


Logic Tensor Networks



Samy Badreddine Artur d'Avila Garcez Luciano Serafini Michael Spranger



Sony AI




 **Sony CSL**



CITY
UNIVERSITY OF LONDON
EST 1894



 **FONDAZIONE
BRUNO KESSLER**



BDD-OIA Dataset, Yiran Xu et al, CVPR'2020


X



Neural
Net
→



Action-inducing objects



Actions

Explain


$\mathbf{x}$



Neural
Net
→



Action-inducing objects



Actions

$p_{\theta}(\mathbf{c}|\mathbf{x})$

red traffic light



0.80

pedestrian



0.95

⋮

stop



0.30


$\mathbf{x}$



Neural
Net
→



Action-inducing objects



Actions

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$




$\mathbf{x}$



Neural
Net
→



Action-inducing objects



Actions

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$

red traffic light




pedestrian



⋮

stop



Option 1: fix using data


$\mathbf{x}$



Neural
Net
→



Action-inducing objects



Actions

$p_{\theta}(\mathbf{c}|\mathbf{x})$



Option 1: fix using data

Option 2: fix using *knowledge*

“If there is a red traffic light or a pedestrian, one should stop”

$\mathbf{x}$



Neural
Net
→

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$



Logical Knowledge

“If there is a red traffic light or a pedestrian, one should stop”

$\mathbf{x}$



Neural
Net
→

$p_{\theta}(\mathbf{c}|\mathbf{x})$



Logical Knowledge

$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$

X




Neural
Net
→

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$



Logical Layer?



Logical Knowledge

$$\phi := (\text{RedTL} \vee \text{Ped}) \Rightarrow \text{Stop}$$

X




Neural
Net
→

$p_{\theta}(\mathbf{c}|\mathbf{x})$ *continuous-discrete*
disconnect



Probs [0,1]

Logical Layer?



Logical Knowledge

$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

X




Neural
Net
→

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$



Probs [0,1]



Logical Knowledge

$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

X



Neural
Net
→

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$



Probs [0,1]

Fuzzy Relaxation

Godel t-conorm

$$\max(a, b)$$

Godel s-implication

$$\max(1 - a, c)$$

0.30

“Truth degree”,
higher is better

$$\mathcal{G}(\phi, \theta)$$

Logical Knowledge


$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

X



Neural
Net
→

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$



Probs [0,1]

$$\nabla_{\theta} \mathcal{G}(\phi, \theta)$$

Differentiable

Logical Knowledge


$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

X



Neural
Net
→

$$p_{\theta}(\mathbf{c}|\mathbf{x})$$



Fuzzy Relaxation

Godel t-conorm

$$\max(a, b)$$

Godel s-implication

$$\max(1 - a, c)$$

0.95 ↑

“Truth degree”,
higher is better



$$\nabla_{\theta} \mathcal{G}(\phi, \theta)$$

Differentiable

Logical Knowledge


$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

Logic Tensor Networks (LTN) relax logical knowledge into neural network layers.




Use Case 1: Learning

Logic Tensor Networks (LTN) relax logical knowledge into neural network layers.



Why fuzzy relaxation?



Use Case 1: Learning

Fuzzy logic operators with symbolic broadcasting.

$$\begin{aligned} &\forall s \text{ } \text{stop}(s) \wedge \text{go}(s) \\ &\forall s \text{ } (\text{redTL}(s) \rightarrow \text{greenTL}(s)) \\ &\forall s \text{ } ((\text{redTL}(s) \vee \text{ped}(s)) \rightarrow \text{stop}(s)) \\ &\forall s \text{ } (\text{stopSign}(s) \rightarrow \text{stop}(s)) \end{aligned}$$

Family	$a \wedge b$	$a \vee b$	Implication $a \rightarrow b$	
			Residuated	S-implication
Gödel	$\min(a, b)$	$\max(a, b)$	$\begin{cases} 1 & a \leq b \\ b & \text{else} \end{cases}$	$\max(1 - a, b)$
Product	$a \cdot b$	$a + b - a \cdot b$	$\text{clip}_{[0,1]}(b/a)$	$1 - a + a \cdot b$
Łukasiewicz	$\text{clip}_{[0,1]}(a + b - 1)$	$\text{clip}_{[0,1]}(a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$
...				

Fuzzy logic operators with symbolic broadcasting.

$\forall s \text{ (stop}(s) \wedge \text{go}(s))$
 $\forall s \text{ (redTL}(s) \multimap \text{greenTL}(s))$
 $\forall s \text{ ((redTL}(s) \vee \text{ped}(s)) \multimap \text{stop}(s))$
 $\forall s \text{ (stopSign}(s) \multimap \text{stop}(s))$

Family	$a \wedge b$	$a \vee b$	Implication $a \rightarrow b$	
			Residuated	S-implication
Gödel	$\min(a, b)$	$\max(a, b)$	$\begin{cases} 1 & a \leq b \\ b & \text{else} \end{cases}$	$\max(1 - a, b)$
Product	$a \cdot b$	$a + b - a \cdot b$	$\text{clip}_{[0,1]}(b/a)$	$1 - a + a \cdot b$
Łukasiewicz	$\text{clip}_{[0,1]}(a + b - 1)$	$\text{clip}_{[0,1]}(a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$
...				

- 1 Agrees with boolean logic on deterministic inputs.

Fuzzy logic operators with symbolic broadcasting.

$\forall s \text{ (stop}(s) \wedge \text{go}(s))$
 $\forall s \text{ (redTL}(s) \multimap \text{greenTL}(s))$
 $\forall s \text{ ((redTL}(s) \vee \text{ped}(s)) \multimap \text{stop}(s))$
 $\forall s \text{ (stopSign}(s) \multimap \text{stop}(s))$

Family	$a \wedge b$	$a \vee b$	Implication $a \rightarrow b$	
			Residuated	S-implication
Gödel	$\min(a, b)$	$\max(a, b)$	$\begin{cases} 1 & a \leq b \\ b & \text{else} \end{cases}$	$\max(1 - a, b)$
Product	$a \cdot b$	$a + b - a \cdot b$	$\text{clip}_{[0,1]}(b/a)$	$1 - a + a \cdot b$
Łukasiewicz	$\text{clip}_{[0,1]}(a + b - 1)$	$\text{clip}_{[0,1]}(a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$
...				

- 1 Agrees with boolean logic on deterministic inputs.
- 2 Useful identities (commutativity, associativity, De Morgan’s laws, ...)



Fuzzy logic operators with symbolic broadcasting.

$$\begin{aligned} &\forall s \text{ (stop}(s) \wedge \text{go}(s)) \\ &\forall s \text{ (redTL}(s) \multimap \text{greenTL}(s)) \\ &\forall s \text{ ((redTL}(s) \vee \text{ped}(s)) \multimap \text{stop}(s)) \\ &\forall s \text{ (stopSign}(s) \multimap \text{stop}(s)) \end{aligned}$$

Family	$a \wedge b$	$a \vee b$	Implication $a \rightarrow b$	
			Residuated	S-implication
Gödel	$\min(a, b)$	$\max(a, b)$	$\begin{cases} 1 & a \leq b \\ b & \text{else} \end{cases}$	$\max(1 - a, b)$
Product	$a \cdot b$	$a + b - a \cdot b$	$\text{clip}_{[0,1]}(b/a)$	$1 - a + a \cdot b$
Łukasiewicz	$\text{clip}_{[0,1]}(a + b - 1)$	$\text{clip}_{[0,1]}(a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$	$\text{clip}_{[0,1]}(1 - a + b)$
...				



- 1 Agrees with boolean logic on deterministic inputs.
- 2 Useful identities (commutativity, associativity, De Morgan’s laws, ...)
- 3 Scalable $\longrightarrow$ For a formula ϕ with N boolean variables, exact probabilistic reasoning (an alternative) is $O(2^N)$.
The relaxation is $O(|\phi|)$.

Logic Tensor Networks (LTN) relax logical knowledge into neural network layers.



Use Case 1: Learning

Logic Tensor Networks (LTN) relax logical knowledge into neural network layers.



Use Case 1: Learning

Use Case 2: Querying

$\mathbf{x}$



$p_{\theta}(\mathbf{c}|\mathbf{x})$

LTN

NN



“Stop for red TL or pedestrian”
 $(\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$ 0.95

Training

x



$p_{\theta}(\mathbf{c}|\mathbf{x})$

LTN

NN



red traffic light



0.80

pedestrian



0.95

⋮

stop



0.95

“Stop for red TL or pedestrian”

$(\text{RedTL} \vee \text{Ped}) \implies \text{Stop} \rightarrow 0.95$

Training

Derive a new “Hazard” concept

$\text{hazard}(s) := \text{ped}(s) \vee \text{redTL}(s) \vee \text{stopSign}(s)$

$\rightarrow 0.95$

Querying

$\mathbf{x}$



$p_{\theta}(\mathbf{c}|\mathbf{x})$

LTN

NN



red traffic light



0.80

pedestrian



0.95

⋮

stop



0.95

“Stop for red TL or pedestrian”

$(\text{RedTL} \vee \text{Ped}) \implies \text{Stop} \rightarrow 0.95$



Training

Derive a new “Hazard” concept



$\text{hazard}(s) := \text{ped}(s) \vee \text{redTL}(s) \vee \text{stopSign}(s)$

$\rightarrow 0.95$


Querying



⋮



⋮



“Do we ever run a red light?”

$\exists s (\text{redTL}(s) \wedge \text{go}(s))$

$\mathbf{x}$



$p_{\theta}(\mathbf{c}|\mathbf{x})$

LTN

NN



“Stop for red TL or pedestrian”

$(\text{RedTL} \vee \text{Ped}) \implies \text{Stop} \rightarrow 0.95$


Training

Derive a new “Hazard” concept

$\text{hazard}(s) := \text{ped}(s) \vee \text{redTL}(s) \vee \text{stopSign}(s)$


Querying

$\rightarrow 0.95$



“Do we ever run a red light?”

$\exists s (\text{redTL}(s) \wedge \text{go}(s))$



Retrieve the riskiest scenes

$\arg \max_s \mathcal{G}(\text{hazard}(s) \wedge \neg \text{stop}(s))$

$\mathbf{x}$



$p_{\theta}(\mathbf{c}|\mathbf{x})$

LTN

NN



red traffic light



0.80

pedestrian



0.95

⋮

stop



0.95

“Stop for red TL or pedestrian”

$(\text{RedTL} \vee \text{Ped}) \implies \text{Stop} \rightarrow 0.95$



Training

Derive a new “Hazard” concept


$\text{hazard}(s) := \text{ped}(s) \vee \text{redTL}(s) \vee \text{stopSign}(s)$

Querying

$\rightarrow 0.95$




⋮



“Do we ever run a red light?”

$\exists s (\text{redTL}(s) \wedge \text{go}(s))$

*Sometimes
useful as biased
logits/rankings*



⋮



Retrieve the riskiest scenes

$\arg \max_s \mathcal{G}(\text{hazard}(s) \wedge \neg \text{stop}(s))$

2016

2018

2021

2024

2026

Seminal work
Serafini and Garcez

FOL syntax wrapper
Spranger

PyTorch version
Carraro et al

Structured extension
Rinaldi and Serafini

AIJ Paper: Formalization and numerically stable version of LTN.
Reference implementation with recipes for many tasks.

Binary and multiclass classification

(semi-)supervised




weakly supervised

$$3 + 9 = 12$$

$$2 + 8 = 10$$


Clustering



Linear regression



Relational learning
and reasoning




Research trends and Applications

Optimization

Which fuzzy operators are best for gradient-based learning?

Family	$a \wedge b$	$a \vee b$	Implication $a \rightarrow b$	
			Residuated	S-implication
Gödel	$\min(a, b)$	$\max(a, b)$	$\begin{cases} 1 & a \leq b \\ b & \text{else} \end{cases}$	$\max(1 - a, b)$
Product	$a \cdot b$	$a + b - a \cdot b$	$\text{clip}(b/a)$	$1 - a + a \cdot b$
Lukasiewicz	$\text{clip}(a + b - 1)$	$\text{clip}(a + b)$	$\text{clip}(1 - a + b)$	$\text{clip}(1 - a + b)$
		...		




- (1) Emile van Krieken et al, *Analyzing differentiable fuzzy logic operators*, AIJ 2022
- (2) Samy Badreddine et al, *logLTN: differentiable fuzzy logic in the logarithm space*, 2023
- (3) Natalia Ślusarz et al, *Logic of differentiable logics*, LPAR 2024
- (4) Jaron Maene et al, *On the hardness of probabilistic NeSy learning*, ICML 2024

Research trends and Applications

Relation to probabilistic semantics

Reinterpret fuzzy relaxations as an estimation of probabilistic semantics.




- (1) Kareem Ahmed et al, *Semantic strengthening of NeSy learning*, AISTATS 2023
- (2) Jaron Maene et al, *On the hardness of probabilistic NeSy learning*, ICML 2024
- (3) Samy Badreddine et al, *Where to approximate in NeSy inference*, CoLorAI 2026

Research trends and Applications

Learnability

“Reasoning shortcuts”: Identifiability issues with logical losses.



$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

Many logical solutions


i.RedTL ✓ Ped ✓ Stop ✓

- (1) Emanuele Marconato et al, *Analysis and mitigation of reasoning shortcuts*, NeurIPS 2023
- (2) Emanuele Marconato et al, *BEARS: Make NeSy models aware of their reasoning shortcuts*, UAI 2024
- (3) Samuele Bortolotti et al, *NeSy benchmark suite for concept quality and reasoning shortcuts*, NeurIPS 2025
- (4) Emanuele Marconato et al, *Symbol Grounding in NeSy AI*, JAIR 2026

Research trends and Applications

Learnability

“Reasoning shortcuts”: Identifiability issues with logical losses.



$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

Many logical solutions


- i. RedTL ✓ Ped ✓ Stop ✓
- ii. RedTL ✗ Ped ✓ Stop ✓
- iii. RedTL ✓ Ped ✗ Stop ✓

- (1) Emanuele Marconato et al, *Analysis and mitigation of reasoning shortcuts*, NeurIPS 2023
- (2) Emanuele Marconato et al, *BEARS: Make NeSy models aware of their reasoning shortcuts*, UAI 2024
- (3) Samuele Bortolotti et al, *NeSy benchmark suite for concept quality and reasoning shortcuts*, NeurIPS 2025
- (4) Emanuele Marconato et al, *Symbol Grounding in NeSy AI*, JAIR 2026

Research trends and Applications

Learnability

“Reasoning shortcuts”: Identifiability issues with logical losses.



$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

Many logical solutions


- i. RedTL ✓ Ped ✓ Stop ✓
- ii. RedTL ✗ Ped ✓ Stop ✓
- iii. RedTL ✓ Ped ✗ Stop ✓
- iv. RedTL ✗ Ped ✗ Stop ✗

- (1) Emanuele Marconato et al, *Analysis and mitigation of reasoning shortcuts*, NeurIPS 2023
- (2) Emanuele Marconato et al, *BEARS: Make NeSy models aware of their reasoning shortcuts*, UAI 2024
- (3) Samuele Bortolotti et al, *NeSy benchmark suite for concept quality and reasoning shortcuts*, NeurIPS 2025
- (4) Emanuele Marconato et al, *Symbol Grounding in NeSy AI*, JAIR 2026

Research trends and Applications

Learnability

“Reasoning shortcuts”: Identifiability issues with logical losses.



$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

Many logical solutions


- i. RedTL ✓ Ped ✓ Stop ✓
- ii. RedTL ✗ Ped ✓ Stop ✓
- iii. RedTL ✓ Ped ✗ Stop ✓
- iv. RedTL ✗ Ped ✗ Stop ✗

- (1) Emanuele Marconato et al, *Analysis and mitigation of reasoning shortcuts*, NeurIPS 2023
- (2) Emanuele Marconato et al, *BEARS: Make NeSy models aware of their reasoning shortcuts*, UAI 2024
- (3) Samuele Bortolotti et al, *NeSy benchmark suite for concept quality and reasoning shortcuts*, NeurIPS 2025
- (4) Emanuele Marconato et al, *Symbol Grounding in NeSy AI*, JAIR 2026

Research trends and Applications

Learnability

“Reasoning shortcuts”: Identifiability issues with logical losses.



$$\phi := (\text{RedTL} \vee \text{Ped}) \implies \text{Stop}$$

Many logical solutions

- i. RedTL ✓ Ped ✓ Stop ✓
- ii. RedTL ✗ Ped ✓ Stop ✓
- iii. RedTL ✓ Ped ✗ Stop ✓
- iv. RedTL ✗ Ped ✗ Stop ✗


→ Can we guarantee that we reach the groundtruth one?

- (1) Emanuele Marconato et al, *Analysis and mitigation of reasoning shortcuts*, NeurIPS 2023
- (2) Emanuele Marconato et al, *BEARS: Make NeSy models aware of their reasoning shortcuts*, UAI 2024
- (3) Samuele Bortolotti et al, *NeSy benchmark suite for concept quality and reasoning shortcuts*, NeurIPS 2025
- (4) Emanuele Marconato et al, *Symbol Grounding in NeSy AI*, JAIR 2026



Research trends and Applications

Reasoning with Knowledge Graphs

[1] Injecting prior knowledge in Recommender Systems




[2,3] Querying knowledge graphs



“Which drugs interact with proteins associated with diseases t_1 or t_2 ?”

$?D : \exists P . \text{interacts}(D, P) \wedge [\text{assoc}(P, t_1) \vee \text{assoc}(P, t_2)]$



$\text{likes_genre}(x, g) \wedge \text{has_genre}(z, g) \implies \text{likes}(x, z)$

(1) Tommaso Carraro, *Overcoming recommendation limitations with NeSy integration*, RecSys 2023

(2) Erik Arakelyan et al, *Complex query answering with neural link predictors*, ICLR 2021


(3) Ouns El Harzli et al, *What’s wrong with gradient-based query answering*, NeSy 2023

Research trends and Applications

Semantic Image Interpretation

Impose ontological constraints on object types and relations.

$$\text{eat}(x, y) \rightarrow \text{person}(x) \wedge \text{food}(y)$$
$$\text{food}(x) \leftrightarrow \text{pizza}(x) \vee \text{ice_cream}(x) \vee \dots$$



- (1) Ivan Donadello et al, *LTN for semantic image interpretation*, IJCAI 2017
- (2) Francesco Manigrasso et al, *Boosting zero-shot learning through NeSy integration*, Pattern Recognition 2026
- (3) Luca Bergamin et al, *Integrating background knowledge in medical semantic segmentation with LTN*, IJCNN 2025
- (4) Stefano Colamonaco et al, *Weakly supervised segmentation as semantic-based regularization*, 2026

Research trends and Applications

Analyzing and Constraining Dynamic Systems

Query dynamic properties expressed in temporal logic and use them for learning.

$$\text{Car}(x) \wedge \text{Ped}(y) \wedge \text{Crossing}(y) \wedge \text{Close}(x, y) \rightarrow \text{X Stop}(x)$$
$$\text{car}(x) \wedge \text{run}(x) \wedge \text{X fstop}(x) \rightarrow \text{XE run}(x)$$



Temporal extensions of LTN


- (1) Elena Umili et al, *Grounding LTLf specifications in image sequences*, KR 2023
- (2) Fabrizio De Santis et al, *NeSy learning for predictive process monitoring via two-stage LTN with rule pruning*, PAKDD 2026
- (3) Devashish Gaikwad et al, *NeSy Process Anomaly Detection*, CAiSE 2026
- (4) Luca Boscarato et al, *First-order Temporal Logic Tensor Networks*, 2026

Research trends and Applications

Interpretable predictions and NeSy probing with LLMs

[1] uses LLMs to write logical rules, and Logic Tensor Networks to compose them and make an interpretable decisions.

[2] uses LTN as a probe to analyze deductions in LLMs.



(1) Genan Dai et al, *LLM enhanced logic tensor networks for stance detection*, Neural Networks 2025

(2) Francesco Manigrasso et al, *Probing LLMs for logical reasoning*, NeSy 2024


Research trends and Applications

World Models

Regularize latent world model dynamics by translating geometric transformation rules into logical formulas.





$$\begin{aligned} (z_t^F = \hat{z}_{t+1}^R) \wedge (\text{RotNet}_\theta(z_t^U, \rightarrow) = \hat{z}_{t+1}^U) \\ (z_t^R = \hat{z}_{t+1}^F) \wedge (\text{RotNet}_\theta(z_t^U, \leftarrow) = \hat{z}_{t+1}^U) \end{aligned}$$



(1) Ajay Ramachandra, Gesina Schwalbe, Zoe Falomir, *Injecting qualitative spatial reasoning into world models via LTNs*,

At IJCAI 2026 !

Logic Tensor Networks relax **logical knowledge** in
loss functions and queries for *learning* and *reasoning*.



LTN repo