

Logic Tensor Networks

Acile Chahboun

Seminar Formal Methods and Automated Reasoning



Introduction

- ▶ Logic Tensor Networks, a neurosymbolic framework combining neural networks with first-order logic
- ▶ Logical Knowledge and data are learned together in a single system
- ▶ Built around a fully differentiable logical language called Real Logic

Reference: Badreddine, S., d'Avila Garcez, A., Serafini, L., Spranger, M. (2021). Logic Tensor Networks. *arXiv:2012.13635*.

Key Terminology

- ▶ **Neurosymbolic AI:** An AI approach that combines neural networks for learning with symbolic logic for reasoning.
- ▶ **Neural Network:** A model that learns patterns from data by adjusting numerical parameters.
- ▶ **Symbolic AI:** An AI approach that represents knowledge explicitly using symbols, rules, and logical relations.
- ▶ **First-Order Logic:** A formal logic system built from objects, variables, predicates, functions, and quantifiers.
- ▶ **Logic Tensor Network:** A framework that represents logical knowledge in a numerical, differentiable form so it can be used with neural networks.

Key Terminology

- ▶ **Real Logic:** The logic used in LTNs. Unlike classical logic, statements can be partly true, with truth values between 0 and 1.
- ▶ **Truth Value:** A number measuring how true a statement is, e.g. $\text{Cat}(\text{image}) = 0.92$.
- ▶ **Grounding:** The process of assigning real numerical meaning to logical symbols. An image becomes a tensor, and $\text{Cat}(x)$ becomes a neural classifier.
- ▶ **Tensor:** A numerical object used by neural networks. Scalars, vectors, matrices, and higher-dimensional arrays are all tensors.
- ▶ **Constant:** A symbol for a specific object, e.g. `Alice`, `Paris`, `image1`.

Key Terminology

- ▶ **Variable:** A placeholder that can range over many objects, e.g. x in $\text{Cat}(x)$.
- ▶ **Domain:** The type or category of objects a symbol belongs to, e.g. Person , City , Image .
- ▶ **Function:** A mapping from one or more objects to another object, e.g. $\text{fatherOf}(x)$.
- ▶ **Predicate:** A statement about objects that returns a truth value, e.g. $\text{Friend}(\text{Alice}, \text{Bob})$ or $\text{Digit}(\text{image}, 8)$.
- ▶ **Formula:** A logical expression built from predicates, variables, connectives, and quantifiers, e.g. $\text{Cat}(x) \rightarrow \text{Animal}(x)$.

Key Terminology

- ▶ **Axiom:** A formula treated as background knowledge that the model should satisfy, e.g. “All cats are animals.”
- ▶ **Connective:** A logical operator such as AND, OR, NOT, or IMPLIES.
- ▶ **Quantifier:** A logical operator that talks about many objects. $\forall x$ means “for all x ” and $\exists x$ means “there exists an x .”
- ▶ **Aggregation:** The process of combining many truth values into one truth value, often used for quantifiers.
- ▶ **Satisfiability:** A score measuring how well the model satisfies all given facts and rules.

Key Terminology

- ▶ **Learning:** Adjusting the model's parameters to maximize satisfiability.
- ▶ **Querying:** Asking the trained model how true a statement is, e.g. $\text{Digit}(\text{image1}, 8) = 0.87$.
- ▶ **Reasoning:** Using known facts and rules to determine whether another statement follows.
- ▶ **Differentiability:** A mathematical property that allows the system to be trained using gradient-based learning.
- ▶ **Gradient:** A signal that tells the model how to change its parameters during training.

Key Terminology

- ▶ **Fuzzy Logic:** A logic system where statements can be partly true rather than only true or false.
- ▶ **Knowledge Base:** A collection of facts and rules expressed as logical formulas that the model must satisfy.
- ▶ **Semantic Loss:** A training loss derived from logical constraints, penalizing the model when formulas in the knowledge base are not satisfied.
- ▶ **t-norm:** A function used to compute the fuzzy AND of two truth values, e.g. product t-norm: $a \cdot b$.
- ▶ **t-conorm:** A function used to compute the fuzzy OR of two truth values, e.g. $a + b - a \cdot b$.

Motivation

- ▶ AI systems need to do more than recognize patterns, they should be able to learn from data, reason about what they have learned, and use that knowledge to make decisions.
- ▶ Modern deep learning is strong at learning from raw data using sub-symbolic representations such as vectors and tensors.
- ▶ However, reasoning often requires a higher-level language that can express rules, relationships, and constraints clearly.

The Core Problem

Learning / Neural

Learns patterns from data

Uses vectors and tensors

Works well with noisy real-world inputs

Statistical generalization

Bottom-up: data \rightarrow patterns

Approximate and probabilistic

Knowledge is implicit

What is likely?

Reasoning / Symbolic

Derives conclusions from facts and rules

Uses symbols and logical formulas

Works well with explicit structured knowledge

Logical inference

Top-down: rules \rightarrow conclusions

Precise and rule-based

Knowledge is explicit and human-readable

What must follow?

The Core Problem

Neural systems struggle with:

Explicit logical rules

Long reasoning chains

Guarantees and truthfulness

Interpretability

Symbolic systems struggle with:

Learning directly from raw data

Perception tasks such as images or signals

Noise, uncertainty, and exceptions

Scalability to large real-world datasets

Core Issue

Modern AI needs both sides: the ability to learn from data and the ability to reason with explicit knowledge. LTN addresses this by combining neural learning with differentiable first-order logic.


NLM Recap

- ▶ **Neural Logic Machines (NLMs)** are neural models designed to perform logic-like reasoning over objects and their relationships.
- ▶ Instead of treating data as isolated examples, they represent objects, predicates, and relations using tensors.
- ▶ This allows the model to reason about structured information such as `friend(x,y)`, `parent(x,y)`, or `connected(x,y)`.
- ▶ NLMs can learn relational patterns and generalize across different numbers of objects, making them useful for tasks involving graphs, rules, and object interactions.

NLM Recap

- ▶ The logic in NLMs is learned **implicitly** from data. The model may discover rule-like patterns, but those rules are not always written explicitly in a human-readable form.
- ▶ **This is where LTNs differ:** LTNs allow logical knowledge to be stated directly as formulas, grounded into tensors, and optimized through differentiable Real Logic.

Logic Tensor Networks Overview



Real Logic: Syntax

- ▶ **Real Logic uses a first-order logical language:**
Built from symbols that represent objects, variables, functions, and relations.
- ▶ **Domains** define types of objects, e.g. `Person`, `City`, `Image`, `Digit`.
- ▶ **Constants** name specific objects, e.g. `Alice`, `Paris`, `img1`.
- ▶ **Variables** range over objects in a domain, e.g. `x` may range over all images.
- ▶ **Functions** map objects to other objects, e.g. `fatherOf(x)` returns a person.
- ▶ **Predicates** express relations or properties, e.g. `Cat(x)`, `Friend(x,y)`, `Digit(img,8)`.

Real Logic: Syntax

- ▶ **Formulas** combine predicates using logical operators, e.g. $\text{Cat}(x) \rightarrow \text{Animal}(x)$ and $\text{Friend}(x,y) \wedge \text{Friend}(y,z)$.
- ▶ **Quantifiers** express rules over many objects, e.g. $\forall x \text{Cat}(x) \rightarrow \text{Animal}(x)$ and $\exists x \text{Digit}(x,8)$.

To summarize

Real Logic syntax defines what counts as a valid logical expression before those expressions are grounded into tensors and evaluated numerically.

Real Logic: Semantics

- ▶ **Semantics gives meaning to the syntax:** It explains how logical symbols and formulas are interpreted numerically.
- ▶ **Symbols are grounded into real-valued objects:** Constants, variables, functions, and predicates are mapped to tensors or tensor operations.
- ▶ **Predicates return truth values in $[0, 1]$:** e.g. $\text{Cat}(\text{img}_1) = 0.92$ means the model strongly believes img_1 is a cat.
- ▶ **Functions become real-valued mappings:** e.g. $\text{fatherOf}(x)$ is interpreted as a function over grounded objects.

Real Logic: Semantics

- ▶ **Logical connectives become fuzzy operators:** AND, OR, NOT, and IMPLIES are computed using differentiable fuzzy-logic operations.
- ▶ **Quantifiers become aggregators:** $\forall x$ and $\exists x$ combine many truth values into one score. e.g. $\forall x \text{ Cat}(x) \rightarrow \text{Animal}(x)$ is evaluated by checking the rule over many grounded examples and aggregating the results.

To summarize

Real Logic semantics turns logical formulas into differentiable numerical computations, so they can be optimized during neural-network training.

Grounding Intuition

- ▶ **Grounding connects logic to data:** Logical symbols are given numerical meaning so a neural model can process them.
- ▶ **Objects become tensors:** e.g. an image becomes a pixel tensor; a person or word may become an embedding vector.
- ▶ **Predicates become functions or neural networks:** e.g. $\text{Cat}(x)$ can be a classifier that returns how true it is that x is a cat.
- ▶ **Relations become truth-valued mappings:** e.g. $\text{Friend}(x,y)$ returns a score between 0 and 1 for each pair of objects.

Grounding Intuition

- ▶ **Rules become differentiable computations:** A formula such as $\text{Cat}(x) \rightarrow \text{Animal}(x)$ is evaluated numerically, not just symbolically.
- ▶ **Learning adjusts the grounding:** The model changes embeddings, predicates, and functions so the logical rules become more satisfied.

To summarize

Grounding is the bridge between symbolic logic and neural learning: it turns symbols and formulas into tensors, functions, and truth scores.

Semantics of Real Logic: Tensor Intuition

Classical first-order logic

A variable assignment gives each variable one value at a time:

$$x = a_i, \quad y = b_j \quad \Rightarrow \quad I(f(x, y)) = I(f)(a_i, b_j)$$

Real Logic / LTNs

Variables are grounded over whole batches:

$$\mathcal{G}(x) = A = [a_1, a_2, \dots] \quad \mathcal{G}(y) = B = [b_1, b_2, \dots]$$

$$\mathcal{G}(f(x, y))_{i,j} = \mathcal{G}(f)(A[i], B[j])$$

Semantics of Real Logic: Tensor Intuition

Tensor view

	b_1	b_2	b_3
a_1	$f(a_1, b_1)$	$f(a_1, b_2)$	$f(a_1, b_3)$
a_2	$f(a_2, b_1)$	$f(a_2, b_2)$	$f(a_2, b_3)$

Each entry: $x = A[i]$, $y = B[j]$.

Main idea

variables over domains \rightarrow all assignments \rightarrow tensor of truth values in $[0, 1]$

Each entry is a fuzzy truth value in $[0, 1]$.

Grounding Terms & Predicates

- ▶ **Grounding** maps logical symbols into numerical objects via a grounding function \mathcal{G} :

\mathcal{G} : logical symbols \rightarrow numerical objects

- ▶ **Constants become tensors:** e.g.

$$\mathcal{G}(\text{img}_1) = v_{\text{img}_1} \in \mathbb{R}^{28 \times 28 \times 1}$$

- ▶ **Variables become batches of tensors:** if

$$x = [\text{img}_1, \text{img}_2, \text{img}_3] \text{ then } \mathcal{G}(x) \in \mathbb{R}^{n \times H \times W \times C}$$

- ▶ **Functions become numerical mappings:**

$$\mathcal{G}(f(x)) = \mathcal{G}(f)(\mathcal{G}(x)), \quad \mathcal{G}(f) : \mathbb{R}^d \rightarrow \mathbb{R}^m$$

Grounding Terms & Predicates

- ▶ **Predicates become truth-valued functions:**

$$\mathcal{G}(\text{Cat}(x)) = \mathcal{G}(\text{Cat})(\mathcal{G}(x)) = \sigma(\text{MLP}(v_{\text{img}_1})) = 0.92$$

- ▶ **Predicates over batches return truth-value tensors:**

$$\mathcal{G}(\text{Cat}(x)) = [0.92, 0.13, 0.77], \quad \mathcal{G}(\text{Friend}(x, y)) \in [0, 1]^{n \times m}$$

- ▶ **Rules become differentiable computations:**

$$\mathcal{G}(\text{Cat}(x) \rightarrow \text{Animal}(x)) = [0.99, 1.00, 0.93]$$

$$\mathcal{G}(\forall x \text{ Cat}(x) \rightarrow \text{Animal}(x)) = \text{Agg}([0.99, 1.00, 0.93]) = 0.97$$

Connectives

Logical connectives combine statements into larger formulas. In Real Logic, truth values lie in $[0, 1]$, so connectives combine numbers, not just Boolean values.

Negation (NOT)

$$\neg P(x) = 1 - P(x) \quad \text{e.g. } 1 - 0.8 = 0.2$$

Conjunction (AND)

$$P(x) \wedge Q(x) = P(x) \cdot Q(x) \quad \text{e.g. } 0.8 \cdot 0.6 = 0.48$$

Disjunction (OR)

$$P(x) \vee Q(x) = P(x) + Q(x) - P(x) \cdot Q(x) \quad \text{e.g. } 0.8 + 0.6 - 0.48 = 0.92$$

Connectives

Implication (IF... THEN)

$$P(x) \rightarrow Q(x)$$

The rule is violated mainly when $P(x)$ is high but $Q(x)$ is low. e.g. $\text{Cat}(x) \rightarrow \text{Animal}(x)$ is satisfied when anything classified as a cat is also classified as an animal.

Why connectives matter

Connectives turn logical rules into differentiable numerical computations:

$$\text{Cat}(x) \wedge \text{Small}(x) \rightarrow \mathcal{G}(\text{Cat}(x)) \cdot \mathcal{G}(\text{Small}(x))$$

Logical formulas can then be evaluated, optimized, and learned through gradient-based training.

Quantifiers & Aggregation

Quantifiers evaluate a rule over many examples and return one final truth score.

Universal quantifier \forall

$\forall x \text{ Cat}(x) \rightarrow \text{Animal}(x)$ — for every x , if x is a cat, then x is an animal.

$$\text{Agg}_{\forall}([0.99, 0.88, 0.94, 0.30]) \rightarrow \text{one score}$$

Behaves like a **soft minimum** — a low score pulls the final result down.

Existential quantifier \exists

$\exists x \text{ Red}(x)$ — there exists at least one red object.

$$\text{Agg}_{\exists}([0.10, 0.20, 0.95, 0.15]) \rightarrow \text{one score}$$

Behaves like a **soft maximum** — a high score raises the final result.

Quantifiers & Aggregation

Aggregation: generalized mean

$$M_p(t_1, \dots, t_n) = \left(\frac{t_1^p + \dots + t_n^p}{n} \right)^{1/p}$$

- ▶ $p = 1 \rightarrow$ normal average
- ▶ large $p \rightarrow$ closer to maximum (\exists)
- ▶ soft-min form \rightarrow closer to minimum (\forall)

$$\text{Agg}_{\forall}([t_1, \dots, t_n]) = 1 - M_p(1 - t_1, \dots, 1 - t_n)$$

Main idea

\forall = soft “all” \exists = soft “at least one”, aggregation makes quantified rules differentiable and trainable.

Quantifier Subtlety

In classical logic, some formulas are logically equivalent. In LTNs, they may **not** be numerically equivalent because quantifiers use **soft aggregation**.

Classical logic equivalence

$$\forall x (\phi(x) \wedge \psi(x)) \equiv (\forall x \phi(x)) \wedge (\forall x \psi(x))$$

Both sides evaluate to **False** when any instance fails, the equivalence holds.

Quantifier Subtlety

In LTNs: soft conjunction $a \wedge b := a \cdot b$, universal aggregation $A_{\forall} =$ average.

Example values

x	$\phi(x)$	$\psi(x)$	$\phi(x) \cdot \psi(x)$
x_1	0.9	0.8	$0.9 \cdot 0.8 = 0.72$
x_2	0.4	0.7	$0.4 \cdot 0.7 = 0.28$

Left-hand side

$$\forall x (\phi(x) \wedge \psi(x)) = A_{\forall}([0.72, 0.28]) = \frac{0.72 + 0.28}{2} = 0.50$$

Right-hand side

$$(\forall x \phi(x)) \wedge (\forall x \psi(x)) = \frac{0.9 + 0.4}{2} \cdot \frac{0.8 + 0.7}{2} = 0.65 \cdot 0.75 = 0.4875$$

Quantifier Subtlety

Final comparison

Setting	Formula 1	Formula 2	Same?
Classical logic	False	False	Yes
LTN / Real Logic	0.50	0.4875	No

Main idea

In classical logic:

$$\forall x(\phi(x) \wedge \psi(x)) \equiv (\forall x \phi(x)) \wedge (\forall x \psi(x))$$

But in LTNs:

$$A_{\forall}(\phi(x) \wedge \psi(x)) \neq A_{\forall}(\phi(x)) \wedge A_{\forall}(\psi(x))$$

because aggregation does **not** generally commute with fuzzy connectives.

Quantifier Subtlety: Variable Groundings

Classical logic

In classical logic, renaming a bound variable does not change the meaning:

$$\forall x p(x) \equiv \forall y p(y)$$

Both are interpreted over the same domain.

Real Logic / LTNs

Each variable name is associated with its own grounding $\mathcal{G}(x)$, $\mathcal{G}(y)$. So the formulas behave like:

$$\forall x \in \mathcal{G}(x) p(x) \quad \forall y \in \mathcal{G}(y) p(y)$$

If $\mathcal{G}(x) \neq \mathcal{G}(y)$, the two formulas are evaluated over different tensors and can have different meanings.

Quantifier Subtlety: Binary Classification

Two variables, two groundings

$\mathcal{G}(x^+) =$ batch of positive points $\mathcal{G}(x^-) =$ batch of negative points

The two axioms:

$$\forall x^+ A(x^+) \quad \forall x^- \neg A(x^-)$$

are **not contradictory** — the quantifiers range over different grounded domains.

Binary classification view

Formula	Grounding	Meaning
$\forall x^+ A(x^+)$	positive examples	positives should satisfy A
$\forall x^- \neg A(x^-)$	negative examples	negatives should not satisfy A

Main idea

Variable names are tied to their groundings:

variable name \rightarrow grounded tensor \rightarrow aggregation range

Quantifiers determine **which tensor or domain is aggregated over**.

Diagonal Quantification

Diagonal quantification is used when two variables are already **paired** in the data.

Example pairing

Image	Correct label
img_1	cat
img_2	dog
img_3	bird

Normal quantification — all combinations

Evaluates $x \times y$: every image with every label $\rightarrow 3 \times 3 = 9$ comparisons. May include irrelevant pairs such as $CorrectLabel(img_1, dog)$.

Diagonal quantification — matched pairs only

Evaluates only $(img_1, cat), (img_2, dog), (img_3, bird) \rightarrow 3$ comparisons.

$$\forall \text{Diag}(x, y) \text{ CorrectLabel}(x, y)$$

Diagonal Quantification

Comparison

Normal quantification

Every x with every y

$n \times n$ pairs

Useful for all-pairs reasoning

May include irrelevant pairs

More computationally expensive

Diagonal quantification

Only x_i with y_i

n pairs

Useful for paired datasets

Preserves intended pairs

More efficient

Main idea

Normal quantification compares everything with everything. Diagonal quantification compares each item only with its paired item. Especially useful for:

image \rightarrow label

question \rightarrow answer

input \rightarrow target

Guarded Quantifiers

Guarded quantifiers apply a rule only to a **specific subset** of examples.

$$\forall x P(x) \longrightarrow \forall x : \text{Guard}(x) P(x)$$

Simple intuition

guarded quantifier = quantifier + filter

$$\forall x : \text{AnimalImage}(x) \text{HasAnimalLabel}(x)$$

For every image that is an animal image, it should have an animal label.

Effect of the guard

Example	AnimalImage(x)	Included?
cat image	True	✓
dog image	True	✓
car image	False	
chair image	False	

Guarded Quantifiers

Guard as a mask

$$P(x) = [0.91, 0.84, 0.30, 0.20] \quad \text{Mask}(x) = [1, 1, 0, 0]$$

The guarded quantifier aggregates only [0.91, 0.84], not the full vector.

$$\text{Mask}(x) = \begin{cases} 1 & \text{if } \text{Guard}(x) \geq \theta \\ 0 & \text{otherwise} \end{cases}$$

Guarded quantifier vs implication

Guarded quantifier

Filters examples first

Ignores out-of-guard examples

Aggregates selected cases only

Better for subset rules

Implication form

Evaluates every example

Still scores them

Aggregates all cases

Affected by irrelevant cases

Main idea

Do not apply this rule to everything, apply it only to examples selected by the guard.
Makes rules more precise when only part of the dataset should contribute.

Objective / Thesis

LTNs are not only a supervised learning method — they are a general **neurosymbolic framework** for expressing different AI tasks as logical constraints.

Core technical idea

Instead of only minimizing prediction error, LTNs maximize knowledge base satisfiability:

$$\max_{\theta} \text{Sat}_{\mathcal{G}_{\theta}}(\mathcal{K})$$

where \mathcal{K} = knowledge base, \mathcal{G}_{θ} = grounding with learnable parameters, Sat = satisfiability score.

Training objective

$$\mathcal{L}(\theta) = 1 - \text{Sat}_{\mathcal{G}_{\theta}}(\mathcal{K}) \quad \Rightarrow \quad \min_{\theta} \mathcal{L}(\theta)$$

The model adjusts θ so the facts and rules in the knowledge base become as true as possible.

Objective / Thesis

Different tasks, one framework

Task	LTN view
Supervised learning	Labeled facts should be satisfied
Semi-supervised learning	Labels plus rules over unlabeled data
Classification	Class predicates output truth scores
Regression	Functions grounded as learnable mappings
Clustering	Similar objects share cluster structure
Relational learning	Predicates model relations between objects
Query answering	Ask for the truth value of a formula
Reasoning	Test whether a conclusion follows from \mathcal{H}

Main thesis

data + rules \rightarrow grounded formulas \rightarrow satisfiability \rightarrow gradient-based learning

LTNs turn machine-learning problems into optimization over logical satisfiability, using one differentiable logical language for learning, reasoning, and querying.

Learning & Reasoning

Learning

Adjust learnable parameters θ so the knowledge base becomes as satisfied as possible:

$$\max_{\theta} \text{Sat}_{g_{\theta}}(\mathcal{K}) \quad \Leftrightarrow \quad \min_{\theta} \mathcal{L}(\theta) = 1 - \text{Sat}_{g_{\theta}}(\mathcal{K})$$

What gets learned?

Component

Example

Predicate grounding

Learn $\text{Cat}(x)$ as a classifier

Function grounding

Learn $\text{price}(x)$ as a regression function

Object embeddings

Learn vector representations for entities

Relation predicates

Learn $\text{Friend}(x,y)$, $\text{Similar}(x,y)$

Learning & Reasoning

Reasoning

Use the learned knowledge base to evaluate or infer new statements.

$$\text{Cat}(x) \rightarrow \text{Animal}(x), \quad \text{Cat}(\text{img}_1) = 0.95$$

$$\text{Query: Animal}(\text{img}_1)? \Rightarrow \text{Animal}(\text{img}_1) = 0.93$$

Learning vs reasoning

Learning

Adjusts model parameters

Maximizes satisfiability

Changes θ

Happens during training

Reasoning

Uses the trained model

Evaluates new formulas

Keeps θ fixed

Happens during evaluation

Main idea

Learning: make the knowledge base true. Reasoning: ask what follows from it. Both are done through grounded formulas with truth values in $[0, 1]$.

Reasoning by Refutation

Reasoning by refutation tests a conclusion by trying to **break it**: can the knowledge base stay satisfied while the conclusion becomes false?

Setup

We want to know whether $\mathcal{K} \models \phi$, i.e. does ϕ follow from the knowledge base?

Refutation idea

Try to find a counterexample where \mathcal{K} is true but ϕ is false:

$Sat(\mathcal{K})$ should be high while $\mathcal{G}(\phi)$ should be low

Optimization view

max $Sat(\mathcal{K})$ while minimizing $\mathcal{G}(\phi)$

Make \mathcal{K} true and ϕ false.

Reasoning by Refutation

Possible outcomes

Outcome	Meaning
Counterexample found	ϕ does not follow from \mathcal{K}
No counterexample found	ϕ is supported by \mathcal{K}
\mathcal{K} unsatisfiable	Knowledge base may be inconsistent

Simple example

$$\text{Cat}(x) \rightarrow \text{Animal}(x), \quad \text{Cat}(\text{img}_1) = 0.95$$

Refutation attempt: set $\text{Animal}(\text{img}_1) \approx 0$. But then $\text{Cat}(\text{img}_1) \rightarrow \text{Animal}(\text{img}_1)$ becomes poorly satisfied. So the conclusion is **supported**.

Main idea

Can the rules be true while the conclusion is false?

Yes \Rightarrow conclusion fails No \Rightarrow conclusion is supported

Reasoning by refutation turns logical reasoning into an optimization problem.

Example: Reasoning by Refutation

Setup

Knowledge base: $\mathcal{K} = \{A \vee B\}$ Query: $(A \vee B) \models_q A?$

Does A necessarily follow if $A \vee B$ is true? **Answer: No.**

Grounding

$$\mathcal{G}(A) = a, \quad \mathcal{G}(B) = b, \quad a, b \in [0, 1]$$

$$\mathcal{G}(A \vee B) = a + b - ab \quad (\text{probabilistic sum})$$

Possible groundings

Grounding	$\mathcal{G}(A)$	$\mathcal{G}(B)$	$\mathcal{G}(A \vee B)$	Interpretation
\mathcal{G}_1	1	1	1	both true
\mathcal{G}_2	1	0	1	only A true
\mathcal{G}_3	0	1	1	only B true

\mathcal{G}_3 is the counterexample: $\mathcal{G}_3(A \vee B) = 1$ but $\mathcal{G}_3(A) = 0$.

Example: Reasoning by Refutation

Step 4 & 5: Try a candidate grounding

Choose $\mathcal{G}(A) = 0$, $\mathcal{G}(B) = 1$. Then:

$$\mathcal{G}(A \vee B) = 0 + 1 - (0)(1) = 1 \geq 0.95 \quad \checkmark$$

The knowledge base is satisfied.

Step 6 & 7: Check the conclusion

$$\mathcal{G}(A) = 0 < 0.95$$

The conclusion is not satisfied. Counterexample found:

$\mathcal{G}(A)$	$\mathcal{G}(B)$	$\mathcal{G}(A \vee B)$	$\mathcal{G}(A) < q?$
0	1	1	yes

$$\therefore (A \vee B) \not\models_q A$$

Example: Reasoning by Refutation

Main idea

$A \vee B$ does not imply A because B alone can make $A \vee B$ true.

Reasoning by refutation searches for:

knowledge base true enough but conclusion false enough

If such a case exists, the conclusion does **not** follow.

Experimental Setup: Binary Classification

Goal

Classify examples as positive or negative using a predicate $A(x)$, where $A(x)$ is the truth value that x belongs to the positive class.

Data & grounding

x^+ = positive examples x^- = negative examples

$$\mathcal{G}(A)(x) = \sigma(\text{MLP}(x)) \in [0, 1]$$

e.g. $A(x) = 0.91$ means the model strongly believes x is positive.

Training objective

$$\max_{\theta} \text{Sat}_{\mathcal{G}_{\theta}}(\mathcal{K}) \quad \Leftrightarrow \quad \min_{\theta} \mathcal{L}(\theta) = 1 - \text{Sat}_{\mathcal{G}_{\theta}}(\mathcal{K})$$

Training pushes $A(x^+) \rightarrow 1$ and $A(x^-) \rightarrow 0$.

Experimental Setup: Binary Classification

Knowledge base

$$\mathcal{K} = \{ \forall x^+ A(x^+), \quad \forall x^- \neg A(x^-) \}$$

Formula	Meaning
$\forall x^+ A(x^+)$	all positive examples should be classified positive
$\forall x^- \neg A(x^-)$	all negative examples should not be classified positive

Results: Binary Classification

Expected behavior after training

Example type	Predicate value	Meaning
Positive x^+	$A(x^+) \approx 1$	classified as positive
Negative x^-	$A(x^-) \approx 0$	classified as negative
Negative x^-	$\neg A(x^-) \approx 1$	satisfies negative rule

Result in LTN terms

$$\forall x^+ A(x^+) \approx 1 \quad \forall x^- \neg A(x^-) \approx 1$$
$$\text{Sat}_{g_\theta}(\mathcal{X}) \rightarrow 1 \quad \mathcal{L}(\theta) \rightarrow 0$$

Accuracy vs satisfiability

Metric	What it asks
Accuracy	Did the prediction cross the threshold?
Satisfiability	How strongly did it satisfy the logical formula?

e.g. $A(x^+) = 0.60$ may count as correct with threshold 0.5, but the rule is only weakly satisfied in LTN terms.

Strengths

Main strengths

Strength	Why it matters
Unified framework	Many tasks written in the same Real Logic language
Uses background knowledge	Rules, constraints, and labels in one knowledge base
Differentiable logic	Formulas trained with gradient-based methods
Handles uncertainty	Truth values in $[0, 1]$, not just true/false
Supports multiple tasks	Classification, regression, clustering, querying
More interpretable	Some knowledge explicitly written as formulas
Works with neural predicates	Predicates grounded as MLPs or classifiers

Strengths

Main strength

A normal neural network learns from data: $x \rightarrow y$. An LTN learns from data **and** rules:

examples + logical constraints \rightarrow satisfiability optimization

LTNs make logic trainable, turning symbolic knowledge into differentiable objectives that guide neural learning.

Challenges

Main challenges

Challenge	Why it matters
Choosing the right formulas	Poorly written rules can mislead the model
Aggregation sensitivity	Formula structure affects satisfaction and gradients
Fuzzy semantics choices	Different t-norms change model behavior
Scalability	Quantifiers over many variables create large products
Optimization difficulty	Nonlinear objectives can be hard to train
Partial interpretability	Neural groundings can still be black boxes
No perfect reasoning guarantee	Optimization may find local optima

Main challenge

Writing a formula is also designing a training objective. $\forall x(\phi(x) \wedge \psi(x))$ and $(\forall x \phi(x)) \wedge (\forall x \psi(x))$ behave differently during training. The question is not only *what* to encode, but *how* to encode it so the optimization behaves correctly.

logical syntax affects numerical learning

Open Questions

Key open questions

Question	Why it matters
Best fuzzy logic operators?	Different t-norms change training behavior
How to write good formulas?	Syntax affects the final loss and gradients
Scalability for large KBs?	Many variables create large tensor computations
Reliability of refutation?	No counterexample found \neq none exists
Handling conflicting rules?	Real KBs may contain noisy or inconsistent rules
Interpretability of groundings?	Learned predicates may still be black boxes
Compete with task-specific models?	Specialized models may outperform on benchmarks

Open Questions

Main idea

LTNs give a powerful framework, but practical success depends on:

good rules + good grounding + stable optimization + scalable computation

The bigger question: can we build AI systems that learn from data, reason with rules, and remain scalable?

Conclusion

What LTNs achieve

Component	Role in LTN
Data	Provides examples and observations
Logic	Provides rules, constraints, and structure
Grounding	Maps symbols to tensors, functions, and predicates
Fuzzy truth values	Allow statements to be partly true
Satisfiability	Measures how well the knowledge base is respected
Gradient learning	Updates neural components to increase satisfiability

Final takeaway

LTNs = neural networks + first-order logic + differentiable satisfiability

LTNs show that logic does not have to remain separate from deep learning:

logical formulas → differentiable computations → trainable AI systems

Logic Tensor Networks are a flexible neurosymbolic framework for **learning**, **reasoning**, and **querying** under logical constraints.

Thank you for listening.

Questions?

Acile Chahboun
Seminar Formal Methods and Automated Reasoning