Vector space embeddings of modal logic formulas: theory and application

Pietro Casavecchia [1,2] Alberto Paparella [1,2] Leonardo Serrentino [1,2]

[1] University of Ferrara, Department of Mathematics and Computer Science
[2] Applied Computational Logic and Artificial Intelligence Lab

Introduction

Vector space embeddings of logic formulas aim to represent symbolic logic as vectors. This enables, for example, neural networks and machine learning models to perform reasoning over logic-based representations [4]. Such embeddings can be given [3] or learned from data [6]. However, there seem to be no literature about embeddings of modal logics. Therefore, a formal framework for vector space embeddings of modal formulas would be of great value, providing an essential tool for faster reasoning.

This work describes a novel approach to embed modal logic formulas over a vector space, with the key idea of interpreting a set of representative modal logic models as dimensions and hence providing a binary encoding of formulas based on the models that satisfy them. This shares some ideas with [5], however following different goals. Furthermore, some examples of both theoretical and applicative interest are described, solving classical problems in an heuristic way somewhat similar to [1, 2]. Finally, we also provide an open-source implementation [a].

Method

Let $\mathcal{P} = \{p, q, r, \ldots\}$ be a set of propositional letters, $\{\wedge, \vee, \neg\}$ a set of propositional logic operators, $\{\Box, \Diamond\}$ a set of normal modal logic operators, $\mathcal{F} = \{\varphi_1, \varphi_2, \ldots, \varphi_n\}$ a set of well-formed formulas, and $\mathcal{M} = \{M_1, M_2, \ldots, M_m\}$ a set of finite Kripke models, where each model $M_i \in \mathcal{M}$ is defined as:

$$M_i = (W_i, R_i, V_i)$$

where $W_i = \{w_1, ..., w_z\}$ is the set of worlds in M_i, $R_i \subseteq W_i \times W_i$ is an accessibility relation between worlds, and $V_i : W_i \times P \rightarrow \{0, 1\}$ is a valuation function assigning a truth value to each propositional letter in each world.

A vector space embedding for formulas in $\mathcal{F}$ is defined through a satisfaction matrix S of dimensions $n \times m$, where each element $S_{i,j} \in \{0, 1\}$ is given by:

$$S_{i,j} = \begin{cases} 0, & \text{if } M_j, w_1 \not\models_e \varphi_i \\ 1, & \text{if } M_j, w_1 \models_e \varphi_i \end{cases} \tag{1}$$

The matrix S summarizes the satisfaction of every formula in $\mathcal{F}$ in every Kripke structure in $\mathcal{M}$. Note that $S \subseteq \mathbb{S}$, where $\mathbb{S}$ is the complete matrix of every possible formula evaluated on every possible Kripke model. Such an embedding can be interpreted as a vector space comprised by Kripke models (the columns of S) as axis where formulas (the rows of S) are represented by binary vectors, i.e., each coordinate is either 0 or 1.

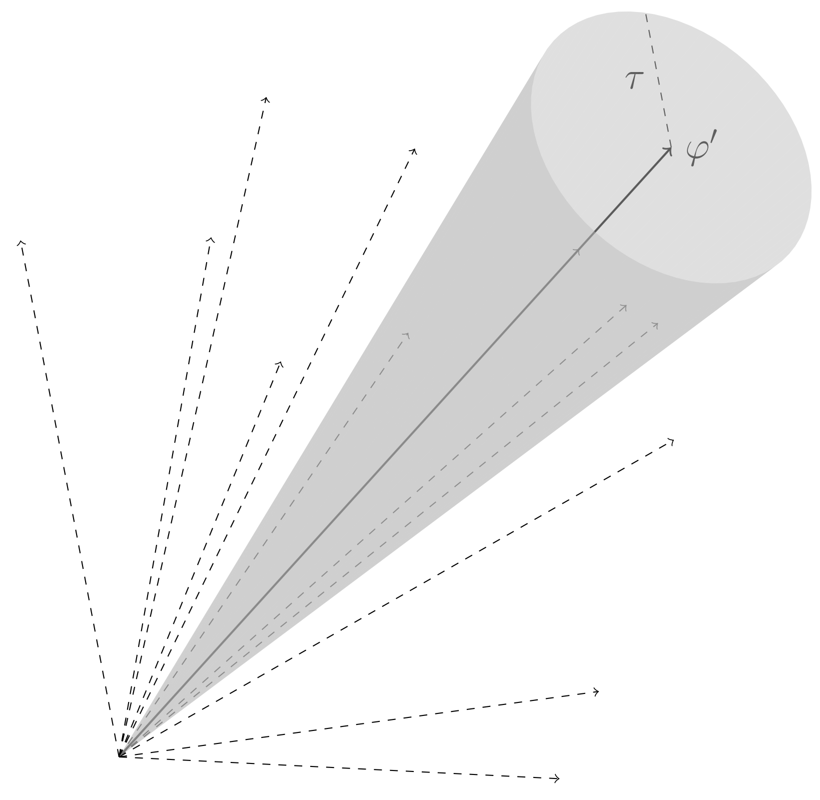

Figure 1. Abstract representation of a high-dimensional vector space; formulas in the gray region are considered similar to φ'.

Given a vector space embedding as described in eq. 1, one can be interested in a notion of "similarity" between formulas: we propose to consider a notion of similarity based on the number of models the formulas share. A figurative example is represented in 1. Such a notion can be grasped by common similarity metrics between vectors, e.g., the rather popular cosine similarity and its variances. In particular, we employ a similarity measure between two fomulas φ_i and φ_j defined as follows:

$$\text{sim}(\varphi_i, \varphi_j) = \frac{S(i,\cdot)^T S(j,\cdot)}{\min(||S(i,\cdot)||_1, ||S(j,\cdot)||_1)} \tag{2}$$

where $||S(i,\cdot)||_1$ represents the number of Kripke structures satisfying formula φ_i; in fact, $S(i,\cdot)$ represents the i-th row of matrix S, while $||S(i,\cdot)||_1$ represents its 1-norm (also called Manhattan distance).

[a] The implementation is open-source and is made available through the following repository: https://github.com/aclai-lab/ESSLLI2025.

Application

Given a vector space embedding as described in eq. 1 and a notion of similarity such as the one in eq. 2, one could easily solve, for example, the problem of stating if a formula entails another, i.e., $\varphi_i \models_e \varphi_j$ where

$$\varphi_i \models_e \varphi_j \iff \forall M \in \mathcal{M} : M, w_1 \models_e \varphi_i \Rightarrow M, w_1 \models_e \varphi_j$$

In fact:

$$\varphi_i \models_e \varphi_j \iff \text{sim}(\varphi_i, \varphi_j) = 1 \wedge \min(||S(i,\cdot)||_1, ||S(j,\cdot)||_1) = ||S(i,\cdot)||_1 \tag{3}$$

An example is given in 2. One should notice that this still an heuristic: there could be a model $\overline{M} \notin \mathcal{M} : \overline{M}, w_1 \models \varphi_i \wedge \overline{M}, w_1 \not\models \varphi_j$. Hence, the choice of $\mathcal{F}$ and $\mathcal{M}$ is crucial to get reliable results.

Another application of interest could be solving the minimization problem, which can be defined as follows: given a formula φ_1, find the formula φ_2 which satisfies exactly the same models (i.e., $\varphi \models_e \varphi_1 \wedge \varphi_1 \models_e \varphi$) and has the least number of symbols. One possible way to solve this problem is considering all formulas:

$$\varphi \in \mathcal{F} : \text{sim}(\varphi_i, \varphi_j) = 1 \wedge ||S(i,\cdot)||_1 = ||S(j,\cdot)||_1 \tag{4}$$

and taking the one with the least number of symbols. This could be taken even a step further if someone is interested in getting a least number of symbols even at the cost of the formulas disagreeing on some models. Should this be the case, one could consider a relaxation of the problem, refered to as quasi-minimization, setting a similarity threshold τ, and considering all formulas $\varphi \in \mathcal{F} : \text{sim}(\varphi_1, \varphi) \geq \tau$ and taking the one with the least number of symbols.

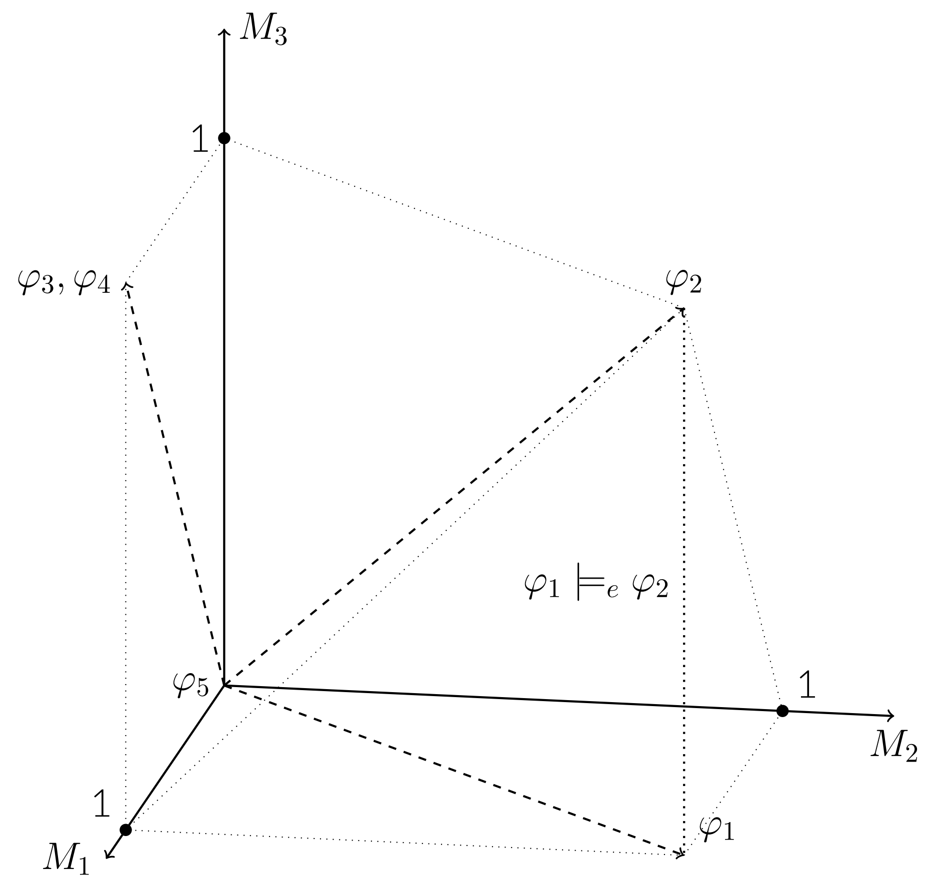

Figure 2. Example of a semantic entailment $\varphi_1 \models_e \varphi_2$ in a 3-dimensional vector space; φ_1 is the projection of φ_2 on the (hyper)-plane formed by the models satisfying φ_1.

Conclusions

In this work, a possible embedding of modal logic formulas has been proposed aiming to solve classical problems, such as minimization and entailment. One should notice how considering only binary encodings is rather constraining, resulting in much sparsity in the vector space. Hence, a natural next step would be generalizing this approach to modal fuzzy logics. Furthermore, the effectiveness of the embedding should be tested, possibly leveraging a modal reasoner checking validity over a broad set of formulas.

References

[1] Davide Bresolin, Fernando Jiménez, Gracia Sánchez, and Guido Sciavicco.
Finite satisfiability of propositional interval logic formulas with multi-objective evolutionary algorithms.
In *Proceedings of the Twelfth Workshop on Foundations of Genetic Algorithms XII*, FOGA XII '13, page 25–36, New York, NY, USA, 2013. Association for Computing Machinery.

[2] Stefan Harmeling.
Solving satisfiability problems with genetic algorithms.
In *Genetic Algorithms and Genetic Programming at Stanford 2000: Book of Student Papers from John Koza's Course at Stanford on Genetic Algorithms and Genetic Programming*, pages 206–213. Stanford Bookstore, 2000.

[3] Taisuke Sato.
Embedding tarskian semantics in vector spaces.
In *The Workshops of the The Thirty-First AAAI Conference on Artificial Intelligence, Saturday, February 4-9, 2017, San Francisco, California, USA*, volume WS-17 of *AAAI Technical Report*. AAAI Press, 2017.

[4] Paul Smolensky.
Tensor product variable binding and the representation of symbolic structures in connectionist systems.
Artificial Intelligence, 46(1):159–216, 1990.

[5] Johan van Benthem and Nick Bezhanishvili.
Modal structures in groups and vector spaces.
Journal of Logic and Computation, 34(1):75–124, 02 2023.

[6] William Yang Wang and William W. Cohen.
Learning first-order logic embeddings via matrix factorization.
In *International Joint Conference on Artificial Intelligence*, 2016.