Report on the Artificial Intelligence journal Special issue on Combining Constraint Solving with Mining and Learning

Tias Guns

IJCAI17 AI & CP workshop

Context

AIJ Special issue: 17 papers, published 03/2017

Related events:

- Dagstuhl 11201 (2011): Constraint Programming meets Machine Learning and Data Mining
- ECAI 2012 workshop CoCoMile
- AAAI 2013 workshop CoCoMile
- Dagstuhl 14411 (2014): Constraints, Optimization and Data

Similar in spirit:

- Euro working group on Data Science meets Opt. (ds-o.org)
- CP 2017 track on CP & Machine Learning

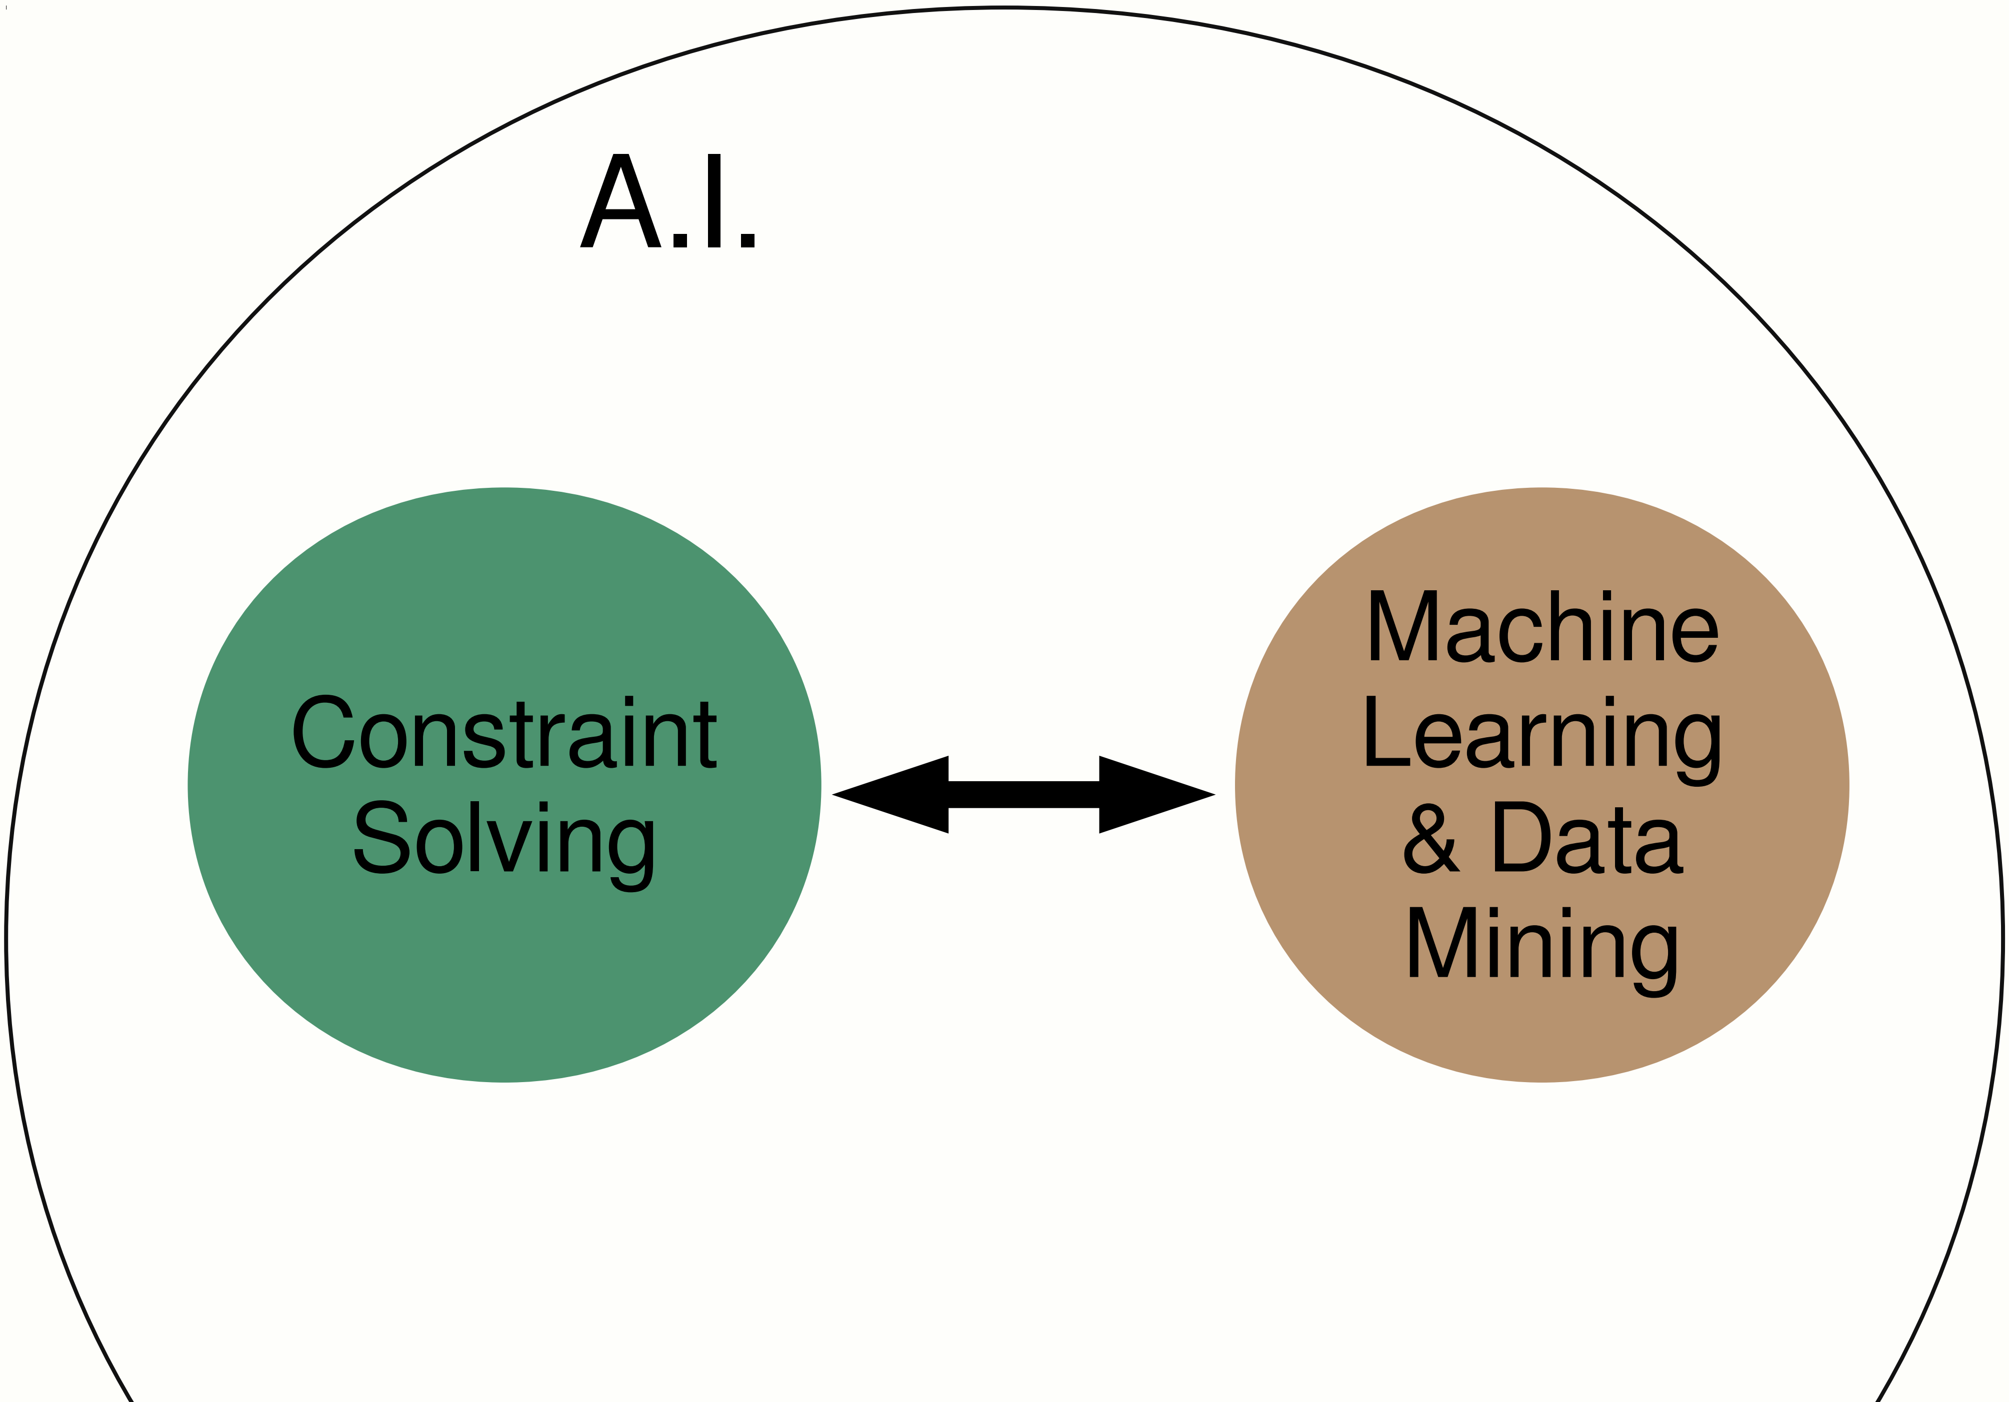
A.I.
Constraint
Solving
Machine
Learning
& Data
Mining

Constraint Solving

?

Machine Learning & Data Mining

Combinatorial problem solving

- constraint satisfaction/enumeration
- constraint optimisation

Extracting regularities from data

- repeating patterns
- function that discriminates classes
- function that estimates a value

Constraint Solving

!

Machine Learning & Data Mining

Combinatorial problem solving

- constraints and objectives are functions
- search generates data

Extracting regularities from data

- often is combinatorial optimisation

  \+ more side constraints
  (in pattern mining, clustering, Bayesian networks)

  \+ more structure
  (in structured output prediction)

AIJ special issue topics:

1) Constraint Solving for ML & DM

2) ML & DM for Constraint Solving

Constraint Solving for ML & DM

Constraint-based itemset mining

Find all *frequent* sets of items in a dataset

- Modeling language for constrained itemset mining
  [MiningZinc: a declarative framework for constraint-based mining. T. Guns, A. Dries, S. Nijssen, G. Tack and L. De Raedt]
- Top-k SAT, for constrained itemsets and sequences
  [Mining Top-k Motifs with a SAT-based Framework. S. Jabbour, L. Sais and Y. Salhi]
- Pareto front, equivalences and DynCSP
  [Skypattern Mining: from Pattern Condensed Representations to Dynamic Constraint Satisfaction Problem. W. Ugarte, P. Boizumault, B. Cremilleux, A. Lepailleur, S. Loudni, M. Plantevit, C. Raissi and A. Soulet]

Constraint Solving for ML & DM

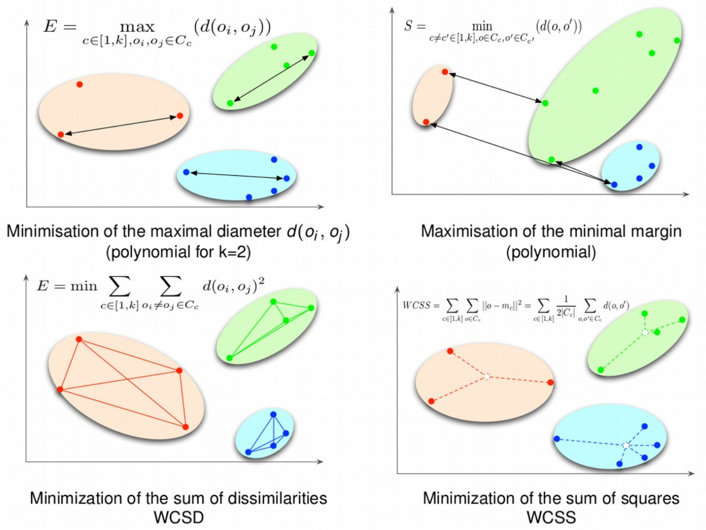

<u>Constrained clustering</u>

Find a partitioning based on similarity

- Centroid-based (like k-means), CP
  [Constrained Clustering by Constraint Programming. T-B-H Dao, K-C Duong and C. Vrain]
- Hierarchical clustering, MIP
  [A Flexible ILP Formulation for Hierarchical Clustering. S. Gilpin and I. Davidson]
- Correlation clustering, MaxSAT
  [Cost-Optimal Constrained Correlation Clustering via weighted Partial Maximum Satisfiability. J. Berg and M. Jarvisalo]

Constraint Solving for ML & DM

ML with hard constraints

Find *f*, *f*(X) approximates y for given (X,y) data

Structured X or structured y: adds structural constraints

- Predict structured output from structured input, OptSMT
  [Structured Learning Modulo Theories. S. Teso, R. Sebastiani and A. Passerini]
- Sampling predictive structures, by learning quality function
  [Learning an efficient constructive sampler for graphs. F. Costa]

Constraint Solving for ML & DM

<u>ML with hard constraints</u>

Find *f*, *f*(X) approximates y for given (X,y) data

Learning <u>with side-constraints</u>

- Kernel machines and logical constraints between predicates
  [Semantic-Based Regularization for Learning and Inference. M. Diligenti, M. Gori and C. Saccà]
- Bayesian Networks, with partial data and known equivalences therein.
  [Learning Bayesian Networks under Equivalence Constraints. T. Yao, A. Choi and A. Darwiche]

Constraint Solving for ML & DM

<u>ML with hard constraints</u>

Find *f*, *f*(X) approximates y for given (X,y) data

Learning <u>as constrained optimisation</u>

- Learn structure of Bayesian Network with MIP
  [Integer Linear Programming for the Bayesian Network Structure Learning Problem. M. Bartlett and J. Cussens]
- Use *relational* LP for linear learning problems
  [Relational Linear Programming. K. Kersting, M. Mladenov and P. Tokmakov]

ML & DM for Constraint Solving

<u>Learning to solve better</u>

- Automatic Construction of Parallel Portfolios via Algorithm Configuration
  [M. Lindauer, H. Hoos, K. Leyton-Brown and T. Schaub]
- Algorithm Recommender System. Application to Algorithm Portfolio Selection
  [M. Misir and M. Sebag]

ML & DM for Constraint Solving

<u>Learning better models</u>

Learning constraints

- Constraint Acquisition
  [C. Bessiere, F. Koriche, N. Lazaar and B. O’Sullivan]

Learning objective functions

- Learn NN/DT, use in CP,LS,MINLP,SMT
  [Empirical Decision Model Learning. M. Lombardi, M. Milano and A. Bartolini]
- Learn DT/RT, use in MIP
  [Auction Optimization using Regression Trees and Linear Models as Integer Programs. S. Verwer, Y. Zhang and Q. Chuan Ye]

Constraint Solving

Overview

Machine Learning & Data Mining

→ Constraint Solving for ML & DM

- Constraint-based itemset mining
- Constrained clustering
- ML with hard constraints
  - structured input/output
  - side-constraints
  - with constraint solver

← ML & DM for Constraint Solving

- Learning to solve better
- Learning to model better (constraints and objectives)

Trends?

<u>In constrained pattern mining and clustering:</u>

- Typical problems:
  - Discrete problems
  - With additional, complex, side-constraints
  - Worthwhile to enumerate all solutions or find a near-optimal one
- Model classic setting so that additional constrains can be added
- Global constraints!
- Novel problem settings (too hard to write algo from scratch)

Trends?

In the machine learning papers:

- Additional constraints that are:
  - Hard
  - Not relaxable
  - Over discrete structures

- Some use a *solver*, some don't
- Harder to keep track of (different focus)

Trends?

Learning for solving

- Most mature topic
- Known under many names in diff. communities!
  - Algorithm selection/configuration [CP, SAT]
  - Hyperparameter Optimisation [ML]
  - Black box optimisation
  - Reinforcement learning...

All are advancing very rapidly and with moderately increasing interaction

Trends?

Learning constraints and objectives

Many open questions:

- hard/soft constraints?
- error/uncertainty of predictions?
- function complexity vs solving complexity? (c.f. *tractable* learning in ML)

Outlook

- More structured problems in ML/DM
- More uncertainty in combinatorial problems
  (from the problem domain or learned functions)
- More reinforcement learning everywhere
  (sequential decision making under uncertainty)