elloumi et al, 2019].

The candidate will explore this direction, develop and implement new algorithms inside toulbar2 and experiment with them on large collections of real problem instances, many of which are not known to be currently solvable.

Candidate profile

The subject is at the intersection of Propositional Satisfiability, Constraint programming, Operations Research. The ideal candidate should be familiar with Constraint Programming or Propositional Satisfiability algorithms and be mathematically comfortable with the theory and algorithms of linear optimization and duality. Some of the algorithms developed during the internship will be implemented and empirically tested on large benchmarking problems sets, https://forgemia.inra.fr/thomas.schiex/cost-function-library, possibly using computing clusters. This requires a good programming ability (in C++, some additional proficiency in scripting language such as python being welcome).

Send by email to the Master supervisors in French or English your detailed CV, a motivation letter, and transcripts of Licence and Master 1 degrees. Reference names (professors to contact) will be a plus.

There will be an opportunity to continue for a PhD inside the ANITI Toulouse institute for AI https://www.univ-toulouse.fr/ANITI.

References

D Allouche, S. de Givry, G. Katsirelos, T. Schiex, M. Zytnicki. Anytime Hybrid Best-First Search with Tree Decomposition for Weighted CSP. In Proc. of CP-15, pp 12-28, Cork, Ireland, 2015.

J Bilmes. Graphical Models and Automatic Speech Recognition. Mathematical Foundations of Speech and Language Processing pp 191-245, 2004.

M Cooper, S. de Givry, M. Sanchez, T. Schiex, M. Zytnicky, T. Werner. Soft Arc-consistency revisited. In Journal of Artificial Intelligence, 174(7-8), pp 449-478, 2010.

M Erdogdu, Y Deshpande, A Montanari. Inference in graphical models via semidefinite programming hierarchies. In Proc. of NeurIPS, pp 417-425, 2017.
Sourour Elloumi, Amélie Lambert, Arnaud Lazare. Solving unconstrained 0-1 polynomial programs through quadratic convex reformulation. Hal-01872996v3, 2019.

Qixing Huang, Yuxin Chen, Leonidas Guibas. Scalable Semidefinite Relaxation for Maximum A Posterior Estimation. In Proc. of ICML, 2014.

J Kappes, M Speth, G Reinelt, and C Schnorr. Towards efficient and exact MAP-inference for large scale discrete computer vision problems via combinatorial optimization. In Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, 2013.

D Koller, N Friedman. Probabilistic graphical models: principles and techniques. The MIT Press, 2009.

Y Nesterov. Introductory Lectures on Convex Programming.

F Rossi, P Van Beek, and T Walsh. Handbook of constraint programming. Elsevier, 2006.

B Savchynskyy, S Schmidt, JH Kappes, and C Schnörr. Efficient MRF energy minimization via adaptive diminishing smoothing. In Proc. of UAI-12, Catalina Island, CA, 2012.

P Wang, Chunhua Shen, Anton van den Hengel, Philip Torr. Efficient Semidefinite Branch-and-Cut for MAP-MRF Inference. 2015.