UCLA Computer Science

STAR AI RESEARCH LAB UCLA

Bridging Data and Knowledge in Neuro-Symbolic Learning

Guy Van den Broeck

Summer school on Neurosymbolic Programming - Jul 11 2022

The AI Dilemma

Pure (Logic) Reasoning ⟷ **Pure Learning**

The AI Dilemma

Pure (Logic) Reasoning **Pure Learning**

- Slow thinking: deliberative, cognitive, model-based, extrapolation
- Amazing achievements until this day

- “*Pure logic is brittle*”

  noise, uncertainty, incomplete knowledge, …

The AI Dilemma

Pure (Logic) Reasoning **Pure Learning**

- Fast thinking: instinctive, perceptive, model-free, interpolation
- Amazing achievements recently
- "*Pure learning is brittle*"

bias, algorithmic fairness, interpretability, explainability, adversarial attacks, unknown unknowns, calibration, verification, missing features, missing labels, data efficiency, shift in distribution, general robustness and safety

fails to incorporate a sensible model of the world

The AI Dilemma

Pure (Logic) Reasoning **Pure Learning**

Integrate reasoning into modern learning algorithms

Today: Deep learning with structured output constraints
Learning monotonic neural networks

Knowledge in Vision, Robotics, NLP, Activity Recognition

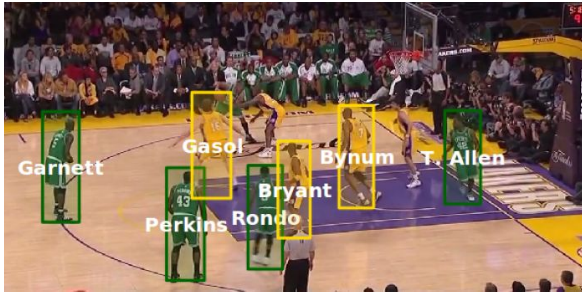

People appear at most
once in a frame

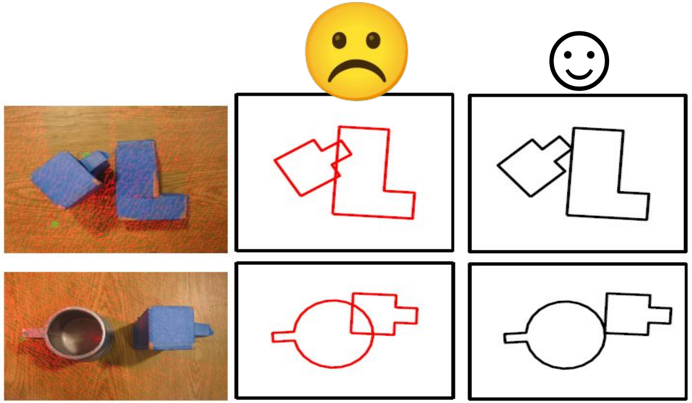

Rigid objects don't overlap

At least one verb
in each sentence.
If X and Y are married,
then they are people.

Cut the orange before
squeezing the orange

[Lu, W. L., Ting, J. A., Little, J. J., & Murphy, K. P. (2013). Learning to track and identify players from broadcast sports videos.], [Wong, L. L., Kaelbling, L. P., & Lozano-Perez, T., Collision-free state estimation. ICRA 2012], [Chang, M., Ratinov, L., & Roth, D. (2008). Constraints as prior knowledge], [Ganchev, K., Gillenwater, J., & Taskar, B. (2010). Posterior regularization for structured latent variable models]… and many many more!

Motivation: Deep Learning

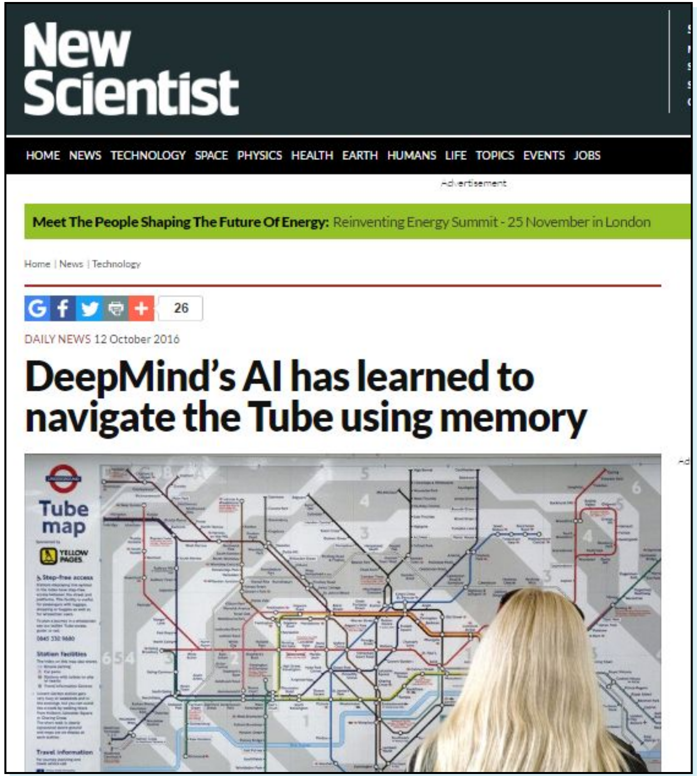
New Scientist

HOME NEWS TECHNOLOGY SPACE PHYSICS HEALTH EARTH HUMANS LIFE TOPICS EVENTS JOBS

Advertisement

Meet The People Shaping The Future Of Energy: Reinventing Energy Summit - 25 November in London

Home | News | Technology

26

DAILY NEWS 12 October 2016

DeepMind's AI has learned to navigate the Tube using memory


nature International weekly journal of science

Home | News & Comment | Research | Careers & Jobs | Current Issue | Archive | Audio & Video | For

News & Comment > News > 2016 > November > Article

NATURE | NEWS

Google's AI reasons its way around the London Underground

DeepMind's latest technique uses external memory to solve tasks that require logic and reasoning — a step toward more human-like AI.

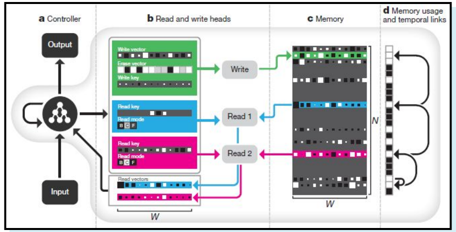

[Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., et al.. (2016). Hybrid computing using a neural network with dynamic external memory. *Nature*, *538*(7626), 471-476.]

Motivation: Deep Learning

DeepMind's latest technique uses external memory to **solve tasks that require logic and reasoning — a step toward more human-like AI.**

… but …

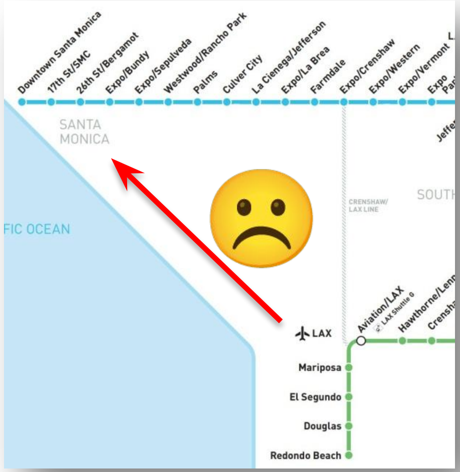

optimal planner recalculating a shortest path to the end node. To ensure that the network always moved to a valid node, the output distribution was renormalized over the set of possible triples outgoing from the current node. The performance

it also received input triples during the answer phase, indicating the actions chosen on the previous time-step. This makes the problem a 'structured prediction'

[Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., et al.. (2016). Hybrid computing using a neural network with dynamic external memory. *Nature*, *538*(7626), 471-476.]

Warcraft Shortest Path

Predicting the minimum-cost path

[Differentiation of Blackbox Combinatorial Solvers, Marin Vlastelica, Anselm Paulus, Vít Musil, Georg Martius, Michal Rolínek, 2019]

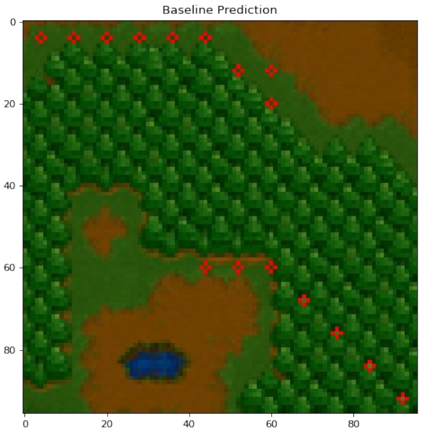
Baseline Prediction

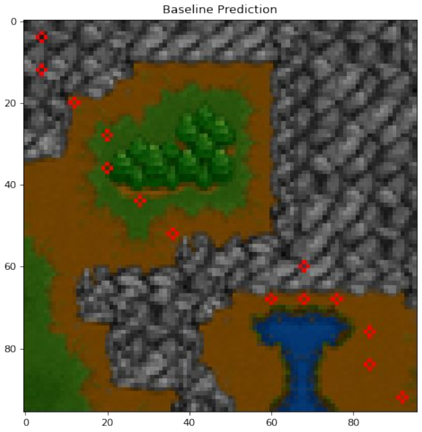
Baseline Prediction

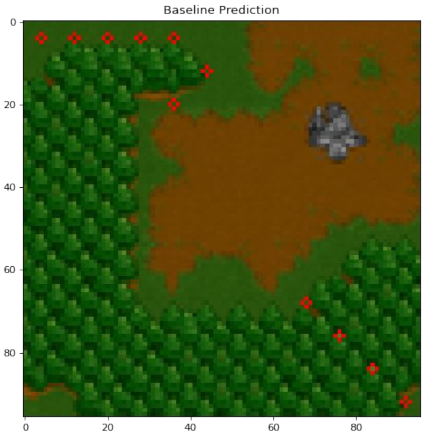
Baseline Prediction

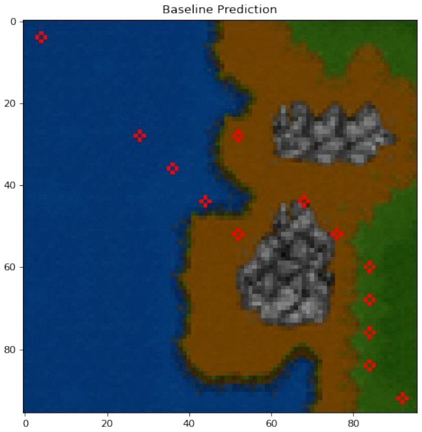
Baseline Prediction

Predict Loan Amount

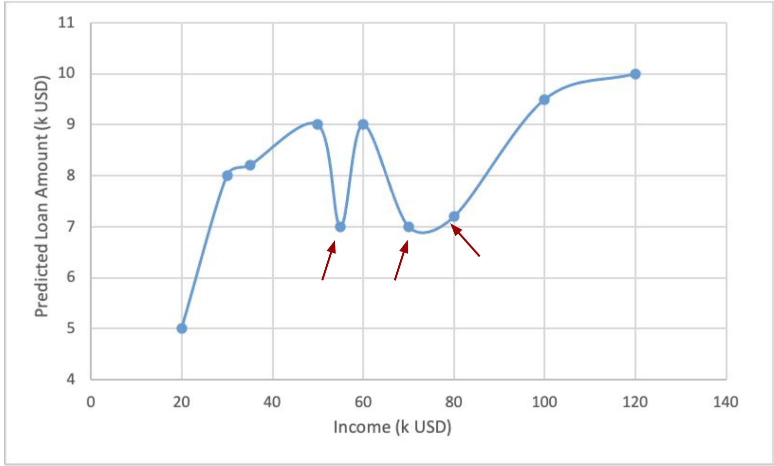

Neural Network Model: Increasing income can decrease the approved loan amount

Monotonicity (Prior Knowledge):
Increasing income should increase the approved loan amount

Knowledge vs. Data

- Where did the world knowledge go?
  - Python scripts
    - Decode/encode/search cleverly
    - Fix inconsistent beliefs
  - Rule-based decision systems
  - Dataset design
  - "a big hack" (with author's permission)
- In some sense we went backwards

  Less principled, scientific, and intellectually satisfying ways of incorporating knowledge

Deep Learning with Constraints

without constraint

without constraint

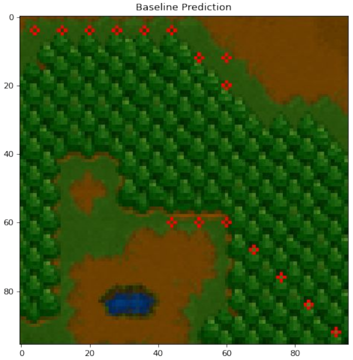

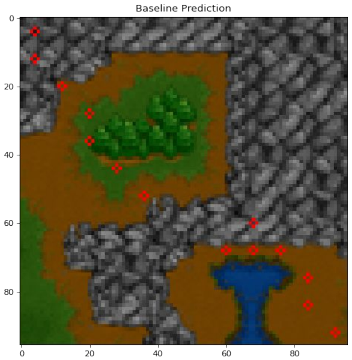

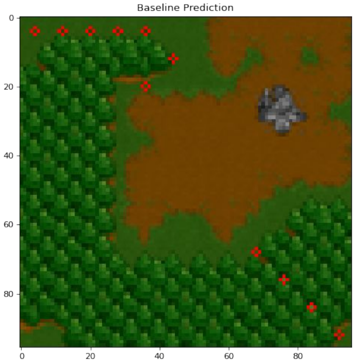

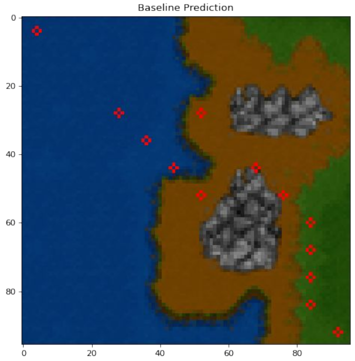

Warcraft min-cost simple-path prediction results

Test accuracy %	Coherent	Incoherent	Constraint
ResNet-18	44.8	**97.7**	56.9

Is prediction the shortest path?
This is the real task!

Are individual edge predictions correct?

Is output a path?

Kareem Ahmed, Eric Wang, Kai-Wei Chang and Guy Van den Broeck. Neuro-Symbolic Entropy Regularization, 2021.

pylon

A PyTorch Framework for Learning with Constraints

Kareem Ahmed Tao Li Thy Ton Quan Guo,
Kai-Wei Chang Parisa Kordjamshidi Vivek Srikumar
Guy Van den Broeck Sameer Singh

http://pylon-lib.github.io

Declarative Knowledge of the Output

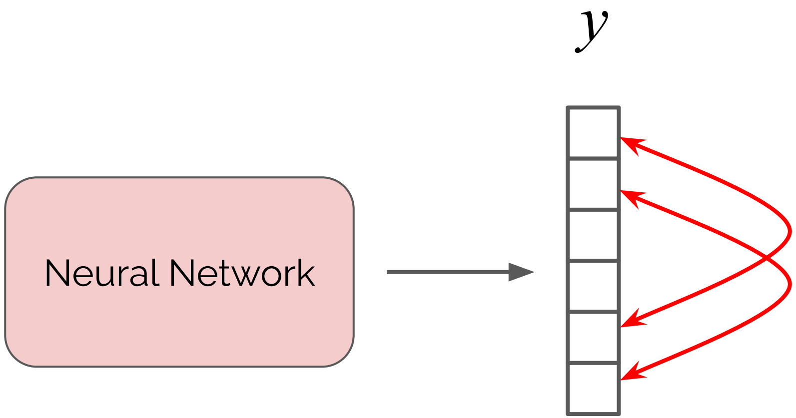

How is the output structured?
Are all possible outputs valid?

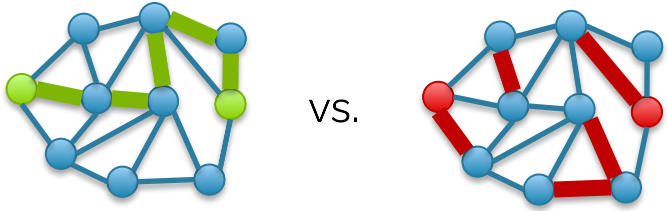

How are the outputs related to each other?

Learning this from data is inefficient
Much easier to express this declaratively

pylon

Library that extends PyTorch to allow injection of declarative knowledge

- **Easy to Express Knowledge:** users write arbitrary constraints on the output
- **Integrates with PyTorch:** minimal change to existing code
- **Efficient Training:** compiles into loss that can be efficiently optimized
    - Exact semantic loss (see later)
    - Monte-carlo estimate of loss
    - T-norm approximation
    - *your solver?*

pylon

PyTorch Code

```
for i in range(train_iters):
    ...
    py = model(x)
    ...
    loss = CrossEntropy(py,...)
```

1 Specify knowledge as a predicate

```
def check(y):
    ...
    return isValid
```

pylon

PyTorch Code

```
for i in range(train_iters):
    ...
    py = model(x)
    ...
    loss = CrossEntropy(py,...)

    loss += constraint_loss(check)(py)
```

1 Specify knowledge as a predicate

```
def check(y):
    ...
    return isValid
```

2 Add as loss to training

```
loss += constraint_loss(check)
```

pylon

PyTorch Code

```
for i in range(train_iters):
    ...
    py = model(x)
    ...
    loss = CrossEntropy(py,...)

    loss += constraint_loss(check)(py)
```

1 Specify knowledge as a predicate

```
def check(y):
    ...
    return isValid
```

2 Add as loss to training

```
loss += constraint_loss(check)
```

3 pylon derives the gradients
(solves a combinatorial problem)

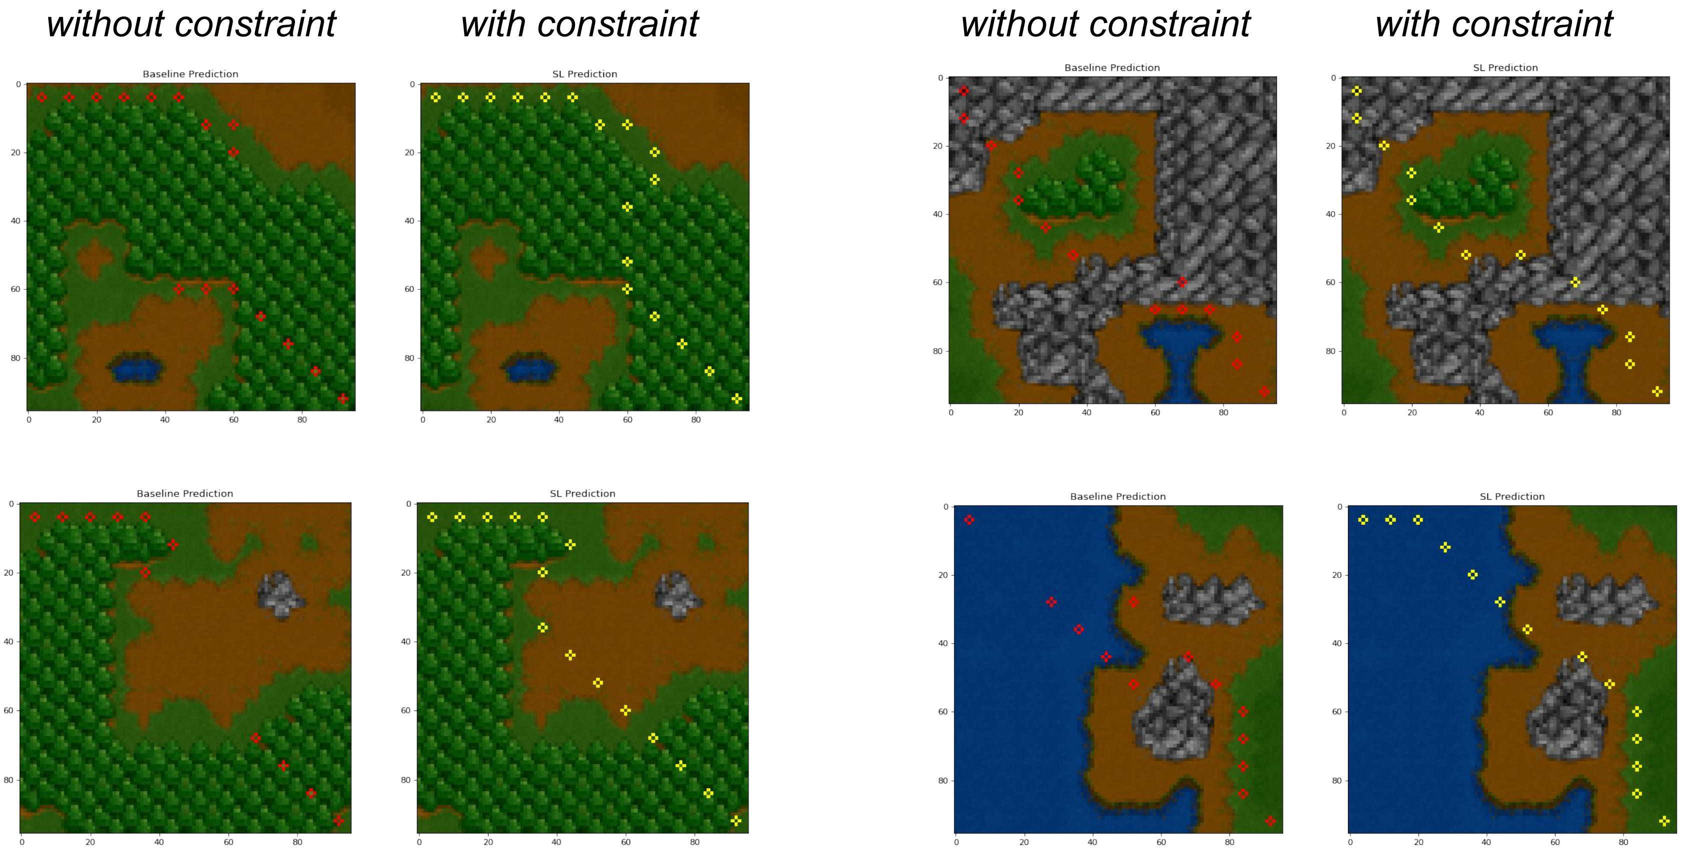
without constraint
with constraint
without constraint
with constraint
Baseline Prediction
SL Prediction
Baseline Prediction
SL Prediction
Baseline Prediction
SL Prediction
Baseline Prediction
SL Prediction

Warcraft min-cost simple-path prediction results

Test accuracy %	Coherent	Incoherent	Constraint
ResNet-18	44.8	**97.7**	56.9
+ Semantic loss	**50.9**	**97.7**	**67.4**

Kareem Ahmed, Eric Wang, Kai-Wei Chang and Guy Van den Broeck. Neuro-Symbolic Entropy Regularization, 2021.

Semantic Loss

*<u>Q</u>: How close is output **p** to satisfying constraint α?*

*<u>A</u>: Semantic loss function L(α,**p**)*

- Axioms, for example:
  - If α constrains to one label, L(α,**p**) is cross-entropy
  - If α implies β then L(α,**p**) ≥ L(β,**p**) (*α more strict*)
- Implied Properties:
  - If α is equivalent to β then L(α,**p**) = L(β,**p**)
  - If **p** is Boolean and satisfies α then L(α,**p**) = 0

SEMANTIC
Loss!

Axioms imply unique semantic loss:

$$\mathrm{L}^{\mathrm{s}}(\alpha, \mathsf{p}) \propto -\log \underbrace{\sum_{\mathbf{x} \models \alpha} \prod_{i:\mathbf{x} \models X_i} \mathsf{p}_i \prod_{i:\mathbf{x} \models \neg X_i} (1 - \mathsf{p}_i)}_{\text{Probability of satisfying constraint } \alpha \text{ after sampling from neural net output layer } \mathbf{p}}$$

Probability of satisfying constraint α after sampling from neural net output layer **p**

In general: #P-hard ☹

Do this probabilistic-logical reasoning during learning in a computation graph

Circuits = Computation Graphs

- Logical circuits that can count solutions (#SAT) also compute semantic loss efficiently in size of circuit

L(α,**p**) = L(, **p**) = - log( )

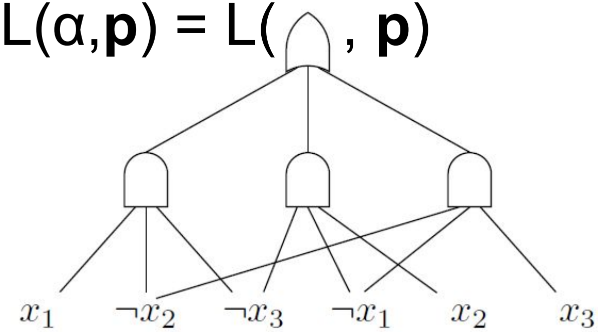

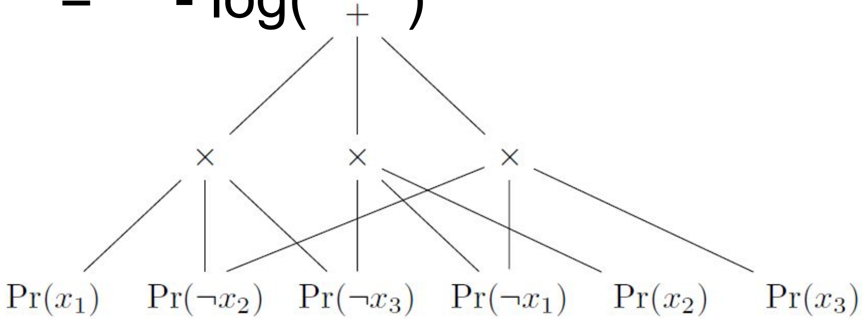

- Compilation into circuit by SAT solvers (once)
- Add circuit to neural network output in pytorch/tensorflow/...

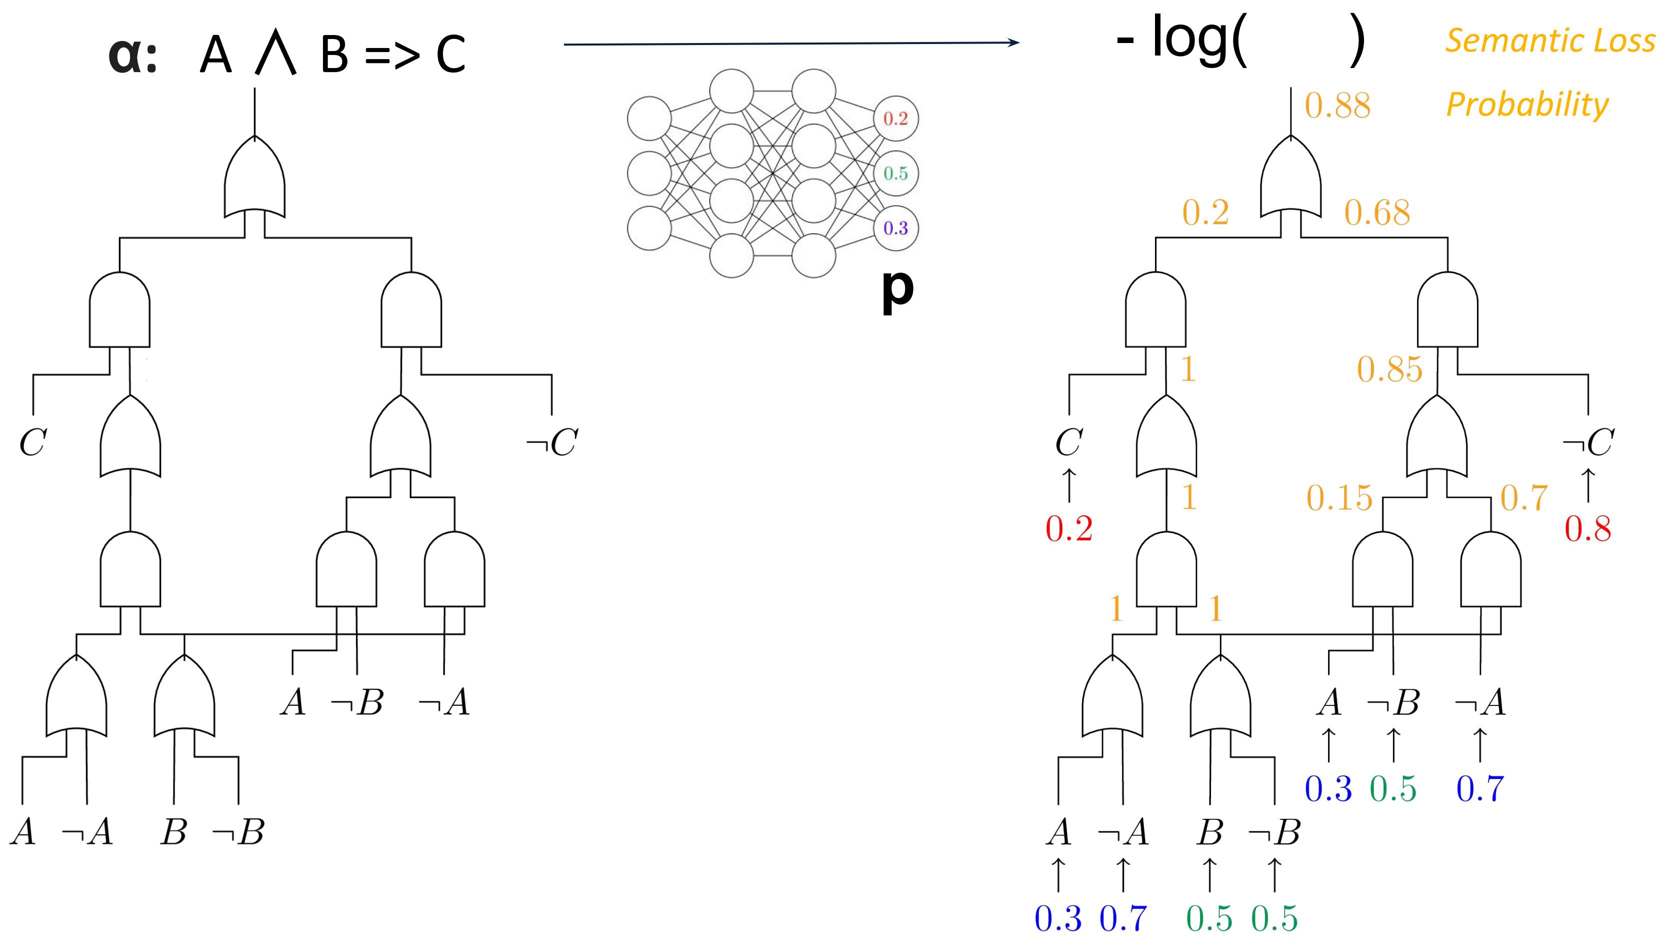
α: A ∧ B => C
- log( )
Semantic Loss
Probability
0.88
0.2
0.68
1
0.85
C
¬C
0.15
0.7
0.2
0.8
1
A
¬B
¬A
0.3
0.5
0.7
A
¬A
B
¬B
0.3
0.7
0.5
0.5
p

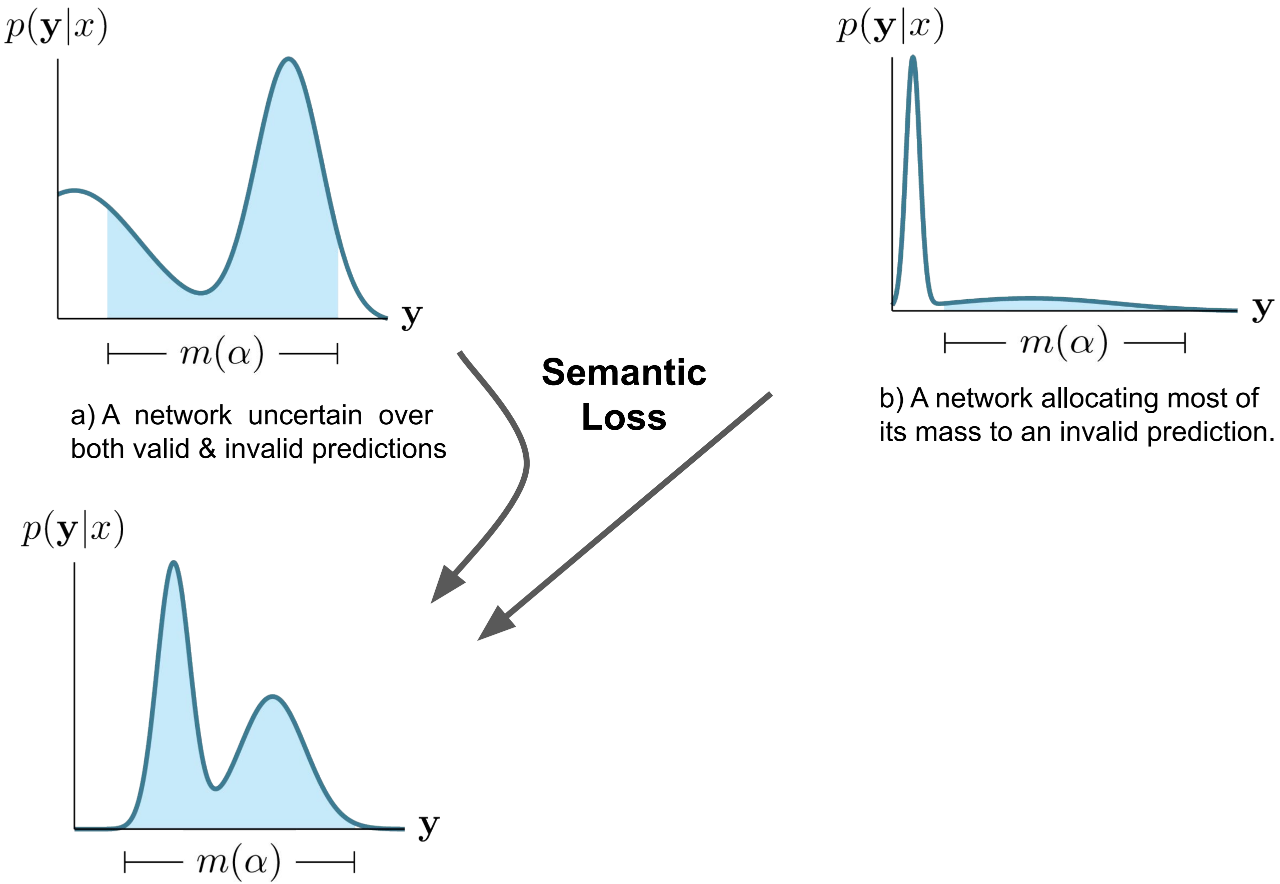

a) A network uncertain over both valid & invalid predictions

b) A network allocating most of its mass to an invalid prediction.

c) A network allocating most of its mass to models of the formula

Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang and Guy Van den Broeck. A Semantic Loss Function for Deep Learning with Symbolic Knowledge, *ICML*, 2018.

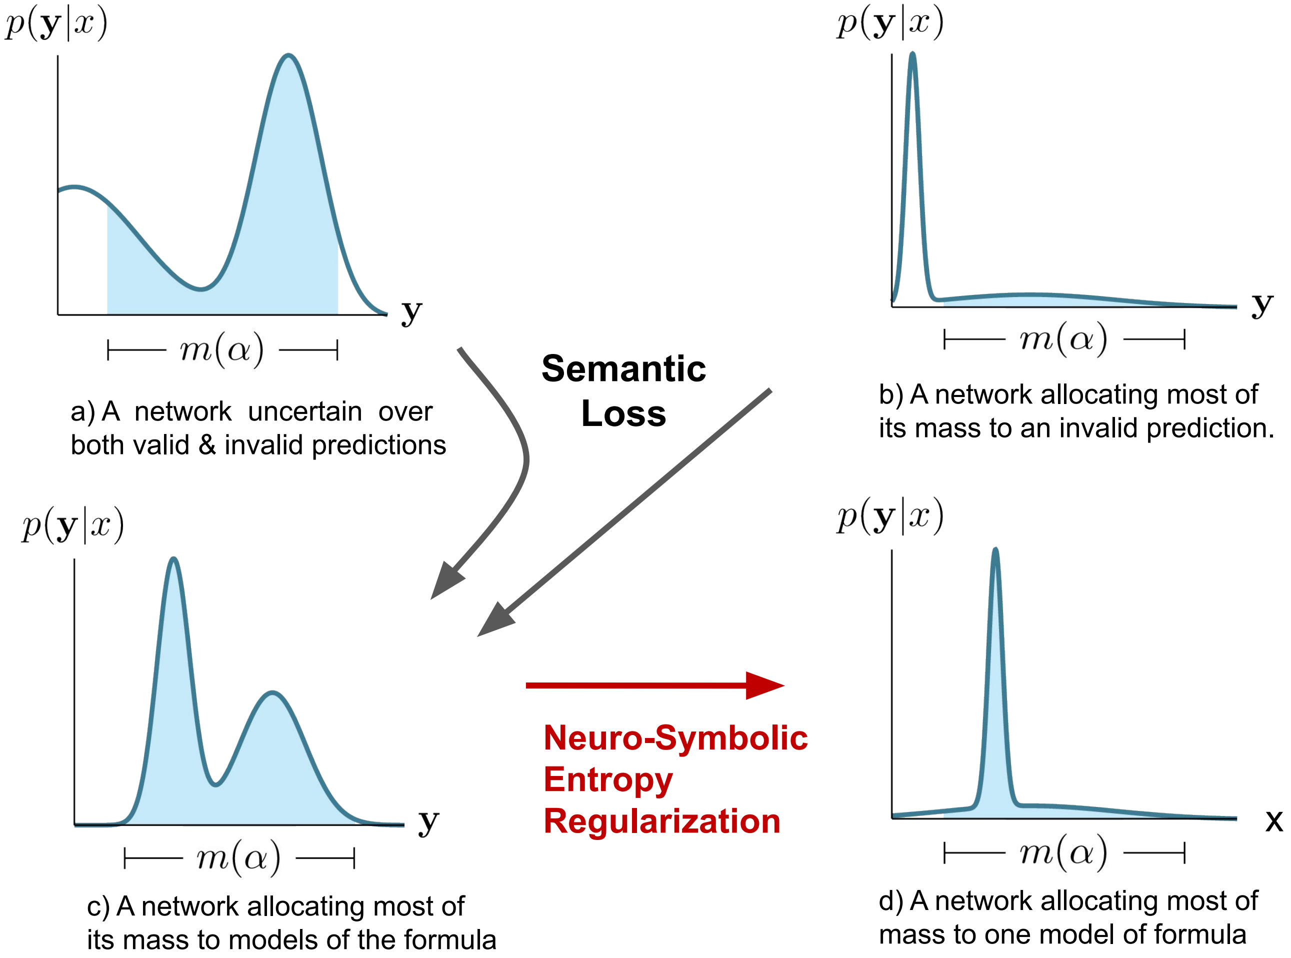

Kareem Ahmed, Eric Wang, Kai-Wei Chang and Guy Van den Broeck. Neuro-Symbolic Entropy Regularization, 2021.

Two complementary neuro-symbolic losses

Neural net $P(\mathbf{Y}|x)$

Constraint α

Semantic Loss $\longrightarrow$ $P(\alpha|x)\uparrow: -\log P(\alpha|x)\downarrow$

Neuro-symbolic Entropy-regularization $\longrightarrow$ $-\mathbb{E}_{P(\mathbf{Y}|x,\alpha)}[\log P(\mathbf{Y}|x,\alpha)]$

Warcraft min-cost simple-path prediction results

Test accuracy %	Coherent	Incoherent	Constraint
ResNet-18	44.8	**97.7**	56.9
Semantic loss	**50.9**	**97.7**	**67.4**
+ Entropy All	51.5	97.6	67.7
+ Entropy Circuit	**55.0**	**97.9**	**69.8**

Kareem Ahmed, Eric Wang, Kai-Wei Chang and Guy Van den Broeck. Neuro-Symbolic Entropy Regularization, 2021.

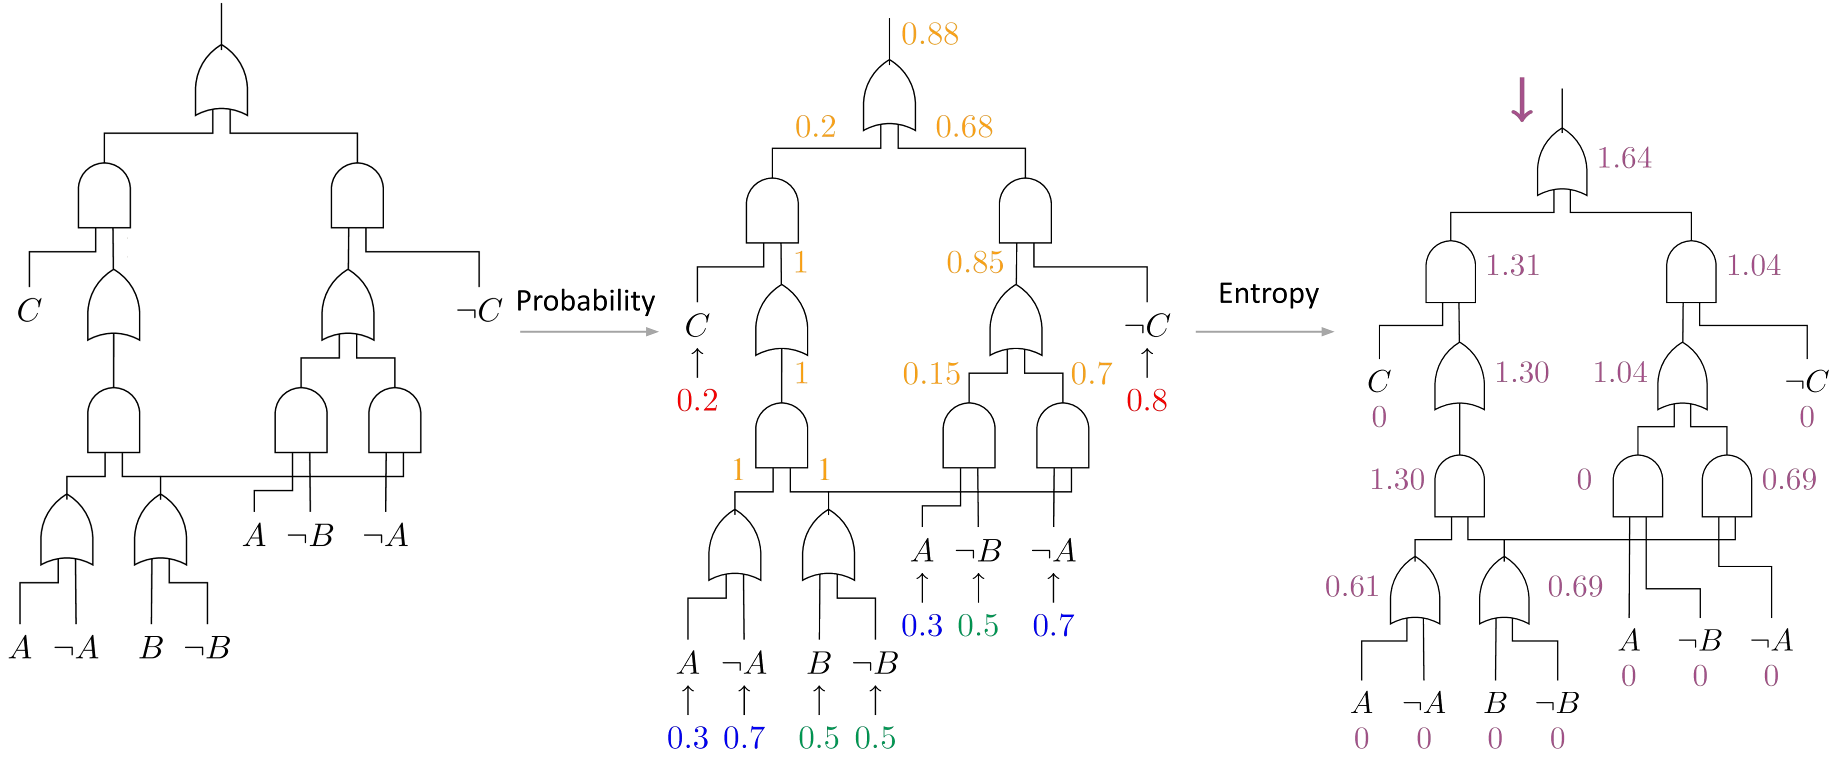

C
$\neg C$
A $\neg B$ $\neg A$
A $\neg A$ B $\neg B$
Probability
0.88
0.2
0.68
1
0.85
C
$\neg C$
0.2
0.8
1
0.15
0.7
1
1
A $\neg B$ $\neg A$
0.3 0.5 0.7
A $\neg A$ B $\neg B$
0.3 0.7 0.5 0.5
Entropy
1.64
1.31
1.04
C
0
1.30
1.04
$\neg C$
0
1.30
0
0.69
0.61
0.69
A $\neg B$ $\neg A$
0 0 0
A $\neg A$ B $\neg B$
0 0 0 0

Joint entity-relation extraction in natural language processing

#		3	5	10	15	25	50	75
ACE05	Baseline	4.92 ± 1.12	7.24 ± 1.75	13.66 ± 0.18	15.07 ± 1.79	21.65 ± 3.41	28.96 ± 0.98	33.02 ± 1.17
	Self-training	7.72 ± 1.21	12.83 ± 2.97	16.22 ± 3.08	17.55 ± 1.41	27.00 ± 3.66	32.90 ± 1.71	37.15 ± 1.42
	Product t-norm	8.89 ± 5.09	14.52 ± 2.13	19.22 ± 5.81	21.80 ± 7.67	30.15 ± 1.01	34.12 ± 2.75	37.35 ± 2.53
	Semantic Loss	12.00 ± 3.81	14.92 ± 3.14	22.23 ± 3.64	27.35 ± 3.10	30.78 ± 0.68	36.76 ± 1.40	38.49 ± 1.74
	+ Full Entropy	**14.80** ± 3.70	15.78 ± 1.90	23.34 ± 4.07	28.09 ± 1.46	31.13 ± 2.26	36.05 ± 1.00	39.39 ± 1.21
	+ NeSy Entropy	14.72 ± 1.57	**18.38** ± 2.50	**26.41** ± 0.49	**31.17** ± 1.68	**35.85** ± 0.75	**37.62** ± 2.17	**41.28** ± 0.46
SciERC	Baseline	2.71 ± 1.10	2.94 ± 1.00	3.49 ± 1.80	3.56 ± 1.10	8.83 ± 1.00	12.32 ± 3.00	12.49 ± 2.60
	Self-training	3.56 ± 1.40	3.04 ± 0.90	4.14 ± 2.60	3.73 ± 1.10	9.44 ± 3.80	14.82 ± 1.20	13.79 ± 3.90
	Product t-norm	**6.50** ± 2.00	8.86 ± 1.20	10.92 ± 1.60	13.38 ± 0.70	13.83 ± 2.90	19.20 ± 1.70	19.54 ± 1.70
	Semantic Loss	6.47 ± 1.02	**9.31** ± 0.76	11.50 ± 1.53	12.97 ± 2.86	14.07 ± 2.33	20.47 ± 2.50	23.72 ± 0.38
	+ Full Entropy	6.26 ± 1.21	8.49 ± 0.85	11.12 ± 1.22	14.10 ± 2.79	17.25 ± 2.75	**22.42** ± 0.43	24.37 ± 1.62
	+ NeSy Entropy	6.19 ± 2.40	8.11 ± 3.66	**13.17** ± 1.08	**15.47** ± 2.19	**17.45** ± 1.52	22.14 ± 1.46	**25.11** ± 1.03

Kareem Ahmed, Eric Wang, Kai-Wei Chang and Guy Van den Broeck. Neuro-Symbolic Entropy Regularization, 2022.

Semantic Probabilistic Layers

- How to give a 100% guarantee that Boolean constraints will be satisfied?
- Bake the constraint into the neural network as a special layer

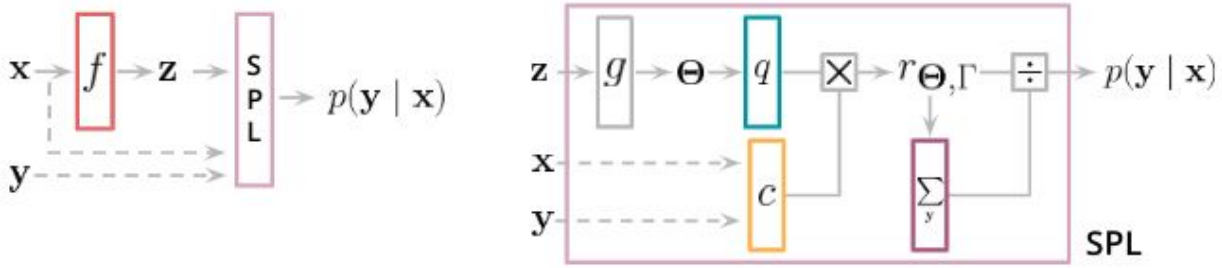

- Secret sauce is again tractable circuits – computation graphs for reasoning

Kareem Ahmed, Stefano Teso, Kai-Wei Chang, Guy Van den Broeck and Antonio Vergari. Semantic Probabilistic Layers for Neuro-Symbolic Learning, 2022.

Warcraft Shortest Path

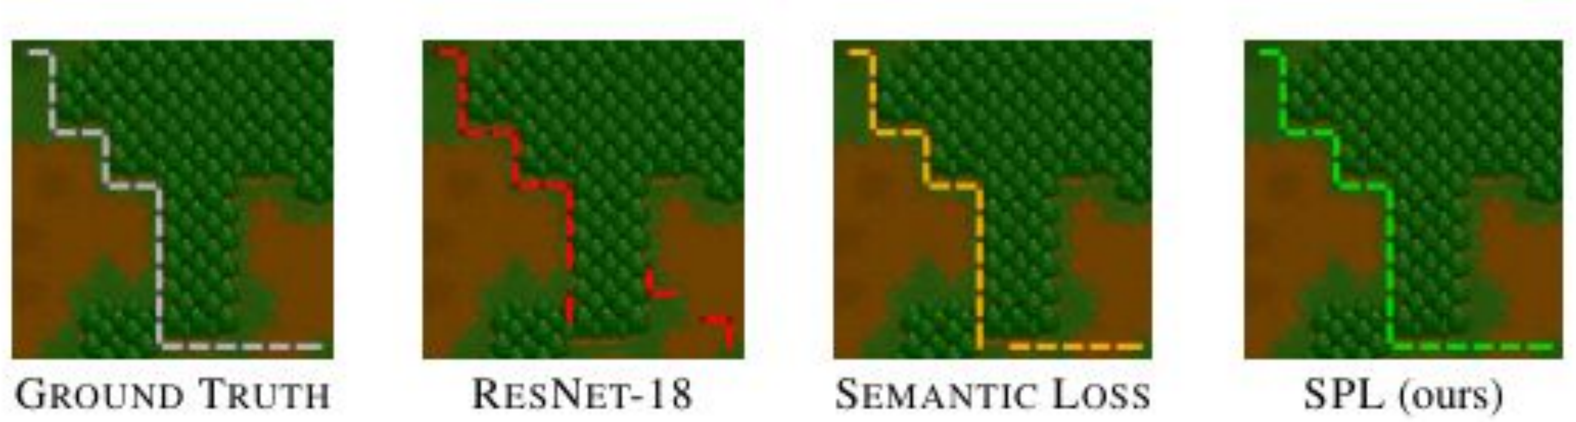

Hierarchical Multi-Label Classification

"if the image is classified as a dog, it must also be classified as an animal"

"if the image is classified as an animal, it must be classified as either cat or dog"

Dataset	Exact Match	
	HMCNN	MLP+SPL
CellCycle	3.05 ± 0.11	$\mathbf{3.79 \pm 0.18}$
Derisi	1.39 ± 0.47	$\mathbf{2.28 \pm 0.23}$
Eisen	5.40 ± 0.15	$\mathbf{6.18 \pm 0.33}$
Expr	4.20 ± 0.21	$\mathbf{5.54 \pm 0.36}$
Gasch1	3.48 ± 0.96	$\mathbf{4.65 \pm 0.30}$
Gasch2	3.11 ± 0.08	$\mathbf{3.95 \pm 0.28}$
Seq	5.24 ± 0.27	$\mathbf{7.98 \pm 0.28}$
Spo	$\mathbf{1.97 \pm 0.06}$	1.92 ± 0.11
Diatoms	48.21 ± 0.57	$\mathbf{58.71 \pm 0.68}$
Enron	5.97 ± 0.56	$\mathbf{8.18 \pm 0.68}$
Imclef07a	79.75 ± 0.38	$\mathbf{86.08 \pm 0.45}$
Imclef07d	76.47 ± 0.35	$\mathbf{81.06 \pm 0.68}$

Neuro-Symbolic Learning Settings

Learn

1. neural network given symbols and constraints and data

2. neural network and constraints given symbols and data

3. neural network and constraints and symbols given data

Everyone is working on 1. Ongoing work on 2.

Neuro-Symbolic Joint Training

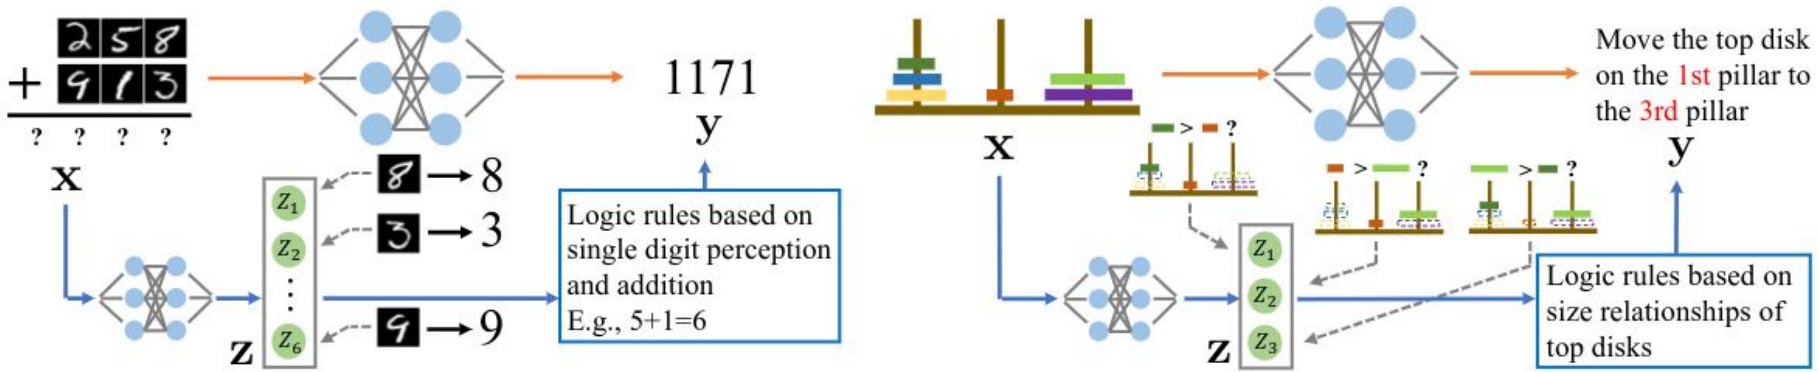

Learn invariant features using neural networks. Learn logic to tie it all together.

Neuro-Symbolic Joint Training

Model	Multi-digit addition [test seq length + train/test img]						Tower of Hanoi		
	5 w/ test	10 w/ test	20 w/ test	5 w/ train	10 w/ train	20 w/ train	Task #1	Task #2	Task #3
DeepProbLog†	88.30	77.46	timeout	94.92	89.74	timeout	89.28	97.96	89.33
LSTM	81.40	56.97	39.05	88.92	77.40	63.23	78.26	**98.32**	74.36
DNC	81.49	59.64	33.83	81.88	59.96	37.85	76.20	97.87	73.87
NToC(ours)	**89.82**	**77.97**	**63.55**	**89.97**	**86.07**	**71.96**	**85.16**	97.94	**85.49**

Learn invariant features using neural networks. Learn logic to tie it all together.

Monotonicity Invariants for Neural Networks

Predict Loan Amount

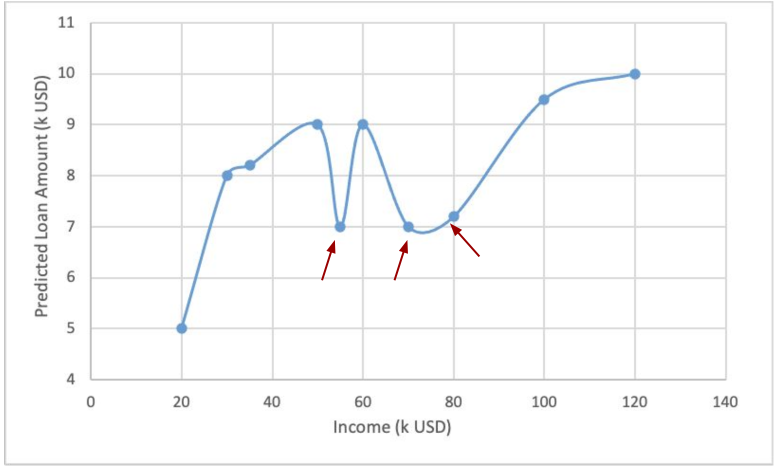

Neural Network Model: Increasing income can decrease the approved loan amount

Monotonicity (Prior Knowledge):
Increasing income should increase the approved loan amount

Counterexamples

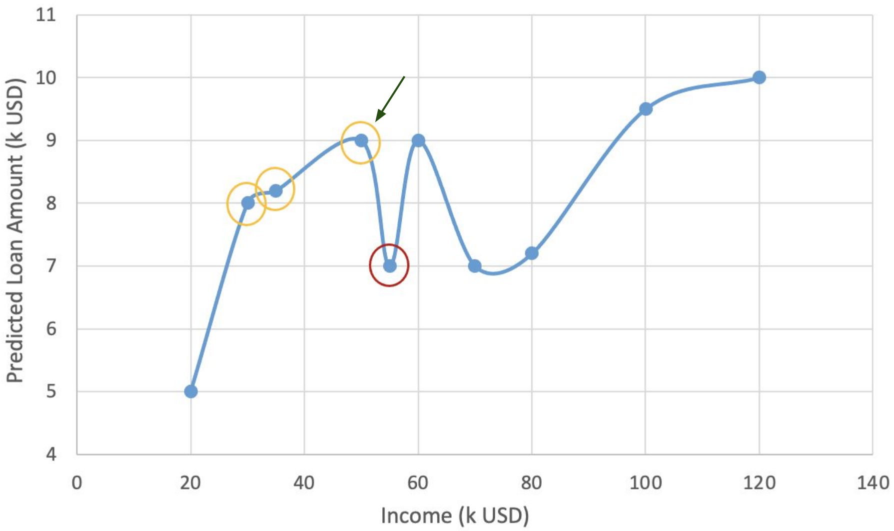

$$\exists x, y \; x \leq y \implies f(x) > f(y)$$

Computed using SMT(LRA) logical reasoning solver

Maximal counterexamples (largest violation) using OMT

Aishwarya Sivaraman, Golnoosh Farnadi, Todd Millstein and Guy Van den Broeck. Counterexample-Guided Learning of Monotonic Neural Networks, *NeurIPS*, 2020.

Counterexample-Guided Predictions

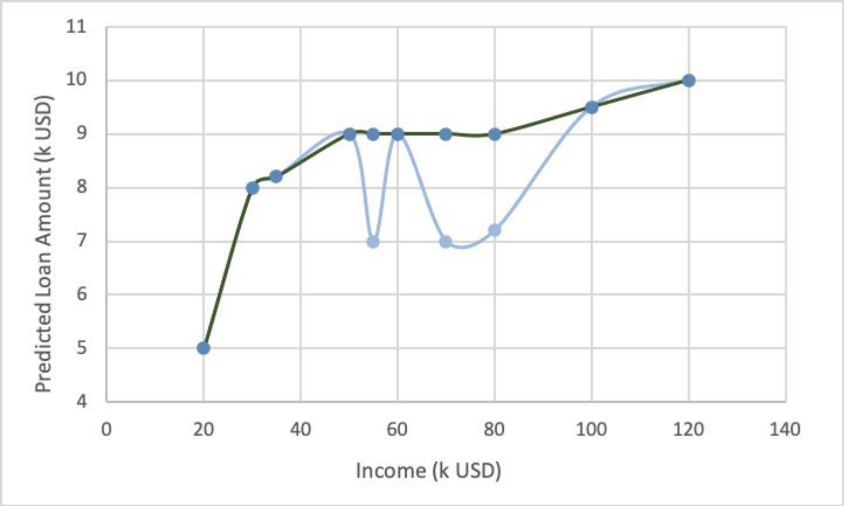

Monotonic Envelope:

- Replace each prediction by its maximal counterexample
- Envelope construction is online (during prediction)
- Guarantees monotonic predictions for any ReLU neural net

- Works for high-dimensional input
- Works for multiple monotonic features

Aishwarya Sivaraman, Golnoosh Farnadi, Todd Millstein and Guy Van den Broeck. Counterexample-Guided Learning of Monotonic Neural Networks, *NeurIPS*, 2020.

Monotonic Envelope: Performance

Dataset	Feature	NN_b	Envelope
Auto-MPG	Weight	9.33±3.22	**9.19±3.41**
	Displ.	9.33±3.22	9.63±2.61
	W,D	9.33±3.22	9.63±2.61
	W,D,HP	9.33±3.22	9.63±2.61
Boston	Rooms	14.37±2.4	**14.19±2.28**
	Crime	14.37±2.4	**14.02±2.17**

Dataset	Feature	NN_b	Envelope
Heart	Trestbps	0.85±0.04	0.85±0.04
	Chol.	0.85±0.04	0.85±0.05
	T,C	0.85±0.04	0.85±0.05
Adult	Cap. Gain	0.84	0.84
	Hours	0.84	0.84

Guaranteed monotonicity at little to no cost

Aishwarya Sivaraman, Golnoosh Farnadi, Todd Millstein and Guy Van den Broeck. Counterexample-Guided Learning of Monotonic Neural Networks, *NeurIPS*, 2020.

Counterexample-Guided Learning

How to use monotonicity to improve model quality?
“Monotonicity as inductive bias”

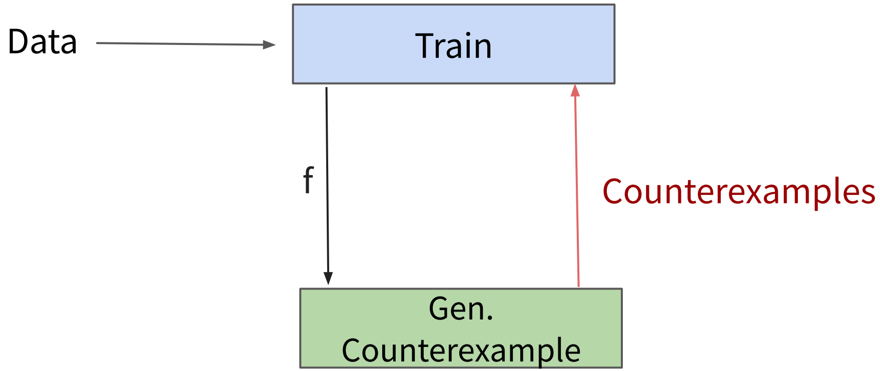

T Epochs

Aishwarya Sivaraman, Golnoosh Farnadi, Todd Millstein and Guy Van den Broeck. Counterexample-Guided Learning of Monotonic Neural Networks, *NeurIPS*, 2020.

Counterexample-Guided Learning: Performance

Dataset	Feature	NN_b	CGL
Auto-MPG	Weight	9.33±3.22	**9.04±2.76**
	Displ.	9.33±3.22	**9.08±2.87**
	W,D	9.33±3.22	**8.86±2.67**
	W,D,HP	9.33±3.22	**8.63±2.21**
Boston	Rooms	14.37±2.4	**12.24±2.87**
	Crime	14.37±2.4	**11.66±2.89**

Dataset	Feature	NN_b	CGL
Heart	Trestbps	0.85±0.04	**0.86±0.02**
	Chol.	**0.85±0.04**	**0.85±0.05**
	T,C	0.85±0.04	**0.86±0.06**
Adult	Cap. Gain	**0.84**	**0.84**
	Hours	**0.84**	**0.84**

Monotonicity is a *great* inductive bias for learning

Aishwarya Sivaraman, Golnoosh Farnadi, Todd Millstein and Guy Van den Broeck. Counterexample-Guided Learning of Monotonic Neural Networks, *NeurIPS*, 2020.

COMET:
Counterexample-Guided Monotonicity Enforced Training

Table 4: Monotonicity is an effective inductive bias. COMET outperforms Min-Max networks on all datasets. COMET outperforms DLN in regression datasets and achieves similar results in classification datasets.

Dataset	Features	Min-Max	DLN	COMET
Auto-MPG	Weight	9.91±1.20	16.77±2.57	**8.92±2.93**
	Displ.	11.78±2.20	16.67±2.25	**9.11±2.25**
	W,D	11.60±0.54	16.56±2.27	**8.89±2.29**
	W,D,HP	10.14±1.54	13.34±2.42	**8.81±1.81**
Boston	Rooms	30.88±13.78	15.93±1.40	**11.54±2.55**
	Crime	25.89±2.47	12.06±1.44	**11.07±2.99**

Dataset	Features	Min-Max	DLN	COMET
Heart	Trestbps	0.75±0.04	0.85±0.02	**0.86±0.03**
	Chol.	0.75±0.04	0.85±0.04	**0.87±0.03**
	T,C	0.75±0.04	**0.86±0.02**	**0.86±0.03**
Adult	Cap. Gain	0.77	**0.84**	**0.84**
	Hours	0.73	**0.85**	0.84

COMET = Provable Guarantees + SotA Results

Aishwarya Sivaraman, Golnoosh Farnadi, Todd Millstein and Guy Van den Broeck. Counterexample-Guided Learning of Monotonic Neural Networks, *NeurIPS*, 2020.

The AI Dilemma

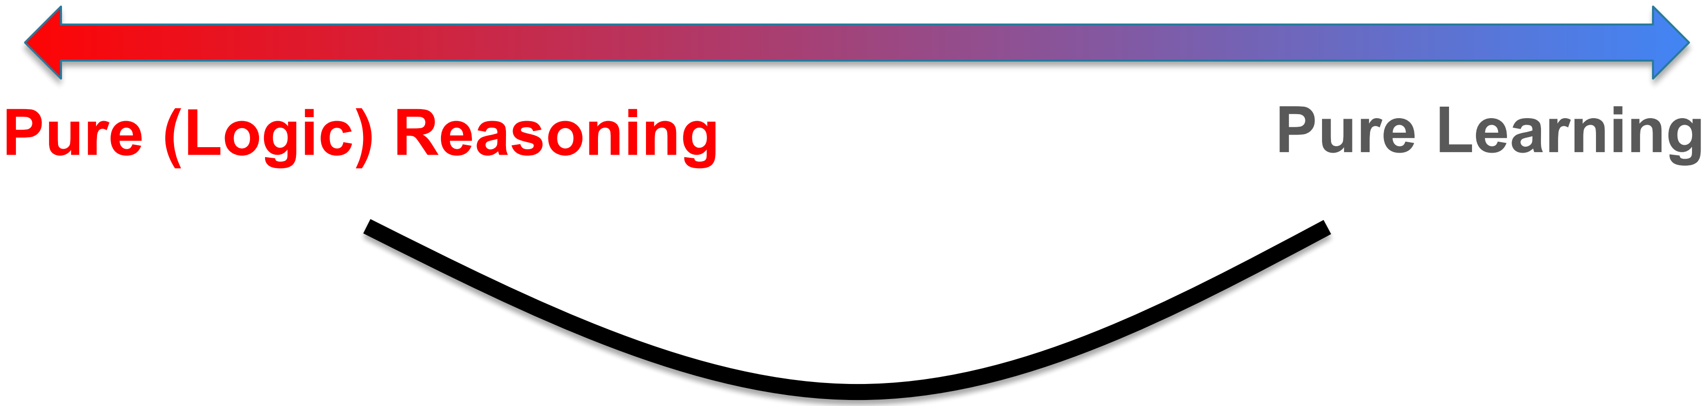

- Knowledge is (hidden) everywhere in ML
- A little bit of reasoning goes a long way!

Deep learning with structured output constraints
Learning monotonic neural networks

Thanks

This was the work of many wonderful students/postdoc/collaborators!

References: http://starai.cs.ucla.edu/publications/