

Learning and Reasoning with Constraint Solving

Prof. Tias Guns
<tias.guns@kuleuven.be>
 @TiasGuns

Joint work with team members:

- Rocs Canoy
- Jayanta Mandi
- Maxime Mulamba
- Victor Bucarey Lopez
- Emilio Gamba
- Ignace Bleukx

And external collaborators:

- Michelangelo Diligenti (Sienna Uni, It)
- Michele Lombardi (Bologna, It)
- Bart Bogaerts (VUB, Be)





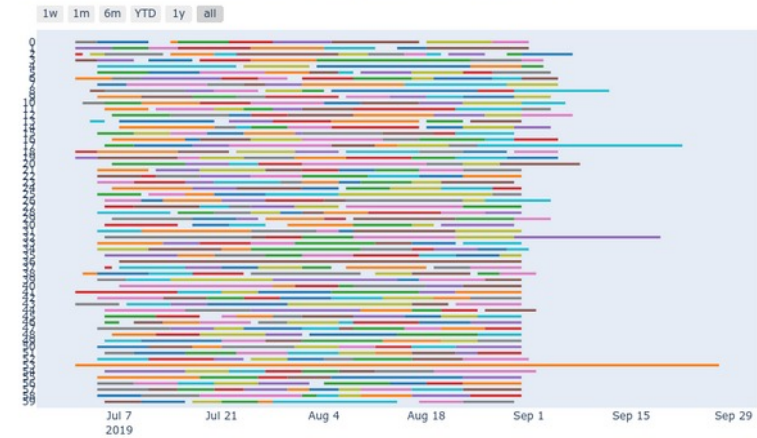
Constraint solving

“Solving combinatorial optimisation problems”

- Vehicle Routing
- Scheduling
- Packing
- Other combinatorial problems



P-Large-02 (59 ROOMS), ExitStatus.OPTIMAL (1558.940814725 seconds)



Example: room scheduling

Demo with CPMpy

https://github.com/CPMpy/cpm.py/blob/master/examples/room_assignment.ipynb

Example: room scheduling (backup slide)

```
def model_rooms(df, max_rooms, verbose=True):
    n_requests = len(df)

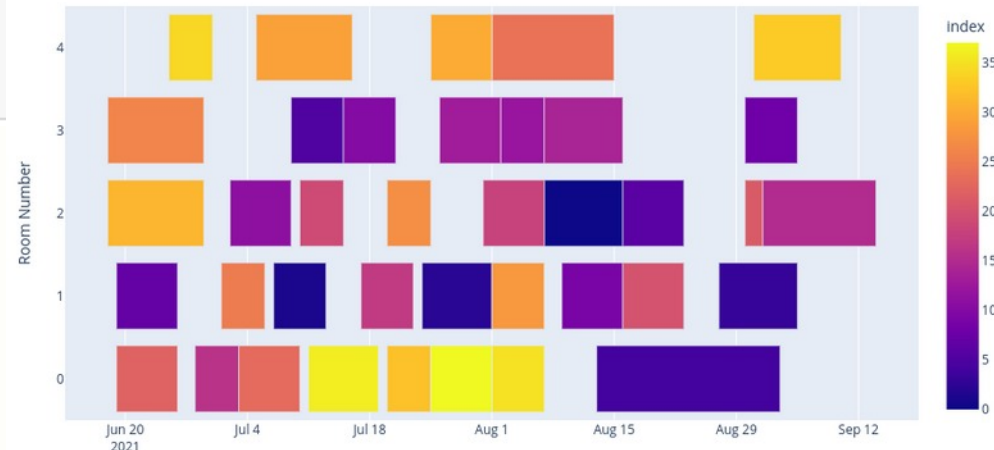
    # All requests must be assigned to one out of the rooms (same room during entire period).
    requestvars = intvar(0, max_rooms-1, shape=(n_requests,))

    m = Model()

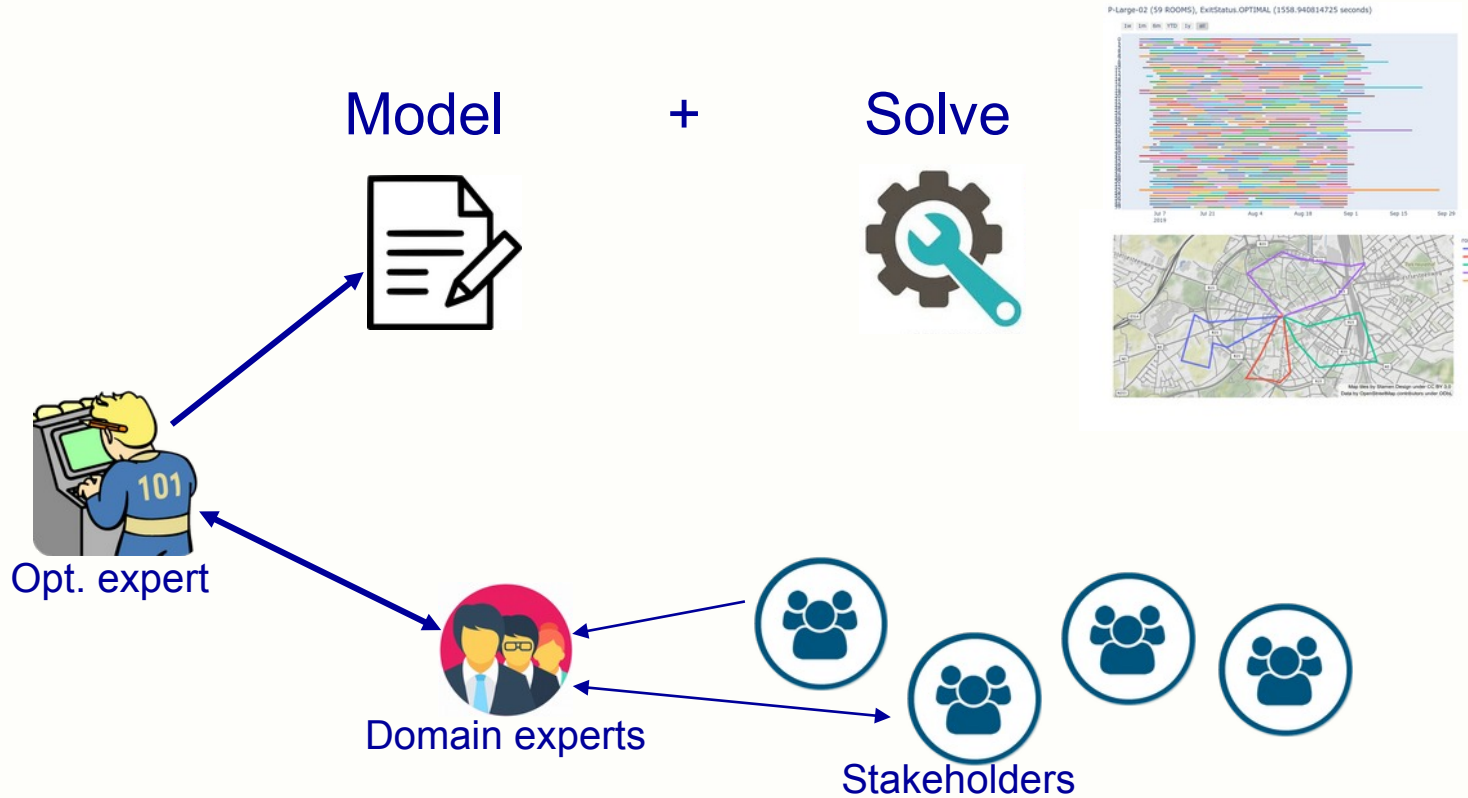
    # Some requests already have a room pre-assigned
    for idx, row in df.iterrows():
        if not pd.isna(row['room']):
            m += (requestvars[idx] == int(row['room']))

    # A room can only serve one request at a time.
    # <=> requests on the same day must be in different rooms
    for day in pd.date_range(min(df['start']), max(df['end'])):
        overlapping = df[(df['start'] <= day) & (day < df['end'])]
        if len(overlapping) > 1:
            m += AllDifferent(requestvars[overlapping.index])

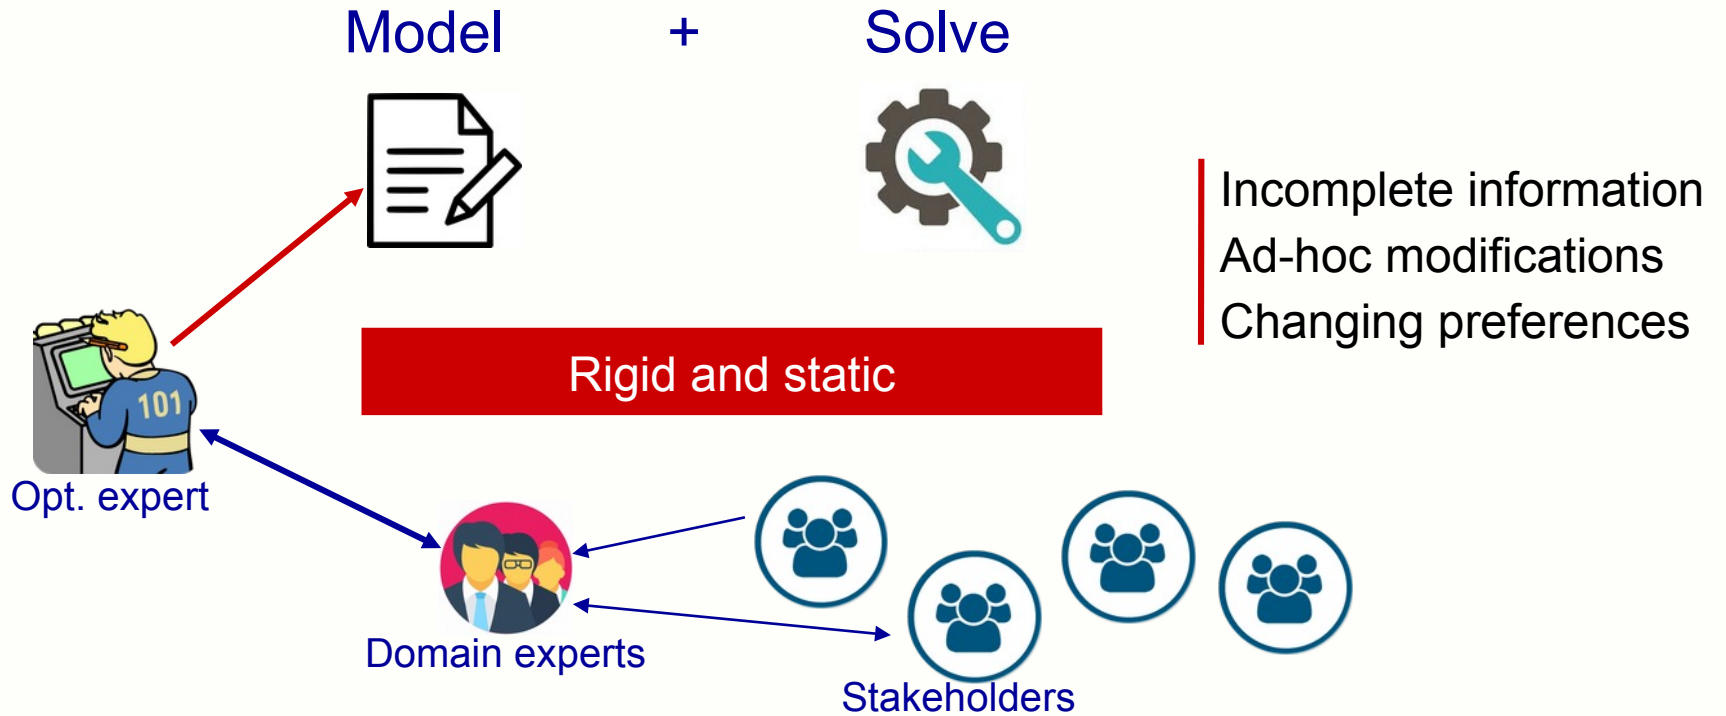
    return (m, requestvars)
```




Current combinatorial optimisation practice



Current combinatorial opt. practice, **problem**




Vision




Vision

- Learning implicit user preferences
- Learning from the environment




Vision

- Learning implicit user preferences
- Learning from the environment
- Explaining constraint solving



Vision


- Learning implicit user preferences
- Learning from the environment
- Explaining constraint solving
- Stateful interaction



CHAT-Opt: Conversational **H**uman-**A**ware **T**echnology for **O**ptimisation



European Research Council
Established by the European Commission




Towards **co-creation** of constraint optimisation solutions

- Solver that learns from user and environment
- Towards conversational: explanations and stateful interaction

<https://people.cs.kuleuven.be/~tias.guns/chat-opt.html>

Hiring post-docs!


Conversational **H**uman-**A**ware Technology for **O**ptimisation



What would the ideal CP system be?

- Easy integration with Machine Learning libraries
=> Python and numpy arrays
- Efficient repeated solving
=> Incremental
- Use CP/SAT/MIP or any combination
=> solver independent and multi-solver

Conversational Human-Aware Technology for Optimisation



What would the ideal CP system be?

- Easy integration with Machine Learning
=> Python and numpy arrays
- Efficient repeated solving
=> Incremental
- Use CP/SAT/MIP or any combination
=> solver independent and multi-solver

<https://github.com/CPMpy/cpmPy>

CPMpy: Constraint Programming and Modeling in Python

CPMpy is a Constraint Programming and Modeling library in Python, based on native solver access.

Constraint Programming is a methodology for solving combinatorial optimisation assignment problems or covering, packing and scheduling problems. Problems that require searching over discrete decision variables.

CPMpy allows to model search problems in a high-level manner, by defining decision variables, constraints and an objective over them (similar to MiniZinc and Essence'). You can then solve the model using a native solver like or-tools, which then compute the optimal answer.

Source code and bug reports at <https://github.com/CPMpy/cpmPy>

Getting started:

- [Installation instructions](#)
- [Getting started with Constraint Programming and CPMpy](#)
- [Quickstart sudoku notebook](#)
- [More examples](#)

User Documentation:

This talk

1. Integrating CP with ML *predictions*

2. Integrating CP with ML *learning*

3. CP Explanations *with implicit hitting sets*

+ High-level overview of how CPMpy enables this

Perception-based constraint solving

What if part of the input is in an image?

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

Pedagogical example: Visual Sudoku

- Every image represents a valid sudoku
- Explicitly know: CP constraints
- Need to predict: image labels


=> test limits of reasoning on learning

Perception data and constraint solving

Other application settings:

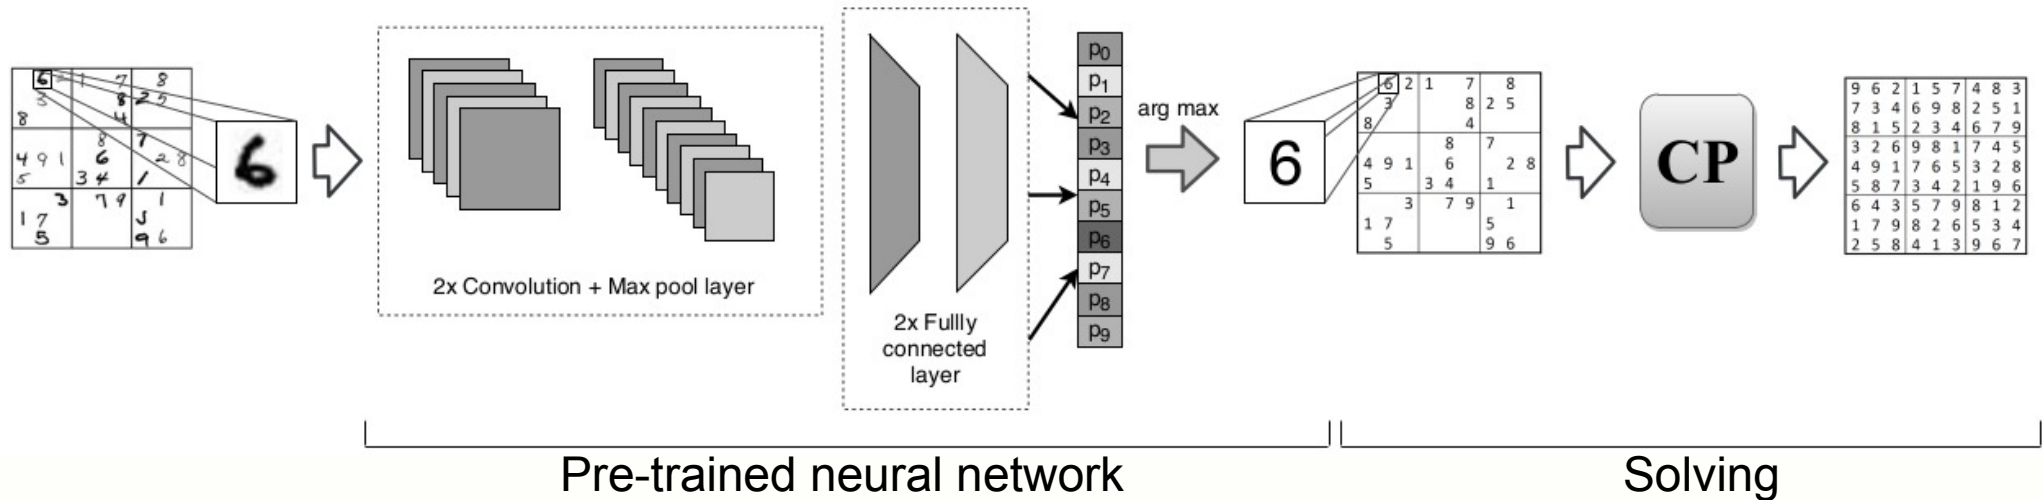
- Document analysis
- Paper-based configuration problems (tax forms)
- Object-detection based reasoning
- Visual relationship detection
- ...

ITR - I SAHAJ INDIVIDUAL INCOME TAX RETURN										Assessment Year			
(For Individuals having Income from Salaries, One House Property, Other Sources Interest etc) Refer to Instructions for eligibility.										2016-17			
(A 1) First Name PRADIP		(A 2) Middl Name KUMARAN			(A 3) Last Name KULKARNI			(A 4) Permanent Account Number A B 1 2 3 4 5 6 7 A					
(A 5) Sex (Tick) <input checked="" type="checkbox"/> Male <input type="checkbox"/> Female		(A 6) Date of Birth (DD/MM/YYYY) 01/04/1963								(A 7) Income Tax Word/Circle 0568			
(A 8) Flat/Door/Building FLAT NO 5		(A 9) Name of Premises/ Building/Village KALISTHANPARA			(A 10) Road / Street NETAJI ROAD			(A 11) Area / Locality HALISAHAR					
(A 12) Town/City/District MUMBAI		(A 13) State MAHARASHTRA		(A 14) Country INDIA			(A 15) Pin Code 0 4 2 1 2 3						
(A 17) Mobile No/Residential/Office Phone No with STD Code 03483270258				(A 18) Mobile No. 2 9474316768			(A 19) Fill only if you belong to <input checked="" type="checkbox"/> Govt <input type="checkbox"/> PSU <input type="checkbox"/> Others						
(A 20) Fill only one <input type="checkbox"/> Tax Refundable <input type="checkbox"/> Tax Payable		<input type="checkbox"/> Nil Tax Balance		<input checked="" type="checkbox"/>									
(A 21) Residential Status (Tick) <input checked="" type="checkbox"/> Resident <input type="checkbox"/> Non Resident <input type="checkbox"/> Resident but not ordinarily resident													
(A 22) Fill only one Filed <input checked="" type="checkbox"/> On or before due date-139(1). <input type="checkbox"/> After due date-139(4). <input type="checkbox"/> Revised Return- 139(5)													



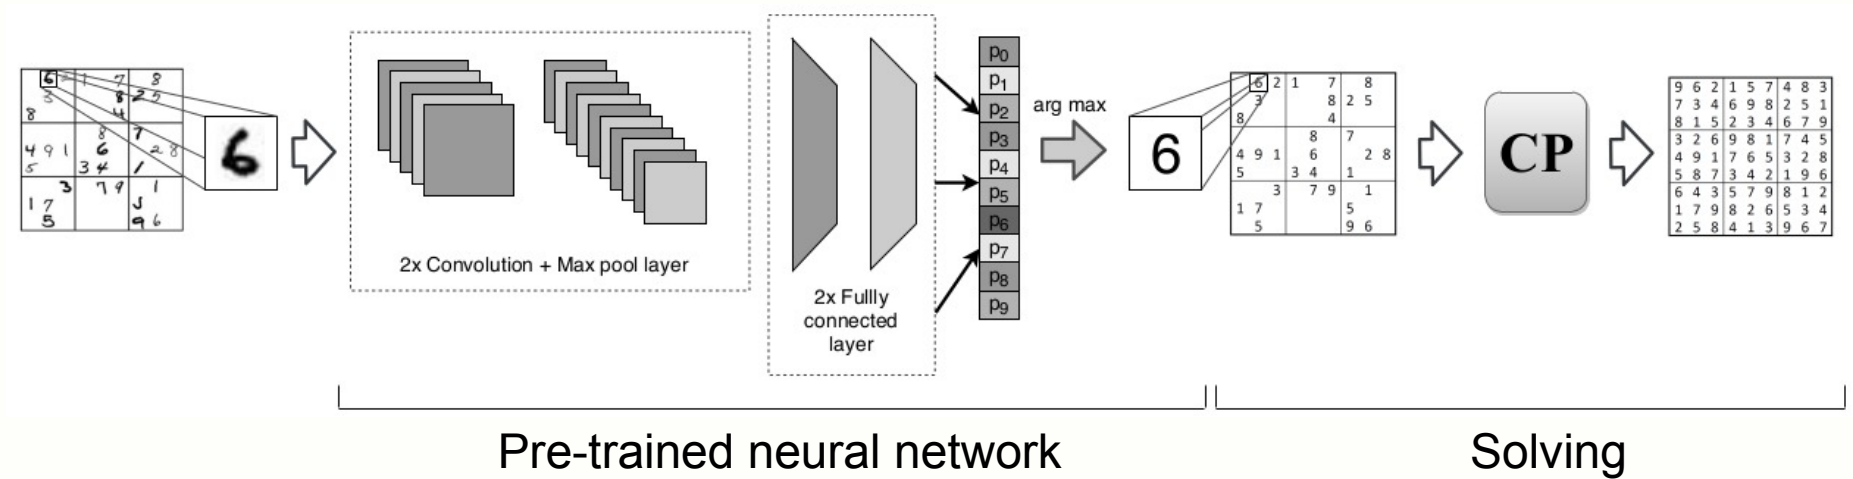
Perception-based constraint solving

Visual sudoku (naïve)



	img	accuracy cell	grid	failure rate grid	time average (s)
baseline	94.75%	15.51%	14.67%	84.43%	0.01

Perception-based constraint solving



What is going on?

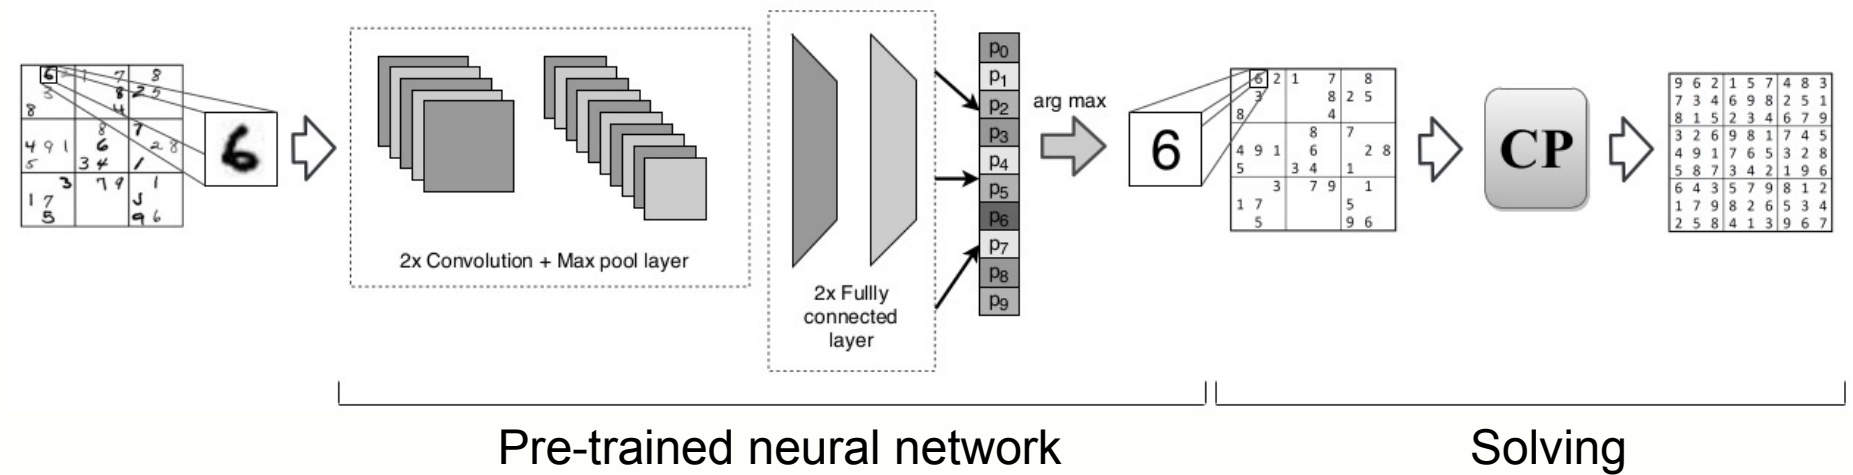
- Each cell predicts the maximum likelihood value:

$$\hat{y}_{ij} = \arg \max P(y_{ij} = k | X_{ij})$$

- This is a multi-output problem (one for each given cell):
together this is the 'maximum likelihood' interpretation
- If $\text{sudoku}(\hat{y}) = \text{False}$: no solution, interpretation is wrong...

What about the *next* most likely interpretation?

Perception-based constraint solving



What about the *next* most likely interpretation?

- Treat prediction as *joint inference* problem:

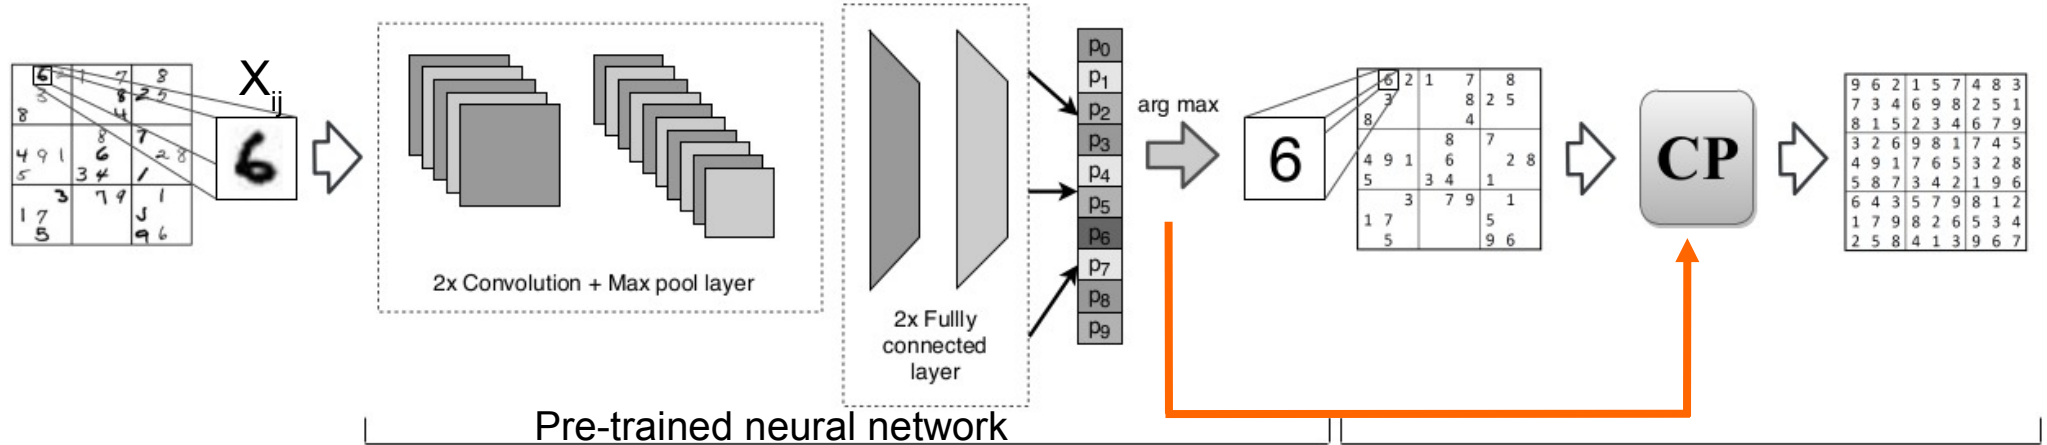
$$\hat{y} = \arg \max \prod_{ij} P(y_{ij} = k | X_{ij}) \quad \text{s.t.} \quad \text{sudoku}(\hat{y})$$

- This is the **constrained** 'maximum likelihood' interpretation

=> Structured output prediction

Used e.g. in NLP: [Punyakank, COLING04]

Perception-based constraint solving



Can we use a constraint solver for that?

$$\hat{y} = \arg \max \prod_{ij} P(y_{ij} = k | X_{ij}) \quad \text{s.t.} \quad \text{sudoku}(\hat{y})$$

- Log-likelihood trick:

$$\min \sum_{\substack{(i,j) \in \\ \text{given } \{1, \dots, 9\}}} \sum_{k \in \{1, \dots, 9\}} \underbrace{-\log(P_{\theta}(y_{ij} = k | X_{ij})) * \mathbb{1}[s_{ij} = k]}_{\text{constant}} \quad \text{s.t.} \quad \text{sudoku}(\hat{y})$$

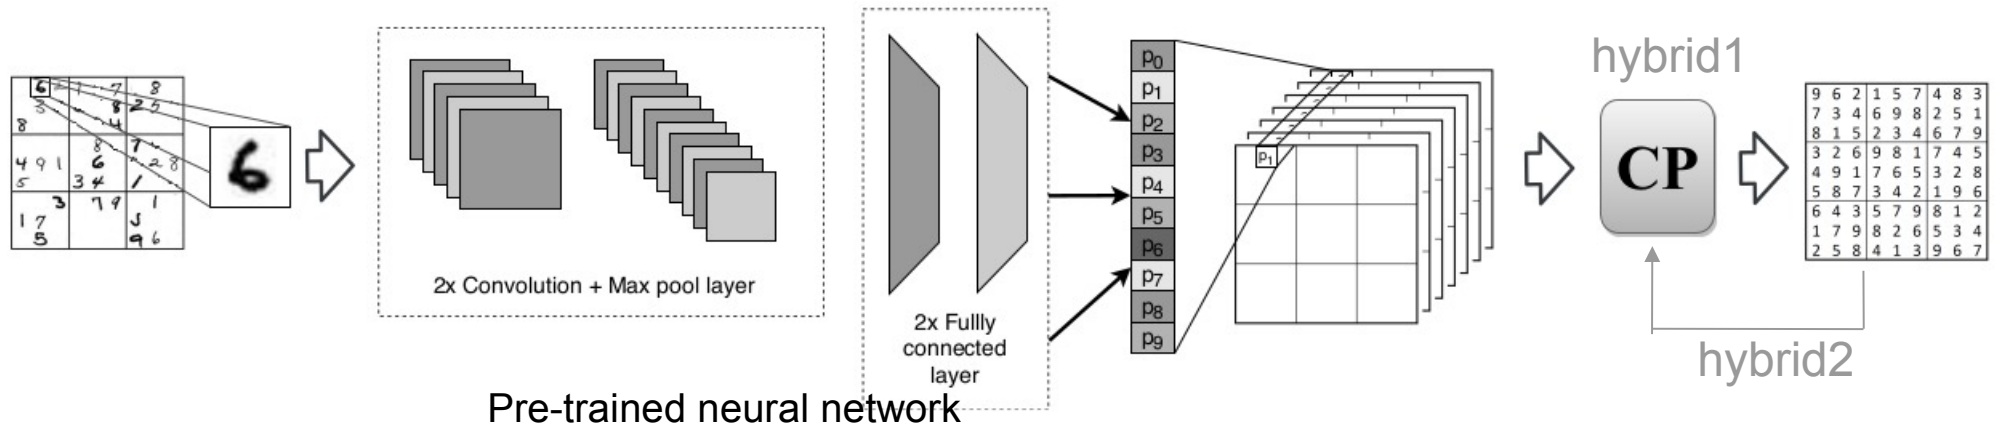
Visual sudoku

Demo with CPMpy

https://github.com/CPMpy/cpm.py/blob/master/examples/advanced/visual_sudoku.ipynb

Perception-based constraint solving

Hybrid: CP solver does *joint inference* over raw probabilities

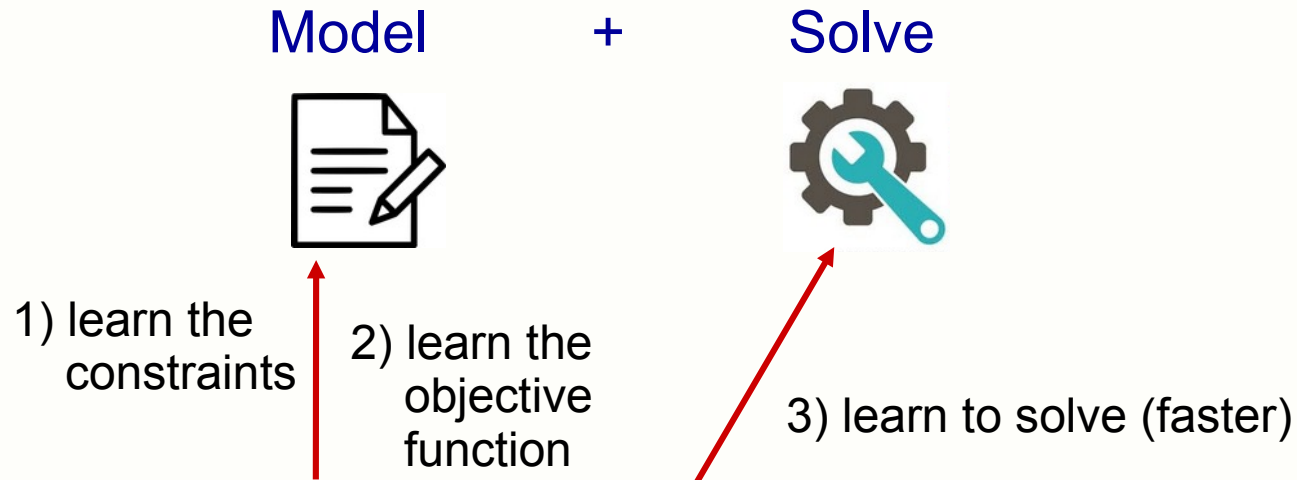


	img	accuracy cell	grid	failure rate grid	time average (s)
baseline	94.75%	15.51%	14.67%	84.43%	0.01
hybrid1	99.69%	99.38%	92.33%	0%	0.79
hybrid2	99.72%	99.44%	92.93%	0%	0.83


This talk

1. Integrating CP with ML *predictions*
 - 2. Integrating CP with ML *learning***
 3. CP Explanations *with implicit hitting sets*
- + High-level overview of how CPMpy enables this

Research trend




Can we *learn* it instead?



Learn the objective function

Prediction + Optimisation (regression of weights)



Predict: hourly energy prices +

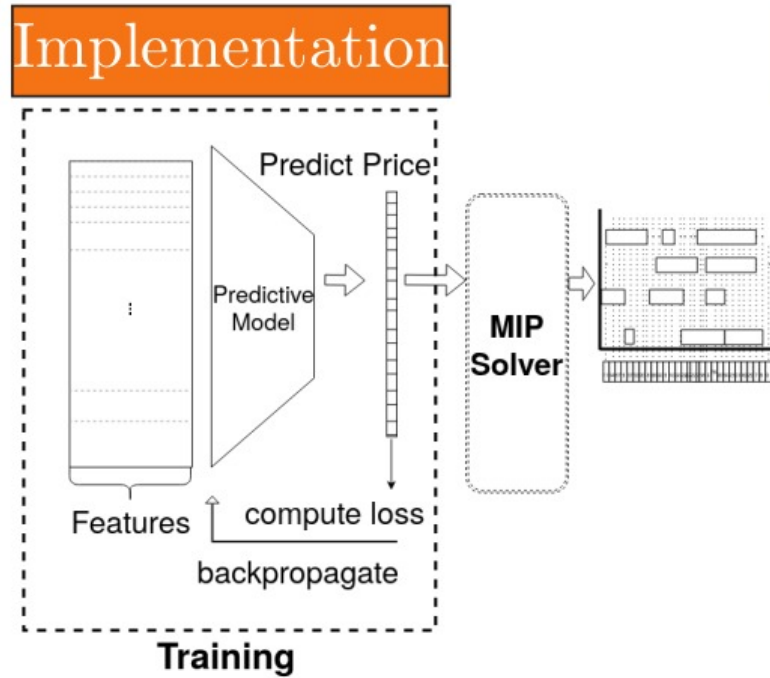
Optimize: energy-aware scheduling

Other examples:

- Optimize steel plant production waste, by predicting steel defects
- Optimize money transport, by predicting value of coins at clients


• ...

prediction-focussed regression



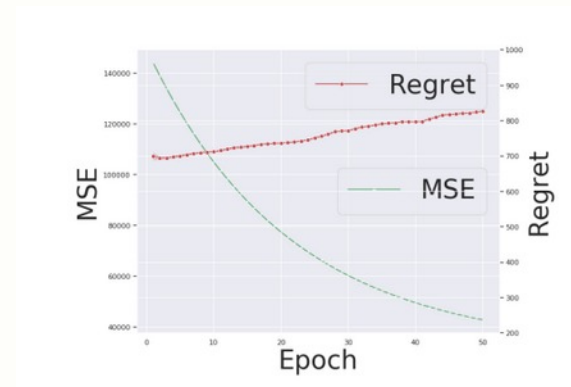
Pre-trained neural network

MSE loss not the best proxy for *task* loss....



Why?

- MSE = average of individual errors of the vector
- Joint inference = *joint* error
→ some errors worse than others!



Which errors worse?

is combinatorial, need to solve to know

$$\operatorname{argmin}_{\omega} \mathbb{E} [\operatorname{regret}(\underbrace{m(\bar{x}_i; \omega)}_{\text{predicted cost vector}}, \underbrace{\bar{c}_i}_{\text{true cost vector}})]$$


network ω params

$$\begin{aligned} \operatorname{regret}(\hat{c}, c) &= f(\hat{v}, c) - f(v^*, c) \\ \text{with } v^* &= \operatorname{argmin}_{v \in V} f(v, c) \\ \hat{v} &= \operatorname{argmin}_{v \in V} f(v, \hat{c}) \end{aligned}$$

Challenges:

- no explicit gradient
- V is implicit, exponential size
- $\operatorname{argmin} f$ may be NP-hard

Learning approaches (gradient descent)



Key challenges:

- 1) suitable loss function? (non-differentiable solver)
- 2) scalability due to repeated solving:
once per instance per epoch

Related work for discrete optimisation

- **Differentiating KKT of a relaxed problem** [Wilder, B., Dilkina, B., & Tambe, M. (2019, July)., Ferber, A., Wilder, B., Dilkina, B., & Tambe, M. (2020, April)]
- **Differentiating HSD of a relaxed problem** [Mandi, J., & Guns, T. (2020)]
- **Subgradient of a surrogate loss** [Elmachtoub, A. N., & Grigas, P. (2022), Mulamba, M. & Mandi, J. & Diligenti, M. & Lombardi, M. & Bucarey, V. & Guns, T.]
- **Differentiation by perturbation** [Pogančić, Marin Vlastelica, et al. (2020), Niepert, M., Minervini, P., & Franceschi, L. (2021)]

Decision-focused learning

Suitable loss function?

Key observation:


“The objective function induces a ranking over feasible solutions”

	Obj with true costs	Obj with predicted
Sol 1 [a,c,b,d,a]	12 (rank: 1)	14 (rank: 3)
Sol 2 [a,b,c,d,a]	15 (rank: 2)	10 (rank: 1)
Sol 3 [a,c,d,b,a]	16 (rank: 3)	11 (rank: 2)
Sol 4 [a,d,b,c,a]	23 (rank: 4)	16 (rank: 4)
Sol 5 [a,d,c,b,a]	28 (rank: 5)	18 (rank: 5)

Decision-focussed learning

Assume a set of feasible solutions S .

“The objective function induces a ranking over feasible solutions”




=> We can now use techniques from the much more mature ‘Learning to Rank’ field in ML!

Listwise Learning 2 Rank for DFL

Discrete exponential distribution
in solution space

$$p(v; c) = \begin{cases} \frac{1}{Z} \exp(-f(v, c)/\tau) & v \in V \\ 0 & v \notin V \end{cases}$$




We obtain 2 distributions (one for true costs, one for predicted costs)
over a finite sample of feasible solutions S


=> Can use the standard Kullback-Leibler Divergence loss!

Results


Shortest Path Problem





Degree 4



Degree 6




Degree 8



Scheduling Problem


Decision-focused learning with L2R

2nd Key bottleneck: repeatedly calling the solver



Decision-focused learning with L2R

2nd Key bottleneck: repeatedly calling the solver

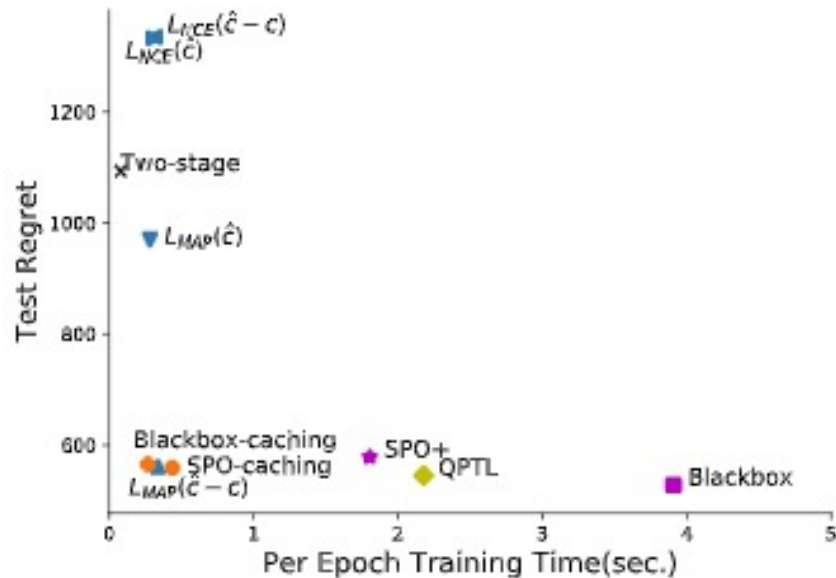


Can use **cached solutions** as approximate solver!!

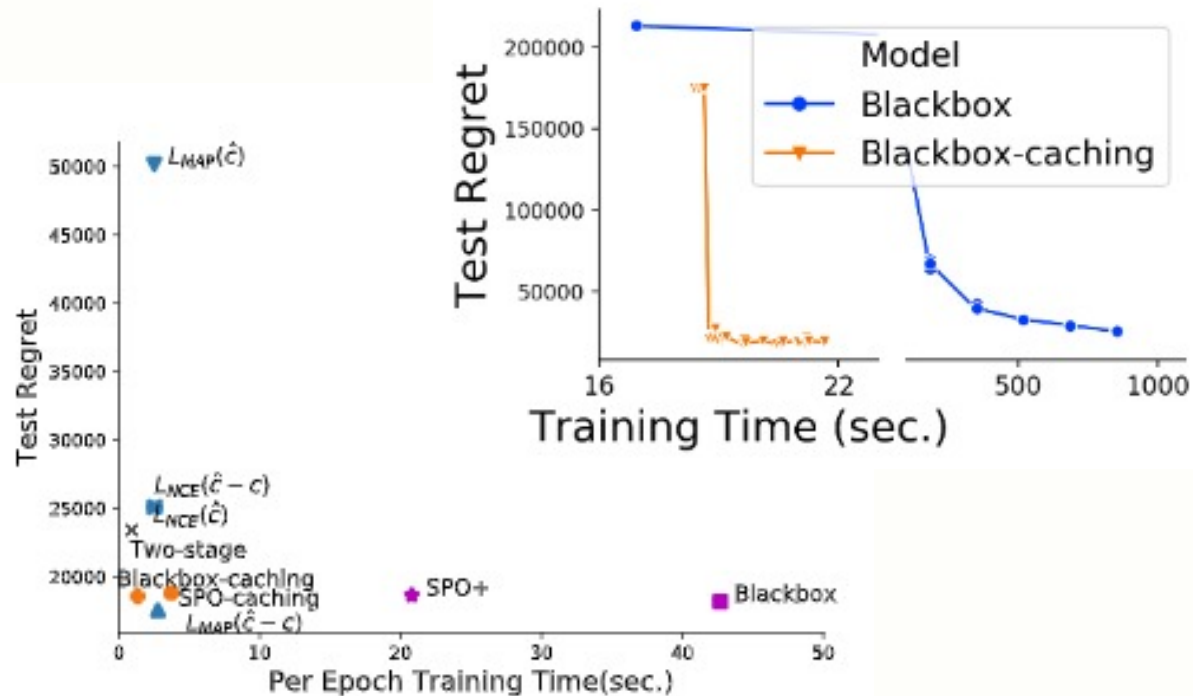
These cached solutions are the feasible set S
(also: sampling schemes: call the solver only 10% of the times)

Results

Caching scheme compatible with **all** methods that call a blackbox solver (*call the cache instead, 90% of time*)



(a) Knapsack-120



(b) Energy-3

Implementation in gradient descent loop

Standard:

Algorithm 1: Stochastic gradient descent

Input : training data $\mathcal{D} = \{X, y\}_{i=1}^n$, learning rate γ

```
1 initialize  $\theta$  (neural network weights)
2 for epochs do
3   for batches do
4     sample batch  $(X, y) \sim \mathcal{D}$ 
5      $\hat{y} \leftarrow g(z, \theta)$  (forward pass: compute predictions)
6     Compute loss  $L(y, \hat{y})$  and gradient  $\frac{\partial L}{\partial \theta}$ 
7     Update  $\theta = \theta - \gamma \frac{\partial L}{\partial \theta}$  through backpropagation (backward pass)
8   end
9 end
```

with Listwise ranking:

Algorithm 3: Stochastic gradient descent with KL on solutions

Input : training data $\mathcal{D} = \{X, y\}_{i=1}^n$, architecture g , learning rate γ , sample rate r

```
1 initialize  $\theta$  (neural network weights of  $g$ )
2  $sols \leftarrow \{solver(y) \mid (X, y) \in \mathcal{D}\}$  (initialize with true solutions)
3 for epochs do
4   for batches do
5     sample batch  $(X, y) \sim \mathcal{D}$ 
6      $\hat{y} \leftarrow g(X, \theta)$  (forward pass: compute predictions)
7     if  $random() \leq r$  then
8        $sols \leftarrow sols \cup \{solver(\hat{y})\}$ 
9     end
10    Compute loss  $L = KL(distr(y, sols), distr(\hat{y}, sols))$  and grad.  $\frac{\partial L}{\partial \theta}$ 
11    Update  $\theta = \theta - \gamma \frac{\partial L}{\partial \theta}$  through backpropagation (backward pass)
12  end
13 end
```

Decision-focused learning

Demo with CPMpy (older method)

https://github.com/CPMpy/cpm.py/blob/master/examples/advanced/predict_plus_optimize.ipynb

This talk

1. Integrating CP with ML *predictions*
 2. Integrating CP with ML *learning*
 3. **CP Explanations *with implicit hitting sets***
- + High-level overview of how CPMpy enables this

Debugging a model

Solver says UNSAT, what now?

→ compute Minimal Unsatisfiable Subset (MUS)

```
def mus(constraints):
    m = Model(constraints)
    assert ~m.solve(), "MUS: model must be UNSAT"

    core = m.get_core() # or all constraints
    i = 0
    while i < len(core):
        subcore = core[:i] + core[i+1:] # check if all but i makes core SAT

        if Model(subcore).solve():
            i += 1 # removing it makes it SAT, must keep
        else:
            core = subcore # overwrite core, so core[i] is next one

    return core
```

(faster if the solver supports unsat core extraction and assumptions)

What if a model is SAT?

- User may not understand all derivations
- Or wants to learn about it

Ex. 2019 Holy Grail Challenge (E. Freuder)

Logic Grid Puzzles (aka Zebra/Einstein puzzles)

- Parse puzzles and translate into CSP
- Solve CSP automatically
- **Explain** in a human-understandable way how to solve this puzzle

ZEBRA TUTOR: EXPLAINING LOGIC GRID PUZZLES

ZebraTutor is an end-to-end solution for solving logic grid puzzles and for explaining, in a human-understandable way, how this solution can be obtained from the clues. Here is an example puzzle. The computer has already solved the problem ! But can you solve it ?

Click on the arrows to navigate in the human-like solving process.

PUZZLE

	the_other_Type2	tehama	zearing	plymouth	shaver_lake	the_other_Type1	oregon	kansas	washington	alaska	the_other_native	mattie	ernesto	roxanne	zachary
110															
111															
112															
113															
the_other_native															
mattie															
ernesto															

CLUES

1. Mattie is 113 years old
2. The person who is 111 years old doesn't live in Plymouth
3. The person who lives in Shaver Lake is 1 year younger than Roxanne
4. The person who lives in Zearing isn't a native of Alaska
5. Roxanne is 2 years younger than the Kansas native
6. If the person who lives in Zeating is a native of Alaska, one is a native of Alaska and the other is from Kansas
7. The centenarian who lives in Plymouth isn't a native of Alaska
8. The Washington native is 1 year older than Ernesto
9. The person who lives in Tehama is a native of either Kansas or Oregon
10. The Oregon native is either Zachary or the person who lives in Tehama

Explanation steps

Let E' & $S' \Rightarrow n$ be one explanation step.

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

E' = a subset of previously derived facts E

S' = a minimal subset of constraints S such that $E' \& S' \Rightarrow n$

n = a newly derived fact

How? $MUS(\sim n \& E \& S)$ is a valid explanation step

The best/easiest explanation step...

Let $f(E, S, n)$ be a cost-function that quantifies how good (e.g. easy to understand) an explanation step is.

Simple MUS-based algo:

```
X_best = None
For n in optimal_propagate(constraints):
    X = MUS(~n & E & S)
    If f(X) < f(X_best):
        X_best = X
return X_best
```

But MUS gives no guarantees on quality, only subset minimal

Optimal unsatisfiable subsets

O(C)US: use an implicit hitting set algorithm (like MaxHS)

Algorithm 4: OCUS(T, f, p)

```
1  $\mathcal{H} \leftarrow \emptyset$ 
2 while true do
3    $\mathcal{S} \leftarrow \text{COST-OPTIMAL-HITTINGSET}(\mathcal{H}, f, p)$   $\longrightarrow$  MIP solver
4   if  $\neg \text{SAT}(\mathcal{S})$  then  $\longrightarrow$  SAT/CP solver
5     return  $\mathcal{S}$ 
6   end
7    $\mathcal{S} \leftarrow \text{GROW}(\mathcal{S}, T)$   $\longrightarrow$  Big efficiency gains
8    $\mathcal{H} \leftarrow \mathcal{H} \cup \{T \setminus \mathcal{S}\}$   $\longrightarrow$  If incremental...
9 end
```

Step-wise explanations

Demo with CPMpy

https://github.com/CPMpy/cpm.py/blob/master/examples/advanced/explain_stepwise_csp.ipynb

This talk

1. Integrating CP with ML *predictions*
 2. Integrating CP with ML *learning*
 3. CP Explanations *with implicit hitting sets*
- + High-level overview of how CPMpy enables this**


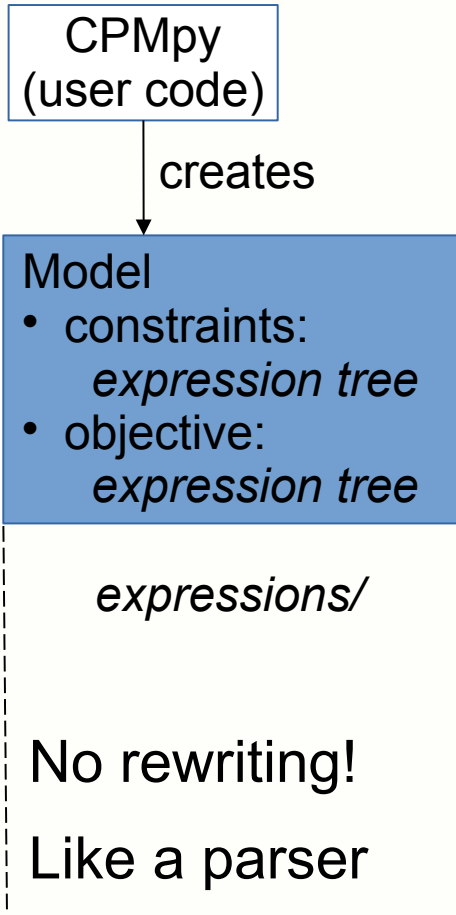
Design

Design principle:

Aim to be a thin layer on top of solver API

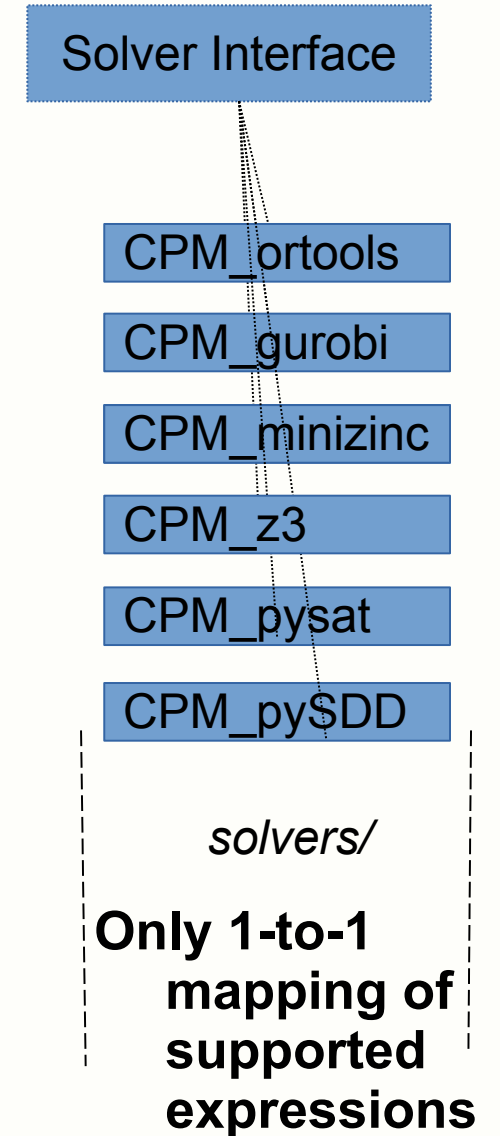
Central concept: CPMpy expression

Design

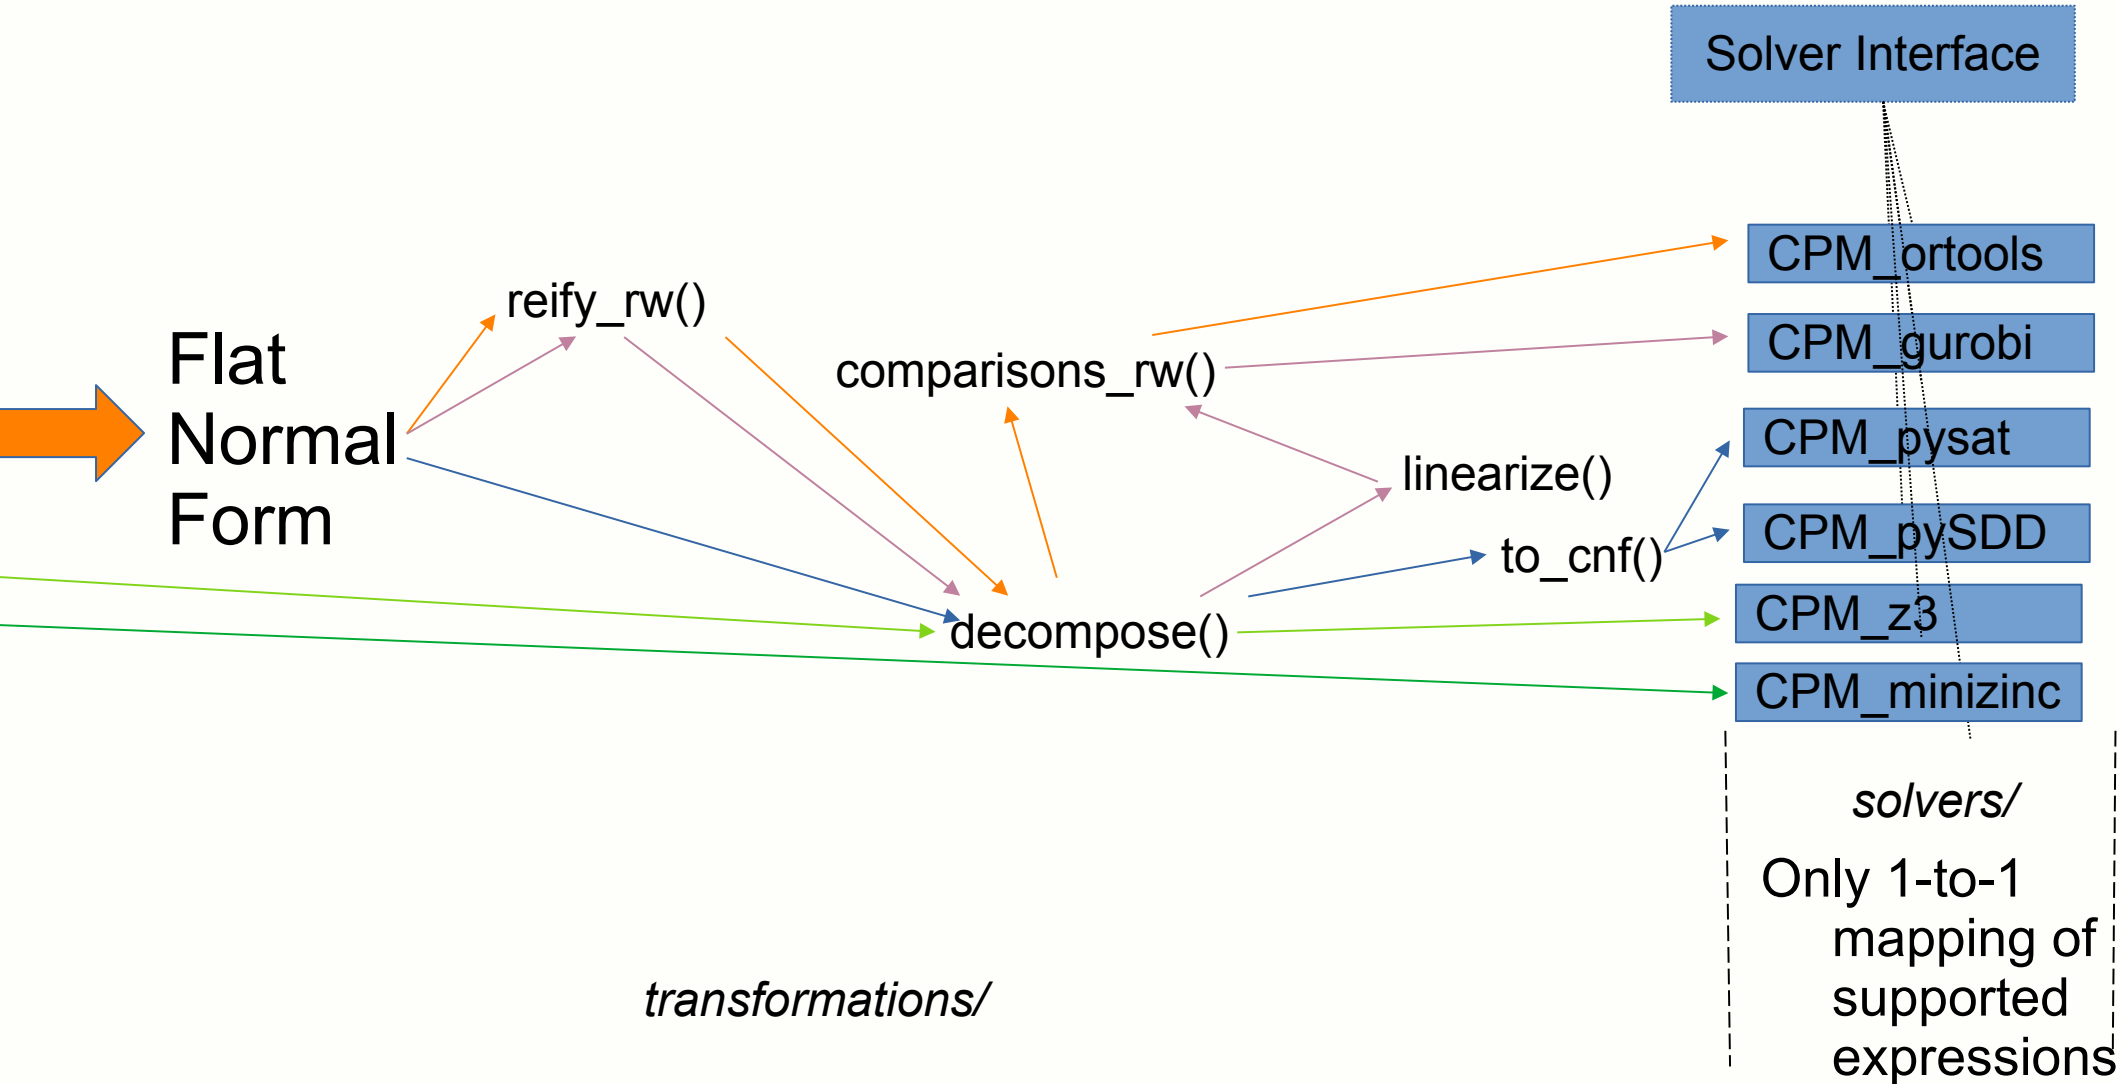


Hardest part

transformations/



Transformations in a nutshell



Transformations are functions

- Flat Normal Form:
 - Removes nested expressions (except reification)
 - All subsequent transformations can assume input is 'normalized'
- All transformations are pure *functions*:
 - Can call them in any order, and indep. of solver objects
 - State can be passed as an argument, but not required

=> they are **incremental**
- Great for debugging too

Solvers

We only interface to Python APIs
(unfortunately, no Common CP solver API : (

Key principle: solver can implement any subset of expressions!

Solvers can also choose to:

- Support assumptions or not
- Be incremental or not
- Expose own solver parameters

Currently:

- ortools
- pysat
- minizinc
- gurobi
- pySDD

Near future: ExactSolver, Z3

Wishlist: Mistral2, Geas,

Gecode

This talk

0. Data/Visualisations

1. Integrating CP with ML *predictions*


2. Integrating CP with ML *learning*

3. CP Explanations *with implicit hitting sets*

4. How does CPMpy enable this?



CHAT-Opt: Conversational **H**uman-**A**ware **T**echnology for **O**ptimisation



Towards **co-creation** of constraint optimisation solutions

- Solver that learns from user and environment
- Towards conversational: explanations and stateful interaction