WORKING PAPER 300/2026

Generalised Geometric Logic: A Logic for Expressing Neural Network Architectures

Ramit Das
Purbita Jana

MADRAS SCHOOL OF ECONOMICS
Gandhi Mandapam Road
Chennai 600 025
India

May 2026

Generalised Geometric Logic: A Logic for Expressing Neural Network Architectures

Ramit Das
Cadence, Bangalore, India
t1i2m3a4r5@gmail.com

and

Purbita Jana
Assistant Professor and
Chair of M.Sc. Data Science Programme,
Madras School of Economics, Chennai, India.
purbita@mse.ac.in

WORKING PAPER 300/2026

May 2026

Price : Rs. 35

MADRAS SCHOOL OF ECONOMICS
Gandhi Mandapam Road
Chennai600 025
India

Phone: 2230 0304/2230 0307/2235 2157
Fax: 2235 4847/2235 2155
Email : info@mse.ac.in
Website: www.mse.ac.in

Generalised Geometric Logic: A Logic for Expressing Neural Network Architectures

Ramit Das
Purbita Jana

Abstract

Neural networks achieve strong empirical performance, yet their architectural semantics and compositional structure remain difficult to analyse formally. This paper develops a logical–topological framework for reasoning about neural network architectures independently of learning dynamics. Focusing on the ART/LAPART family as a canonical testbed in which bidirectional interaction and stability are explicit, we provide a semantic interpretation of excitation relations using geometric logic and topological systems. We extend the classical setting to fuzzy and frame-valued semantics in order to capture graded and potentially incomparable activation strengths. With this extension we show that we *express* SHAP - a Neural Network Explainability methodology. The contribution is foundational: it clarifies how architectural causal structure can be represented, compared, and composed. While the technical development centres on ART and LAPART, the framework isolates structural principles—compositionality, graded influence, and continuity—that can be extend to modern deep neural network architectures.

Keywords: Geometric logic, Topological systems, Neural network semantics, Adaptive Resonance Theory, LAPART, Fuzzy topology, Explainable AI, SHAP values, Frame-valued semantics, Neural architectures
JEL Codes: C02, C45, C63, C65, D83

Acknowledgement

The authors are grateful to the anonymous reviewers of the KR Conference and RTLG 2026 for their valuable comments and suggestions. A preliminary version of this work was presented at RTLG 2026, and the feedback received there has helped improve the present manuscript.

Ramit Das
Purbita Jana

1 Introduction

Formal explanations of neural networks typically focus either on optimisation dynamics or on post-hoc explanation techniques. While these approaches are essential, they often leave implicit the question of *architectural semantics*: how the structure of a neural network itself supports stable causal relationships between activations, independently of training trajectories and specific parameter values.

In this work, we study architectural semantics through a logical and topological lens. Our goal is not to analyse predictive performance or learning algorithms, but to provide a mathematically precise language for expressing and reasoning about patterns of excitation, compositional structure, and causal propagation within neural architectures.

We focus on the ART/LAPART family of neural networks introduced by Carpenter and Grossberg and later extended by Healy. Although these architectures predate modern deep learning models, they make explicit two features that are central to semantic analysis: stability under repeated inputs and bidirectional interaction until equilibrium. These properties make LAPART a natural benchmark for connecting neural computation with logic and topology.

Our technical tool is geometric logic, the logic of finite observations and affirmative assertions. Its semantics is given by topological systems, which unify syntax, algebra, and interpretation within a single framework. This perspective is particularly suitable for neural settings, where information is partial, graded, and continuously varying, and where negation is not operationally observable. We extend the classical setting to fuzzy and frame-valued semantics in order to model graded and non-comparable modes of activation.

We show how ART and LAPART architectures induce (generalised) topological systems, and how geometric sequents can be interpreted as stable causal propagation of excitation. Although our constructions are developed in detail for resonance-based architectures, the abstractions are not architecture specific. We also show that in the fuzzy setting we gain the ability to express SHAP a neural network explainability paradigm.

Why Geometric and Topological Semantics?

Geometric logic is the logic of finite observations and affirmative assertions. Its semantics is given by topological systems, unifying syntax, algebra, and interpretation. This makes it particularly suitable for reasoning about systems with partial, graded, or continuously varying information.

Neural networks naturally exhibit these features: activations are real-valued, influences may not be totally ordered, and small perturbations often lead to small output changes.

2 Definitions

Before presenting our main results, we recall the basic notions required in this paper. In particular, we review (i) the neural network architectures ART and LAPART, and (ii) topological systems, which serve as semantic models for Geometric Logic.

Once this preparatory material is in place, we develop our main constructions in Section 3.

2.1 Neural Network Models

The objective of a machine learning algorithm is to infer a functional relationship that maps input observations to outputs. With the availability of large-scale data and fast computation, such algorithms provide a general framework for approximating hidden functions, with parameter tuning enabling adaptation to different contexts.

Carpenter and Grossberg (1) introduced a neural network architecture known as ART (Adaptive Resonance Theory). ART is designed to address the problem of *classification*, namely, the partitioning of input data into categories. The algorithm exhibits several notable properties:

- **Plasticity:** The system adapts to new input patterns and can form new categories as needed.
- **Stability:** Once a category is learned, it is retained; repeated presentation of the same input yields the same classification.
- **Attention:** The learning process is regulated by attentional mechanisms, implemented via gain control and vigilance parameters, which ensure self-stabilising behaviour.

Building on the ART framework, Healy (4) proposed the LAPART (Laterally Primed Adaptive Resonance Theory) architecture. LAPART consists of two coupled ART systems and is designed to recognise familiar sequences of patterns by learning associations between pattern pairs inferred from prior experience.

2.2 ART and LAPART

ART is an *unsupervised clustering* architecture: inputs are grouped into equivalence classes determined by a vigilance-controlled similarity criterion, rather than supervised class labels. Throughout this paper, the term "classification" is used only in this equivalence-class sense.

For simplicity, we assume binary input patterns. The connection weights are taken to lie in $[0, 1]$, which aligns naturally with the fuzzy and frame-valued semantics introduced

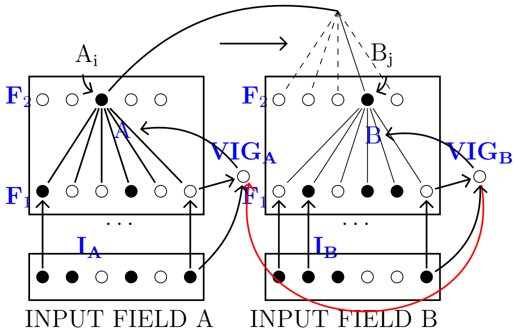

Figure 1: A LAPART system

later. Although ART algorithm takes the floor of the real values we can simply imagine an extension which will be relevant in our discourse later where the ceiling is not used. The operator $\wedge$ denotes componentwise conjunction on binary patterns (after feedforward processing), and should not be interpreted as XOR on real-valued weights.

The vigilance parameter is a hyperparameter controlling the granularity of the induced clusters. The size of the F_2 layer is therefore data-dependent and not fixed *a priori*.

LAPART couples two ART modules and propagates activity bidirectionally until a stable pair of templates is accepted. A learned association $A_i \mapsto B_j$ is interpreted semantically as a stable inference between categories.

We present the details of the algorithms in Appendix A. Having established the ART and LAPART frameworks, we next introduce the necessary notions from topology and topological systems in order to connect these architectures with models of Geometric Logic.

2.3 Topology and Motivation

Geometric logic is interpreted in *topological systems*, a point-free (or "pointless") notion of topology in which opens form a frame and points are related to opens via a satisfaction relation. Although this differs from classical point-set topology, the two notions are equivalent in the spatial case.

An L-topological space generalises ordinary topology by allowing membership values in a frame L rather than $\{0, 1\}$. Classical topological spaces arise when $L = \{0, 1\}$, while fuzzy topologies arise when $L = [0, 1]$. Such structures are well suited to neural semantics, where activation is graded and not necessarily totally ordered.

The motivation for using topology is twofold. First, geometric logic precisely matches the closure properties of observable activation patterns. Second, continuity provides a principled notion of semantic preservation between architectures, allowing comparisons at the level of structure rather than numerical parameters. A brief primer is accessible from Appendix B.

3 LAPART and Topological Systems

With the above definitions in place, we now explain how the ART components of a LAPART architecture may be viewed as topological systems. Our approach proceeds by assigning a logical language of varying expressive strength to capture node excitation within the neural network. Each extension of the logic induces a corresponding refinement of the associated semantics.

Geometric Logic. The logic is intended to express the cause of excitation of a particular node in the neural network. To this end, we introduce a primitive predicate $r(x)$, intended to mean that the node represented by the variable x excites the node r.

For a given ART_A system, we define the syntax of the logic as

$$\phi \in \mathcal{L}_A ::= \top \mid \bot \mid r(x) \mid \phi_1 \wedge \phi_2 \mid \bigvee_{i \in S} \phi_i,$$

where $r \in ptA$, the set of nodes of the ART_A network.

An analogous logic $\mathcal{L}_B$ is defined for the ART_B system.

Semantics. Formulas are interpreted over the neural network model $\mathcal{M}_A = (ptA, E_A)$, where E_A denotes the directed excitation relation. The satisfaction relation is defined as follows:

- $a \models r(x)$ iff $(a, r) \in E_A$ and node a excites node r;
- $a \models \phi_1 \wedge \phi_2$ iff $a \models \phi_1$ and $a \models \phi_2$;
- $a \models \bigvee S$ iff $a \models \phi$ for some $\phi \in S$.

Note that $a \not\models a(x)$, since nodes do not excite themselves, and if $(p, q) \in E_A$, then $q \not\models p(x)$.

We present the **Associated topological system** in Appendix C. Overall idea is to establish that a ART system is a topological system and the connections between the two are like continuous functions between two topological systems.

4 Extending the Logic

In the zeroth-order fragment of geometric logic without equality, one gains the ability to express whether a node causes excitation in a neighbouring node. However, when we examine neural network algorithms more closely, it becomes evident that nodes excite other nodes at varying excitation levels. Moreover, two distinct nodes may excite the same node in qualitatively different ways. Such nuances naturally arise in ambiguous systems, where the underlying knowledge is certain, but the information admits gradation—either quantitatively or qualitatively.

This is true for nearly all contemporary AI algorithms. Let us revisit our ART systems with this perspective. Imagine the ART algorithm modified to allow for graded weights in place of the boolean setting.

Accordingly, we refine our model of an ART_A system as

$$\mathcal{M}_A = (ptA, E_A, wt_A),$$

where $wt_A : E_A \to \mathbb{R}$ assigns a real-valued weight to each directed edge.

4.1 Fuzzy Extension

In fuzzy logic, the satisfaction relation is replaced by a graded one of the form

$$\models: \mathfrak{M}_A \times \mathcal{L}_A \to [0,1].$$

In our setting, we do not modify the *syntax* of the logic $\mathcal{L}_A$, but instead employ the zeroth-order fragment (without equality) of fuzzy geometric logic (2). In particular, the one-variable predicate fragment suffices for our purposes.

The predicate $r(x)$ now carries a graded meaning. For instance, if

$$\models (a, r(x)) = 0.5,$$

then node a excites node r with excitation level 0.5.

Formally, we define the satisfaction relation by

$$\models (a, r(x)) = \begin{cases} wt_A(a,r), & \text{if } (a,r) \in E_A, \\ 0, & \text{otherwise.} \end{cases}$$

This fuzzy satisfaction relation naturally yields a connection with fuzzy topological systems (2). Consequently, the resulting framework allows the analysis of neural network semantics along the lines of (4), while providing strictly richer expressive power.

4.2 Shapley-Based Semantics in the $[0, 1]$ Case

Neural networks are often described as black boxes because it is difficult for humans to understand or justify why a trained model produces a particular output for a given input. This opacity arises from the networks' large scale, deep layered structure, and the highly non-linear interactions among a vast number of interconnected parameters.

Explainability of neural networks is a young research area focused on developing tools that provide insight into model behaviour, with particular interest in model-agnostic methods that operate independently of network architecture. A prominent example is SHAP (SHapley Additive exPlanations) explored first in paper (7). It is based on ideas from coalitional game theory. It offers a principled approach to attributing predictions to input features.

We imagine a SHAP based explainability of ART systems and propose in this subsection how SHAP (SHapley Additive exPlanations) values can be used to instantiate such degrees in neural network explanatory semantics.

Let $(X, \models, A)$ be a $[0,1]$-topological system, where the satisfaction relation

$$\models : X \times A \to [0,1]$$

assigns to each formula a graded degree of support.

Fix a point $x \in X$ and an output node r. Let $F = \{f_1, \dots, f_n\}$ denote the set of input features or lower-level activations that influence r. Following (7), we associate to r a cooperative game

$$v_r : \mathcal{P}(F) \to [0,1],$$

where $v_r(S)$ is the normalised activation of r when only the features in S are present. The Shapley value $\phi_i(r)$ of feature f_i represents its average marginal contribution to the activation of r.

We interpret atomic predicates $p_i(x)$, expressing the activation of feature f_i, by defining

$$\models (x, p_i) := \phi_i(r).$$

The satisfaction relation is then extended inductively to compound formulae using the standard clauses of fuzzy geometric logic:

$$\models (x, \varphi \wedge \psi) = \min(\models (x, \varphi), \models (x, \psi)), \qquad \models \left(x, \bigvee_j \varphi_j \right) = \sup_j \models (x, \varphi_j).$$

This construction yields a $[0,1]$-valued satisfaction relation compatible with fuzzy geometric logic and hence with $[0,1]$-topological systems. Under this interpretation, Shapley values quantify the degree to which atomic observations support a given outcome, while conjunction and disjunction correspond to joint and alternative modes of support.

Remark. The resulting semantics is explanatory rather than purely architectural, as the Shapley values depend on a background distribution used to evaluate marginal contributions. Nevertheless, it provides a principled way to relate feature attribution methods to the graded satisfaction relation $\models$ in $[0,1]$-topological systems.

We move the details of Geometric Logic to the Appendix D.

5 Conclusion

We have developed a logical–topological framework for reasoning about neural network architectures using geometric logic and its generalisations. Focusing on ART and LAPART systems, we provided a semantic interpretation of excitation relations and showed how fuzzy and frame-valued extensions give rise to spatial L-topological systems equivalent to L-topological spaces. Within this setting, geometric sequents admit a natural interpretation as stable causal relationships between activations.

The contribution of this work is architectural and foundational. Rather than addressing optimisation or learning dynamics, we clarify how compositional structure and graded causal influence are supported by network architecture itself. Allowing truth values to range over frames enables the representation of incomparable modes of activation that arise naturally in neural computation.

Although ART and LAPART serve as our primary case studies, the framework is not limited to resonance-based models.

For future direction it would be interesting if we could show that deep convolutional architectures, such as AlexNet, fit naturally into the same semantic framework when analysed at the level of architectural composition. This would suggest that generalised geometric logic provides a unifying semantic perspective across both classical and modern neural architectures.

In addition, a closer integration with learning dynamics, a systematic study of learned inference rules, and extensions to attention-based and transformer architectures would be warranted. More broadly, this work represents a step toward logics tailored to the structural and explanatory requirements of contemporary machine learning systems.

References

[1] Carpenter, Gail A, Grossberg, S.: A massively parallel architecture for a self-organizing neural pattern recognition machine, *Computer vision, graphics, and image processing*, 37(1), 54–115 (1987).

[2] Chakraborty, M.K., Jana, P.: Fuzzy topology via fuzzy geometric logic. *Int. J. Approx. Reasoning* 80, 334–347 (2017)

[3] Denniston, J.T., Melton, A., Rodabaugh, S.E.: Interweaving algebra and topology: Lattice-valued topological systems, *Fuzzy Sets Syst.*, 192, 58–103 (2012)

[4] Healy, M.J.: Continuous functions and neural network semantics, *Nonlinear Analysis* 30(3), 1335–1341 (1997)

[5] Jana, P: Generalised geometric logic, *Categories and General Algebraic Structures with Applications*, 22(1), 43–58 (2025)

[6] Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25, 1097–1105 (2012)

[7] Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems* 30, 4765–4774 (2017)

[8] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, Aidan N, Kaiser, Ł., Polosukhin, I.: Attention is all you need, *Advances in neural information processing systems*, 30 (2017).

[9] Vickers, S.: *Topology via Logic*. Cambridge University Press (1989)

A Appendix

A ART AND LAPART ALGORITHM

A.1 ART algorithm

We now describe the ART algorithm at a detail sufficient for our purposes, with particular emphasis on features relevant to the LAPART architecture. For simplicity, we assume that inputs are binary patterns. Let I denote an input vector drawn from the set of all binary strings of length m.

- The ART network consists of three layers: the input layer I, the comparison layer F_1, and the category layer F_2. Each input node is connected one-to-one with nodes in F_1, so F_1 contains m nodes.

- The number of nodes in F_2, denoted by $|F_2| = m_2$, may vary over time t, corresponding to the iteration number of the learning process. This variation is controlled by the vigilance parameter.

- Each node $j \in F_2$ is connected to all nodes in F_1 by weighted edges. At time t, these weights form a vector $\vec{w}_j(t) \in \mathbb{R}^m$.

- The collection of weights between F_1 and F_2 at time t can be represented by the incidence matrix

$$M_t = \begin{pmatrix} w_{1,1} & w_{1,2} & \cdots & w_{1,m} \\ w_{2,1} & w_{2,2} & \cdots & w_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ w_{m_2,1} & w_{m_2,2} & \cdots & w_{m_2,m} \end{pmatrix}.$$

- At each time step t, a node $j \in F_2$ satisfies the vigilance criterion if

$$\frac{|I \wedge \vec{w}_j|}{|I|} \geq \rho,$$

where ρ is the vigilance parameter. If no existing node satisfies this condition, a new node is added to F_2 with weight vector $\vec{\mathbf{1}}$.

- A winner-take-all mechanism selects a unique node $j \in F_2$ for learning at time t. The winner is chosen by maximising

$$\max_{j \in F_2} \frac{|I \wedge \vec{w}_j|}{|\alpha + \vec{w}_j|},$$

where $\alpha > 0$ is a small constant.

- Only the weight vector of the winning node is updated, according to

$$\vec{w}_j(t+1) = \beta(I \wedge \vec{w}_j(t)) + (1-\beta)\vec{w}_j(t),$$

where β is the learning rate.

- The operation $\wedge$ may be interpreted as bitwise conjunction (or equivalently, addition modulo 2 for binary inputs).

After processing an input $I \in \mathcal{I}$, the ART system assigns I to an equivalence class E, which admits two equivalent representations:

1. A unique category node $F_{2,i}$.
2. A binary template pattern T^i, determined by the weights connecting F_1 to $F_{2,i}$, such that

$$T^i_k = 1 \iff \forall I \in E,\ I_k = 1,$$

for $k \in [m]$.

A.2 LAPART

With the ART architecture in place, we now describe the LAPART system. Our choice of LAPART is motivated by its role as a natural interface between logical inference and neural network semantics (4).

LAPART couples two ART systems, denoted ART_A and ART_B, and may be viewed as simulating bidirectional propagation until a stable equilibrium is reached.

- **Connections:**
  1. Each node $i \in F_2^A$ is adaptively connected to each node $j \in F_2^B$.
  2. A single additional fixed connection $VIG_B \to VIG_A$ is imposed between the vigilance nodes, enforcing learned inferences.
- **Algorithm:**
  1. An input I_A is presented to ART_A, yielding a winning node $F_{2,i}^A$ and a corresponding template $T^{A,i}$.
  2. Activations propagate along the adaptive connections from F_2^A to F_2^B, inducing activity in some node $F_{2,j}^B$.
  3. This activation generates a template pattern $T^{B,j}$ in the F_1^B layer.
  4. If the induced template does not sufficiently match the input I_B, the vigilance node VIG_B is activated, which in turn triggers VIG_A and causes a reset in ART_A.
  5. The process repeats until both ART systems accept the induced templates.

The LAPART system takes a pair of inputs (I_A, I_B) and classifies them as (A_i, B_j). We interpret this outcome as a learned inference

$$A_i \mapsto B_j.$$

B TOPOLOGICAL SYSTEMS

B.1 Topological Space

Definition 1 (Frame)**.** *A **frame** $(A, \wedge, \bigvee, \top, \bot)$ is a complete lattice such that for all $x \in A$ and all subsets $Y \subseteq A$,*

$$x \wedge \bigvee Y \;=\; \bigvee \{\, x \wedge y \mid y \in Y \,\}.$$

That is, finite meets distribute over arbitrary joins.

Throughout this section, we will denote a frame simply by its underlying set A, with the operations understood.

Definition 2 (Frame homomorphism)**.** *Let A and B be frames. A **frame homomorphism** $f : A \to B$ is a function preserving arbitrary joins and finite meets, that is,*

$$f\Big(\bigvee S\Big) = \bigvee_{x\in S} f(x) \quad \textit{and} \quad f(x \wedge y) = f(x) \wedge f(y),$$

for all $S \subseteq A$ and $x, y \in A$.

Definition 3 (L-topological space)**.** *Let X be a set and τ a collection of L-valued subsets of X, that is, functions $\tilde{A} : X \to L$, satisfying:*

1. *$\tilde{\emptyset}, \tilde{X} \in \tau$, where $\tilde{\emptyset}(x) = 0_L$ and $\tilde{X}(x) = 1_L$ for all $x \in X$;*
2. *if $\{\tilde{A}_i\}_{i\in I} \subseteq \tau$, then $\bigcup_{i\in I} \tilde{A}_i \in \tau$, where*

$$\Big(\bigcup_{i\in I} \tilde{A}_i\Big)(x) = \sup_{i\in I} \tilde{A}_i(x);$$

3. *if $\tilde{A}_1, \tilde{A}_2 \in \tau$, then $\tilde{A}_1 \cap \tilde{A}_2 \in \tau$, where*

$$(\tilde{A}_1 \cap \tilde{A}_2)(x) = \tilde{A}_1(x) \wedge \tilde{A}_2(x).$$

*Then (X, τ) is called a L-**topological space**. The elements of τ are called L-**open sets**.*

B.2 Topological Systems

Definition 4. *(9) A **topological system** is a triple $(X, \models, A)$, where X is a non-empty set, A is a frame, and $\models \subseteq X \times A$ is a satisfaction relation such that:*

1. *for any* finite *subset $S \subseteq A$,*

$$x \models \bigwedge S \ \textit{ iff } \ x \models a \textit{ for all } a \in S;$$

2. *for any subset $S \subseteq A$,*

$$x \models \bigvee S \ \textit{ iff } \ x \models a \textit{ for some } a \in S.$$

Definition 5 (L-topological system)**.** *(5) A L-**topological system** is a triple $(X, \models, A)$, where X is a non-empty set, A is a frame, and $\models$ is a L-valued relation*

$$\models : X \times A \to L,$$

satisfying:

1. for any finite subset $S \subseteq A$,

$$\models (x, \bigwedge S) = \inf\{\models (x, s) \mid s \in S\};$$

2. for any subset $S \subseteq A$,

$$\models (x, \bigvee S) = \sup\{\models (x, s) \mid s \in S\}.$$

Definition 6 (Continuous function)**.** *Let $(X, \models, A)$ and $(Y, \models', B)$ be L-topological systems. A **continuous function** between them is a pair (f_1, f_2), where $f_1 : X \to Y$ is a function and $f_2 : B \to A$ is a frame homomorphism, such that*

$$\models (x, f_2(b)) \ = \models' (f_1(x), b),$$

for all $x \in X$ and $b \in B$.

Definition 7 (Spatial)**.** *A L-topological system $(X, \models, A)$ is said to be **spatial** if, for all $a, b \in A$,*

$$\big(\forall x \in X,\ \models (x, a) = \models (x, b)\big) \ \Rightarrow\ a = b.$$

Theorem 1. *(3) The category of spatial L-topological systems is equivalent to the category of L-topological spaces.*

C Associated Topological System

We now construct a topological system from ART_A. Let $X = ptA$, and define the extent of a formula as follows.

Definition 8. *The **extent** function*

$$\mathrm{ext}_A : \mathcal{L}_A \to \wp(ptA)$$

is defined by

$$\mathrm{ext}_A(r(x)) = \{a \in ptA \mid a \models r(x)\},$$

and extended inductively over the logical connectives.

It follows immediately that

$$\mathrm{ext}_A(\top) = ptA, \quad \mathrm{ext}_A(\bot) = \emptyset,$$

$$\mathrm{ext}_A(\phi_1 \wedge \phi_2) = \mathrm{ext}_A(\phi_1) \cap \mathrm{ext}_A(\phi_2), \quad \mathrm{ext}_A\Big(\bigvee S\Big) = \bigcup_{\phi \in S} \mathrm{ext}_A(\phi).$$

Hence, the collection of all extents forms a topology on ptA.

Definition 9. *Define*

$$\Omega A = \{\mathrm{ext}_A(\phi) \mid \phi \in \mathcal{L}_A\}.$$

One may verify that $(ptA, \in, \Omega A)$ and $(ptB, \in, \Omega B)$ are topological systems in the sense of Definition 4.

Continuous functions. A continuous function between $(ptA, \in, \Omega A)$ and $(ptB, \in, \Omega B)$ is a pair (f_1, f_2), where $f_1 : ptA \to ptB$ and $f_2 : \Omega B \to \Omega A$ is a frame homomorphism satisfying

$$f_1(a) \models r(x) \ \text{ iff } \ a \models f_2(r(x)),$$

for all $a \in ptA$ and $r(x) \in \Omega B$.

Since $(ptA, \Omega A)$ and $(ptB, \Omega B)$ are topological spaces, any continuous function $f : ptA \to ptB$ induces a continuous function of topological systems via the pair (f, f^{-1}). In this sense, the logic of ART_A is capable of expressing the logic of ART_B through the notion of continuity.

D GEOMETRIC LOGIC

We now proceed to further generalise the underlying logic to capture even more refined neural network semantics. This is achieved in the next subsection by introducing the propositional fragment of generalised geometric logic and studying its models and interconnections.

D.1 Generalisation of Fuzzy Extension

We now propose a more general class of neural networks in which edge weights take values in a **frame**. This allows us to model situations where excitation levels may be incomparable, even though they originate from the same value set.

Intuitively, different nodes may influence a given node in fundamentally different ways. To incorporate this feature, we define

$$wt_A : E_A \to L,$$

where L is a **frame**.

To interpret such networks logically, we employ generalised fuzzy logic. In this setting, the satisfaction relation takes the form

$$\models : \mathfrak{M}_A \times \mathcal{L}_A \to L,$$

where L is a frame. Accordingly, we define

$$\models (a, r(x)) = \begin{cases} wt_A(a,r), & \text{if } (a,r) \in E_A, \\ 0_L, & \text{otherwise.} \end{cases}$$

D.2 Generalised Geometric Logic: A Logic of Finite Observations

In this subsection, we consider the propositional fragment of generalised geometric logic (5), which we call the *generalised logic of finite observations*. The motivation for this terminology will become clear below.

As observed in (9), the logic of finite observations coincides with the logic of affirmative assertions. An assertion is affirmative if and only if it is true precisely in those circumstances where it can be affirmed. Such assertions are grounded in finite time, finite effort, and observable evidence. This explains why the logic admits the connectives $\wedge$, $\bigvee$, $\top$, and $\bot$, but excludes $\neg$ and $\rightarrow$.

To bring this notion closer to real-world reasoning, it is often preferable to measure the *degree* to which an affirmative assertion holds, rather than merely determining its binary validity. This motivates the use of graded valuation functions.

Let Φ be a set of propositional variables. The language $\mathcal{L}_G$ of generalised logic of finite observations is given by

$$\phi ::= \top \mid \bot \mid p \mid \phi_1 \wedge \phi_2 \mid \bigvee \{\phi_i\}_{i \in I},$$

where $p \in \Phi$ and I is an index set.

The inference rules are as follows:

1. $\phi \vdash \phi$,
2. $\dfrac{\phi \vdash \psi \qquad \psi \vdash \chi}{\phi \vdash \chi}$,
3. (i) $\phi \vdash \top$, (ii) $\phi \wedge \psi \vdash \phi$, (iii) $\phi \wedge \psi \vdash \psi$,
   (iv) $\dfrac{\phi \vdash \psi \qquad \phi \vdash \chi}{\phi \vdash \psi \wedge \chi}$,
4. (i) $\phi \vdash \bigvee S$ ($\phi \in S$), (ii) $\dfrac{\phi \vdash \psi \qquad \text{all } \phi \in S}{\bigvee S \vdash \psi}$,
5. $\phi \wedge \bigvee S \vdash \bigvee \{\phi \wedge \psi \mid \psi \in S\}$.

Proposition 1. *$\bigvee \{\phi \wedge \psi \mid \psi \in S\} \vdash \phi \wedge \bigvee S$ is derivable.*

Proof. The derivation follows by standard applications of the inference rules. □

A valuation function $v : \Phi \to L$ extends uniquely to $\hat{v} : \mathcal{L}_G \to L$ defined by

1. $\hat{v}(\top) = 1_L$,
2. $\hat{v}(\bot) = 0_L$,
3. $\hat{v}(\phi \wedge \psi) = \hat{v}(\phi) \wedge \hat{v}(\psi)$,
4. $\hat{v}(\bigvee\{\phi_i\}) = \sup_i \hat{v}(\phi_i)$.

Allowing L to be an arbitrary frame accommodates incomparable truth values, a phenomenon common in real-world reasoning. In this way, we generalise the logic of finite observations to handle graded and non-comparable information.

Definition 10. *$\phi \vdash \psi$ is* valid *if and only if $\hat{v}(\phi) \leq \hat{v}(\psi)$ for all $\hat{v} : \mathcal{L}_G \to L$.*

Proposition 2. *All inference rules are valid.*

D.3 L-Topological Systems via Generalised Geometric Logic

Let $(X, \models, A)$ be defined as follows: X is the set of all extended valuation functions $\hat{v}$, A is the set of geometric formulae, and

$$\models(\hat{v}, \phi) = \hat{v}(\phi).$$

Proposition 3. *For all $\hat{v} \in X$,*

1. *$\models(\hat{v}, \phi \wedge \psi) = \models(\hat{v}, \phi) \wedge \models(\hat{v}, \psi)$,*
2. *$\models(\hat{v}, \bigvee\{\phi_i\}) = \sup_i \models(\hat{v}, \phi_i)$.*

Define $\phi \approx \psi$ if $\models(\hat{v}, \phi) = \models(\hat{v}, \psi)$ for all $\hat{v} \in X$. This is an equivalence relation, yielding the quotient $A/\approx$.

Proposition 4. *$(A/\approx, \leq, \wedge, \bigvee)$ is a frame.*

Theorem 2. *$(X, \models', A/\approx)$ is an L-topological system.*

Proposition 5. *The above L-topological system is spatial.*

D.4 L-Topology via Generalised Geometric Logic

From the L-topological system $(X, \models', A/\approx)$, we construct an L-topological space $(X, \mathrm{ext}(A/\approx))$, where

$$\mathrm{ext}([\phi])(\hat{v}) = \models'(\hat{v}, [\phi]).$$

One verifies that this collection is closed under finite intersection and arbitrary union, and hence forms an L-topology.

Since the system is spatial, Theorem 1 implies that the associated L-topological system and L-topological space are equivalent.

Finally, given an L-topological space (X, τ), one may recover a corresponding generalised geometric theory by associating propositions to L-open sets and axioms encoding inclusion, unions, and finite intersections. Each point $x \in X$ then determines a model in which the truth value of a proposition is given by its membership degree.

Thus, L-topology and generalised geometric logic provide mutually equivalent perspectives for analysing graded neural network semantics.

D.5 Interpretation of Sequents

An expression of the form $\phi \vdash \psi$, where ϕ and ψ are geometric formulae, is called a *sequent*. The sequent $\phi \vdash \psi$ is read as "ϕ turnstile ψ" and intuitively expresses that ψ follows from ϕ.

In this section, we explain how sequents arising in geometric logic and its generalisations can be interpreted within the context of LAPART. A rule of inference in geometric logic and its extensions (including fuzzy geometric logic and generalised geometric logic) is typically of the form $\dfrac{S_1, S_2, \ldots, S_i}{S}$ where each of $S_1, S_2, \ldots, S_i$ and S is a sequent. The sequents $S_1, S_2, \ldots, S_i$ are called the *premises*, and S is the *conclusion*. Note that the set of premises may be empty, in which case the rule represents an axiom.

Within the LAPART framework, a sequent of the form

$$p(x) \vdash q(x)$$

is interpreted as follows: the activation of node p is consistently followed by the activation of node q, due to an excitatory connection whose strength exceeds a prescribed threshold. In other words, the implication encoded by the sequent captures a reliable causal relationship between node activations.

Under this interpretation, the inference rule $\dfrac{p(x) \vdash q(x) \qquad q(x) \vdash r(x)}{p(x) \vdash r(x)}$ expresses the following reasoning pattern: if activation of node p is always followed by activation of node q through a strong excitatory connection, and activation of node q is always followed by activation of node r through a strong excitatory connection, then activation of node p is always followed by activation of node r through a strong excitatory connection.

In this way, the transitivity of entailment in geometric logic corresponds naturally to the propagation of excitation through successive layers of the LAPART architecture. Other inference rules of geometric logic and its extensions admit analogous interpretations in terms of neural activation patterns and their structural dependencies.

MSE Monographs

* Monograph 36/2017
  Underlying Drivers of India's Potential Growth
  C.Rangarajan and D.K. Srivastava

* Monograph 37/2018
  India: The Need for Good Macro Policies *(4th Dr. Raja J. Chelliah Memorial Lecture)*
  Ashok K. Lahiri

* Monograph 38/2018
  Finances of Tamil Nadu Government
  K R Shanmugam

* Monograph 39/2018
  Growth Dynamics of Tamil Nadu Economy
  K R Shanmugam

* Monograph 40/2018
  Goods and Services Tax: Revenue Implications and RNR for Tamil Nadu
  D.K. Srivastava, K.R. Shanmugam

* Monograph 41/2018
  Medium Term Macro Econometric Model of the Indian Economy
  D.K. Srivastava, K.R. Shanmugam

* Monograph 42/2018
  A Macro-Econometric Model of the Indian Economy Based on Quarterly Data
  D.K. Srivastava

* Monograph 43/2019
  The Evolving GST
  Indira Rajaraman

* Monograph 44/2025 Landscape Analysis of the Labour Market of the Freight Logistics Sector in India
  Gopal Krishna Roy, Brinda Viswanathan, Ashrita. B, Madhuritha Murali and Mohit Sharma

* Monograph 45/2025 The Fisc and India's Energy Transition
  Laveesh Bhandari

MSE Working Papers
Recent Issues

* Working Paper 291/2025
  Bribing and Vulnerability of the Informal Sector in India
  Devlina and Santosh Kumar Sahu

* Working Paper 292/2026
  Patent Valuation under Fragile Institutional Enforcement: A Continuous-Time Markov Approach
  Srikanth Pai, Akila Hariharan, and Naveen Srinivasan

* Working Paper 293/2026
  Early Detection of ESG Policy Violations Using Machine Learning Techniques
  Gautami Parate and Arpita Choudhary

* Working Paper 294/2026
  Stock Market Reactions to COP26 and Climate Change Exposures of Indian Firms
  Saumitra N Bhaduri, Ekta Selarka and Alankrti Aggrwal

* Working Paper 295/2026
  European Union Regulations in Indian Tyre Industry
  Ekta Selarka and Subrata Sarkar

* Working Paper 296/2026
  Women Director Networks and Corporate Social Responsibility of Indian firms
  Kavitha Nambiar and Ekta Selarka

* Working Paper 297/2026
  Does Perception Matter? The Role of Monetary Policy Uncertainty in Policy Transmission
  Aariya Sen

* Working Paper 298/2026
  Reimagining Gender Budgeting Framework in India: Linking Fiscal Outlays and Gendered Outcomes
  N R Bhanumurthy, Bhabesh Hazarika and Aritri Chakravarty

* Working Paper 299/2026
  Explainable Decision Support in Multi-Agent AI Systems Using L-Valued Information Flow and Shapley Aggregation
  Purbita Jana
