

# PUSHING DATA INTO CP MODELS USING GRAPHICAL MODEL LEARNING & SOLVING

*CP 2020*



CÉLINE BROUARD<sup>1</sup>, S. DE GIVRY<sup>2</sup> & T. SCHIEX<sup>2</sup>

---

<sup>1</sup> Université Fédérale de Toulouse, INRAE MIAT, UR 875, Toulouse, France

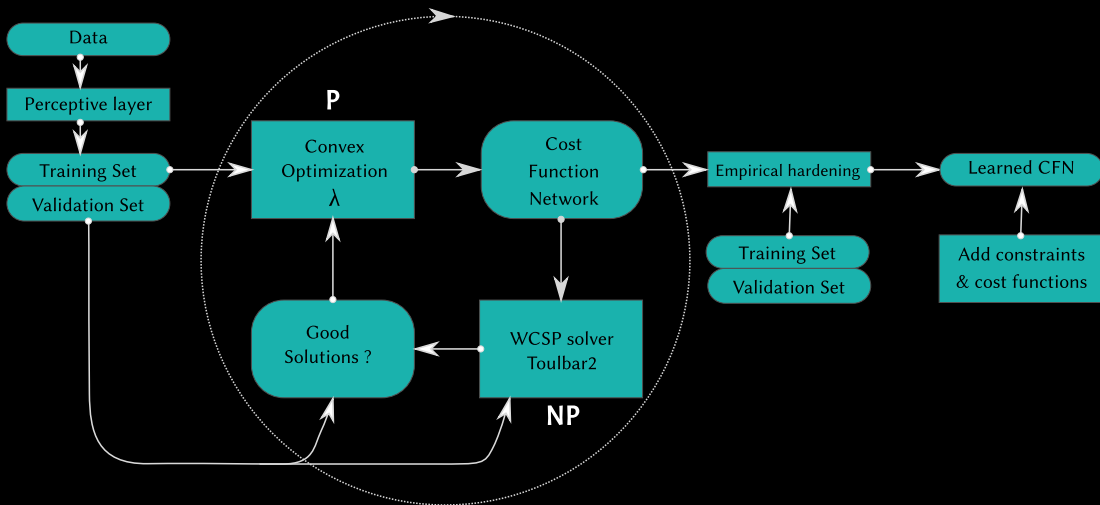
<sup>2</sup> Université Fédérale de Toulouse, ANITI, INRAE MIAT, UR 875, Toulouse, France

CP AND ML TRACK

SEPTEMBER 2020



# LEARNING A COST FUNCTION NETWORK FROM HIGH-QUALITY SOLUTIONS



## You'll learn

- how we use graphical models to connect CP with probabilistic Machine Learning
- how the NP-hard regularization loop can be made practical
- how we learn playing the Sudoku from images (without rules)
- how it compares with DL architectures that “learn to reason”
- how we can combine learned user preferences with (car) configuration constraints

## You'll learn

- how we use graphical models to connect CP with probabilistic Machine Learning
- how the NP-hard regularization loop can be made practical
- how we learn playing the Sudoku from images (without rules)
- how it compares with DL architectures that “learn to reason”
- how we can combine learned user preferences with (car) configuration constraints

# PLEASE, STAY WITH US AND...

## You'll learn

- how we use graphical models to connect CP with probabilistic Machine Learning
- how the NP-hard regularization loop can be made practical
- how we learn playing the Sudoku from images (without rules)
- how it compares with DL architectures that “learn to reason”
- how we can combine learned user preferences with (car) configuration constraints

				8		7		
4	9	1		6			2	8
5			3	4		1		
		3		7	9		1	
1	7					5		
	5					9	6	
	6	2	1		7		8	
	3				8	2	5	
8					4			

## You'll learn

- how we use graphical models to connect CP with probabilistic Machine Learning
- how the NP-hard regularization loop can be made practical
- how we learn playing the Sudoku from images (without rules)
- how it compares with DL architectures that “learn to reason”
- how we can combine learned user preferences with (car) configuration constraints

				8		7		
4	9	1		6			2	8
5			3	4		1		
		3		7	9		1	
1	7					5		
	5					9	6	
	6	2	1		7		8	
	3				8	2	5	
8					4			

## You'll learn

- how we use graphical models to connect CP with probabilistic Machine Learning
- how the NP-hard regularization loop can be made practical
- how we learn playing the Sudoku from images (without rules)
- how it compares with DL architectures that “learn to reason”
- how we can combine learned user preferences with (car) configuration constraints

				8		7		
4	9	1		6			2	8
5			3	4		1		
		3		7	9		1	
1	7					5		
	5					9	6	
	6	2	1		7		8	
	3				8	2	5	
8					4			

## What is it ?

A description of a multivariate function as the combination of small functions

Cost Function Network  $\mathcal{M}$

(unbounded)

- a set  $V$  of variables
- variable  $X \in V$  has domain  $D^X$
- a set  $C$  of cost functions
- $c_S \in C : \prod_{X \in S} D^X \rightarrow \bar{\mathbb{Z}}$

$n$  variables

max. size  $d$

$(\infty)$

Joint cost function

Weighted Constraint Satisfaction Problem

$$C_{\mathcal{M}}(v) = \sum_{c_S \in C} c_S(v[S])$$



## What is it ?

A description of a multivariate function as the combination of small functions

### Cost Function Network $\mathcal{M}$

(unbounded)

- a set  $V$  of variables
- variable  $X \in V$  has domain  $D^X$
- a set  $C$  of cost functions
- $c_S \in C : \prod_{X \in S} D^X \rightarrow \bar{\mathbb{Z}}$

$n$  variables

max. size  $d$

$(\infty)$

Joint cost function

Weighted Constraint Satisfaction Problem

$$C_{\mathcal{M}}(v) = \sum_{c_S \in C} c_S(v[S])$$

What is it ?

A description of a multivariate function as the combination of small functions

Cost Function Network  $\mathcal{M}$

(unbounded)

- a set  $V$  of variables
- variable  $X \in V$  has domain  $D^X$
- a set  $C$  of cost functions
- $c_S \in C : \prod_{X \in S} D^X \rightarrow \bar{\mathbb{Z}}$

$n$  variables

max. size  $d$

$(\infty)$

Joint cost function

Weighted Constraint Satisfaction Problem

$$C_{\mathcal{M}}(v) = \sum_{c_S \in C} c_S(v[S])$$

# WHAT DO WE WANT TO LEARN ?

## Definition (Learning a pairwise CFN from high quality solutions)

Given:

- a set of variables  $V$ ,
- a set of assignments  $E$  i.i.d. from an unknown distribution of high-quality solutions

Find a pairwise CFN  $\mathcal{M}$  that can be solved to produce high-quality solutions

## Pairwise CFN with cost-tables

- $\frac{n(n-1)}{2}$  tables of  $d^2$  costs +  $n$  tables of  $d$  costs
- A constant table can be ignored.

# WHAT DO WE WANT TO LEARN ?

## Definition (Learning a pairwise CFN from high quality solutions)

Given:

- a set of variables  $V$ ,
- a set of assignments  $E$  i.i.d. from an unknown distribution of high-quality solutions

Find a pairwise CFN  $\mathcal{M}$  that can be solved to produce high-quality solutions

## Pairwise CFN with cost-tables

- $\frac{n(n-1)}{2}$  tables of  $d^2$  costs +  $n$  tables of  $d$  costs
- A constant table can be ignored.

# STOCHASTIC GRAPHICAL MODELS

## Markov Random Field $\mathcal{M}$

- a set  $V$  of domain variables
- a set  $\Phi$  of potential functions
- $\varphi_S \in \Phi : \prod_{X \in S} D^X \rightarrow \mathbb{R}^+$

## Joint function and probability distribution

$$\Phi_{\mathcal{M}}(v) = \prod_{\varphi_S \in \Phi} \varphi_S(v[S])$$

$$P_{\mathcal{M}}(v) \propto \Phi_{\mathcal{M}}(v)$$

## From products to sum and back

(up to some precision)

MRF  $\mathcal{M}$

$$\xrightarrow{-\log(x)}$$

CFN  $\mathcal{M}^\ell$

$$\xrightarrow{\exp(-x)}$$

MRF  $\mathcal{M}$

# STOCHASTIC GRAPHICAL MODELS

## Markov Random Field $\mathcal{M}$

- a set  $V$  of domain variables
- a set  $\Phi$  of potential functions
- $\varphi_S \in \Phi : \prod_{X \in S} D^X \rightarrow \mathbb{R}^+$

## Joint function and probability distribution

$$\Phi_{\mathcal{M}}(v) = \prod_{\varphi_S \in \Phi} \varphi_S(v[S])$$

$$P_{\mathcal{M}}(v) \propto \Phi_{\mathcal{M}}(v)$$

## From products to sum and back

(up to some precision)

$$\text{MRF } \mathcal{M} \xrightarrow{-\log(x)} \text{CFN } \mathcal{M}^{\ell} \xrightarrow{\exp(-x)} \text{MRF } \mathcal{M}$$

## Markov Random Field $\mathcal{M}$

- a set  $V$  of domain variables
- a set  $\Phi$  of potential functions
- $\varphi_S \in \Phi : \prod_{X \in S} D^X \rightarrow \mathbb{R}^+$

## Joint function and probability distribution

$$\Phi_{\mathcal{M}}(v) = \prod_{\varphi_S \in \Phi} \varphi_S(v[S])$$

$$P_{\mathcal{M}}(v) \propto \Phi_{\mathcal{M}}(v)$$

## From products to sum and back

(up to some precision)

MRF  $\mathcal{M}$

$$\xrightarrow{-\log(x)}$$

CFN  $\mathcal{M}^\ell$

$$\xrightarrow{\exp(-x)}$$

MRF  $\mathcal{M}$

# MAXIMUM LOGLIKELIHOOD FOR CFN LEARNING

## Maximum likelihood estimation from i.i.d. sample $E$

- Likelihood of  $\mathcal{M}$ : probability of  $E$  under  $\mathcal{M}$
- Maximum likelihood  $\mathcal{M}$ : a MRF  $\mathcal{M}$  that gives maximum probability to  $E$ .

## Maximum loglikelihood $\mathcal{M}$ on $\mathcal{M}_\ell$

$$\begin{aligned}\mathcal{L}(\mathcal{M}, E) &= \log(\prod_{v \in E} P_{\mathcal{M}}(v)) = \sum_{v \in E} \log(P_{\mathcal{M}}(v)) \\ &= \sum_{v \in E} \log(\Phi_{\mathcal{M}}(v)) - \log(Z_{\mathcal{M}}) \\ &= \underbrace{\sum_{v \in E} (-C_{\mathcal{M}^\ell}(v))}_{\text{-costs of } E \text{ samples}} - \underbrace{\log\left(\sum_{t \in \prod_{x \in V} D^X} \exp(-C_{\mathcal{M}^\ell}(t))\right)}_{\text{Soft-Min of all assignment costs}}\end{aligned}$$



# MAXIMUM LOGLIKELIHOOD FOR CFN LEARNING

## Maximum likelihood estimation from i.i.d. sample $E$

- Likelihood of  $\mathcal{M}$ : probability of  $E$  under  $\mathcal{M}$
- Maximum likelihood  $\mathcal{M}$ : a MRF  $\mathcal{M}$  that gives maximum probability to  $E$ .

## Maximum loglikelihood $\mathcal{M}$ on $\mathcal{M}_\ell$

$$\begin{aligned}\mathcal{L}(\mathcal{M}, E) &= \log(\prod_{v \in E} P_{\mathcal{M}}(v)) = \sum_{v \in E} \log(P_{\mathcal{M}}(v)) \\ &= \sum_{v \in E} \log(\Phi_{\mathcal{M}}(v)) - \log(Z_{\mathcal{M}}) \\ &= \underbrace{\sum_{v \in E} (-C_{\mathcal{M}^\ell}(v))}_{\text{-costs of } E \text{ samples}} - \underbrace{\log\left(\sum_{t \in \prod_{x \in V} D^X} \exp(-C_{\mathcal{M}^\ell}(t))\right)}_{\text{Soft-Min of all assignment costs}}\end{aligned}$$

## Regularized Log-Likelihood estimation

- penalizes log-likelihood proportionally to the  $L_1$  norm of the costs learned ( $\lambda$ )
- avoids over-fitting by pushing non essential costs to zero: learns scopes.

## PE MRF: ADMM optimized convex approximation of regularized loglikelihood<sup>1</sup>

- avoids #P-completeness using a concave approximation of  $Z_{\mathcal{M}}$
- statistically sparsistent
- provides a CFN as output

---

<sup>1</sup>Youngsuk Park et al. "Learning the network structure of heterogeneous data via pairwise exponential Markov random fields". In: *Proceedings of machine learning research* 54 (2017), p. 1302.

## Regularized Log-Likelihood estimation

- penalizes log-likelihood proportionally to the  $L_1$  norm of the costs learned ( $\lambda$ )
- avoids over-fitting by pushing non essential costs to zero: learns scopes.

## PE MRF: ADMM optimized convex approximation of regularized loglikelihood<sup>1</sup>

- avoids #P-completeness using a concave approximation of  $Z_{\mathcal{M}}$
- statistically sparsistent
- provides a CFN as output

---

<sup>1</sup>Youngsuk Park et al. “Learning the network structure of heterogeneous data via pairwise exponential Markov random fields”. In: *Proceedings of machine learning research* 54 (2017), p. 1302.

# SELECTING A SUITABLE VALUE OF $\lambda$

## Using empirical risk minimization

- for each sample  $v$  in the validation set
- assign a fraction of  $v$  and solve with a WCSP solver
- prefer  $\lambda$  that gives solutions close to  $v$

## Controlling PyToulbar2 NP-hard optimization effort

- bounded optimization effort (backtrack, time, gap. Here: 50,000 backtracks)
- controllable fraction of  $v$  assigned

## Empirical hardening

Set positive costs that are never violated in the training/validation sets to  $\infty$ .

# SELECTING A SUITABLE VALUE OF $\lambda$

## Using empirical risk minimization

- for each sample  $v$  in the validation set
- assign a fraction of  $v$  and solve with a WCSP solver
- prefer  $\lambda$  that gives solutions close to  $v$

## Controlling PyToulbar2 NP-hard optimization effort

- bounded optimization effort (backtrack, time, gap. Here: 50,000 backtracks)
- controllable fraction of  $v$  assigned

## Empirical hardening

Set positive costs that are never violated in the training/validation sets to  $\infty$ .

# SELECTING A SUITABLE VALUE OF $\lambda$

## Using empirical risk minimization

- for each sample  $v$  in the validation set
- assign a fraction of  $v$  and solve with a WCSP solver
- prefer  $\lambda$  that gives solutions close to  $v$

## Controlling PyToulbar2 NP-hard optimization effort

- bounded optimization effort (backtrack, time, gap. Here: 50,000 backtracks)
- controllable fraction of  $v$  assigned

## Empirical hardening

Set positive costs that are never violated in the training/validation sets to  $\infty$ .

## An exemplar of reasoning for benchmarking


- Recurrent Relational Neural Net<sup>2</sup>:  $18 \times (10,000 + 1,000 + 1,000)$  training, validation and test samples of variable difficulty (17 to 34 hints).
- SAT-Net<sup>3</sup> (DL friendly convex Max-SAT relaxation):  $(9,000 + 1,000)$  easy training and test samples (36.2 hints average).

---

<sup>2</sup>Rasmus Berg Palm, Ulrich Paquet, and Ole Winther. “Recurrent Relational Networks”. In: *Advances in Neural Information Processing Systems, Montréal, Canada*. 2018, pp. 3372–3382.

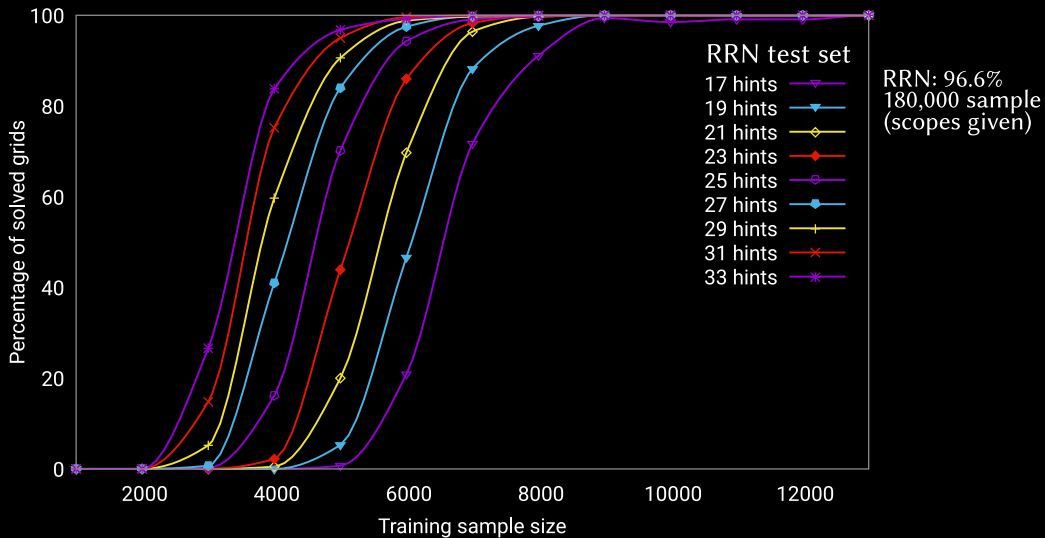
<sup>3</sup>Po-Wei Wang et al. “SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver”. In: *Proc. of ICML-19, Long Beach, California, USA*. vol. 97. PMLR, 2019, pp. 6545–6554.

# BETTER WITH LESS DATA AND COMPARABLE BIASES






# BETTER WITH LESS DATA AND COMPARABLE BIASES



# BETTER WITH LESS DATA AND COMPARABLE BIASES



# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net hints as images (8,000 samples): ..... 74.7% max. accuracy
- PE MRF+Toulbar2 (8,000+1024 samples): ..... 78.4% max. accuracy
- On hard 15 hints test RNN problems: ..... 61.2% max. accuracy
- Empirical hardening: ..... 79.8% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net (hints as images), 9,000 samples ..... 63.2%
- PE MRF+Toulbar2, 8,000+1,024 samples ..... 78.1%
- On hard 17 hints test RRN problems ..... 81.2%
- Empirical hardening ..... > 99%
- Hints and solutions as images ..... 76.3%

# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net (hints as images), 9,000 samples ..... 63.2%
- PE MRF+Toulbar2, 8,000+1,024 samples ..... 78.1%
- On hard 17 hints test RRN problems ..... 81.2%
- Empirical hardening ..... > 99%
- Hints and solutions as images ..... 76.3%

# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net (hints as images), 9,000 samples ..... 63.2%
- PE MRF+Toulbar2, 8,000+1,024 samples ..... 78.1%
- On hard 17 hints test RRN problems ..... 81.2%
- Empirical hardening ..... > 99%
- Hints and solutions as images ..... 76.3%

# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net (hints as images), 9,000 samples ..... 63.2%
- PE MRF+Toulbar2, 8,000+1,024 samples ..... 78.1%
- On hard 17 hints test RRN problems ..... 81.2%
- Empirical hardening ..... > 99%
- Hints and solutions as images ..... 76.3%

# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net (hints as images), 9,000 samples ..... 63.2%
- PE MRF+Toulbar2, 8,000+1,024 samples ..... 78.1%
- On hard 17 hints test RRN problems ..... 81.2%
- Empirical hardening ..... > 99%
- Hints and solutions as images ..... 76.3%



# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net (hints as images), 9,000 samples ..... 63.2%
- PE MRF+Toulbar2, 8,000+1,024 samples ..... 78.1%
- On hard 17 hints test RRN problems ..... 81.2%
- Empirical hardening ..... > 99%
- Hints and solutions as images ..... 76.3%

# BENEFITS FROM NUMERICAL LEARNING AND SOLVING

## Sudoku digits can be LeNet decoded and fed to PE MRF/Toulbar2

- LeNet has 99.2% accuracy on handwritten digits
- SAT-Net test set, hints as images (36.2 avg): ..... 74.7% max. accuracy
- Hints + solutions as images: ..... 52% max. accuracy

## Performances on SAT-Net test set

- SAT-Net (hints as images), 9,000 samples ..... 63.2%
- PE MRF+Toulbar2, 8,000+1,024 samples ..... 78.1%
- On hard 17 hints test RRN problems ..... 81.2%
- Empirical hardening ..... > 99%
- Hints and solutions as images ..... 76.3%

## Renault “big” dataset

[irit.fr/ Helene.Fargier/BR4CP/benches.html](http://irit.fr/~Helene.Fargier/BR4CP/benches.html)

- 268 variables (87 decision variables) with 324 values at most
- 332 constraints (max. arity 12)
- 24,566,537,954,855,758,069,760 possible configurations ( $\approx 2^{74}$ )<sup>4</sup>
- sample of 8,337 user configurations

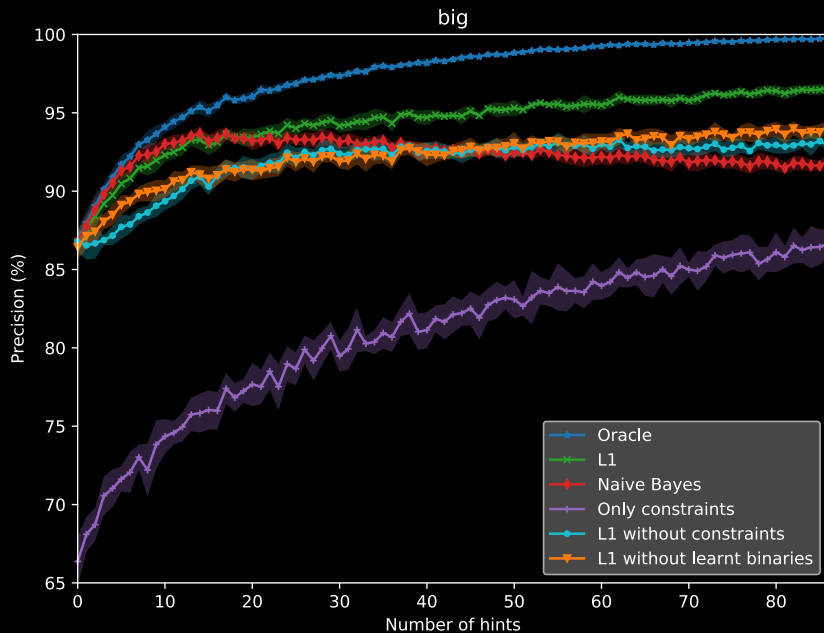
## Configuration assistant: complete ongoing user configuration on the next variable

- Learned CFN (unary + binary CFs, 10 fold CV, 2' / fold ) + configuration constraints
- Naive Bayes (knowing the partial assignment)
- Oracle (optimal stochastic choice for the test set given the partial assignment)

---

<sup>4</sup>Counted in 1.8” by Toulbar2 treewidth aware solution counter.

# ACCURACY



## Conclusions

- Flexible learning framework for CP with understandable and editable output
- Numerical (even integer) weights are enough to interact with DL output
- Anytime NP-hard prediction: more powerful than differentiable convex relaxation
- Convex relaxations may be the best we can do in P (Unique Game Conjecture)
- One should try anytime NP-hard learning too

## Conclusions

- Flexible learning framework for CP with understandable and editable output
- Numerical (even integer) weights are enough to interact with DL output
- Anytime NP-hard prediction: more powerful than differentiable convex relaxation
- Convex relaxations may be the best we can do in P (Unique Game Conjecture)
- One should try anytime NP-hard learning too

## Conclusions

- Flexible learning framework for CP with understandable and editable output
- Numerical (even integer) weights are enough to interact with DL output
- Anytime NP-hard prediction: more powerful than differentiable convex relaxation
- Convex relaxations may be the best we can do in P (Unique Game Conjecture)
- One should try anytime NP-hard learning too

## Conclusions

- Flexible learning framework for CP with understandable and editable output
- Numerical (even integer) weights are enough to interact with DL output
- Anytime NP-hard prediction: more powerful than differentiable convex relaxation
- Convex relaxations may be the best we can do in P (Unique Game Conjecture)
- One should try anytime NP-hard learning too



## Conclusions

- Flexible learning framework for CP with understandable and editable output
- Numerical (even integer) weights are enough to interact with DL output
- Anytime NP-hard prediction: more powerful than differentiable convex relaxation
- Convex relaxations may be the best we can do in P (Unique Game Conjecture)
- One should try anytime NP-hard learning too