

Learning from user and environment in combinatorial optimisation

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

Joint work with team members:

- Rocs Canoy
- Jayanta Mandi
- Maxime Mulamba
- Victor Bucarey Lopez
- Ahmed KA Abdullah
- Emilio Gamba

And external collaborators:



- Peter Stuckey (Monash Uni, Au)
- Emir Demirovic (TU Delft, NL)
- Michelangelo Diligenti (Sienna Uni, It)
- Michele Lombardi (Bologna, It)
- Bart Bogaerts (VUB, Be)



Combinatorial optimisation

“Solving *constrained* optimisation problems”

- Vehicle Routing
- Scheduling
- Configuration
- Graph problems



Constraint solving paradigm

Model

+

Solve

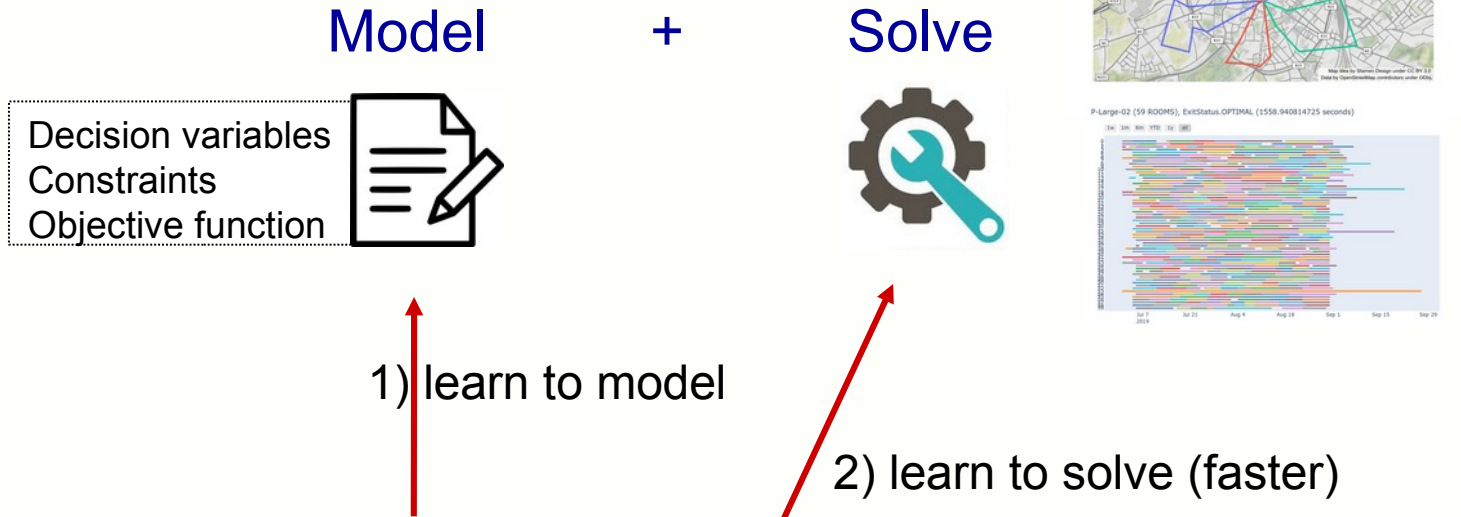
Decision variables
Constraints
Objective function




P-Range-02 (59 ROOMS), ExitStatus:OPTIMAL (1558.940814725 seconds)



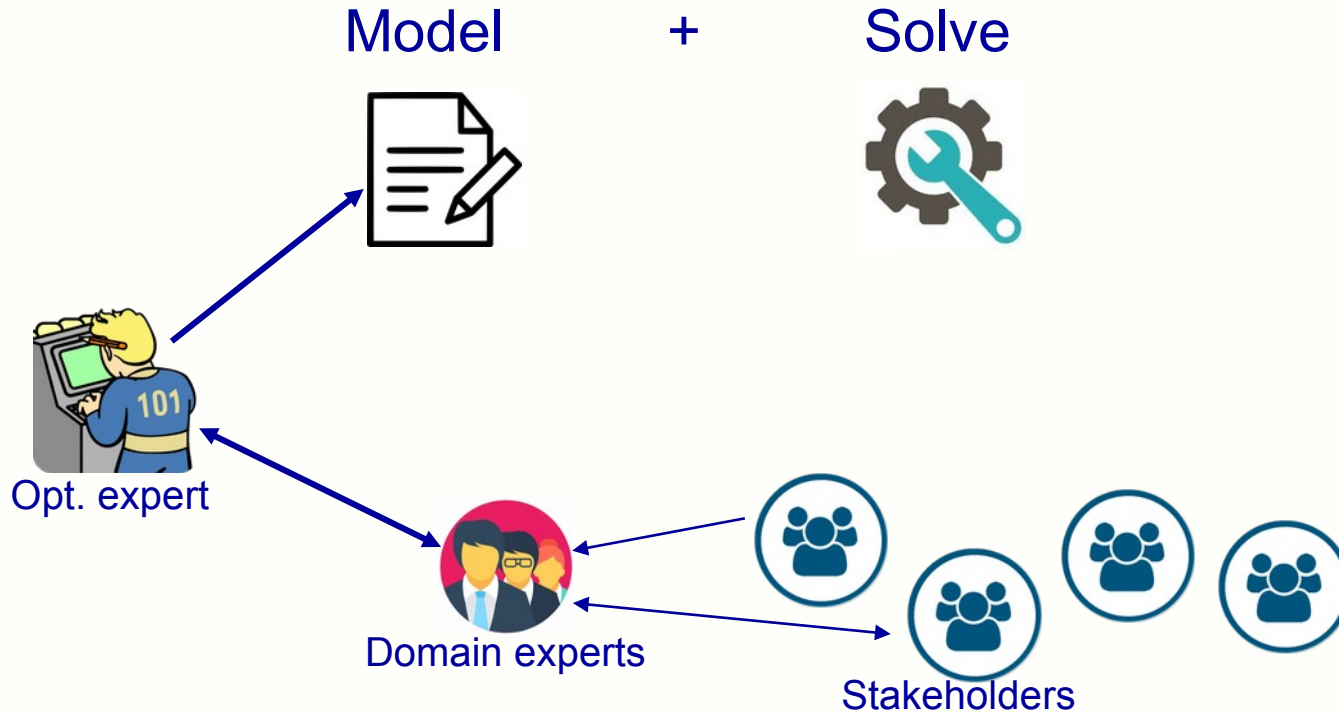
Research trend



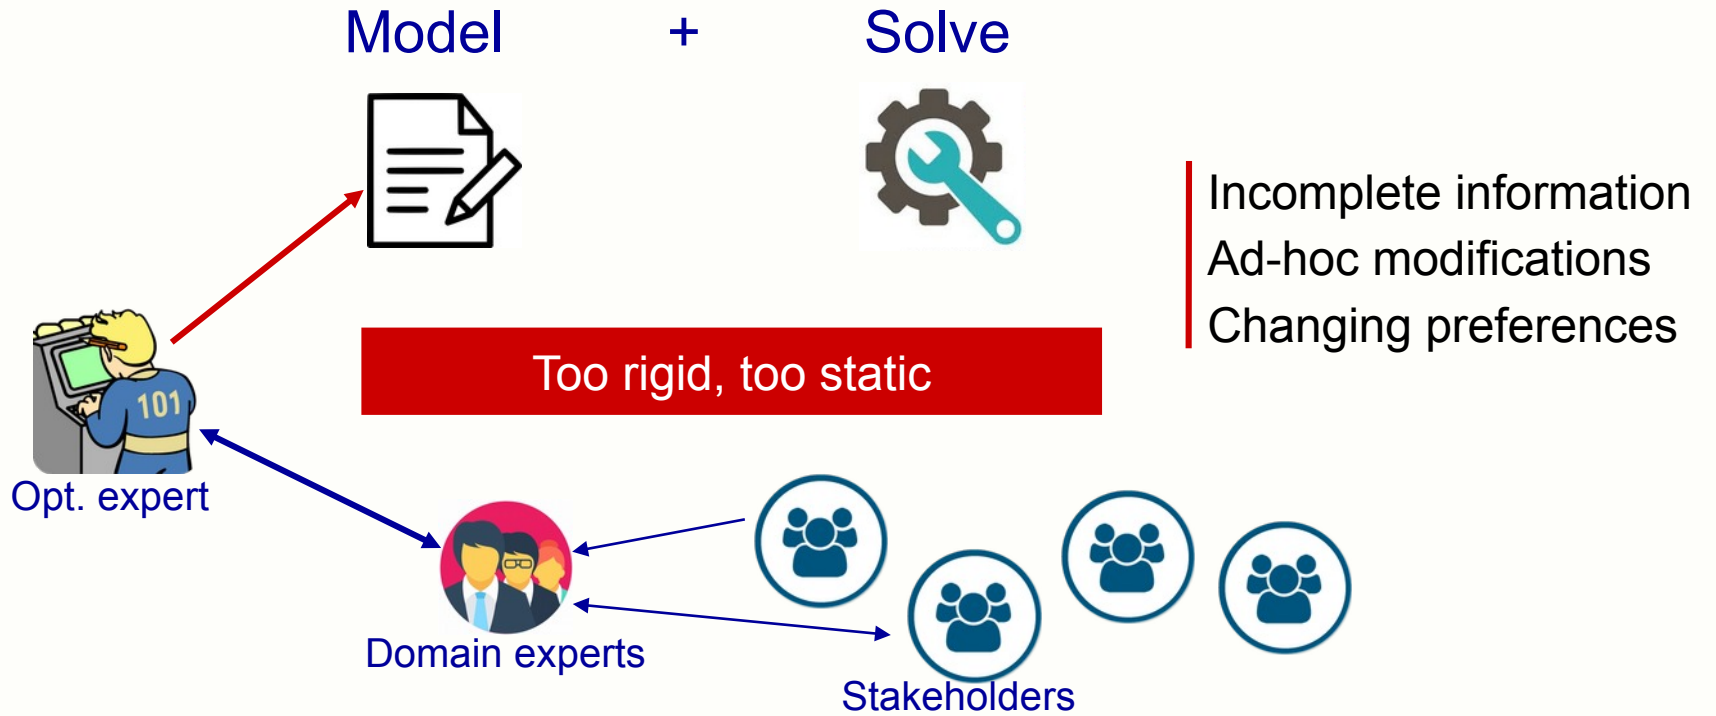
Can we *learn* it instead?



Current combinatorial optimisation practice



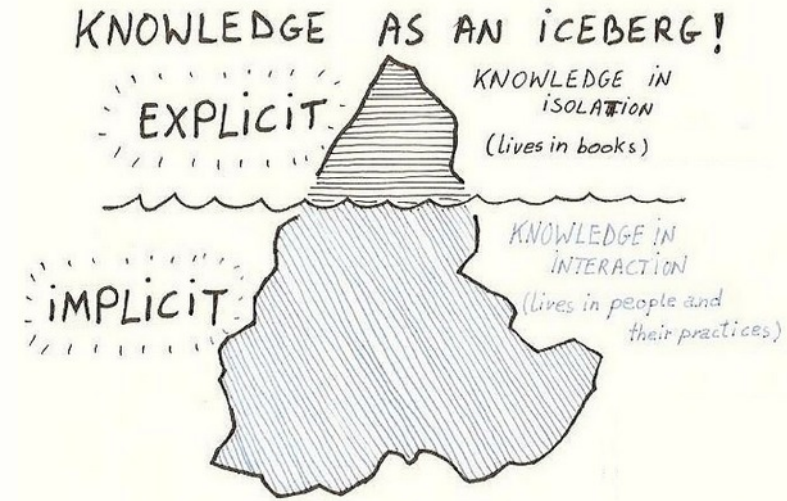
Current combinatorial opt. practice, **problem**





Prediction + constraint solving


- Part explicit knowledge:
in a formal language
- Part implicit knowledge:
learned from data





Prediction + constraint solving

- Part explicit knowledge:
in a formal language
- Part implicit knowledge:
learned from data



- » **perception** (*vision, natural language, ...*)
- » tacit knowledge (*user preferences, social, ...*)
- » complex environment (*demand, prices, ...*)

Perception-based Constraint Solving: a demo application



The logo for 'Sudoku Assistant' is a 3x3 grid. The top-left cell contains the number '7'. The top-right cell contains a gear icon. The middle-left cell contains a neural network icon. The middle-middle cell contains the number '4'. The middle-right cell contains the number '6'. The bottom-middle cell contains a camera icon.

Sudoku Assistant

Tias Guns, Milan Pesa, Maxime Mulamba,
Ignace Bleux, Emilio Gamba, Senne Berden




Logos of funding organizations: European Union flag, ERC (European Research Council), AI FLANDERS (BUILDING OUR DIGITAL FUTURE), KU LEUVEN, VUB, and ARTIFICIAL INTELLIGENCE RESEARCH GROUP.



<https://visualsudoku.cs.kuleuven.be>

Perception-based constraint solving

Visual sudoku (naïve)



Pre-trained neural network

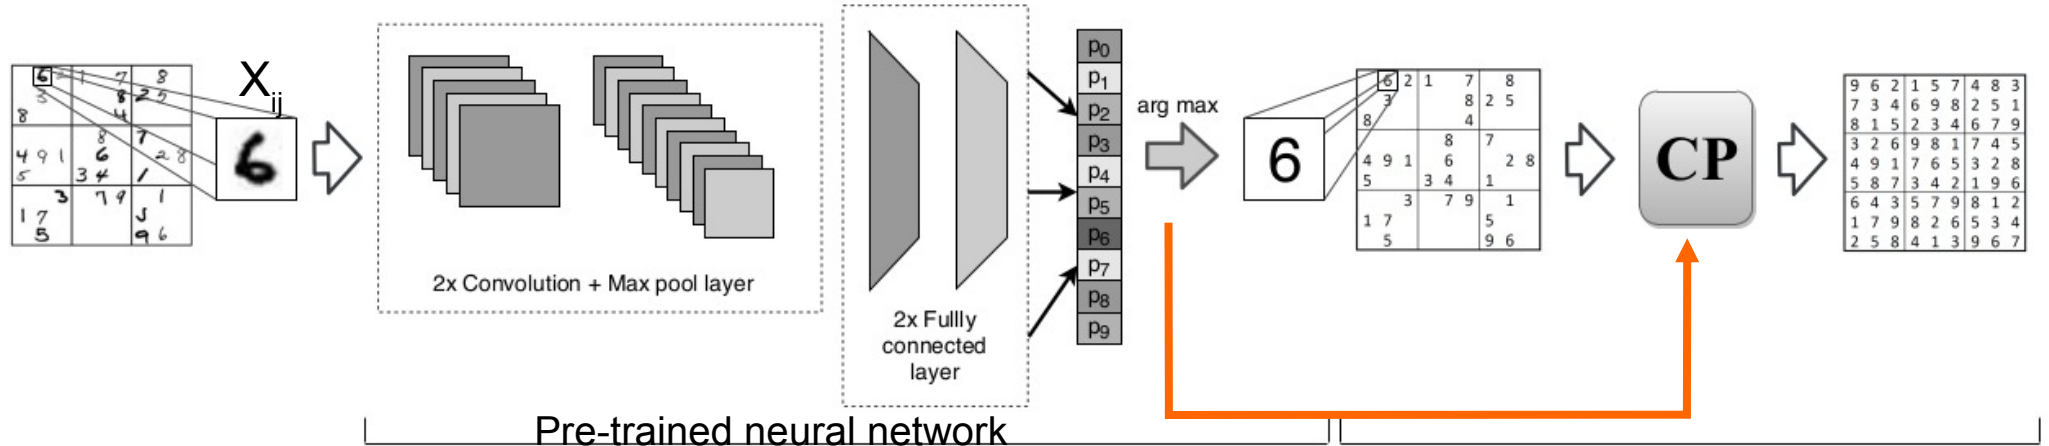
Solving

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



Check Alignment

Perception-based constraint solving



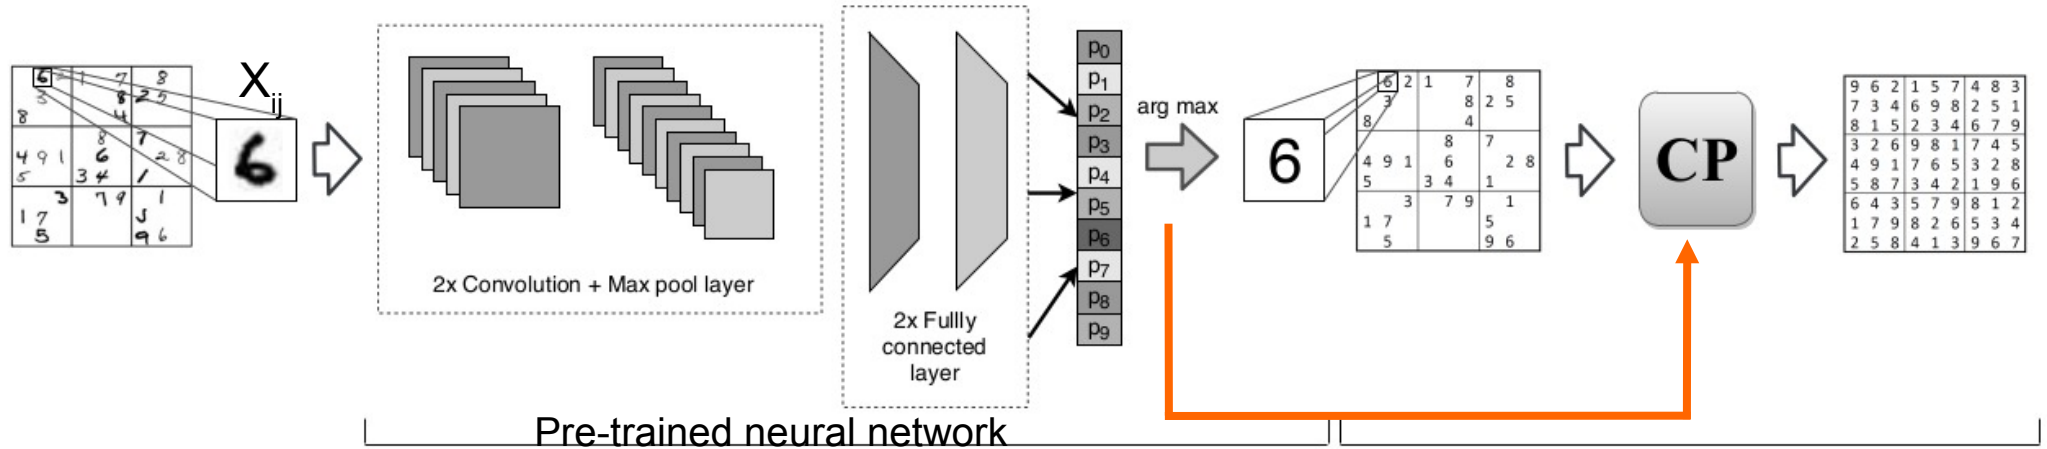
What about the *next* most likely interpretation?

- 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

Perception-based constraint solving



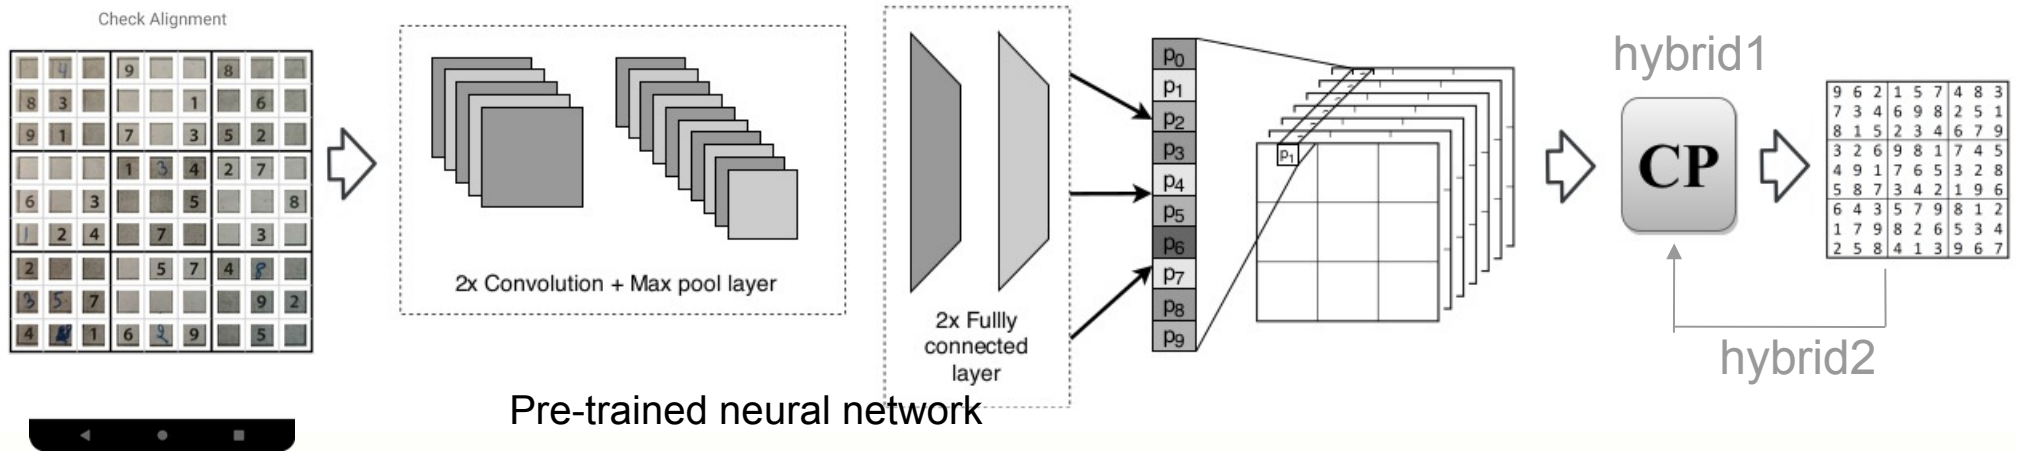
$$\hat{y} = \arg \max \prod_{ij} P(y_{ij} = k | X_{ij}) \quad \text{s.t.} \quad \text{sudoku}(\hat{y})$$

- Log-likelihood trick:

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

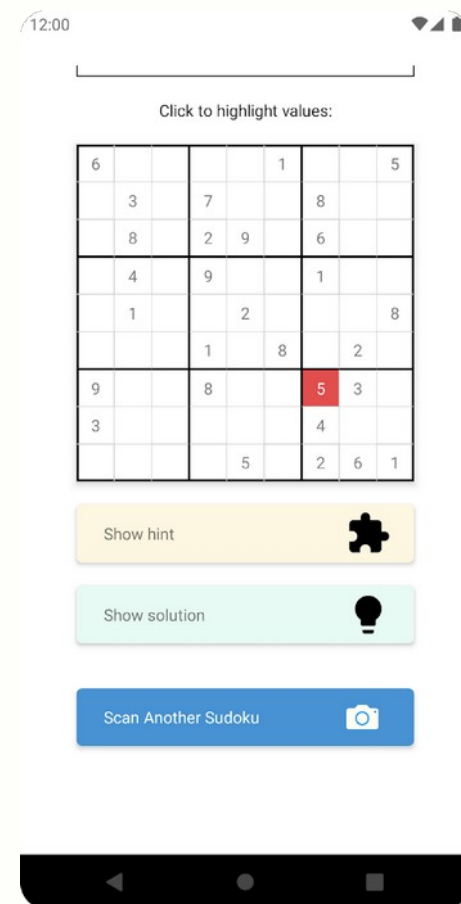
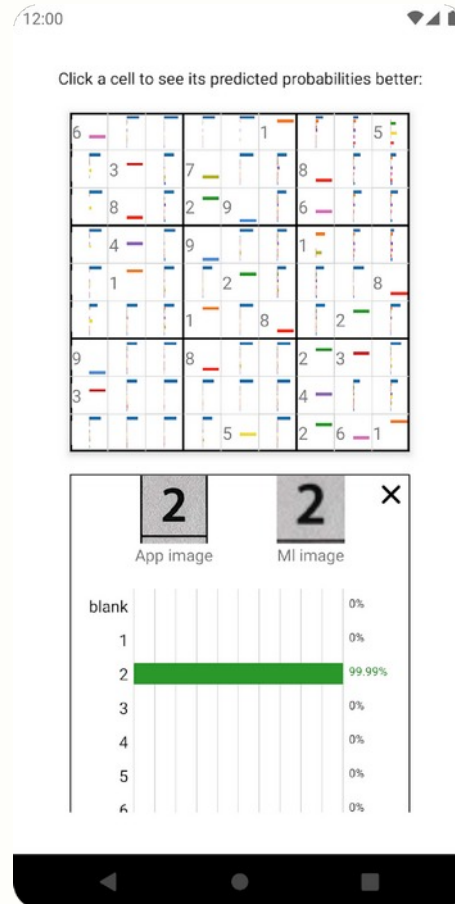

Perception-based constraint solving

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	<u>99.72%</u>	99.44%	92.93%	0%	0.83

Sudoku Assistant demo, continued




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			
<small>(For Individuals having Income from Salaries, One House Property, Other Sources Interest etc) Refer to Instructions for eligibility.</small>										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 16] Email Address pranab.banejee83@gmail.com			
[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		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)													





Prediction + constraint solving

- Part explicit knowledge:
in a formal language
- Part implicit knowledge:
learned from data



- » perception (*vision, natural language, ...*)
- » **tacit knowledge** (*user preferences, social, ...*)
- » complex environment (*demand, prices, ...*)

Tacit knowledge (user preferences)

Example:

“in real-life operations, the quality of a route is not exclusively defined by its theoretical length, duration, or cost”

- Data:
 - stop list, zones
 - TSP solutions with “good, average, bad” labels

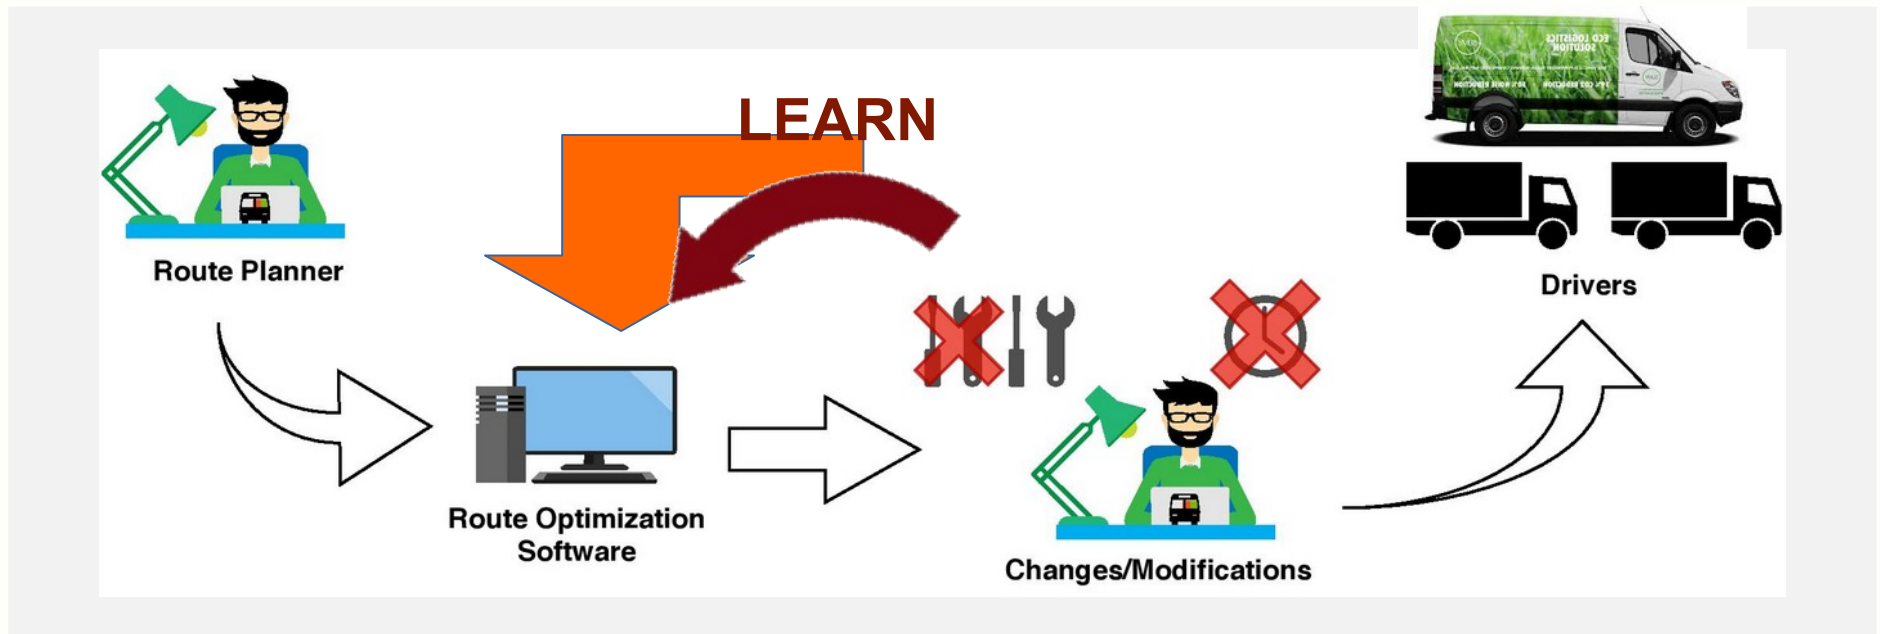
→ learn and route



Tacit knowledge (user preferences)

“Vehicle routing by learning from historical solutions”

[Rocsildes Canoy and Tias Guns, CP19], **Best student paper award**



GOAL: Learn preferences, reduce manual effort, adapt to changes over time!

Tacit knowledge (user preferences)

Small data: 6 months = 26 weeks = 130 week days (instances)



Tacit knowledge (user preferences)


For single vehicles, in mobility data mining literature:

- ▶ Driver turn prediction [Krumm, 2008]
- ▶ Prediction of remainder of route early in the trip [Ye et al., 2015]
- ▶ Prediction of route given origin and destination [Wang et al., 2015]



Can we use similar techniques
to learn preferences across routings of multiple vehicles?

And can we optimize over them with constraint solving?



Learning and prediction part


1st order Markov approximation:

$$P([s_1, s_2, s_3, \dots]) = P(s_1) * P(s_2 | s_1) * P(s_3 | s_2) * \dots$$

→ estimate the $P(s_y | s_x)$ by observing the transitions in the actually driven routes

probability of transition = relative nr of observations in the data

$$t_{ij} = \frac{f_{ij} + \alpha}{N_i + \alpha d},$$




Concept drift: new/gone clients

When 'counting' the probabilities:

- can include a *prior* on each historic instance wrt. current day
- e.g. weighing of the instance:

$$\mathbf{F} = \sum_t w_t \mathbf{A}^t.$$

- ♦ uniform = unit weight
- ♦ by similarity = how much overlap in clients with current day
- ♦ by time = more recent instances get higher weight
incl. exponential smoothing



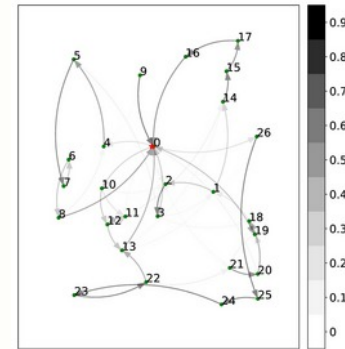
Constrained optimisation: what now?



Goal: find constrained maximum likelihood solution:

$$\text{maximize } P([s_1, s_2, s_3, \dots]) = P(s_1) * P(s_2 | s_1) * P(s_3 | s_2) * \dots$$


$$\text{s.t. VRP}([s_1, s_2, s_3, \dots])$$



Standard probability computation trick: log-likelihood

$$\max \prod_{(i,j) \in X} \Pr(\text{next stop} = j \mid \text{current stop} = i),$$

$$= \max \sum_{(i,j) \in A} \log(t_{ij}) x_{ij}.$$



→ *VRP*: replace *distance matrix* by *negative log-likelihood matrix*!

$$\min \sum_{(i,j) \in A} c_{ij} x_{ij} \quad \Rightarrow \quad \min \sum_{(i,j) \in A} -\log(t_{ij}) x_{ij}.$$

Compatible with ALL vrp solvers

Concept drift, quality AFTER solving

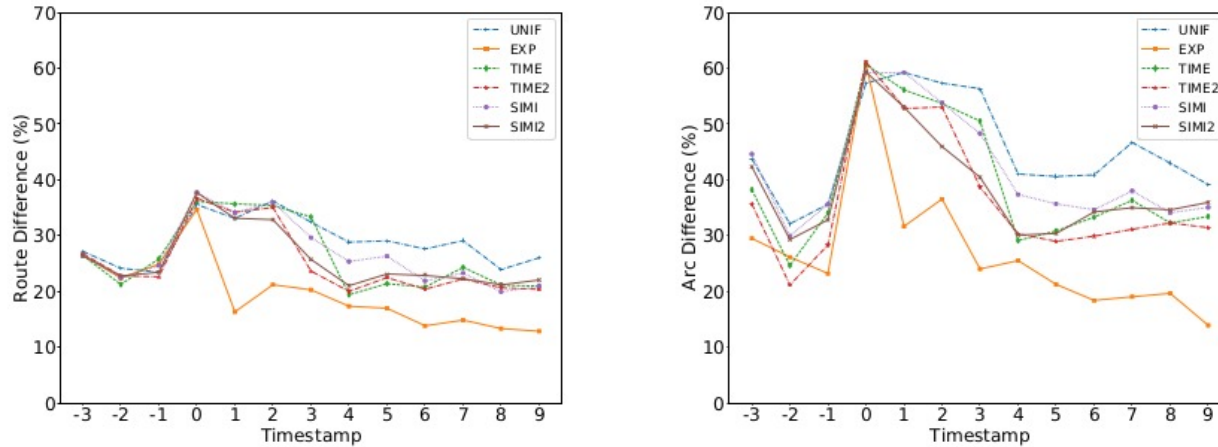


Fig. 7 Route and arc difference during concept drift (drop in number of stops)




Fig. 8 Route and arc difference during concept drift (rise in number of stops)


Extension: Neural instead of Markov?

Opportunities: contextual *features* (day of week, nr vehicles...)

Challenges:

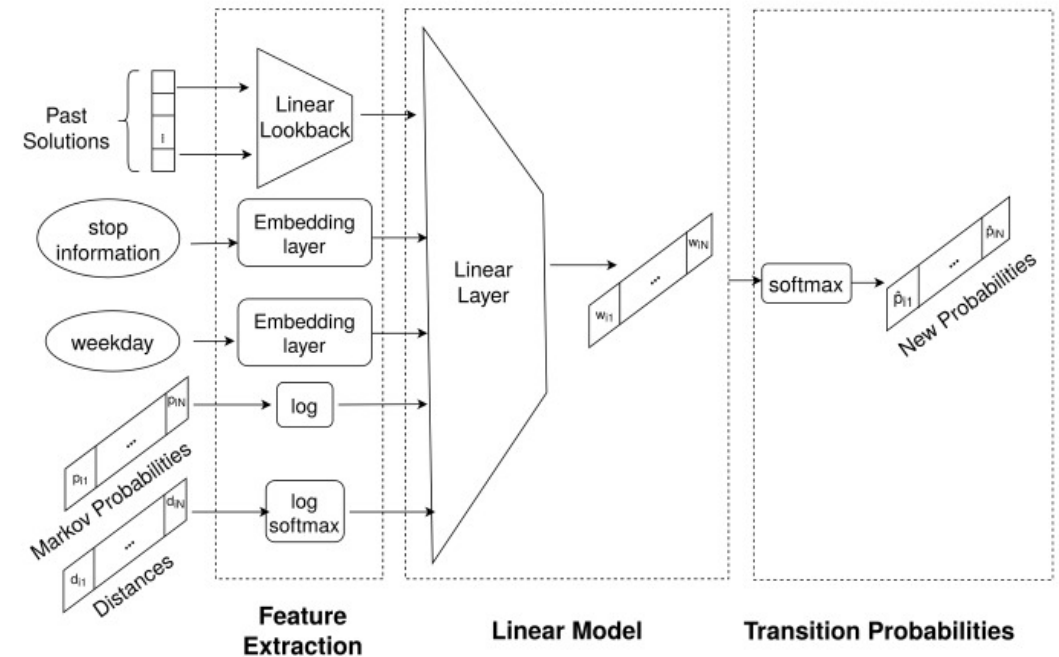
- Need to predict $n \times n$ outputs
- What input representation?
 - Encoding the (variable amount) of stops
 - Encoding all $n \times n$ distances
 - Encoding the temporal historic instances
- What loss function?
- Only few data (large networks will overfit)

Neural representation



- Key ideas:
 - domain-specific architecture,
 - 1 source → all stops

- Contextual features:



Loss function: Classification loss as proxy for arc-difference
(log likelihood of arcs in 'preferred' solution)

Summary comparison:

	CE	AD (%)	RD(%)	Distance
Markov Counting	2.44	18.55	17.26	418
Neural Net (Without linear lookback)	1.04	18.04	17.02	414
Distance based VRP	11.90	73.14	46.93	366

Learning the preferences

= imitation of user choices → copying, not *intelligence*?

Optimisation software is meant to do *better* than a user
(by considering larger nr of candidates and better resolving of conflicts)



I prefer route X even if it is 2 kilometers longer
→ trade's off distance versus preference


Give control to combine both:

$$t'_{ij} = \beta t_{ij} + (1 - \beta) d_{ij}.$$



Prediction + constraint solving

- Part explicit knowledge:
in a formal language
- Part implicit knowledge:
learned from data




- » perception (*vision, natural language, ...*)
- » tacit knowledge (*user preferences, social, ...*)
- » **complex environment** (*demand, prices, ...*)

Time for end-to-end training!

Learn the objective function

Prediction + Optimisation (regression of weights)



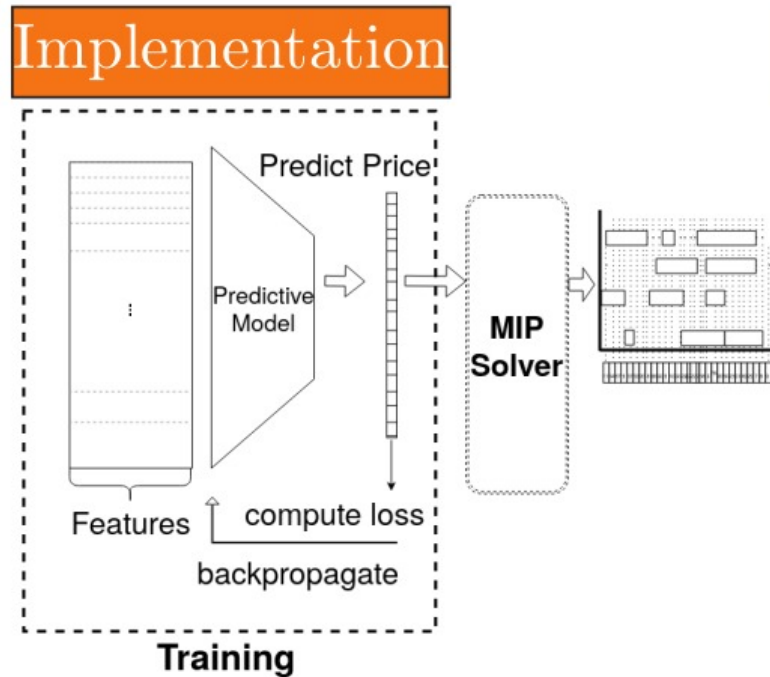
Predict: hourly energy prices +

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

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

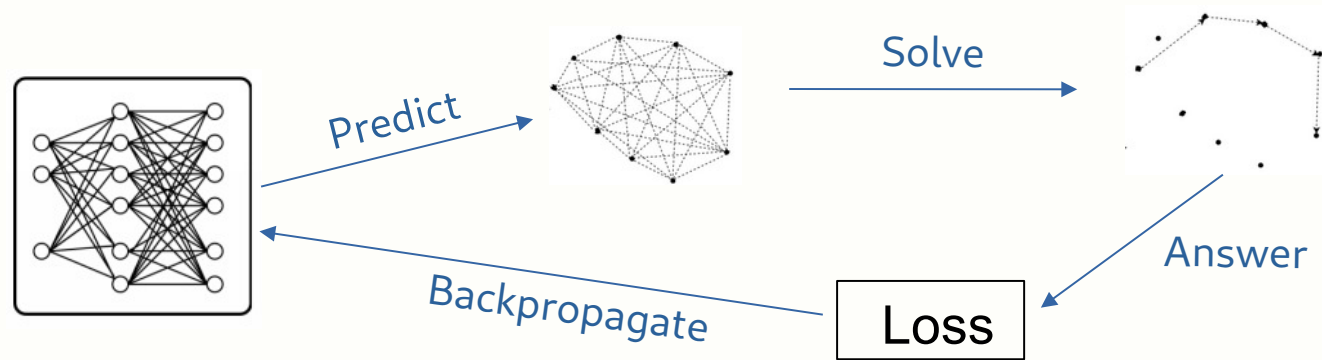
network params

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

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 (QP) problem** [Wilder, B., Dilkina, B., & Tambe, M. (2019, July)., Ferber, A., Wilder, B., Dilkina, B., & Tambe, M. (2020, April)]
- **Differentiating HSD of a relaxed (LP) 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-focused 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 empirical distributions (one for true costs, one for predicted)
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



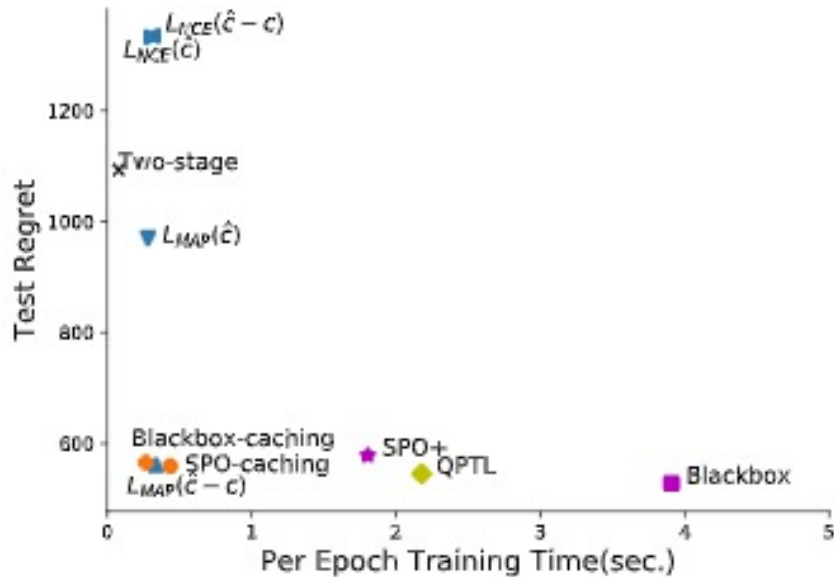
Can use **cached solutions** as approximate solver!!

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 θ (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

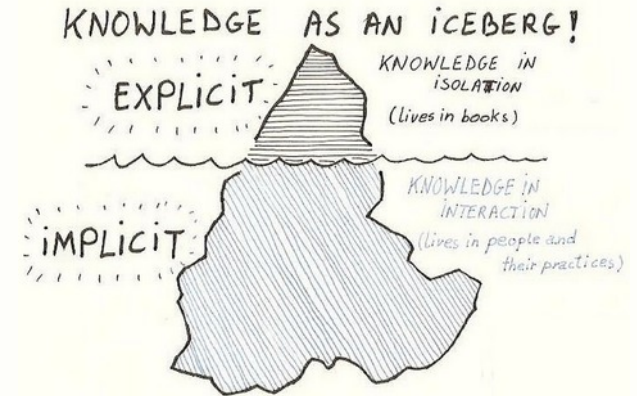
Input : training data $\mathcal{D} = \{X, y\}_{i=1}^n$, architecture g , learning rate γ , sample rate r

- 1 initialize θ (neural network weights of g)
- 2 $\text{sols} \leftarrow \{\text{solver}(y) \mid (X, y) \in \mathcal{D}\}$ (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** $\text{random}() \leq r$ **then**
- 8 $\text{sols} \leftarrow \text{sols} \cup \{\text{solver}(\hat{y})\}$
- 9 **end**
- 10 Compute loss $L = \text{KL}(\text{distr}(y, \text{sols}), \text{distr}(\hat{y}, \text{sols}))$ 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**




Key take-aways:

- Explicit knowledge: use solver
- Implicit knowledge: do learning
- Joint inference / collective classification:
maximize log likelihood!
- Keep revisiting the solving AND the learning,
hybridize and use properties of one in the other!
- Prediction + Optimisation with decision-focussed learning possible




Bigger picture




Bigger picture

- Learning implicit user preferences
- Learning from the environment




Bigger picture

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




Bigger picture

- 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

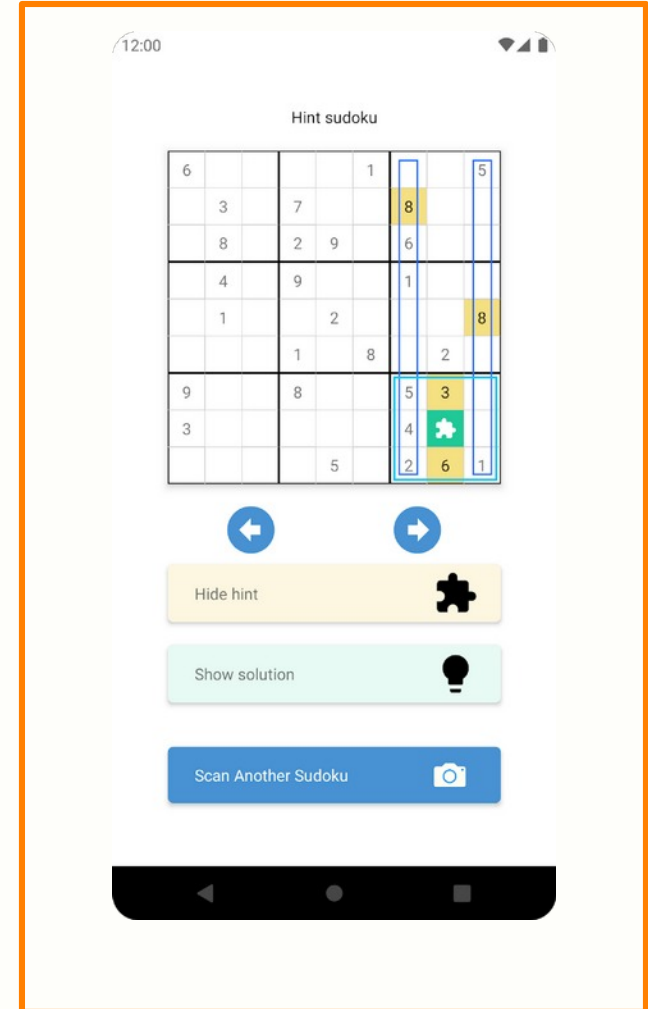
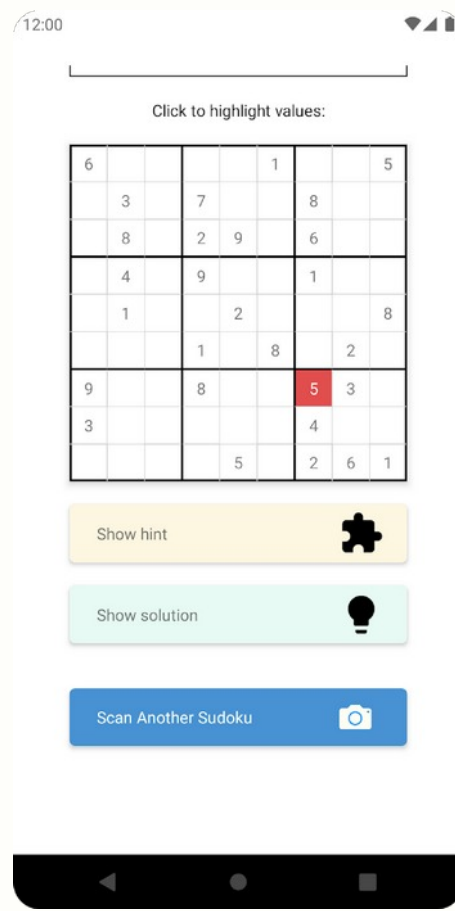
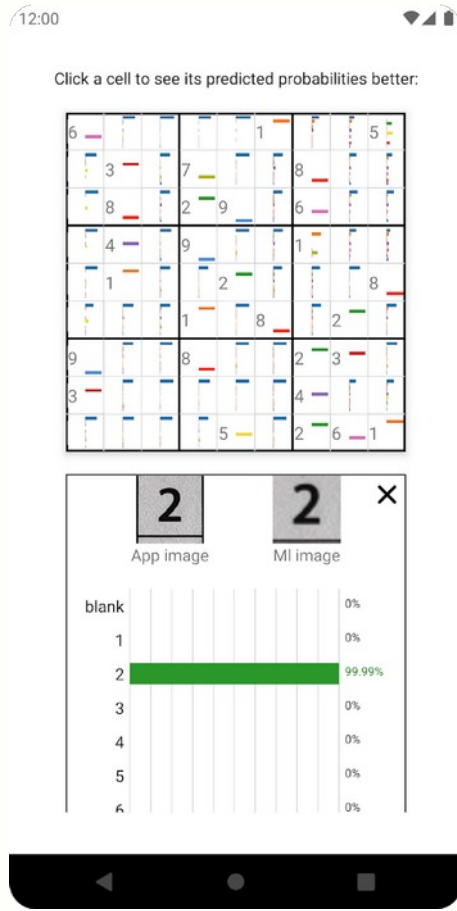


Towards **co-creation** of constraint optimisation solutions


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

Hiring post-docs!

Sudoku Assistant, explanation steps




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



Easy to integrate (Python) libraries

- 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



Easy to integrate (Python) libraries

- Easy integration with Machine Learning
=> Python and numpy arrays
- Efficient repeated solving
=> Incremental
- Use CP/SAT/MIP or any commercial solver
=> 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 searching over discrete decision variables.

CPMpy allows to model search problems in a high-level manner, by defining decision constraints and an objective over them (similar to MiniZinc and Essence²). You can define functions and indexing while doing so. This model is then automatically translated to a 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:



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

<https://people.cs.kuleuven.be/~tias.guns>

 @TiasGuns

Hiring post-docs!