


Contents lists available at [ScienceDirect](#)

International Journal of Approximate Reasoning

journal homepage: www.elsevier.com/locate/ijar

Embeddings as epistemic states: Limitations on the use of pooling operators for accumulating knowledge

Steven Schockaert

Cardiff University, School of Computer Science & Informatics, Abacws building, Senghennydd Road, CF24 4AG, Cardiff, UK

ARTICLE INFO

Article history:

Received 11 November 2022
Received in revised form 3 July 2023
Accepted 5 July 2023
Available online xxxx

Keywords:

Learning and reasoning
Knowledge representation
Propositional logic

ABSTRACT

Various neural network architectures rely on pooling operators to aggregate information coming from different sources. It is often implicitly assumed in such contexts that vectors encode epistemic states, i.e. that vectors capture the evidence that has been obtained about some properties of interest, and that pooling these vectors yields a vector that combines this evidence. We study, for a number of standard pooling operators, under what conditions they are compatible with this idea, which we call the epistemic pooling principle. While we find that all the considered pooling operators can satisfy the epistemic pooling principle, this only holds when embeddings are sufficiently high-dimensional and, for most pooling operators, when the embeddings satisfy particular constraints (e.g. having non-negative coordinates). We furthermore show that these constraints have important implications on how the embeddings can be used in practice. In particular, we find that when the epistemic pooling principle is satisfied, in most cases it is impossible to verify the satisfaction of propositional formulas using linear scoring functions, with two exceptions: (i) max-pooling with embeddings that are upper-bounded and (ii) Hadamard pooling with non-negative embeddings. This finding helps to clarify, among others, why Graph Neural Networks sometimes under-perform in reasoning tasks. Finally, we also study an extension of the epistemic pooling principle to weighted epistemic states, which are important in the context of non-monotonic reasoning, where max-pooling emerges as the most suitable operator.

© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

One of the key challenges in many sub-areas of Machine Learning is to learn suitable vector space embeddings of the objects of interest (e.g. graphs, images or sentences). A question which is usually left implicit is what the embedding of an object represents. We can take at least two different views on this. First, we may consider that embeddings essentially serve as a compact representation of a *distance metric*, which intuitively captures some form of similarity. What matters, then, is that objects which are similar, in some sense, are represented by similar vectors, while objects which are dissimilar are not. This intuition provides the foundation, for instance, for the use of contrastive pre-training strategies [1,2]. Second, we may consider that embeddings are essentially compact encodings of *epistemic states*. In other words, the embedding of an object encodes what we know about that object. Embeddings then essentially play a similar role as formulas in propositional logic. This view implicitly underpins most strategies that combine neural network learning with aspects of symbolic reasoning,

E-mail address: schockaerts1@cardiff.ac.uk.

<https://doi.org/10.1016/j.ijar.2023.108981>

0888-613X/© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

e.g. when using a semantic loss function to encourage neural network predictions to satisfy certain constraints [3] or when using neural network predictions as input to a probabilistic logic program [4]. In this paper, we focus on this second view.

In practice, the embedding of an object is often obtained by combining the embeddings of related objects using some kind of pooling operator. For instance, in the context of Computer Vision, convolutional feature extractors such as ResNet [5] provide an embedding for each sub-region of the image. An embedding of the overall image is then typically obtained by averaging these sub-region embeddings. Along similar lines, a common setting in Natural Language Processing consists in using a transformer based language model such as BERT [6] to obtain paragraph embeddings, and to average these embeddings to obtain an embedding for a full document. In multi-modal settings, it is common to obtain embeddings for the individual modalities first, and to subsequently aggregate these embeddings [7]. Graph Neural Networks [8,9] also crucially rely on pooling operators, learning node representations by aggregating embeddings derived from neighbouring nodes. Essentially, in all these cases we have an embedding \mathbf{e} which is obtained by aggregating embeddings $\mathbf{e}_1, \dots, \mathbf{e}_k$ using some pooling operator \diamond :

$$\mathbf{e} = \diamond(\mathbf{e}_1, \dots, \mathbf{e}_k) \quad (1)$$

Under the epistemic¹ view, this pooling operator is implicitly assumed to aggregate the knowledge that is captured by the embeddings $\mathbf{e}_1, \dots, \mathbf{e}_k$. For instance, the embeddings $\mathbf{e}_1, \dots, \mathbf{e}_k$ may encode which objects are present in different parts of the image. After pooling these embeddings, we should end up with an embedding \mathbf{e} that captures which objects are present throughout the entire image. Let us write $\Gamma(\mathbf{e}_i)$ for the knowledge that is captured by the embedding \mathbf{e}_i . More precisely, we will think of $\Gamma(\mathbf{e}_i)$ as a set of properties that are known to be satisfied. If we view pooling as the process of accumulating knowledge from different sources (e.g. from different regions of the image, or different neighbours in a graph neural network), then we would expect the following to be true: $\Gamma(\mathbf{e}) = \Gamma(\mathbf{e}_1) \cup \dots \cup \Gamma(\mathbf{e}_n)$. We will refer to this principle as the *epistemic pooling principle*. As an important special case, we will consider the case where the properties of interest correspond to possible worlds (i.e. propositional interpretations). The set $\Gamma(\mathbf{e}_i)$ then contains the set of possible worlds that can be excluded based on the knowledge encoded in the embedding \mathbf{e}_i . In other words, $\Gamma(\mathbf{e}_i)$ can then be characterised as a propositional formula. The epistemic pooling principle then states that the formula corresponding to $\diamond(\mathbf{e}_1, \dots, \mathbf{e}_k)$ should be equivalent to the conjunction of the formulas corresponding to $\mathbf{e}_1, \dots, \mathbf{e}_k$.

The main aim of this paper is to study under which conditions the epistemic pooling principle can be satisfied. Analysing a number of standard pooling operators, we find that the epistemic pooling principle can be satisfied for all of them, but with several important caveats:

- We need at least as many dimensions as there are properties. In settings where we want to model propositional formulas, the properties of interest correspond to possible worlds. Without further restrictions, this means that we need embeddings with as many dimensions as there are possible worlds.
- For most of the pooling operators, we find that embeddings need to be constrained in a particular way, e.g. by only allowing vectors with non-negative coordinates.
- We also identify important restrictions on how embeddings can be linked to the formulas they capture. In particular, when summation or averaging is used for pooling, we find that the satisfaction of propositional formulas can, in general, not be predicted by a linear classifier when the epistemic pooling principle is satisfied.

The fundamental question which we want to answer is whether a vector-based representation can act as a formal knowledge representation framework, or whether reasoning with neural networks is inevitably approximate in nature. While we focus on a theoretical analysis of pooling operators in this paper, our results provide a number of important insights for the design of neural network models for tasks that require reasoning. For instance, two operators emerge from our analysis as being particularly suitable for applications where we need to reason about propositional formulas: max-pooling, with the constraint that the coordinates of all embeddings are upper-bounded by some constant z , and the Hadamard operator (i.e. the component-wise product), with the constraint that all coordinates are non-negative.

Our results also help to explain the outperformance of neuro-symbolic methods [10] over Graph Neural Networks (GNN) in certain applications. GNNs are, in theory, capable of reasoning in the two-variable fragment of first-order logic with counting quantifiers (FOC₂) [11]. Nonetheless, in practice, GNNs often perform relatively poorly in tasks that require reasoning, even when staying within FOC₂. Crucially, GNNs typically use averaging (or summation) for pooling the messages coming from adjacent nodes. As we will see in Section 5.4, in such cases, GNNs can only capture logical reasoning if the node embeddings are essentially binary. The discrete nature of the required embeddings makes them difficult to learn, which helps to explain why GNNs are often outperformed by models that are more specifically tailored towards reasoning [12,13].

Furthermore, our lower bounds on the required dimensionality of epistemic embeddings suggest that, beyond toy problems, using neural networks for reasoning inevitably requires some kind of modularity. For instance, an important challenge

¹ The word *epistemic*, in this paper, merely refers to the idea that vectors capture what we know about some properties of interest. In particular, note that our setting is not directly related to epistemic logics.




Fig. 1. A graph containing four types of nodes. The nodes c_1, c_2 represent committees, r_1, r_2, r_3 represent researchers, a_1, a_2, a_3 represent articles, and t_1, t_2, t_3, t_4, t_5 represent research topics. Topic nodes are connected to the articles that discuss them. Article nodes are connected to their authors. Researchers are connected to the committees they belong to.

in Natural Language Processing is to design models that can reason about information that comes from different sources (e.g. different news sources). Combining different text fragments by pooling their embeddings would require a prohibitively high dimensionality, if we want these embeddings to capture epistemic states in a faithful way. Instead, it is more common to use vectors to encode what we know about a given entity, or about the relationship between two entities. Reasoning about the overall story then requires us to combine these different entity and relation vectors, by using the structure of the problem domain in some way, e.g. by using GNNs or neuro-symbolic approaches. By exploiting this structure, embeddings can be used to capture more focused knowledge, which means that the prohibitively high dimensionality that is otherwise needed can be avoided.

The remainder of this paper is structured as follows. In the next section, we first discuss a number of simple examples to illustrate the considered setting. Subsequently, in Section 3, we formalise the problem setting and introduce the notations that will be used throughout the paper. In Section 4, we then analyse under which conditions the epistemic pooling principle can be satisfied. One of the main findings from this section is that the requirement to satisfy the epistemic pooling principle fundamentally constrains which embeddings can be allowed, and how these embeddings encode knowledge. In Section 5, we then investigate how this impacts our ability to use vector embeddings for propositional reasoning. In particular, we focus on the problem of verifying whether a given propositional formula is satisfied in the epistemic state encoded by a given vector. Subsequently, in Section 6, we look at a generalisation of the epistemic pooling principle for dealing with weighted epistemic states. In this case, vectors encode the strength with which we believe a given property to be satisfied. Section 7 presents a discussion of our results in the context of related work, after which we summarise our conclusions.

2. Illustrative examples

In this section, we consider a number of toy examples to illustrate our problem setting. These examples involve message-passing GNNs, where pooling arises because we need to aggregate the messages coming from adjacent nodes. In Section 2.1, we discuss a simple example to illustrate the basic principles. This example involves aggregating sets of properties, without any further reasoning. We then build on this example in Section 2.2 to illustrate how propositional reasoning can be carried out in this framework. Section 2.3 illustrates a setting where we need to reason in the presence of background knowledge. Finally, Section 2.4 shows how this can be extended to situations involving non-monotonic reasoning with defeasible knowledge.

2.1. Basic setting

Fig. 1 displays a graph with four types of nodes, referring to research topics (t_1, t_2, t_3, t_4, t_5), scientific articles (a_1, a_2, a_3), researchers (r_1, r_2, r_3) and committees (c_1, c_2), respectively. In this graph, research topics are connected to the articles that discuss them, articles are connected to their authors and researchers are connected to the committees they belong to. We say that a researcher is an expert on a topic if they have published at least one article which discusses this topic. We say that a committee is complete if it contains an expert on each of the five topics. For instance, in the case of Fig. 1, we can see that researcher r_1 is an expert on topics t_1, t_2 ; researcher r_2 is an expert on topics t_1, t_2, t_3, t_4 ; and researcher r_3 is an expert on topics t_4, t_5 . It follows that committee c_2 is complete, as each of the five topics are covered by its two members, but committee c_1 is not. We are now interested in designing a message-passing GNN which can predict whether a given committee is complete or not. A message-passing GNN learns embeddings of the nodes of a given graph, by iteratively updating each node's embedding based on the embeddings of its neighbouring nodes. To specify a message-passing GNN, we thus need to specify (i) the initial embedding of each node and (ii) the update mechanism which is used.

A straightforward solution is as follows. Let us write $\mathbf{x}^{(l)} \in \mathbb{R}^5$ for the representation of a node x in layer l of the GNN. In particular, $\mathbf{x}^{(0)}$ represents the input embedding of node x , which we define as follows:

$$\mathbf{t}_1^{(0)} = (1, 0, 0, 0, 0) \quad \mathbf{t}_2^{(0)} = (0, 1, 0, 0, 0) \quad \mathbf{t}_3^{(0)} = (0, 0, 1, 0, 0) \quad \mathbf{t}_4^{(0)} = (0, 0, 0, 1, 0) \quad \mathbf{t}_5^{(0)} = (0, 0, 0, 0, 1)$$

In other words, we use a one-hot encoding for representing the different topics. The input embeddings for all the other nodes are set to $\mathbf{x}^{(0)} = (0, 0, 0, 0, 0)$. In the subsequent layers, the node representations are updated as follows:

$$\mathbf{x}^{(l+1)} = \max(\{\mathbf{x}^{(l)}\} \cup \{\mathbf{y}^{(l)} \mid (y, x) \in \mathcal{E}\}) \quad (2)$$

where \mathcal{E} represents the set of edges, i.e. $(y, x) \in \mathcal{E}$ if there is an edge from node y to node x , and the maximum is applied component-wise. It is easy to verify that an article a_i covers topic t_j if the j^{th} coordinate of $\mathbf{a}_i^{(1)}$ is 1, for $l \geq 1$. Similarly, a researcher r_i is an expert on topic j if the j^{th} coordinate of $\mathbf{r}_i^{(1)}$ is 1, for $l \geq 2$; and a committee c_i contains expertise on topic j if the j^{th} coordinate of $\mathbf{c}_i^{(1)}$ is 1, for $l \geq 3$. We thus have that the committee c_i is complete iff $\mathbf{c}_i^{(3)} = (1, 1, 1, 1, 1)$. Note that the latter condition can be checked using a linear scoring function, since it is equivalent with $\mathbf{c}_i^{(3)} \cdot \mathbf{1} \geq 5$, where $\mathbf{1} = (1, 1, 1, 1, 1)$.

In the proposed construction, a number of particular design decisions were made. For instance, we used 5-dimensional embeddings for representing nodes and we used the maximum for aggregating the evidence from neighbouring nodes. Our analysis in this paper is aimed at studying the importance of such decisions. For instance, we may wonder whether it is possible to devise an encoding that relies on four-dimensional embeddings. If we are allowed to replace the maximum with an arbitrarily complex function (and we allow non-linear scoring functions), then the answer is clearly positive.² However, as we will see in Section 4, when we restrict ourselves to standard pooling operators, such as max-pooling, summation or averaging, the answer is negative. Another question is whether a variant of the construction above can be found which relies on summation or averaging, rather than max-pooling. Here, the answer depends on what assumptions we make on the final node classifier, i.e. the function that maps the embedding $\mathbf{c}^{(l)}$ of a committee onto a decision. The most common approach is to rely on a linear classifier to make such predictions. In that case, we can show that no encoding of the problem can be found that relies on summation or averaging, as we will see in Section 5.

2.2. Propositional reasoning

The aforementioned example simply required us to aggregate sets of features, which corresponded to research topics in that example. Throughout this paper, we will refer to such features as *properties*. At first glance, it may seem like this setting only involves rather basic forms of reasoning. However, by identifying the considered properties with possible worlds, we can in fact design GNNs which perform propositional reasoning. We illustrate this with an example. Let us consider the same setting as before, but now we only focus on research topics, articles and researchers. We say that a researcher is a generalist if they have worked on at least two sufficiently distinct topics. Let us write t_i to denote that some researcher has worked on topic t_i . Then we define:

$$\text{generalist} \equiv (t_1 \wedge t_4) \vee (t_1 \wedge t_5) \vee (t_2 \wedge t_5)$$

where the idea is that the other topic combinations are too similar to each other (e.g. working on t_1 and t_2 would not make someone a generalist because these are related research topics). We want to predict whether a given researcher is a generalist by applying a linear classifier to the corresponding node embedding. As before, we assume that available knowledge about researchers is encoded as a graph with topic, article and researcher nodes (although we no longer consider committees). Let $\omega_1, \dots, \omega_{32}$ be an enumeration of all the possible worlds of the propositional logic over the set of atoms $\{t_1, \dots, t_5\}$. The input embedding $\mathbf{t}_i^{(0)}$ of the node t_i is now defined as a 32-dimensional vector, where the j^{th} component is 0 if $\omega_j \models \neg t_i$ and -1 otherwise. Note how each component now corresponds to a possible world. The j^{th} coordinate in the embedding of node n is 0 if that node captures the knowledge that the world ω_j can be excluded. The input embeddings of the article and research nodes are initialised as $(-1, \dots, -1)$. The embeddings of the nodes in the subsequent layers are again computed using (2). For $l \geq 1$ we then have that $\mathbf{a}_i^{(1)} = (x_1, \dots, x_{32})$ where $x_j = 0$ iff article a_i has some topic t_k such that $\omega_j \models \neg t_k$, and $x_j = -1$ otherwise. For $l \geq 2$, we have that $\mathbf{r}_i^{(1)} = (x_1, \dots, x_{32})$ where $x_j = 0$ iff researcher r_i has written an article that has some topic t_k such that $\omega_j \models \neg t_k$. In other words, we have that $x_j = 0$ if we can exclude ω_j as a possible representation of the expertise of r_i based on the knowledge encoded in the graph. To test whether we can entail that r_i is a generalist, it suffices to check whether every countermodel of $(t_1 \wedge t_4) \vee (t_1 \wedge t_5) \vee (t_2 \wedge t_5)$ can be excluded. Let $\mathbf{c} \in \mathbb{R}^{32}$ be the vector whose j^{th} component is 0 if ω_j is a model of $(t_1 \wedge t_4) \vee (t_1 \wedge t_5) \vee (t_2 \wedge t_5)$ and 1 otherwise. Then we have that researcher r_i is a generalist iff

$$\mathbf{r}_i^{(2)} \cdot \mathbf{c} \geq 0$$

² In fact, we can even use one-dimensional embeddings. Let n be a number which is larger than the number of topics k . Then we can initialise topic i as n^{i-1} . To aggregate two node representations x and y , we first decompose them as $x = x_0 + x_1n + x_2n^2 + \dots + x_{k-1}n^{k-1}$ and $y = y_0 + y_1n + y_2n^2 + \dots + y_{k-1}n^{k-1}$, with $x_0, \dots, x_{k-1}, y_0, \dots, y_{k-1} \in \{0, 1\}$. Then the aggregated embedding can be defined as $\max(x_0, y_0) + \max(x_1, y_1)n + \dots + \max(x_{k-1}, y_{k-1})n^{k-1}$.

By associating properties with possible worlds, we can thus study settings in which GNNs need to aggregate knowledge encoded using propositional formulas. While the kind of knowledge that we had to consider in this example was rather simple, in general the formulas that need to be aggregated can be arbitrarily complex propositional formulas (including e.g. formulas with negation). The topic of propositional reasoning by pooling embeddings, as illustrated in the aforementioned example, will be studied in Section 5. Among others, we will see that the choice of the pooling operator plays an important role, where faithful reasoning is, in general, not possible with averaging and summation, when linear scoring functions are used.

2.3. Background knowledge

Background knowledge can straightforwardly be taken into account by restricting the set of possible worlds to the models of a given knowledge base. Embeddings then capture which models of the knowledge base are still possible. To illustrate this, let us consider a similar setting as before, but with four topics: *Artificial Intelligence*, *Machine Learning*, *Knowledge Representation* and *Databases*. Let us denote these topics as t_{AI} , t_{ML} , t_{KR} and t_{DB} and suppose we have a knowledge base K containing the following two rules:

$$t_{ML} \rightarrow t_{AI} \tag{3}$$

$$t_{KR} \rightarrow t_{AI} \tag{4}$$

In other words, if an article is about Machine Learning, then it is also about Artificial Intelligence, and the same holds for Knowledge Representation. This knowledge base has the following models (where we denote an interpretation by the set of atoms it makes true):

$$\begin{aligned} \omega_1 = \{\} & \quad \omega_2 = \{t_{AI}\} & \quad \omega_3 = \{t_{AI}, t_{ML}\} & \quad \omega_4 = \{t_{AI}, t_{KR}\} & \quad \omega_5 = \{t_{AI}, t_{ML}, t_{KR}\} \\ \omega_6 = \{t_{DB}\} & \quad \omega_7 = \{t_{AI}, t_{DB}\} & \quad \omega_8 = \{t_{AI}, t_{ML}, t_{DB}\} & \quad \omega_9 = \{t_{AI}, t_{KR}, t_{DB}\} & \quad \omega_{10} = \{t_{AI}, t_{ML}, t_{KR}, t_{DB}\} \end{aligned}$$

Accordingly, we can represent nodes using 10-dimensional embeddings. The input embeddings of the topic nodes would then be given as:

$$\begin{aligned} \mathbf{t}_{AI}^{(0)} &= (0, -1, -1, -1, -1, 0, -1, -1, -1, -1) \\ \mathbf{t}_{ML}^{(0)} &= (0, 0, -1, 0, -1, 0, 0, -1, 0, -1) \\ \mathbf{t}_{KR}^{(0)} &= (0, 0, 0, -1, -1, 0, 0, 0, -1, -1) \\ \mathbf{t}_{DB}^{(0)} &= (0, 0, 0, 0, 0, -1, -1, -1, -1, -1) \end{aligned}$$

Note, for instance, how $\max(\mathbf{t}_{AI}^{(0)}, \mathbf{t}_{ML}^{(0)}) = \mathbf{t}_{ML}^{(0)}$. As a result, an article node that is connected to both t_{AI} and t_{ML} would have the same embedding as an article node that is only connected to t_{ML} , which reflects the fact that the rule $t_{ML} \rightarrow t_{AI}$ has been taken into account as background knowledge.

2.4. Non-monotonic reasoning

Our focus on accumulating knowledge, and the use of the union in the formulation of the epistemic pooling principle, may suggest that the problem setting we consider in this paper is only suitable for monotonic reasoning. In contrast, many applications require some kind of non-monotonic inference, where we can defeasibly infer something to be true based on the absence of evidence to the contrary. To illustrate this idea, let us again consider the topics t_{AI} , t_{ML} , t_{KR} and t_{DB} and the background knowledge expressed in (3)-(4). Let us now additionally assume that we want to implement the following behaviour:

- If we know that the topic of an article is AI, and we have no evidence that the article is about KR, then we will assume that the topic is ML (i.e. “AI articles are typically about ML”).
- However, if we know that the article is about AI and KR, then this inference should not be made.

This intuition cannot be implemented using standard propositional entailment, given the non-monotonic nature of the desired inferences: we cannot at the same time have $K \cup \{t_{AI}\} \models t_{ML}$ and $K \cup \{t_{AI}, t_{KR}\} \not\models t_{ML}$ for any knowledge base K . The standard solution is to assume that each possible world ω has some plausibility degree $\pi(\omega)$ and to interpret “if α then typically β ” as the constraint that $\max\{\pi(\omega) \mid \omega \models \alpha \wedge \beta\} > \max\{\pi(\omega) \mid \omega \models \alpha \wedge \neg\beta\}$ [14–17]. In other words, such a default rule is satisfied if β is true in the most plausible models of α . Using possibilistic logic, we can conveniently encode default knowledge using weighted logical formulas [16]. In particular, the available knowledge can be encoded as follows:

$$\begin{aligned}
 1 &: & t_{ML} &\rightarrow t_{AI} \\
 1 &: & t_{KR} &\rightarrow t_{AI} \\
 0.5 &: & t_{KR} &\rightarrow \perp \\
 0.5 &: & t_{AI} &\rightarrow t_{ML}
 \end{aligned}$$

Let K_λ be the set of formulas whose weight is at least λ . To perform defeasible reasoning in possibilistic logic, we proceed as follows. Let the available evidence be encoded as a propositional formula α . We first find the lowest $\lambda \geq 0$ for which $K_\lambda \cup \{\alpha\}$ is logically consistent. Then we say that β is defeasibly entailed from the evidence α if it is classically entailed from $K_\lambda \cup \{\alpha\}$. It can be straightforwardly verified that the weighted formulas above capture the intended default knowledge. For instance, if $\alpha = t_{AI}$, we have $\lambda = 0$. Since K_0 contains the rule $t_{AI} \rightarrow t_{ML}$ we can defeasibly infer t_{ML} . However, if the available evidence is $t_{AI} \wedge t_{KR}$, then we find $\lambda = 1$ and we can no longer infer t_{ML} . Note how the rule $t_{KR} \rightarrow \perp$ intuitively serves the purpose of blocking the rule $t_{AI} \rightarrow t_{ML}$ in situations where t_{KR} is known to be true.

Let $\omega_1, \dots, \omega_{10}$ be defined as before and let us consider, for this example, three plausibility levels. The main intuition is as follows: the higher the value of a given coordinate in the embeddings, the more strongly we can exclude the corresponding possible world. In particular, if the available evidence is violated in a possible world, then its corresponding coordinate is set to 1. Otherwise, if a formula with a weight of 0.5 is violated, the corresponding coordinate is set to 0.5. Note that the formulas of weight 1 are taken into account implicitly, as we only consider the models of these formulas as our 10 possible worlds. The remaining coordinates are set to 0. The initial embeddings of the topic nodes are then defined as follows:

$$\begin{aligned}
 \mathbf{t}_{AI}^{(0)} &= (1, 0.5, 0, 0.5, 0.5, 1, 0.5, 0, 0.5, 0.5) \\
 \mathbf{t}_{ML}^{(0)} &= (1, 1, 0, 1, 0.5, 1, 1, 0, 1, 0.5) \\
 \mathbf{t}_{KR}^{(0)} &= (1, 1, 1, 0.5, 0.5, 1, 1, 1, 0.5, 0.5) \\
 \mathbf{t}_{DB}^{(0)} &= (1, 1, 1, 1, 1, 0, 0.5, 0, 0.5, 0.5)
 \end{aligned}$$

To test whether a propositional formula α can be defeasibly inferred in the epistemic state encoded by a vector $\mathbf{x} = (x_1, \dots, x_n)$, we need to check whether $\min\{x_i \mid \omega_i \models \alpha\} < \min\{x_i \mid \omega_i \models \neg\alpha\}$. Indeed, if this inequality holds, there is a model of α which is strictly more plausible than any of the models of $\neg\alpha$. For instance, going back to the example, let a_1 be an article node which is only connected to the topic node t_{AI} . The embedding of this article $\mathbf{a}_1^{(1)}$ is given by $(1, 0.5, 0, 0.5, 0.5, 1, 0.5, 0, 0.5, 0.5)$. Note that t_{ML} is satisfied in worlds $\omega_3, \omega_5, \omega_8, \omega_{10}$. We find that:

$$\begin{aligned}
 \min\{x_i \mid \omega_i \models t_{ML}\} &= \min(x_3, x_5, x_8, x_{10}) = 0 \\
 \min\{x_i \mid \omega_i \models \neg t_{ML}\} &= \min(x_1, x_2, x_4, x_6, x_7, x_9) = 0.5
 \end{aligned}$$

Hence we can indeed defeasibly infer that the article is about Machine Learning. Let a_2 be an article which connected to both t_{AI} and t_{KR} . The embedding $\mathbf{a}_2^{(1)}$ is then given by

$$\begin{aligned}
 \mathbf{a}_2^{(1)} &= \max((1, 0.5, 0, 0.5, 0.5, 1, 0.5, 0, 0.5, 0.5), (1, 1, 1, 0.5, 0.5, 1, 1, 1, 0.5, 0.5)) \\
 &= (1, 1, 1, 0.5, 0.5, 1, 1, 1, 0.5, 0.5)
 \end{aligned}$$

We then find:

$$\begin{aligned}
 \min\{x_i \mid \omega_i \models t_{ML}\} &= \min(x_3, x_5, x_8, x_{10}) = 0.5 \\
 \min\{x_i \mid \omega_i \models \neg t_{ML}\} &= \min(x_1, x_2, x_4, x_6, x_7, x_9) = 0.5
 \end{aligned}$$

It is thus not possible to defeasibly infer that a_2 is about Machine Learning. As this example illustrates, non-monotonic reasoning essentially requires that we can model weighted epistemic states. This will be studied in Section 6, where we will see that max-pooling is uniquely suitable for this setting.

3. Problem setting

In this section, we introduce our considered problem setting more formally. We assume that epistemic states are represented using sets of elementary properties. Let us write \mathcal{P} for the set of all these properties. An epistemic state \mathcal{Q} then simply corresponds to a subset of \mathcal{P} . Intuitively, we think of these elementary properties as atomic pieces of evidence. The properties in an epistemic state \mathcal{Q} then correspond to the evidence that is available, whereas the properties in $\mathcal{P} \setminus \mathcal{Q}$ correspond to evidence that has not been encountered. For instance, in the context of image processing, we can think of the properties from \mathcal{P} as elementary visual features, whose presence may be detected in an image. In applications where more intricate forms of reasoning are needed than simply aggregating sets of detected features, we can relate the properties in

\mathcal{P} to possible worlds, as we have seen in Section 2.2. For each possible world ω , we then consider a property p_ω , corresponding to the knowledge that ω can be excluded, i.e. that ω is not a model of the world. In this way, subsets of \mathcal{P} can be used to represent arbitrary propositional knowledge bases. This link with logical reasoning will be developed in Section 5. For now, however, it will suffice to simply think of epistemic states as subsets of \mathcal{P} .

Taking the view that embeddings encode epistemic states, each $\mathbf{e} \in \mathbb{R}^n$ will be associated with a set of properties from \mathcal{P} . Formally, we assume that a scoring function $\gamma_p : \mathbb{R}^n \rightarrow \mathbb{R}$ is available for each property $p \in \mathcal{P}$. We consider two variants of our setting, which differ in whether strict or weak inequalities are used to determine which properties are satisfied. As we will see, this choice has a material impact on the theoretical properties of the resulting framework.

Strict semantics Under the strict semantics, we say that an embedding $\mathbf{e} \in \mathbb{R}^n$ satisfies the property $p \in \mathcal{P}$ if $\gamma_p(\mathbf{e}) > 0$. Let us write $\Gamma(\mathbf{e})$ for the epistemic state encoded by \mathbf{e} , i.e. the set of properties satisfied by \mathbf{e} :

$$\Gamma(\mathbf{e}) = \{p \in \mathcal{P} \mid \gamma_p(\mathbf{e}) > 0\} \quad (5)$$

Let $\diamond : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ represent a pooling operator. The epistemic pooling principle can then be formalised as follows:

$$\Gamma(\mathbf{e} \diamond \mathbf{f}) = \Gamma(\mathbf{e}) \cup \Gamma(\mathbf{f}) \quad (6)$$

If we want to specify that the strict semantics is used, we will also refer to (6) as the *strict epistemic pooling principle*. Intuitively, the embeddings \mathbf{e} and \mathbf{f} capture information coming from two different sources, e.g. two different regions of an image or two different modalities. The principle captured by (6) is that the pooling operator \diamond should merely combine this information: the total evidence that is available is the union of the evidence provided by the two sources. Note that (6) is equivalent to:

$$\forall p \in \mathcal{P}. (\gamma_p(\mathbf{e}) > 0) \vee (\gamma_p(\mathbf{f}) > 0) \Leftrightarrow \gamma_p(\mathbf{e} \diamond \mathbf{f}) > 0 \quad (7)$$

In the following, we will assume that all embeddings are taken from some set $X \subseteq \mathbb{R}^n$. One possibility would be to choose $X = \mathbb{R}^n$, but as we will see, it is sometimes necessary to make a more restrictive choice. For instance, we may have $X = [0, +\infty]^n$ if we want to restrict the discussion to vectors with non-negative coordinates. Regardless of how X is chosen, an important consideration is that the embeddings in X should allow us to capture every possible epistemic state, in the following sense:

$$\forall Q \subseteq \mathcal{P}. \exists \mathbf{e} \in X. \Gamma(\mathbf{e}) = Q \quad (8)$$

If (8) is satisfied, we say that X satisfies *exhaustiveness*. Finally, for the ease of presentation, we introduce the following notations:

$$\text{Pos}_p = \{\mathbf{e} \in X \mid \gamma_p(\mathbf{e}) > 0\}$$

$$\text{Neg}_p = \{\mathbf{e} \in X \mid \gamma_p(\mathbf{e}) \leq 0\} = X \setminus \text{Pos}_p$$

We will refer to Pos_p and Neg_p as the positive and negative regions for property p . Indeed, we have that $\mathbf{e} \in \text{Pos}_p$ iff $p \in \Gamma(\mathbf{e})$, and $\mathbf{e} \in \text{Neg}_p$ otherwise.

Weak epistemic pooling principle Under the weak semantics, we say that an embedding $\mathbf{e} \in \mathbb{R}^n$ satisfies the property $p \in \mathcal{P}$ if $\gamma_p(\mathbf{e}) \geq 0$. Epistemic states are then determined as follows:

$$\Gamma'(\mathbf{e}) = \{p \in \mathcal{P} \mid \gamma_p(\mathbf{e}) \geq 0\} \quad (9)$$

This definition gives rise to the following counterpart of (6), which we will refer to as the *weak epistemic pooling principle*:

$$\Gamma'(\mathbf{e} \diamond \mathbf{f}) = \Gamma'(\mathbf{e}) \cup \Gamma'(\mathbf{f}) \quad (10)$$

We will furthermore require that *exhaustiveness* is satisfied:

$$\forall Q \subseteq \mathcal{P}. \exists \mathbf{e} \in X. \Gamma'(\mathbf{e}) = Q \quad (11)$$

Finally, the positive and negative regions are also defined analogously as before:

$$\text{Pos}'_p = \{\mathbf{e} \in X \mid \gamma_p(\mathbf{e}) \geq 0\}$$

$$\text{Neg}'_p = \{\mathbf{e} \in X \mid \gamma_p(\mathbf{e}) < 0\} = X \setminus \text{Pos}'_p$$

Pooling operators Whether the (strict or weak) epistemic pooling principle can be satisfied for every $\mathbf{e}, \mathbf{f} \in X$ depends on the choice of the scoring functions γ_p , the set $X \subseteq \mathbb{R}^n$ and the pooling operator \diamond . In our analysis, we will focus on the following standard pooling operators:

$$\begin{aligned} \textbf{Average:} \quad & (e_1, \dots, e_n) \diamond_{\text{avg}} (f_1, \dots, f_n) = \frac{\mathbf{e} + \mathbf{f}}{2} \\ \textbf{Summation:} \quad & (e_1, \dots, e_n) \diamond_{\text{sum}} (f_1, \dots, f_n) = \mathbf{e} + \mathbf{f} \\ \textbf{Max-pooling:} \quad & (e_1, \dots, e_n) \diamond_{\text{max}} (f_1, \dots, f_n) = (\max(e_1, f_1), \dots, \max(e_n, f_n)) \\ \textbf{Hadamard:} \quad & (e_1, \dots, e_n) \diamond_{\text{had}} (f_1, \dots, f_n) = (e_1 \cdot f_1, \dots, e_n \cdot f_n) \end{aligned}$$

Note that the epistemic pooling principles in (6) and (10) are defined w.r.t. two arguments. This focus on binary pooling operators simplifies the formulation, while any negative results we obtain naturally carry over to pooling operators with more arguments. Moreover, most of the considered pooling operators are associative, with the exception of \diamond_{avg} . Furthermore, even though \diamond_{avg} itself is not associative, if it satisfies (6) or (10), its effect on the epistemic states encoded by the embeddings will nonetheless be associative, given that we have e.g. $\Gamma(\mathbf{e}_1 \diamond (\mathbf{e}_2 \diamond \mathbf{e}_3)) = \Gamma((\mathbf{e}_1 \diamond \mathbf{e}_2) \diamond \mathbf{e}_3) = \Gamma(\mathbf{e}_1) \cup \Gamma(\mathbf{e}_2) \cup \Gamma(\mathbf{e}_3)$, due to the associativity of the union. We now illustrate the key concepts with a simple example.

Example 1. Let $\mathcal{P} = \{a, b\}$ and suppose embedding are taken from \mathbb{R}^2 . Let the scoring functions γ_a and γ_b be defined as follows:

$$\begin{aligned} \gamma_a(x_1, x_2) &= 1 - d((x_1, x_2), (0, 0)) = 1 - \sqrt{x_1^2 + x_2^2} \\ \gamma_b(x_1, x_2) &= 1 - d((x_1, x_2), (1, 1)) = 1 - \sqrt{(1 - x_1)^2 + (1 - x_2)^2} \end{aligned}$$

Now let $\mathbf{e} = (\frac{1}{4}, 0)$ and $\mathbf{f} = (\frac{3}{4}, 1)$. Then we have $\mathbf{e} \diamond_{\text{avg}} \mathbf{f} = (\frac{1}{2}, \frac{1}{2})$. We find:

$$\begin{aligned} \gamma_a(\mathbf{e}) &= \frac{3}{4} & \gamma_a(\mathbf{f}) &= -\frac{1}{4} & \gamma_a(\mathbf{e} \diamond_{\text{avg}} \mathbf{f}) &= 1 - \frac{1}{\sqrt{2}} \\ \gamma_b(\mathbf{e}) &= -\frac{1}{4} & \gamma_b(\mathbf{f}) &= \frac{3}{4} & \gamma_b(\mathbf{e} \diamond_{\text{avg}} \mathbf{f}) &= 1 - \frac{1}{\sqrt{2}} \end{aligned}$$

and thus

$$\Gamma(\mathbf{e}) = \{a\} \quad \Gamma(\mathbf{f}) = \{b\} \quad \Gamma(\mathbf{e} \diamond_{\text{avg}} \mathbf{f}) = \{a, b\}$$

This means that the epistemic pooling principle (6) is satisfied for \mathbf{e} and \mathbf{f} . On the other hand, for $\mathbf{g} = (10, 10)$, we have $\Gamma(\mathbf{g}) = \Gamma(\mathbf{e} \diamond_{\text{avg}} \mathbf{g}) = \emptyset$, hence the epistemic pooling principle is not satisfied for \mathbf{e} and \mathbf{g} .

Notations Throughout this paper, we write $\delta(A)$ for the boundary of a set $A \subseteq \mathbb{R}^n$. Similarly, we will write $\text{int}(A)$ and $\text{cl}(A)$ for the interior and closure:

$$\begin{aligned} \text{int}(A) &= \{\mathbf{e} \in A \mid \exists \varepsilon > 0. \forall \mathbf{f} \in \mathbb{R}^n. d(\mathbf{e}, \mathbf{f}) < \varepsilon \Rightarrow \mathbf{f} \in A\} \\ \text{cl}(A) &= \{\mathbf{e} \in \mathbb{R}^n \mid \forall \varepsilon > 0. \exists \mathbf{f} \in A. d(\mathbf{e}, \mathbf{f}) < \varepsilon\} \\ \delta(A) &= \text{cl}(A) \setminus \text{int}(A) \end{aligned}$$

4. Realizability of the epistemic pooling principle

In this section we study, for each of the considered pooling operators, whether they can satisfy the strict and weak epistemic pooling principles, and if so, under which conditions this is the case. In all cases, we find that the epistemic pooling principles can only be satisfied if $n \geq |\mathcal{P}|$, with n the dimensionality of the embeddings. For \diamond_{avg} and \diamond_{sum} , we also have to make assumptions on the set X , i.e. the epistemic pooling principles cannot be satisfied for $X = \mathbb{R}^n$ with these pooling operators. We furthermore find that the epistemic pooling principles can only be satisfied if γ_p satisfies some particular conditions. Most significantly, we find that \diamond_{avg} and \diamond_{sum} cannot satisfy the weak epistemic pooling principle with continuous scoring functions γ_p , and that \diamond_{had} cannot satisfy the strict epistemic pooling principle with continuous scoring functions. These results are summarised in Table 1.

Table 1
Summary of results about the realizability of the epistemic pooling principles.

Pooling Operator	Semantics	$X = \mathbb{R}^n$ possible?	Continuous γ_p possible?
Average	Strict	✗	✓
	Weak	✗	✗
Summation	Strict	✗	✓
	Weak	✗	✗
Max-pooling	Strict	✓	✓
	Weak	✓	✓
Hadamard	Strict	✓	✗
	Weak	✓	✓

4.1. Average

Strict semantics The first question we look at is whether the strict epistemic pooling principle (6) can be satisfied for all $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$. The following result shows that this is only possible in the trivial case where every embedding encodes the same epistemic state.

Proposition 1. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, with $\diamond = \diamond_{\text{avg}}$. For any given $p \in \mathcal{P}$ we have*

$$(\forall \mathbf{e} \in \mathbb{R}^n . p \in \Gamma(\mathbf{e})) \vee (\forall \mathbf{e} \in \mathbb{R}^n . p \notin \Gamma(\mathbf{e}))$$

Proof. Suppose there exists some $\mathbf{e} \in \mathbb{R}^n$ such that $p \in \Gamma(\mathbf{e})$. We show that we then have $p \in \Gamma(\mathbf{f})$ for every $\mathbf{f} \in \mathbb{R}^n$. Noting that $\mathbf{f} = \mathbf{e} \diamond_{\text{avg}} (2\mathbf{f} - \mathbf{e})$, we know from (6) that $\Gamma(\mathbf{f}) = \Gamma(\mathbf{e}) \cup \Gamma(2\mathbf{f} - \mathbf{e})$, and thus in particular that $p \in \Gamma(\mathbf{f})$. \square

We will thus have to assume that embeddings are restricted to some subset $X \subset \mathbb{R}^n$. To ensure that X is closed under the pooling operator \diamond_{avg} , we will assume that X is convex. The following three lemmas explain how (6) constrains the scoring functions γ_p .

Lemma 1. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{avg}}$ and X convex. Suppose there exists some $\mathbf{e} \in X$ such that $p \in \Gamma(\mathbf{e})$. It holds that $\text{Neg}_p \subseteq \delta(X)$.*

Proof. Suppose $\mathbf{f} \in \text{int}(X)$. We show that $\mathbf{f} \in \text{Pos}_p$. Let us define ($\lambda \in \mathbb{R}$):

$$\mathbf{x}_\lambda = \lambda \mathbf{e} + (1 - \lambda) \mathbf{f}$$

Note that because we assumed that X is convex, it holds that $\mathbf{x}_\lambda \in X$ for all $\lambda \in [0, 1]$. We have $p \in \Gamma(\mathbf{x}_1)$, as $\mathbf{x}_1 = \mathbf{e}$. By repeatedly applying (6) we find that p belongs to $\Gamma(\mathbf{x}_{\frac{1}{2}})$, $\Gamma(\mathbf{x}_{\frac{1}{4}})$, $\Gamma(\mathbf{x}_{\frac{1}{8}})$, etc. In the limit, we find that for every $\lambda \in]0, 1]$ it holds that $p \in \Gamma(\mathbf{x}_\lambda)$. Since $\mathbf{f} \in \text{int}(X)$, there exists some $0 < \varepsilon \leq 1$ such that $\mathbf{x}_{-\varepsilon} \in X$. Using (6), we find $\Gamma(\mathbf{f}) = \Gamma(\mathbf{x}_{-\varepsilon}) \cup \Gamma(\mathbf{x}_\varepsilon)$. Since we established $p \in \Gamma(\mathbf{x}_\varepsilon)$ we thus find $p \in \Gamma(\mathbf{f})$, which means $\mathbf{f} \in \text{Pos}_p$. \square

Corollary 1. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{avg}}$ and X convex. Let $p \in \mathcal{P}$. It holds that $\dim(\text{Neg}_p) \leq n - 1$*

Lemma 2. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{avg}}$ and X convex. Let $p \in \mathcal{P}$. It holds that Pos_p is convex.*

Proof. Suppose $\mathbf{e}, \mathbf{f} \in \text{Pos}_p$ and define ($\lambda \in [0, 1]$):

$$\mathbf{x}_\lambda = \lambda \mathbf{e} + (1 - \lambda) \mathbf{f}$$

We show that $\mathbf{x}_\lambda \in \text{Pos}_p$ for every $\lambda \in]0, 1[$. By applying (6) to $\mathbf{f} = \mathbf{x}_0$ and $\mathbf{e} = \mathbf{x}_1$ we find that $\mathbf{x}_{\frac{1}{2}} \in \text{Pos}_p$. By applying (6) to \mathbf{x}_0 and $\mathbf{x}_{\frac{1}{2}}$, we find $\mathbf{x}_{\frac{1}{4}} \in \text{Pos}_p$. Similarly, by applying (6) to $\mathbf{x}_{\frac{1}{2}}$ and \mathbf{x}_1 , we find $\mathbf{x}_{\frac{3}{4}} \in \text{Pos}_p$. Continuing in this way, we find $\mathbf{x}_\lambda \in \text{Pos}_p$ for every λ of the form $\frac{j}{2^i}$ with $i \in \mathbb{N}$ and $j \in \{0, 1, \dots, 2^i\}$. Now let $\lambda \in]0, 1[$. We can approximate λ arbitrary well using a value of the form $\frac{j}{2^i}$. In particular, we can always find some $i \in \mathbb{N}$ and $j \in \{0, \dots, 2^i\}$ such that $0 < \frac{j}{2^i} < \lambda$ and $\lambda < 2\lambda - \frac{j}{2^i} < 1$. By (6), we have:

$$\Gamma(\mathbf{x}_\lambda) = \Gamma\left(\mathbf{x}_{\frac{j}{2^i}}\right) \cup \Gamma\left(\mathbf{x}_{2\lambda - \frac{j}{2^i}}\right)$$

Since we already know that $\mathbf{x}_{\frac{j}{2^i}} \in \text{Pos}_p$ we thus also find $\mathbf{x}_\lambda \in \text{Pos}_p$. \square

Lemma 3. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{avg}}$ and X convex. Let $p \in \mathcal{P}$. It holds that Neg_p is convex.*

Proof. Suppose $\mathbf{e}, \mathbf{f} \in \text{Neg}_p$ and define ($\lambda \in [0, 1]$):

$$\mathbf{x}_\lambda = \lambda \mathbf{e} + (1 - \lambda) \mathbf{f}$$

Entirely analogously as in the proof of Lemma 2, we find that $\mathbf{x}_\lambda \in \text{Neg}_p$ for every λ of the form $\frac{j}{2^i}$, with $i \in \mathbb{N}$ and $j \in \{0, \dots, 2^i\}$. Let $\lambda \in]0, 1[$. Let us consider values $\varepsilon_1, \varepsilon_2$ satisfying $0 < \varepsilon_1 < \varepsilon_2 < \lambda < 2\lambda - \varepsilon_2 < 2\lambda - \varepsilon_1 < 1$. Suppose $\mathbf{x}_\lambda \notin \text{Neg}_p$. Then from (6) we know that:

- $\mathbf{x}_{\varepsilon_1} \in \text{Pos}_p$ or $\mathbf{x}_{2\lambda - \varepsilon_1} \in \text{Pos}_p$ needs to hold; and
- $\mathbf{x}_{\varepsilon_2} \in \text{Pos}_p$ or $\mathbf{x}_{2\lambda - \varepsilon_2} \in \text{Pos}_p$ needs to hold.

This means that there are at least two distinct values $\lambda_1, \lambda_2 \in \{\varepsilon_1, \varepsilon_2, 2\lambda - \varepsilon_1, 2\lambda - \varepsilon_2\}$ such that $\mathbf{x}_{\lambda_1} \in \text{Pos}_p$ and $\mathbf{x}_{\lambda_2} \in \text{Pos}_p$. Let us assume w.l.o.g. that $\lambda_1 < \lambda_2$. We can always find some $i \in \mathbb{N}$ and $j \in \{0, \dots, 2^i\}$ such that $\lambda_1 < \frac{j}{2^i} < \lambda_2$. From the preceding discussion we already know that $\mathbf{x}_{\frac{j}{2^i}} \in \text{Neg}_p$. However, from $\mathbf{x}_{\lambda_1} \in \text{Pos}_p$ and $\mathbf{x}_{\lambda_2} \in \text{Pos}_p$, using Lemma 2 we find $\mathbf{x}_{\frac{j}{2^i}} \in \text{Pos}_p$, a contradiction. It follows that $\mathbf{x}_\lambda \in \text{Neg}_p$. \square

From Lemmas 1, 2 and 3, we obtain the following corollary using the hyperplane separation theorem.

Corollary 2. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{avg}}$ and X convex. For any $p \in \mathcal{P}$, there exists a hyperplane H_p such that $\text{Neg}_p \subseteq \delta(X) \cap H_p$.*

The next proposition reveals that the dimensionality of the embeddings needs to be at least $|\mathcal{P}|$ if we want the epistemic pooling principle to be satisfied and at the same time ensure that every epistemic state is modelled by some vector.

Proposition 2. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{avg}}$ and $X \subseteq \mathbb{R}^n$ convex. Suppose furthermore that exhaustiveness (8) is satisfied. It holds that $n \geq |\mathcal{P}|$.*

Proof. Let $p_1, \dots, p_{|\mathcal{P}|}$ be an enumeration of the properties in \mathcal{P} . Note that because of (8), we have that $\text{Neg}_{p_1} \neq \emptyset$ and $\text{Pos}_{p_1} \neq \emptyset$. Moreover, from Lemmas 2 and 3 we know that these regions are both convex. It follows from the hyperplane separation theorem that there exists a hyperplane H_1 which separates Neg_{p_1} and Pos_{p_1} . From Lemma 1, we furthermore know that $\text{Neg}_{p_1} \subseteq \text{cl}(\text{Pos}_{p_1})$, which implies that $\text{Neg}_{p_1} \subseteq H_1$.

Note that $H_1 \cap \text{Neg}_{p_2}$ and $H_1 \cap \text{Pos}_{p_2}$ are convex regions. Moreover, since $\text{Neg}_{p_1} \subseteq H_1$, we find from (8) that $H_1 \cap \text{Neg}_{p_2} \neq \emptyset$ and $H_1 \cap \text{Pos}_{p_2} \neq \emptyset$. It follows from the hyperplane separation theorem that there exists some hyperplane H_2 separating $H_1 \cap \text{Neg}_{p_2}$ and $H_1 \cap \text{Pos}_{p_2}$. Moreover, it holds that $H_1 \cap \text{Neg}_{p_2} \subseteq H_2$. Indeed, suppose there was some $\mathbf{e} \in (H_1 \cap \text{Neg}_{p_2}) \setminus H_2$ and let $\mathbf{f} \in H_1 \cap \text{Pos}_{p_2}$. For $i \in \mathbb{N} \setminus \{0\}$ we define $\mathbf{f}_i = \mathbf{e} \diamond_{\text{avg}} \mathbf{f}_{i-1}$, with $\mathbf{f}_0 = \mathbf{f}$. Then there must be some $i \in \mathbb{N} \setminus \{0\}$ such that \mathbf{f}_i is on the same side of hyperplane H_2 as \mathbf{e} , which implies $\mathbf{f}_i \in \text{Neg}_{p_2}$ since H_2 was chosen as a separating hyperplane. However, using (6) we also find that $p_2 \in \Gamma(\mathbf{f}_i)$ and thus $\mathbf{f}_i \in \text{Pos}_{p_2}$, which is a contradiction. This means that $H_1 \cap \text{Neg}_{p_2} \subseteq H_2$ and thus in particular also that $\text{Neg}_{p_1} \cap \text{Neg}_{p_2} \subseteq H_2$.

We can repeat this argument as follows. For each $j \in \{1, \dots, |\mathcal{P}|\}$, we let H_j be a hyperplane separating Pos_{p_j} and Neg_{p_j} . Let $i \in \{2, \dots, |\mathcal{P}| - 1\}$ and suppose we have already shown for every $j \leq i$ that $H_1 \cap \dots \cap H_{j-1} \cap \text{Neg}_{p_j} \subseteq H_j$. Note that this in particular implies $\text{Neg}_{p_1} \cap \dots \cap \text{Neg}_{p_i} \subseteq H_i$. We know from (8) that $\text{Neg}_{p_1} \cap \dots \cap \text{Neg}_{p_i} \cap \text{Neg}_{p_{i+1}} \neq \emptyset$ and $\text{Neg}_{p_1} \cap \dots \cap \text{Neg}_{p_i} \cap \text{Pos}_{p_{i+1}} \neq \emptyset$. It follows that $H_1 \cap \dots \cap H_i \cap \text{Neg}_{p_{i+1}} \neq \emptyset$ and $H_1 \cap \dots \cap H_i \cap \text{Pos}_{p_{i+1}} \neq \emptyset$. This implies that $H_1 \cap \dots \cap H_i \not\subseteq H_{i+1}$ and thus that $\dim(H_1 \cap \dots \cap H_{i+1}) = \dim(H_1 \cap \dots \cap H_i) - 1$. Similar as for the base case above, we furthermore find that $H_1 \cap \dots \cap H_i \cap \text{Neg}_{p_{i+1}} \subseteq H_{i+1}$.

It follows that $\dim(H_1 \cap \dots \cap H_{|\mathcal{P}|}) = n - |\mathcal{P}|$, which is only possible if $n \geq |\mathcal{P}|$. \square

It is easy to see that the bound from this proposition cannot be strengthened, i.e. that it is possible to satisfy (6) and (8) while $n = |\mathcal{P}|$. One possible construction is as follows. Let p_1, \dots, p_n be an enumeration of the properties in \mathcal{P} . We define:

$$X = \{(x_1, \dots, x_n) \mid x_1, \dots, x_n \geq 0\}$$

and for $i \in \{1, \dots, n\}$ we define:

$$\gamma_{p_i}(e_1, \dots, e_n) = e_i$$




Fig. 2. Construction satisfying (6) and (8) for \diamond_{avg} .

To see why this choice satisfies (8), let $\mathcal{Q} \subseteq \mathcal{P}$ and define $\mathbf{e} = (e_1, \dots, e_n)$ as follows:

$$e_i = \begin{cases} 1 & \text{if } p_i \in \mathcal{Q} \\ 0 & \text{otherwise} \end{cases}$$

Then it is straightforward to verify that $\Gamma(\mathbf{e}) = \mathcal{Q}$. Moreover, it is also clear that (6) is satisfied. Indeed, the i^{th} coordinate of $(e_1, \dots, e_n) \diamond_{\text{avg}} (f_1, \dots, f_n)$ is $\frac{e_i + f_i}{2}$. Hence we have $\frac{e_i + f_i}{2} > 0$ as soon as $e_i > 0$ or $f_i > 0$ (noting that $e_i, f_i \geq 0$), meaning $p_i \in \Gamma((e_1, \dots, e_n) \diamond_{\text{avg}} (f_1, \dots, f_n))$ iff $p_i \in \Gamma(e_1, \dots, e_n) \cup \Gamma(f_1, \dots, f_n)$. A visualisation of this construction is shown in Fig. 2. Note how the construction aligns with the common practice of learning sparse high-dimensional embeddings with non-negative coordinates.

Weak semantics Let us now consider whether the weak epistemic pooling principle can also be satisfied for \diamond_{avg} . Without any restrictions on the scoring functions γ_p , this is clearly the case. In particular, suppose for each $p \in \mathcal{P}$, a function γ_p is defined such that the strict epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, for some convex set X , and suppose furthermore that exhaustiveness (8) is satisfied. Then we can define modified scoring functions as follows:

$$\gamma'_p(\mathbf{e}) = \begin{cases} 1 & \text{if } \gamma_p(\mathbf{e}) > 0 \\ -1 & \text{otherwise} \end{cases} \quad (12)$$

In particular, we have $\gamma_p(\mathbf{e}) > 0$ iff $\gamma'_p(\mathbf{e}) \geq 0$. The fact that the strict epistemic pooling principle is satisfied for the scoring functions γ_p thus implies that the weak epistemic pooling principle is satisfied for the modified scoring functions γ'_p , while we can still model every epistemic state, i.e. exhaustiveness (11) is also satisfied.

The discontinuous nature of the scoring function γ'_p defined in (12) is clearly undesirable in practice. Hence the question arises whether it is possible to satisfy the weak epistemic pooling principle when only continuous scoring functions can be used. The answer to this question is negative. In particular, as the following result shows, if the weak epistemic pooling principle is satisfied with continuous scoring functions, all embeddings encode the same epistemic state.

Proposition 3. *Suppose the epistemic pooling principle (10) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, for $\diamond = \diamond_{\text{avg}}$ and X convex. Suppose γ_p is continuous. It holds that either $\text{Neg}'_p = \emptyset$ or $\text{Pos}'_p = \emptyset$.*

Proof. Suppose there exists some $\mathbf{e} \in \text{Neg}'_p$ and $\mathbf{f} \in \text{Pos}'_p$. We then have $\gamma_p(\mathbf{e}) < 0$ and $\gamma_p(\mathbf{f}) \geq 0$. For $\lambda \in [0, 1]$ we define:

$$\mathbf{x}_\lambda = \lambda \mathbf{f} + (1 - \lambda) \mathbf{e}$$

By repeatedly applying (10), we find that $\gamma_p(\mathbf{x}_{\frac{1}{2^i}}) \geq 0$ for every $i \in \mathbb{N}$. In particular, this means that for every $\varepsilon > 0$ there exists some $\mathbf{x} \in X$ such that $d(\mathbf{e}, \mathbf{x}) < \varepsilon$ and $\gamma_p(\mathbf{x}) \geq 0$. If γ_p is continuous this implies $\gamma_p(\mathbf{e}) = 0$, which is a contradiction since we assumed $\mathbf{e} \in \text{Neg}'_p$. \square

4.2. Summation

Strict semantics We first show that the epistemic pooling principle (6) cannot be satisfied (in a non-trivial way) for all $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$ when using \diamond_{sum} , as we also found for \diamond_{avg} in Proposition 1.

Proposition 4. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, with $\diamond = \diamond_{\text{sum}}$. For any given $p \in \mathcal{P}$ we have*

$$(\forall \mathbf{e} \in \mathbb{R}^n . p \in \Gamma(\mathbf{e})) \vee (\forall \mathbf{e} \in \mathbb{R}^n . p \notin \Gamma(\mathbf{e}))$$

Proof. Suppose $p \in \Gamma(\mathbf{e})$ and let $\mathbf{f} \in \mathbb{R}^n$. Using (6), we find

$$\Gamma(\mathbf{f}) = \Gamma(\mathbf{e} \diamond_{\text{sum}} (\mathbf{f} - \mathbf{e})) = \Gamma(\mathbf{e}) \cup \Gamma(\mathbf{f} - \mathbf{e})$$

and thus $p \in \Gamma(\mathbf{f})$. \square

We thus again need to define a suitable subset $X \subseteq \mathbb{R}^n$. To ensure that X is closed under \diamond_{sum} it is not sufficient that X is convex. For this reason, we will assume that X is conically closed, in particular:

$$\forall \mathbf{e}, \mathbf{f} \in X . \forall \alpha, \beta \geq 0 . \alpha \mathbf{e} + \beta \mathbf{f} \in X \tag{13}$$

We now show that whenever (6) is satisfied for \diamond_{sum} , for all $\mathbf{e}, \mathbf{f} \in X$, it also holds that (6) is satisfied for \diamond_{avg} , meaning that the results we have established for \diamond_{avg} carry over to \diamond_{sum} . We first show the following lemma.

Lemma 4. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{sum}}$ and X satisfying (13). Let $p \in \mathcal{P}$ and $\mathbf{e} \in X$. If $\gamma_p(\mathbf{e}) > 0$ then it holds that $\gamma_p(\lambda \mathbf{e}) > 0$ for every $\lambda > 0$.*

Proof. First note that $\gamma_p(\frac{\mathbf{e}}{2}) > 0$. Indeed, $\gamma_p(\frac{\mathbf{e}}{2}) \leq 0$ would imply $\gamma_p(\mathbf{e}) \leq 0$, given that (6) implies $\Gamma(\mathbf{e}) = \Gamma(\frac{\mathbf{e}}{2}) \cup \Gamma(\frac{\mathbf{e}}{2}) = \Gamma(\frac{\mathbf{e}}{2})$. Repeating the same argument, we find $\gamma_p(\frac{\mathbf{e}}{2^i}) > 0$ for any $i \in \mathbb{N}$. Hence, for every $\lambda > 0$, there exists some $\lambda' \in]0, \lambda[$ such that $\gamma_p(\lambda' \mathbf{e}) > 0$. Since $(\lambda - \lambda')\mathbf{e} \in X$, from (6) we then also find that $\gamma_p(\lambda \mathbf{e}) > 0$. \square

Proposition 5. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$ with $\diamond = \diamond_{\text{sum}}$ and X satisfying (13). Then it also holds that (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$ with $\diamond = \diamond_{\text{avg}}$.*

Proof. Suppose (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$ with $\diamond = \diamond_{\text{sum}}$. Let $\mathbf{e}, \mathbf{f} \in X$ be such that $p \in \Gamma(\mathbf{e} \diamond_{\text{avg}} \mathbf{f})$. In other words $\gamma_p(\frac{1}{2}(\mathbf{e} + \mathbf{f})) > 0$. Using Lemma 4 we then find $\gamma_p(\mathbf{e} + \mathbf{f}) > 0$, and thus $p \in \Gamma(\mathbf{e} + \mathbf{f})$. Since (6) is satisfied for \diamond_{sum} we find $p \in \Gamma(\mathbf{e}) \cup \Gamma(\mathbf{f})$. Conversely, assume that $p \in \Gamma(\mathbf{e}) \cup \Gamma(\mathbf{f})$. Since (6) is satisfied for \diamond_{sum} , this implies $p \in \Gamma(\mathbf{e} + \mathbf{f})$, and using Lemma 4 we find $p \in \Gamma(\frac{1}{2}(\mathbf{e} + \mathbf{f}))$. \square

Among others, it follows from Proposition 5 that whenever (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$ with $\diamond = \diamond_{\text{sum}}$, we have that Pos_p and Neg_p are convex for every $p \in \mathcal{P}$, and $\text{Neg}_p \subseteq \delta(X)$. In particular, we also have the following corollary.

Corollary 3. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{sum}}$ and $X \subseteq \mathbb{R}^n$ satisfying (13). Suppose that exhaustiveness (8) is satisfied. It holds that $n \geq |\mathcal{P}|$.*

Weak semantics In entirely the same way as Proposition 5, we can show the following result.

Proposition 6. *Suppose the epistemic pooling principle (10) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$ with $\diamond = \diamond_{\text{sum}}$ and X satisfying (13). Then it also holds that (10) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$ with $\diamond = \diamond_{\text{avg}}$.*

This means that we have the same negative result as we found for \diamond_{avg} . In particular, from Propositions 3 and 6, we immediately obtain the following corollary.

Corollary 4. *Suppose the epistemic pooling principle (10) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, for $\diamond = \diamond_{\text{sum}}$ and X satisfying (13). Suppose γ_p is continuous. It holds that either $\text{Neg}'_p = \emptyset$ or $\text{Pos}'_p = \emptyset$.*

4.3. Max-pooling

Strict semantics In contrast to what we found for \diamond_{avg} and \diamond_{sum} , when using \diamond_{max} it is possible to satisfy the epistemic pooling principle (6) for all $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$ in a non-trivial way. The main idea is illustrated in Fig. 3. For reasons that will become clear in Section 5, in addition to the case where $X = \mathbb{R}^n$, we also consider the case where $X =]-\infty, z]^n$ for some $z \in \mathbb{R}$. We now first show the following characterisation: whenever (6) is satisfied in a non-trivial way, we always have that Neg_p is of the form $X \cap (Y_1 \times \dots \times Y_n)$, where each Y_i is of the form $] - \infty, b_i[$, $] - \infty, b_i]$ or $] - \infty, +\infty[$. Before we show this result, we show a number of lemmas.

Lemma 5. *Let $X = \mathbb{R}^n$ or $X =]-\infty, z]^n$. Suppose that the epistemic pooling principle (6) is satisfied for every $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{max}}$. Let $\mathbf{x} = (x_1, \dots, x_n) \in X$ and $\mathbf{y} = (y_1, \dots, y_n) \in X$ be such that $\forall i \in \{1, \dots, n\} . x_i \leq y_i$. Then the following implication holds for every $p \in \mathcal{P}$:*




Fig. 3. Construction satisfying (6) and (8) for \diamond_{\max} .

$$\mathbf{y} \in \text{Neg}_p \Rightarrow \mathbf{x} \in \text{Neg}_p \tag{14}$$

Proof. Given that $\mathbf{x} \diamond_{\max} \mathbf{y} = \mathbf{y}$, it follows from (6) that $\Gamma(\mathbf{x}) \subseteq \Gamma(\mathbf{y})$, from which we immediately find (14). \square

Lemma 6. Let $X = \mathbb{R}^n$ or $X =]-\infty, z]^n$. Suppose that the epistemic pooling principle (6) is satisfied for every $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\max}$. Let $\mathbf{x} = (x_1, \dots, x_n) \in X$ and $\mathbf{y} = (y_1, \dots, y_n) \in X$ be such that $\forall i \in \{1, \dots, n\} . x_i \leq y_i$. Then the following implication holds for every $p \in \mathcal{P}$:

$$\mathbf{y} \in cl(\text{Neg}_p) \Rightarrow \mathbf{x} \in cl(\text{Neg}_p)$$

Proof. Suppose $\mathbf{y} \in cl(\text{Neg}_p)$. We need to show that for every $\varepsilon > 0$ there exists some $\mathbf{x}' \in \text{Neg}_p$ such that $d(\mathbf{x}, \mathbf{x}') < \varepsilon$. Let $\varepsilon > 0$. Since $\mathbf{y} \in cl(\text{Neg}_p)$, there exists some $\mathbf{y}' = (y'_1, \dots, y'_n) \in \text{Neg}_p$ such that $d(\mathbf{y}, \mathbf{y}') < \varepsilon$. Due to the assumption that $\forall i \in \{1, \dots, n\} . x_i \leq y_i$, we also have $\forall i \in \{1, \dots, n\} . y'_i - y_i + x_i \leq y'_i$. Using Lemma 5, this implies $\mathbf{y}' - \mathbf{y} + \mathbf{x} \in \text{Neg}_p$. Moreover, we have $d(\mathbf{x}, \mathbf{y}' - \mathbf{y} + \mathbf{x}) = d(\mathbf{y}', \mathbf{y}) < \varepsilon$. It thus follows that $\mathbf{x} \in cl(\text{Neg}_p)$. \square

Lemma 7. Let $X = \mathbb{R}^n$ or $X =]-\infty, z]^n$. Suppose that the epistemic pooling principle (6) is satisfied for every $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\max}$. If $\mathbf{x}, \mathbf{y} \in cl(\text{Neg}_p)$ then we also have $\mathbf{x} \diamond_{\max} \mathbf{y} \in cl(\text{Neg}_p)$.

Proof. Let $\mathbf{x}, \mathbf{y} \in cl(\text{Neg}_p)$. We show that for every $\varepsilon > 0$, there is some $\mathbf{e} \in \text{Neg}_p$ such that $d(\mathbf{x} \diamond_{\max} \mathbf{y}, \mathbf{e}) < \varepsilon$. Since $\mathbf{x}, \mathbf{y} \in cl(\text{Neg}_p)$, there exist $\mathbf{e}_x, \mathbf{e}_y \in \text{Neg}_p$ such that $d(\mathbf{e}_x, \mathbf{x}) < \frac{\varepsilon}{\sqrt{n}}$ and $d(\mathbf{e}_y, \mathbf{y}) < \frac{\varepsilon}{\sqrt{n}}$. Since $\mathbf{e}_x, \mathbf{e}_y \in \text{Neg}_p$, by (6) we also have $\mathbf{e}_x \diamond_{\max} \mathbf{e}_y \in \text{Neg}_p$. Moreover, we have, for $\mathbf{e}_x = (e_{x,1}, \dots, e_{x,n})$, $\mathbf{e}_y = (e_{y,1}, \dots, e_{y,n})$, $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$:

$$d^2(\mathbf{e}_x \diamond_{\max} \mathbf{e}_y, \mathbf{x} \diamond_{\max} \mathbf{y}) = \sum_{i=1}^n (\max(e_{x,i}, e_{y,i}) - \max(x_i, y_i))^2 < \sum_{i=1}^n \left(\frac{\varepsilon}{\sqrt{n}}\right)^2 = \varepsilon^2$$

and thus $d(\mathbf{e}_x \diamond_{\max} \mathbf{e}_y, \mathbf{x} \diamond_{\max} \mathbf{y}) < \varepsilon$. \square

Proposition 7. Let $X = \mathbb{R}^n$ or $X =]-\infty, z]^n$. Suppose that the epistemic pooling principle (6) is satisfied for every $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\max}$. Let $p \in \mathcal{P}$. It holds that $cl(\text{Neg}_p) = X \cap (Y_p^1 \times \dots \times Y_p^n)$ where for every $i \in \{1, \dots, n\}$, we have $Y_p^i =]-\infty, b_i]$ for some $b_i \in \mathbb{R}$ or $Y_p^i =]-\infty, +\infty[$.

Proof. For each $i \in \{1, \dots, n\}$, we can consider two cases:

- Assume that the i^{th} coordinate of the elements from $cl(\text{Neg}_p)$ is bounded, i.e. there exists some $b_i \in \mathbb{R}$ such that for each $(x_1, \dots, x_n) \in cl(\text{Neg}_p)$ it holds that $x_i \leq b_i$. Suppose (x_1, \dots, x_n) and (y_1, \dots, y_n) are elements from $cl(\text{Neg}_p)$ which are maximal in the i^{th} coordinate, i.e. for any $\varepsilon > 0$ we have $(x_1, \dots, x_{i-1}, x_i + \varepsilon, x_{i+1}, \dots, x_n) \notin cl(\text{Neg}_p)$ and $(y_1, \dots, y_{i-1}, y_i + \varepsilon, y_{i+1}, \dots, y_n) \notin cl(\text{Neg}_p)$. Assume furthermore that $x_i < y_i$. We have that $(\max(x_1, y_1), \dots, \max(x_n, y_n)) \in cl(\text{Neg}_p)$ by Lemma 7, which implies $(x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n) \in cl(\text{Neg}_p)$ by Lemma 6. However, this is in contradiction with the assumption we made about the i^{th} coordinate of (x_1, \dots, x_n) . It follows that there is a constant $b_i \in \mathbb{R}$ such that for every $(x_1, \dots, x_n) \in cl(\text{Neg}_p)$, we have

$$b_i = \max\{y \mid (x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n) \in cl(\text{Neg}_p)\}$$

Moreover, from Lemma 6 we find that $(x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n) \in cl(\text{Neg}_p)$ for every $y \in]-\infty, b_i]$.

- Now we consider the case where the i^{th} coordinate of the elements from $cl(\text{Neg}_p)$ is unbounded. Let $(y_1, \dots, y_n) \in cl(\text{Neg}_p)$. We show that for any $z \in \mathbb{R}$ it holds that $(y_1, \dots, y_{i-1}, z, y_{i+1}, \dots, y_n) \in cl(\text{Neg}_p)$. Since we assumed the i^{th} coordinate of the elements from $cl(\text{Neg}_p)$ is unbounded, there exists some $(x_1, \dots, x_n) \in cl(\text{Neg}_p)$ such that $x_i \geq z$. From $(x_1, \dots, x_n) \in cl(\text{Neg}_p)$ and $(y_1, \dots, y_n) \in cl(\text{Neg}_p)$ we find $(\max(x_1, y_1), \dots, \max(x_n, y_n)) \in cl(\text{Neg}_p)$, using Lemma 7, which in turn implies $(y_1, \dots, y_{i-1}, z, y_{i+1}, \dots, y_n) \in cl(\text{Neg}_p)$, using Lemma 6.

Putting these two cases together, find that $cl(\text{Neg}_p)$ is of the form $X \cap (Y_p^1 \times \dots \times Y_p^n)$. \square

Using the characterisation from Proposition 7, we now show that embeddings with a minimum of $|\mathcal{P}|$ dimensions are needed to satisfy the epistemic pooling principle.

Proposition 8. *Let $X = \mathbb{R}^n$ or $X =] - \infty, z]^n$. Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\max}$. Suppose that exhaustiveness (8) is satisfied. It holds that $n \geq |\mathcal{P}|$.*

Proof. From Proposition 7 we know that for each property $p \in \mathcal{P}$, it holds that $cl(\text{Neg}_p)$ is of the form $X \cap (Y_p^1 \times \dots \times Y_p^n)$, where $Y_p^i =] - \infty, b_i]$ or $Y_p^i =] - \infty, +\infty[$. Let us write $p <_i q$ for $p, q \in \mathcal{P}$ to denote that one of the following cases holds:

- $Y_p^i =] - \infty, b_i^p]$ and $Y_q^i =] - \infty, b_i^q]$ with $b_i^p < b_i^q$; or
- $Y_p^i =] - \infty, b_i^p]$ and $Y_q^i =] - \infty, b_i^q]$ with $b_i^p = b_i^q = b$ and there exists an element $(x_1, \dots, x_n) \in \text{Neg}_q$ such that $x_i = b$ while no such element exists in Neg_p . In other words, the upper bound b for the i^{th} coordinate is strict for Neg_p but not for Neg_q .
- $Y_p^i =] - \infty, b_i^p]$ and $Y_q^i =] - \infty, +\infty[$.

For each $i \in \{1, \dots, n\}$, we can choose a property p_i from \mathcal{P} which is minimal w.r.t. the relation $<_i$. Suppose there was some property $q \in \mathcal{P} \setminus \{p_1, \dots, p_n\}$. Let $\mathbf{x} = (x_1, \dots, x_n) \in \text{Pos}_q$. Then for some coordinate i , it must be the case that $Y_q^i =] - \infty, b_i^q]$ and either (i) $x_i > b_i^q$ or (ii) $x_i = b_i^q$ but b_i^q represents a strict upper bound for the i^{th} coordinate. Because p_i was chosen as a minimal element w.r.t. $<_i$ it follows that $(x_1, \dots, x_n) \notin \text{Neg}_{p_i}$. We thus find that for every $\mathbf{x} \in \text{Pos}_q$ it holds that $\mathbf{x} \in \text{Pos}_{p_1} \cup \dots \cup \text{Pos}_{p_n}$. It follows that there is no $\mathbf{x} \in \mathbb{R}^n$ such that $\Gamma(\mathbf{x}) = \{q\}$, meaning that (8) is not satisfied. \square

To show that the epistemic pooling principle (6) and exhaustiveness (8) can indeed be satisfied with $|\mathcal{P}|$ dimensions, let $\mathcal{P} = \{p_1, \dots, p_n\}$ and define:

$$\gamma_{p_i}(e_1, \dots, e_n) = e_i \tag{15}$$

Then we have that the i^{th} coordinate of $(e_1, \dots, e_n) \diamond_{\max} (f_1, \dots, f_n)$ is strictly positive iff $e_i > 0$ or $f_i > 0$, hence we indeed have $\Gamma(\mathbf{e} \diamond_{\max} \mathbf{f}) = \Gamma(\mathbf{e}) \cup \Gamma(\mathbf{f})$, meaning that the epistemic pooling principle (6) is satisfied for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$. To see why exhaustiveness (8) is satisfied, let $\mathcal{Q} \subseteq \mathcal{P}$. We define $\mathbf{q} = (q_1, \dots, q_n)$ as follows:

$$q_i = \begin{cases} 1 & \text{if } p_i \in \mathcal{Q} \\ -1 & \text{otherwise} \end{cases}$$

Weak semantics As before, the main question is whether it is possible to satisfy the weak epistemic pooling principle (10) in a non-trivial way using continuous scoring functions γ_p , since the results from the strict semantics trivially carry over to the weak semantics if non-continuous scoring functions are allowed. This is indeed the case. In fact, with the scoring functions defined in (15), the weak epistemic pooling principle is also satisfied. Moreover, in the same way as for the strict semantics, we find that exhaustiveness (8) is satisfied for this choice. Finally, note that the lower bound $n \geq |\mathcal{P}|$ still applies for the weak semantics, which can be shown in exactly the same way as Proposition 8.

4.4. Hadamard product

Strict semantics Similar as we found for max-pooling, with the Hadamard product \diamond_{had} , it is possible to satisfy the epistemic pooling principle (6) for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, while also satisfying exhaustiveness (8). In addition to the choice $X = \mathbb{R}^n$, we also consider the case where $X = [0, +\infty[^n$. As we will see, the results we establish in this section are valid regardless of whether $X = \mathbb{R}^n$ or $X = [0, +\infty[^n$. The reason why we specifically include the case $X = [0, +\infty[^n$ will become clear in Section 5. Fig. 4 illustrates how the epistemic pooling principle (6) can be satisfied for the case where $X = [0, +\infty[^n$. Note the dual nature of this construction and the construction for \diamond_{avg} from Fig. 2, where the positive regions in Fig. 4 correspond to the negative regions in Fig. 2. Let us write H_i for the following hyperplane:

$$H_i = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_i = 0\}$$




Fig. 4. Construction satisfying (6) and (8) for \diamond_{had} with $X = [0, +\infty[^n$.

These hyperplanes play a particular role in the characterisation of the positive regions Pos_p , as was already illustrated in Fig. 4. The following results make this role explicit.

Lemma 8. Let $X = \mathbb{R}^n$ or $X = [0, +\infty[^n$. Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{had}}$. Let $p \in \mathcal{P}$ and $\mathbf{e} = (e_1, \dots, e_n) \in \text{Pos}_p$. Let $I = \{i \in \{1, \dots, n\} \mid e_i = 0\}$. It holds that

$$X \cap \bigcap_{i \in I} H_i \subseteq \text{Pos}_p$$

Proof. Let $\mathbf{f} = (f_1, \dots, f_n) \in X \cap \bigcap_{i \in I} H_i$. We show that $\mathbf{f} \in \text{Pos}_p$. Let $\mathbf{x} = (x_1, \dots, x_n)$ be defined as follows:

$$x_i = \begin{cases} \frac{f_i}{e_i} & \text{if } e_i \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

Since $(f_1, \dots, f_n) \in \bigcap_{i \in I} H_i$ it holds that $f_i = 0$ whenever $e_i = 0$. We thus have $\mathbf{f} = \mathbf{e} \diamond_{\text{had}} \mathbf{x}$. Using (6), it then follows from $\mathbf{e} \in \text{Pos}_p$ that $\mathbf{f} \in \text{Pos}_p$. \square

It follows that Pos_p is a finite union of regions of the form $X \cap \bigcap_{i \in I} H_i$. In particular, it also follows that $\dim(\text{Pos}_p) \leq n - 1$ if the epistemic pooling principle (6) is satisfied in a non-trivial way. For a given index set $I \subseteq \{1, \dots, n\}$, let us define:

$$\mathcal{P}_I = \{p \in \mathcal{P} \mid X \cap \bigcap_{i \in I} H_i \subseteq \text{Pos}_p\}$$

Lemma 9. Let $X = \mathbb{R}^n$ or $X = [0, +\infty[^n$. Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{had}}$. Let $I, J \subseteq \{1, \dots, n\}$. It holds that:

$$\mathcal{P}_{I \cup J} = \mathcal{P}_I \cup \mathcal{P}_J$$

Proof. We clearly have $\mathcal{P}_I \subseteq \mathcal{P}_{I \cup J}$, since $X \cap \bigcap_{i \in I} H_i \subseteq \text{Pos}_p$ implies $X \cap \bigcap_{i \in I \cup J} H_i \subseteq \text{Pos}_p$. Similarly, we have $\mathcal{P}_J \subseteq \mathcal{P}_{I \cup J}$ and thus we find $\mathcal{P}_I \cup \mathcal{P}_J \subseteq \mathcal{P}_{I \cup J}$. We now show the other direction. Suppose $p \in \mathcal{P}_{I \cup J}$. Let $\mathbf{e} = (e_1, \dots, e_n)$ where $e_i = 0$ if $i \in I \cup J$ and $e_i = 1$ otherwise. Since $p \in \mathcal{P}_{I \cup J}$ we have $\mathbf{e} \in \text{Pos}_p$. Now let $\mathbf{f} = (f_1, \dots, f_n)$ where $f_i = 1$ if $i \in I$ and $f_i = 0$ otherwise. Similarly, let $\mathbf{g} = (g_1, \dots, g_n)$ where $g_i = 1$ if $i \in J$ and $g_i = 0$ otherwise. Then $\mathbf{e} = \mathbf{f} \odot \mathbf{g}$. From (6) it follows that $\mathbf{f} \in \text{Pos}_p$ or $\mathbf{g} \in \text{Pos}_p$. Thus, using Lemma 8 we find that $p \in \mathcal{P}_I$ or $p \in \mathcal{P}_J$. \square

We can now show that at least $|\mathcal{P}|$ dimensions are again needed to satisfy (6) in a nontrivial way.

Proposition 9. Let $X = \mathbb{R}^n$ or $X = [0, +\infty[^n$. Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{had}}$. Suppose that exhaustiveness (8) is satisfied. It holds that $n \geq |\mathcal{P}|$.

Proof. Given (8), for each $p \in \mathcal{P}$, there must exist some $(e_1^p, \dots, e_n^p) \in X$ such that $\Gamma(e_1^p, \dots, e_n^p) = \{p\}$. Let us fix such vectors (e_1^p, \dots, e_n^p) for each $p \in \mathcal{P}$ and define $I_p = \{i \in \{1, \dots, n\} \mid e_i^p = 0\}$. Note that by Lemma 8, we have $X \cap \bigcap_{i \in I_p} H_i \subseteq \text{Pos}_p$. Moreover, since $\Gamma(e_1^p, \dots, e_n^p) = \{p\}$, we have $X \cap \bigcap_{i \in I_p} H_i \not\subseteq \text{Pos}_q$ for any $q \neq p$. In other words, we have $\mathcal{P}_{I_p} = \{p\}$.

For $p \neq q$ we clearly have $I_p \not\subseteq I_q$, since $I_p \subseteq I_q$ would imply $\mathcal{P}_{I_q} = \{p, q\}$. This implies in particular that $I_p \neq \emptyset$ for every $p \in \mathcal{P}$. Now let us consider k distinct properties p_1, p_2, \dots, p_k . Then we cannot have $I_{p_k} \subseteq I_{p_1} \cup I_{p_2} \cup \dots \cup I_{p_{k-1}}$. Indeed

$I_{p_k} \subseteq I_{p_1} \cup \dots \cup I_{p_{k-1}}$ would imply $p_k \in \mathcal{P}_{I_{p_1} \cup \dots \cup I_{p_{k-1}}}$ whereas from Lemma 9 we know that $\mathcal{P}_{I_{p_1} \cup \dots \cup I_{p_{k-1}}} = \mathcal{P}_{I_{p_1}} \cup \dots \cup \mathcal{P}_{I_{p_{k-1}}}$ and we know that the latter is equal to $\{p_1, \dots, p_{k-1}\}$. In other words, there is at least one element in I_{p_k} which does not occur in $I_{p_1}, \dots, I_{p_{k-1}}$, a contradiction. Since this needs to hold for every $k \in \{2, \dots, |\mathcal{P}|\}$, there need to be at least $|\mathcal{P}|$ distinct elements in $I_{p_1} \cup \dots \cup I_{p_{|\mathcal{P}|}}$. This means that $n \geq |\mathcal{P}|$. \square

To show that the epistemic pooling principle (6) and exhaustiveness (8) can indeed be satisfied with $|\mathcal{P}|$ dimensions, let $\mathcal{P} = \{p_1, \dots, p_n\}$ and define:

$$\gamma_{p_i}(e_1, \dots, e_n) = \begin{cases} 1 & \text{if } e_i = 0 \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

Clearly, the i^{th} coordinate of $(e_1, \dots, e_n) \diamond_{\text{had}} (f_1, \dots, f_n)$ is 0 iff $e_i = 0$ or $f_i = 0$, hence we indeed have $\Gamma(\mathbf{e} \diamond_{\text{had}} \mathbf{f}) = \Gamma(\mathbf{e}) \cup \Gamma(\mathbf{f})$, meaning that (6) is satisfied for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$. It is also straightforward to verify that exhaustiveness (8) is satisfied. Note, however, that the scoring function γ_{p_i} defined in (16) is not continuous. As the following result shows, this is not a coincidence.

Proposition 10. *Let $X = \mathbb{R}^n$ or $X = [0, +\infty[^n$. Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, for $\diamond = \diamond_{\text{had}}$. Suppose γ_p is continuous. It holds that either $\text{Neg}_p = \emptyset$ or $\text{Pos}_p = \emptyset$.*

Proof. Suppose $\text{Pos}_p \neq \emptyset$. Then from Lemma 8, we know that there is some $I \subseteq \{1, \dots, n\}$ such that $X \cap \bigcap_{i \in I} H_i \subseteq \text{Pos}_p$. Let us assume that I is minimal, i.e. for any $I' \subset I$ we have $X \cap \bigcap_{i \in I'} H_i \not\subseteq \text{Pos}_p$. Now suppose we also have $\text{Neg}_p \neq \emptyset$. Then we have $I \neq \emptyset$. Let $i \in I$ and $\mathbf{e}_0 \in X \cap \bigcap_{i \in I} H_i$. For $\varepsilon > 0$, let \mathbf{e}_ε be the vector obtained from \mathbf{e}_0 by replacing the i^{th} coordinate by ε . Since we assumed I was minimal, for any $\varepsilon > 0$ we have that $\mathbf{e}_\varepsilon \in \text{Neg}_p$, or equivalently $\gamma_p(\mathbf{e}_\varepsilon) \leq 0$. However, if γ_p is continuous, this would imply $\gamma_p(\mathbf{e}_0) \leq 0$ and thus $\mathbf{e}_0 \in \text{Neg}_p$, a contradiction. \square

The fact that only discontinuous scoring functions can be used is an important limitation in practice. One solution is to make a different choice for the set X . For instance, if we choose $X = [1, +\infty[$, we end up with embeddings where the bounding hyperplanes, of the form $H = \{(e_1, \dots, e_n) \in \mathbb{R}^n \mid e_i = 1\}$, correspond to negative regions. The resulting embeddings thus have similar characteristics as the ones we obtained with \diamond_{avg} . For this reason, we will not consider this option in further detail.

Weak semantics Under the weak semantics, it is possible to satisfy the epistemic pooling principle using continuous scoring functions. For instance, if $\mathcal{P} = \{p_1, \dots, p_n\}$ we can define:

$$\gamma_{p_i}(e_1, \dots, e_n) = -e_i^2 \quad (17)$$

It is straightforward to verify that, with this choice of scoring function, the epistemic pooling principle (10) is satisfied for all $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, while exhaustiveness (11) is also satisfied. Moreover, the same argument as in the proof of Proposition 9 can be used for the weak semantics as well, meaning that we still need at least $|\mathcal{P}|$ dimensions to satisfy (10) with $\diamond = \diamond_{\text{had}}$ in a non-trivial way.

5. Propositional reasoning with epistemic embeddings

Throughout this paper, we model epistemic states as subsets of \mathcal{P} . In general, we can simply think of the elements of \mathcal{P} as atomic pieces of evidence. Crucially, however, this setting is expressive enough to capture propositional reasoning. In particular, suppose each embedding \mathbf{e} is associated with a set $\Lambda(\mathbf{e})$ of propositional formulas. Now suppose the embedding \mathbf{g} is obtained by pooling the embeddings \mathbf{e} and \mathbf{f} , e.g. representing the information we have obtained from two different sources:

$$\mathbf{g} = \mathbf{e} \diamond \mathbf{f}$$

Then we want \mathbf{g} to combine the knowledge captured by \mathbf{e} and \mathbf{f} . In other words, we want $\Lambda(\mathbf{g})$ to be logically equivalent to $\Lambda(\mathbf{e}) \cup \Lambda(\mathbf{f})$. This gives rise to the following variant of the epistemic pooling principle:

$$\Lambda(\mathbf{e} \diamond \mathbf{f}) \equiv \Lambda(\mathbf{e}) \cup \Lambda(\mathbf{f}) \quad (18)$$

where we write \equiv to denote logical equivalence between sets of propositional formulas. In other words, in this setting, we want to be able to do propositional reasoning by pooling embeddings. This can be achieved using our considered setting as follows. Let \mathcal{W} be the set of all possible worlds, i.e. the set of all propositional interpretations over some set of atoms At . We associate one property with each possible world:

$$\mathcal{P} = \{p_\omega \mid \omega \in \mathcal{W}\} \tag{19}$$

Intuitively, p_ω means that ω can be excluded, i.e. that we have evidence that ω is not the true world. Let *Clauses* be the set of all clauses over the considered set of propositional atoms *At*. We can define $\Lambda(\mathbf{e})$ in terms of the scoring functions γ_p as follows:

$$\Lambda(\mathbf{e}) = \{\alpha \in \text{Clauses} \mid \forall \omega \in \mathcal{W}. (\gamma_{p_\omega}(\mathbf{e}) \leq 0) \Rightarrow (\omega \models \alpha)\} \tag{20}$$

In other words, we have that $\alpha \in \Lambda(\mathbf{e})$ if α is true in all the worlds ω that cannot be excluded based on the evidence encoded by \mathbf{e} . Note that $\Lambda(\mathbf{e})$ is deductively closed, in the sense that $\alpha \in \Lambda(\mathbf{e})$ iff $\Lambda(\mathbf{e}) \models \alpha$. We can also straightforwardly show the following characterisation:

Lemma 10. *Suppose \mathcal{P} is defined by (19) and Λ is defined by (20). Let $p_\omega \in \mathcal{P}$. It holds that $\gamma_{p_\omega}(\mathbf{e}) \leq 0$ iff $\omega \models \Lambda(\mathbf{e})$.*

Proof. Suppose $\gamma_{p_\omega}(\mathbf{e}) \leq 0$, then by definition of Λ , for each $\alpha \in \Lambda(\mathbf{e})$ we must have $\omega \models \alpha$. In particular, we must thus have $\omega \models \Lambda(\mathbf{e})$. Conversely, suppose $\gamma_{p_\omega}(\mathbf{e}) > 0$. Let $\{\alpha_1, \dots, \alpha_k\}$ be a set of clauses whose set of models is given by $\{\omega' \mid \gamma_{p_{\omega'}}(\mathbf{e}) \leq 0\}$. Note that such a set of clauses always exists. Moreover, we then have $\{\alpha_1, \dots, \alpha_k\} \subseteq \Lambda(\mathbf{e})$, by definition of Λ , while by construction we have $\omega \not\models \{\alpha_1, \dots, \alpha_k\}$. It follows that $\omega \not\models \Lambda(\mathbf{e})$. \square

We can now prove the following result, which shows that (18) arises as a special case of the (strict) epistemic pooling principle (6).

Proposition 11. *Suppose the epistemic pooling principle (6) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, for some $X \subseteq \mathbb{R}^n$ and some pooling operator \diamond such that X is closed under \diamond . Suppose \mathcal{P} is defined by (19) and Λ is defined by (20). It holds that (18) is satisfied for all $\mathbf{e}, \mathbf{f} \in X$.*

Proof. Let $\mathbf{e}, \mathbf{f} \in X$. By (6) we have that $\Gamma(\mathbf{e} \diamond \mathbf{f}) \supseteq \mathbf{e}$, hence for every ω we have:

$$(\gamma_{p_\omega}(\mathbf{e} \diamond \mathbf{f}) \leq 0) \Rightarrow (\gamma_{p_\omega}(\mathbf{e}) \leq 0)$$

From the definition of Λ , it follows that $\Lambda(\mathbf{e} \diamond \mathbf{f}) \supseteq \Lambda(\mathbf{e})$. Since we similarly have $\Lambda(\mathbf{e} \diamond \mathbf{f}) \supseteq \Lambda(\mathbf{f})$, we find $\Lambda(\mathbf{e} \diamond \mathbf{f}) \supseteq \Lambda(\mathbf{e}) \cup \Lambda(\mathbf{f})$ and in particular:

$$\Lambda(\mathbf{e} \diamond \mathbf{f}) \models \Lambda(\mathbf{e}) \cup \Lambda(\mathbf{f})$$

Conversely, suppose $\alpha \in \Lambda(\mathbf{e} \diamond \mathbf{f})$. Then we have $\forall \omega \in \mathcal{W}. (\gamma_{p_\omega}(\mathbf{e} \diamond \mathbf{f}) \leq 0) \Rightarrow (\omega \models \alpha)$. From (6) we know that $\gamma_{p_\omega}(\mathbf{e} \diamond \mathbf{f}) \leq 0$ iff $\gamma_{p_\omega}(\mathbf{e}) \leq 0$ and $\gamma_{p_\omega}(\mathbf{f}) \leq 0$, hence we obtain:

$$\forall \omega \in \mathcal{W}. (\gamma_{p_\omega}(\mathbf{e}) \leq 0) \wedge (\gamma_{p_\omega}(\mathbf{f}) \leq 0) \Rightarrow (\omega \models \alpha)$$

Using Lemma 10 we find:

$$\forall \omega \in \mathcal{W}. (\omega \models \Lambda(\mathbf{e})) \wedge (\omega \models \Lambda(\mathbf{f})) \Rightarrow (\omega \models \alpha)$$

or, equivalently, $\forall \omega \in \mathcal{W}. (\omega \models \Lambda(\mathbf{e}) \cup \Lambda(\mathbf{f})) \Rightarrow (\omega \models \alpha)$. We thus find $\Lambda(\mathbf{e}) \cup \Lambda(\mathbf{f}) \models \alpha$. Since this holds for every $\alpha \in \Lambda(\mathbf{e} \diamond \mathbf{f})$, we find

$$\Lambda(\mathbf{e}) \cup \Lambda(\mathbf{f}) \models \Lambda(\mathbf{e} \diamond \mathbf{f}) \quad \square$$

We can similarly model propositional reasoning using the weak semantics, by defining the set of formulas associated with an embedding \mathbf{e} as follows

$$\Lambda'(\mathbf{e}) = \{\alpha \in \text{Clauses} \mid \forall \omega \in \mathcal{W}. (\gamma_{p_\omega}(\mathbf{e}) < 0) \Rightarrow (\omega \models \alpha)\} \tag{21}$$

The counterpart to Proposition 11 for the weak semantics is shown in entirely the same way. We can thus use the framework that was studied in Section 4 to combine, and reason about propositional knowledge. However, when we focus on knowledge that is encoded using propositional formulas, we also need an effective way to check whether a given formula α is entailed by the knowledge base $\Lambda(\mathbf{e})$, i.e. whether the knowledge encoded by \mathbf{e} is sufficient to conclude that α holds. To this end, for each propositional formula, we need a scoring function $\phi_\alpha : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $\psi_\alpha(\mathbf{e})$ indicates whether $\Lambda(\mathbf{e}) \models \alpha$. We now study such scoring functions.

5.1. Checking the satisfaction of propositional formulas

Let us consider scoring functions $\psi_\alpha : \mathbb{R}^n \rightarrow \mathbb{R}$, for arbitrary propositional formulas α , which satisfy the following condition:

$$(\psi_\alpha(\mathbf{e}) > 0) \Leftrightarrow (\Lambda(\mathbf{e}) \models \alpha) \quad (22)$$

where we assume that Λ is defined as in (20). In other words, $\psi_\alpha(\mathbf{e}) > 0$ holds if α is (known to be) true in the epistemic state encoded by \mathbf{e} . Let us write $\mathcal{M}(\alpha)$ for the models of a formula α , where $\mathcal{M}(\alpha) \subseteq \mathcal{W}$. We find from the definition of Λ that $\Lambda(\mathbf{e}) \models \alpha$ is equivalent with

$$\forall \omega \in \mathcal{M}(\neg\alpha). \gamma_{p_\omega}(\mathbf{e}) > 0$$

In other words, to check the satisfaction of a propositional formula α , we need a scoring function that allows us to check whether $\gamma_p(\mathbf{e}) > 0$ for every p in some subset of properties $\mathcal{Q} \subseteq \mathcal{P}$. In particular let us define a scoring function $\gamma_{\mathcal{Q}}$ for every $\mathcal{Q} \subseteq \mathcal{P}$ such that

$$(\gamma_{\mathcal{Q}}(\mathbf{e}) > 0) \equiv (\forall p \in \mathcal{Q}. \gamma_p(\mathbf{e}) > 0) \quad (23)$$

If this equivalence is satisfied, we will say that the scoring functions are *faithful*. Then we have:

$$(\psi_\alpha(\mathbf{e}) > 0) \Leftrightarrow (\gamma_{\mathcal{M}(\neg\alpha)}(\mathbf{e}) > 0)$$

The scoring functions of the form ψ_α are thus a special case of scoring functions of the form $\gamma_{\mathcal{Q}}$. For generality, we will study the latter type of scoring functions in the remainder of this section. This also has the advantage that we can stay closer to the setting from Section 4. Analogously to Pos_p and Neg_p , we now define the following regions:

$$\text{Pos}_{\mathcal{Q}} = \{\mathbf{e} \in X \mid \gamma_{\mathcal{Q}}(\mathbf{e}) > 0\}$$

$$\text{Neg}_{\mathcal{Q}} = \{\mathbf{e} \in X \mid \gamma_{\mathcal{Q}}(\mathbf{e}) \leq 0\}$$

Similarly, under the weak semantics, we can consider scoring functions of the form $\gamma'_{\mathcal{Q}}$, with $\mathcal{Q} \subseteq \mathcal{P}$, which are faithfully linked to the scoring functions γ_p as follows:

$$(\gamma'_{\mathcal{Q}}(\mathbf{e}) \geq 0) \equiv (\forall p \in \mathcal{Q}. \gamma_p(\mathbf{e}) \geq 0) \quad (24)$$

The corresponding positive and negative regions are defined as follows

$$\text{Pos}'_{\mathcal{Q}} = \{\mathbf{e} \in X \mid \gamma'_{\mathcal{Q}}(\mathbf{e}) \geq 0\}$$

$$\text{Neg}'_{\mathcal{Q}} = \{\mathbf{e} \in X \mid \gamma'_{\mathcal{Q}}(\mathbf{e}) < 0\}$$

Clearly, if the scoring functions γ_p are continuous, then continuous scoring functions of the form $\gamma_{\mathcal{Q}}$ must also exist, as we can simply define $\gamma_{\mathcal{Q}}(\mathbf{e}) = \min_{q \in \mathcal{Q}} \gamma_q(\mathbf{e})$, and similar for the weak semantics. Our main focus in the remainder of this section is on the following question: is it possible for the scoring functions $\gamma_{\mathcal{Q}}$ and $\gamma'_{\mathcal{Q}}$ to be linear? In other words, can we use linear scoring functions to check the satisfaction of a propositional formula in the epistemic state encoded by a vector \mathbf{e} ? This question is important because of the prevalence of linear scoring functions in the classification layer of neural networks. Our results are summarised in Table 2. One important finding is that linear scoring functions of the form $\gamma_{\mathcal{Q}}$ are not compatible with the strict semantics, regardless of how the embeddings were obtained (and thus also regardless of which pooling operator is considered). For the weak semantics, we find that linear scoring functions are possible with \diamond_{\max} and \diamond_{had} , but crucially, this is only the case if the set of embeddings X is bounded in a suitable way.

5.2. Linear scoring functions under the strict semantics

In the following, we assume that embeddings belong to some region $X \subseteq \mathbb{R}^n$, which we assume to be convex. We make no assumptions about how the embeddings are obtained, requiring only that exhaustiveness (8) is satisfied, i.e. for every epistemic state, there exists some embedding $\mathbf{e} \in X$ which encodes it. We show that scoring functions of the form $\gamma_{\mathcal{Q}}$ can then not be linear. Before showing the main result, we first prove three lemmas.

Lemma 11. *Suppose exhaustiveness (8) is satisfied. Let $\mathcal{Q} = \{p_1, \dots, p_k\}$ be a subset of \mathcal{P} . Assume that $\gamma_{\mathcal{Q}}$ is linear and suppose $\mathbf{e} \in X$ is such that $\gamma_{\mathcal{Q}}(\mathbf{e}) = 0$. For every $\varepsilon > 0$ there exists some $\mathbf{f} \in X$ such that $d(\mathbf{e}, \mathbf{f}) < \varepsilon$ while $\gamma_{\mathcal{Q}}(\mathbf{f}) > 0$.*

Table 2
Summary of results about the possibility of checking entailment of propositional formulas using linear scoring functions.

Pooling Operator		Semantics	Linear γ_Q possible?
Average		Strict	\times
		Weak	\times
Summation		Strict	\times
		Weak	\times
Max-pooling	$X = \mathbb{R}^n$	Strict	\times
		Weak	\times
	$X =]-\infty, z]^n$	Strict	\times
		Weak	\checkmark
Hadamard	$X = \mathbb{R}^n$	Strict	\times
		Weak	\times
	$X = [0, +\infty[^n$	Strict	\times
		Weak	\checkmark

Proof. Note that the linearity of γ_Q means that there exists some hyperplane H which separates the regions Pos_Q and Neg_Q , where $\gamma_Q(\mathbf{e}) = 0$ means that $\mathbf{e} \in H$. If $\mathbf{e} \in \text{int}(X)$ then the claim is trivially satisfied, as there exist vectors \mathbf{f} on either side of the hyperplane which are arbitrarily close to \mathbf{e} . Now assume $\mathbf{e} \in \delta(X)$. Suppose there were an $\varepsilon > 0$ such that for every $\mathbf{f} \in X$ satisfying $d(\mathbf{e}, \mathbf{f}) < \varepsilon$, it holds that $\gamma_Q(\mathbf{f}) \leq 0$. It then follows, given the convexity of X , that H is a bounding hyperplane of X , and in particular that $X \subseteq \text{Neg}_Q$. This is a contradiction, given that we assumed that (8) is satisfied. \square

Lemma 12. Suppose exhaustiveness (8) is satisfied. Let $Q = \{p_1, \dots, p_k\}$ be a subset of \mathcal{P} . Assume that γ_Q is linear and that $\gamma_{p_1}, \dots, \gamma_{p_k}$ are all continuous. For every $\mathbf{e} \in X$ we have:

$$(\gamma_Q(\mathbf{e}) = 0) \Rightarrow (\forall i \in \{1, \dots, k\} \mid \gamma_{p_i}(\mathbf{e}) \geq 0)$$

Proof. Suppose $\gamma_Q(\mathbf{e}) = 0$ and suppose $\gamma_{p_i}(\mathbf{e}) < 0$ for some $i \in \{1, \dots, k\}$. Because γ_{p_i} is continuous, there must be some $\varepsilon > 0$ such that $\gamma_{p_i}(\mathbf{f}) < 0$ for every \mathbf{f} satisfying $d(\mathbf{e}, \mathbf{f}) < \varepsilon$. However, by Lemma 11 we know that there must be such an \mathbf{f} for which $\gamma_Q(\mathbf{f}) > 0$, which implies $\gamma_{p_i}(\mathbf{f}) > 0$, a contradiction. \square

Lemma 13. Suppose exhaustiveness (8) is satisfied. Let $Q = \{p_1, \dots, p_k\}$ be a subset of \mathcal{P} . Assume that γ_Q is linear and that $\gamma_{p_1}, \dots, \gamma_{p_k}$ are all continuous. For every $\mathbf{e} \in X$ we have:

$$(\gamma_Q(\mathbf{e}) = 0) \Rightarrow (\exists i \in \{1, \dots, k\} . \gamma_{p_i}(\mathbf{e}) = 0)$$

Proof. Suppose $\gamma_Q(\mathbf{e}) = 0$. Then there must be some $i \in \{1, \dots, k\}$ such that $\gamma_{p_i}(\mathbf{e}) \leq 0$, by definition of γ_Q . From Lemma 12 we furthermore find $\gamma_{p_i}(\mathbf{e}) \geq 0$. It follows that $\gamma_{p_i}(\mathbf{e}) = 0$. \square

Proposition 12. Suppose exhaustiveness (8) is satisfied. Let $Q = \{p_1, \dots, p_k\}$ be a subset of \mathcal{P} . Suppose $\gamma_{p_1}, \dots, \gamma_{p_k}$ and γ_Q are all linear. Then we have $|Q| = 1$.

Proof. Fig. 5 illustrates the argument provided in this proof. Suppose $k > 1$. Let H_Q be the hyperplane defined by $H_Q = \{\mathbf{e} \mid \gamma_Q(\mathbf{e}) = 0\}$. Let H_i similarly be the hyperplane corresponding to γ_{p_i} . Because of (8) there exists some $\mathbf{e} \in X$ such that $\Gamma(\mathbf{e}) = \{p_1, \dots, p_k\}$. Moreover, for each $i \in \{1, \dots, k\}$ there exists some $\mathbf{f}_i \in X$ such that $\Gamma(\mathbf{f}_i) = \{p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_k\}$. Then we have $\gamma_Q(\mathbf{e}) > 0$ and $\gamma_Q(\mathbf{f}_i) \leq 0$ for every $i \in \{1, \dots, k\}$. For each i , we let \mathbf{g}_i be the point on the intersection between H_Q and the line defined by \mathbf{e} and \mathbf{f}_i . Note that \mathbf{g}_i must exist, since $\gamma_Q(\mathbf{e}) > 0$ and $\gamma_Q(\mathbf{f}_i) \leq 0$. Moreover, we have that $\mathbf{g}_i \in X$ since X is convex. Since $\gamma_{p_1}, \dots, \gamma_{p_k}$ are linear, the fact that $\gamma_{p_j}(\mathbf{e}) > 0$ and $\gamma_{p_j}(\mathbf{f}_i) > 0$, for $j \neq i$, implies that $\gamma_{p_j}(\mathbf{g}_i) > 0$. Since this holds for every $j \neq i$ while $\gamma_Q(\mathbf{g}_i) = 0$ it follows from Lemma 13 that $\gamma_{p_i}(\mathbf{g}_i) = 0$. Now consider $\mathbf{g}^* = \frac{1}{k}(\mathbf{g}_1 + \dots + \mathbf{g}_k)$. By the convexity of X , we have $\mathbf{g}^* \in X$. Moreover, since $\gamma_{p_i}(\mathbf{g}_i) > 0$ if $i \neq j$ and $\gamma_{p_i}(\mathbf{g}_i) = 0$ otherwise, if $k > 1$ we have that $\gamma_{p_i}(\mathbf{g}^*) > 0$ for every $i \in \{1, \dots, k\}$. This would mean that $\Gamma(\mathbf{g}^*) = \{p_1, \dots, p_k\}$ and thus that we should have $\gamma_Q(\mathbf{g}^*) > 0$. However, since $\mathbf{g}_1, \dots, \mathbf{g}_k \in H_Q$, we also have $\mathbf{g}^* \in H_Q$ and thus $\gamma_Q(\mathbf{g}^*) = 0$, a contradiction. It follows that $k = 1$. \square

5.3. Linear scoring functions under the weak semantics

In Section 5.2, we found that the strict semantics is not compatible with the use of linear scoring functions of the form γ_Q . We now explore whether linear scoring functions can be more successful under the weak semantics. We already




Fig. 5. Illustration of the construction from the proof of Proposition 12.

know from Section 4 that \diamond_{avg} and \diamond_{sum} are not compatible with the use of continuous scoring functions under the weak semantics. We therefore focus on the remaining pooling operators, although we will return to \diamond_{avg} in Section 5.4.

5.3.1. Max-pooling

For the max-pooling operator \diamond_{max} , it is possible to satisfy faithfulness (24) using linear scoring functions when choosing $X =] - \infty, z]^n$. For instance, let us fix $X =] - \infty, 0]^n$. Let $\mathcal{P} = \{p_1, \dots, p_n\}$. For $p_i \in \mathcal{P}$ we define:

$$\gamma_{p_i}(e_1, \dots, e_n) = e_i$$

Note that, because we assumed $X =] - \infty, 0]^n$, it holds that $\gamma_{p_i}(e_1, \dots, e_n) \geq 0$ iff $e_i = 0$. As in Section 4.3, with this choice we find that the epistemic pooling principle (10) is satisfied for all $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$ while exhaustiveness (11) is also satisfied. For $\mathcal{Q} = \{p_{i_1}, \dots, p_{i_k}\}$ we now define:

$$\gamma_{\mathcal{Q}}(e_1, \dots, e_n) = e_{i_1} + \dots + e_{i_k}$$

Then we have $\gamma_{\mathcal{Q}}(e_1, \dots, e_n) \geq 0$ iff $e_{i_1} = \dots = e_{i_k} = 0$, which is equivalent to requiring $\gamma_{p_{i_1}}(e_1, \dots, e_n) \geq 0, \dots, \gamma_{p_{i_k}}(e_1, \dots, e_n) \geq 0$. We thus find that faithfulness (24) is indeed satisfied.

We now show that faithfulness (24) cannot be satisfied with linear scoring functions if we choose $X = \mathbb{R}^n$. In Proposition 7 we found that $cl(\text{Neg}_p)$ is of the form $Y_p^1 \times \dots \times Y_p^n$ (for $X = \mathbb{R}^n$), where for every $i \in \{1, \dots, n\}$ we have $Y_p^i =] - \infty, b_i]$ for some $b_i \in \mathbb{R}$ or $Y_p^i =] - \infty, +\infty[$. While this result was shown for the strict semantics, the same argument can be made for the weak semantics, i.e. if the epistemic pooling principle (10) is satisfied for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, then we have that $cl(\text{Neg}'_p)$ is of the form $Y_p^1 \times \dots \times Y_p^n$. Clearly, a region Neg'_p of this form can only arise from a linear scoring function γ_p if $Y_p^i =] - \infty, +\infty[$ for all but one i from $\{1, \dots, n\}$. Note that there must exist some $i \in \{1, \dots, n\}$ such that Y_p^i is of the form $] - \infty, b_i]$ to avoid the trivial case where $\text{Neg}'_p = \mathbb{R}^n$, which would imply that epistemic states in which p is known cannot be modelled. If γ_p is linear, for $p \in \mathcal{P}$, we thus find that there must exist some $i \in \{1, \dots, n\}$ and some $b_i \in \mathbb{R}$ such that

$$\text{Neg}'_p = \{(e_1, \dots, e_n) \mid e_i < b_i\} \tag{25}$$

Note that we have a strict inequality in (25) since we know that Neg'_p is an open set if γ_p is continuous. Given this characterisation of the regions Neg'_p , we now easily find that scoring functions of the form $\text{Neg}'_{\mathcal{Q}}$ can only be linear in the trivial case where $|\mathcal{Q}| = 1$.

Proposition 13. *Suppose that the epistemic pooling principle (10) is satisfied for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, with $\diamond = \diamond_{\text{max}}$. Suppose exhaustiveness (11) is satisfied. Let $\mathcal{Q} = \{p_1, \dots, p_k\}$ be a subset of \mathcal{P} . Suppose $\gamma_{p_1}, \dots, \gamma_{p_k}$ and $\gamma_{\mathcal{Q}}$ are all linear. Then we have $|\mathcal{Q}| = 1$.*

Proof. We know that for each property p_i , there exists an index $\ell_i \in \{1, \dots, n\}$ and a corresponding threshold $b_i \in \mathbb{R}$ such that:

$$\text{Neg}'_{p_i} = \{(e_1, \dots, e_n) \mid e_{\ell_i} < b_i\}$$

and thus we also have:

$$\text{Pos}'_{p_i} = \{(e_1, \dots, e_n) \mid e_{\ell_i} \geq b_i\}$$

Note that $\text{Pos}'_{\mathcal{Q}} = \text{Pos}'_{p_1} \cap \dots \cap \text{Pos}'_{p_n}$. This means:

$$\text{Pos}'_{\mathcal{Q}} = \{(e_1, \dots, e_n) \mid e_{\ell_1} \geq b_1, \dots, e_{\ell_k} \geq b_k\}$$

Clearly, $\text{Pos}'_{\mathcal{Q}}$ can only be linearly separated from $\text{Neg}'_{\mathcal{Q}}$ if $\ell_1 = \dots = \ell_k$. Given (11), this is only possible if $p_1 = \dots = p_k$. \square

Going beyond linear scoring functions, it is possible to satisfy faithfulness (24) even with $X = \mathbb{R}^n$. To see this, let $\mathcal{P} = \{p_1, \dots, p_n\}$ and let us define:

$$\gamma_{p_i}(e_1, \dots, e_n) = -\text{ReLU}(-e_i)$$

with $\text{ReLU}(a) = \max(0, a)$ the rectified linear unit. Note that $\gamma_{p_i}(e_1, \dots, e_n) \geq 0$ iff $\text{ReLU}(-e_i) = 0$ iff $e_i \geq 0$. With this choice, we thus have that Neg'_{p_i} is of the form (25). It is easy to verify that (10) is satisfied for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, and that (11) is also satisfied. For $\mathcal{Q} = \{p_{i_1}, \dots, p_{i_k}\}$ we define:

$$\gamma_{\mathcal{Q}}(e_1, \dots, e_n) = -\text{ReLU}(-e_{i_1}) - \dots - \text{ReLU}(-e_{i_k})$$

Then we have $\gamma_{\mathcal{Q}}(e_1, \dots, e_n) \geq 0$ iff $\text{ReLU}(-e_{i_1}) = \dots = \text{ReLU}(-e_{i_k}) = 0$, which is the case if and only if $e_{i_1} \geq 0, \dots, e_{i_k} \geq 0$, which is equivalent with $\gamma_{p_{i_1}}(e_1, \dots, e_n) \geq 0, \dots, \gamma_{p_{i_k}}(e_1, \dots, e_n) \geq 0$. We thus find that faithfulness (24) is indeed satisfied. Note that through the use of the ReLU activation function, we essentially end up in the setting where $X =]-\infty, 0]^n$, where linear scoring functions are possible.

5.3.2. Hadamard product

For the Hadamard product \diamond_{had} we find that linear scoring functions can be used under the weak semantics, provided that we choose $X = [0, +\infty[^n$. To see this, let $\mathcal{P} = \{p_1, \dots, p_n\}$ and let the scoring functions γ_{p_i} be defined as follows:

$$\gamma_{p_i}(e_1, \dots, e_n) = -e_i$$

It is straightforward to verify that the epistemic pooling principle (10) is then satisfied for every $\mathbf{e}, \mathbf{f} \in [0, +\infty[^n$, while exhaustiveness (11) is also satisfied. For $\mathcal{Q} = \{p_{i_1}, \dots, p_{i_k}\}$ we define:

$$\gamma_{\mathcal{Q}}(e_1, \dots, e_n) = -e_{i_1} - \dots - e_{i_k}$$

Then, for $(e_1, \dots, e_n) \in [0, +\infty[^n$, we have $\gamma_{\mathcal{Q}}(e_1, \dots, e_n) \geq 0$ iff $e_{i_1} = \dots = e_{i_k} = 0$ iff $\gamma_{p_{i_\ell}} \geq 0$ for every $\ell \in \{1, \dots, k\}$. Thus we find that faithfulness (24) is indeed satisfied.

In Section 4.4, we also considered the case where $X = \mathbb{R}^n$. Unfortunately, for this choice of X , scoring functions of the form $\gamma_{\mathcal{Q}}$ can only be linear in the trivial case where $|\mathcal{Q}| = 1$.

Proposition 14. *Suppose that the epistemic pooling principle (10) is satisfied for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, with $\diamond = \diamond_{\text{had}}$. Suppose exhaustiveness (11) is satisfied. Let $\mathcal{Q} = \{p_{i_1}, \dots, p_{i_k}\}$ be a subset of \mathcal{P} . Suppose $\gamma_{p_{i_1}}, \dots, \gamma_{p_{i_k}}$ and $\gamma_{\mathcal{Q}}$ are all linear. Then we have $|\mathcal{Q}| = 1$.*

Proof. Let $H_{\mathcal{Q}}$ be the hyperplane associated with $\gamma_{\mathcal{Q}}$, i.e. $H_{\mathcal{Q}} = \{\mathbf{e} \mid \gamma_{\mathcal{Q}}(\mathbf{e}) = 0\}$, and let H_1, \dots, H_k similarly be the hyperplanes associated with $\gamma_{p_{i_1}}, \dots, \gamma_{p_{i_k}}$. Clearly, for each $\mathbf{e} \in H_{\mathcal{Q}}$ we have $\gamma_{p_{i_1}}(\mathbf{e}) \geq 0, \dots, \gamma_{p_{i_k}}(\mathbf{e}) \geq 0$. This is only possible if the hyperplanes $H_{\mathcal{Q}}, H_1, \dots, H_k$ are all parallel. Given that we assumed that (11) is satisfied, this is only possible if $p_{i_1} = \dots = p_{i_k}$. \square

Similar as we found for max-pooling, it is possible to satisfy faithfulness (24) for $X = \mathbb{R}^n$ by using a non-linear activation function. In this case, the aim of this activation function is to map all vectors $\mathbf{e} \in \mathbb{R}^n$ to some vector in $[0, +\infty[^n$, where we know that linear scoring functions are possible. In particular, we can use the following:

$$\begin{aligned} \gamma_{p_i}(e_1, \dots, e_n) &= -e_i^2 \\ \gamma_{\mathcal{Q}}(e_1, \dots, e_n) &= -e_{i_1}^2 - \dots - e_{i_k}^2 \end{aligned}$$

It is easy to verify that faithfulness (24) is indeed satisfied for this choice, while the epistemic pooling principle (10) is satisfied for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$ and exhaustiveness (11) is also satisfied.

5.4. Reasoning with average pooling

The results we have established so far suggest that neural networks are inherently limited in their reasoning abilities when averaging (or summation) is used. Graph Neural Networks often rely on this pooling operator, even when they are used in applications where they are expected to learn to carry out reasoning tasks. In this section, we analyse whether the empirical success of GNNs in such settings can be explained by weakening the epistemic pooling principle.

In Section 5.2 we found that linear scoring functions cannot be used under the strict semantics. Note that this limitation holds regardless of how the embeddings were obtained. Moreover, the limitation arises as soon as we have two properties p_1 and p_2 , such that we need to be able to check whether p_1 holds, whether p_2 holds, and whether p_1 and p_2 both hold. In Section 4.1 we found that continuous scoring functions cannot be used under the weak semantics for \diamond_{avg} . Since the limitation holds for arbitrary continuous scoring functions, adding a (continuous) non-linearity after pooling cannot resolve

the issue. Hence, there does not appear to exist a straightforward mechanism to use average pooling while also satisfying the epistemic pooling principle, which seems to be at odds with the prevalence and empirical success of models that rely on average pooling. One possible argument is that the epistemic pooling principle is perhaps too restrictive. While it should not be the case that $\mathbf{e} \diamond_{\text{avg}} \mathbf{f}$ captures properties that are not captured by \mathbf{e} or \mathbf{f} , we only need this principle to hold for the embeddings \mathbf{e}, \mathbf{f} that we are likely to encounter in practice. In particular, due to the way neural networks are trained, embeddings satisfying some property p are typically separated by *some margin* from embeddings which do not. Formally, let $\Delta > 0$ represent a given margin. Then we can expect that embeddings \mathbf{e} will either satisfy $\gamma_p(\mathbf{e}) \leq 0$ or $\gamma_p(\mathbf{e}) \geq \Delta$. If this is the case for every property p , we can think of \mathbf{e} as representing a clear-cut epistemic state. In contrast, when $0 < \gamma_p(\mathbf{e}) < \Delta$, we can think of \mathbf{e} as being ambiguous regarding the satisfaction of p . Let us write $X^* \subseteq X$ for the clear-cut epistemic states, i.e.

$$X^* = \{\mathbf{e} \in X \mid \forall p \in \mathcal{P}. \gamma_p(\mathbf{e}) \leq 0 \vee \gamma_p(\mathbf{e}) \geq \Delta\} \quad (26)$$

We show that it is possible to satisfy the epistemic pooling principle for all embeddings from X , while using scoring functions that are linear after the application of a ReLU activation, provided that we only require the scoring functions to satisfy faithfulness (23) for embeddings \mathbf{e} in X^* .

Let $X = [0, +\infty]^n$, $\Delta > 0$ and $\mathcal{P} = \{p_1, \dots, p_n\}$. Let us define the scoring functions γ_p as follows.

$$\gamma_{p_i}(e_1, \dots, e_n) = e_i$$

As in Section 4.1, we find that with this choice, the epistemic pooling principle (6) is satisfied for all $\mathbf{e}, \mathbf{f} \in X$ while exhaustiveness (8) is also satisfied. Note furthermore that $X^* = (\{0\} \cup [\Delta, +\infty])^n$. For $\mathcal{Q} = \{p_{i_1}, \dots, p_{i_k}\}$ we define:

$$\gamma_{\mathcal{Q}}(e_1, \dots, e_n) = \Delta - \sum_{\ell=1}^k \text{ReLU}(\Delta - e_{i_\ell})$$

Then (23) is satisfied for any $(e_1, \dots, e_n) \in X^*$. Indeed, for $(e_1, \dots, e_n) \in X^*$ we find:

$$\begin{aligned} \gamma_{\mathcal{Q}}(e_1, \dots, e_n) > 0 &\Leftrightarrow \sum_{\ell=1}^k \text{ReLU}(\Delta - e_{i_\ell}) < \Delta \\ &\Leftrightarrow \forall \ell \in \{1, \dots, k\}. e_{i_\ell} \geq \Delta \\ &\Leftrightarrow \forall \ell \in \{1, \dots, k\}. \gamma_{p_{i_\ell}}(e_1, \dots, e_n) > 0 \end{aligned}$$

where we used the fact that either $\text{ReLU}(\Delta - e_{i_\ell}) = 0$ or $\text{ReLU}(\Delta - e_{i_\ell}) = \Delta$, given that $(e_1, \dots, e_n) \in X^*$. Note that the ReLU transformation converts (e_1, \dots, e_n) into a vector that is binary, in the sense that each coordinate is either 0 or Δ . Using the sigmoid function σ we can similarly ensure that faithfulness (23) is satisfied for any $(e_1, \dots, e_n) \in X^*$, by constructing vectors that are approximately binary in the aforementioned sense. For instance, we could define:

$$\gamma_{\mathcal{Q}}(e_1, \dots, e_n) = \mu - \sum_{\ell=1}^k \sigma\left(\lambda \left(\frac{\Delta}{2} - e_{i_\ell}\right)\right)$$

where μ is an arbitrary constant satisfying $0 < \mu < 1$ and $\lambda > 0$ is chosen in function of μ . In particular, by choosing a sufficiently large value for λ , we can always ensure that

$$\sigma\left(\lambda \cdot \frac{\Delta}{2}\right) \geq \mu$$

and

$$\sigma\left(-\lambda \cdot \frac{\Delta}{2}\right) < \frac{\mu}{k}$$

We have $\gamma_{\mathcal{Q}}(e_1, \dots, e_n) > 0$ iff

$$\sum_{\ell=1}^k \sigma\left(\lambda \left(\frac{\Delta}{2} - e_{i_\ell}\right)\right) < \mu$$

Given the assumption that λ is sufficiently large, and given that $(e_1, \dots, e_n) \in X^*$, this inequality is satisfied iff $e_{i_\ell} \geq \Delta$ for every $\ell \in \{1, \dots, k\}$. The rest of the argument then follows as before.

We can also use linear scoring functions, by defining X^* such that embeddings are approximately binary. In particular, let $X = [0, 1]^n$ with $\Delta = 1 - \varepsilon$ for some $\varepsilon \in]0, \frac{1}{k}[$. Let $\gamma_{p_j}(e_1, \dots, e_n) = e_j$ as before. For $\mathcal{Q} = \{p_{i_1}, \dots, p_{i_k}\}$, we define:

$$\gamma_{\mathcal{Q}}(e_1, \dots, e_n) = \left(\sum_{\ell=1}^k e_{i_\ell} \right) - k + 1$$

We then have $\gamma_{\mathcal{Q}}(e_1, \dots, e_n) > 0$ iff

$$\sum_{\ell=1}^k e_{i_\ell} > k - 1$$

Given our assumption that $\varepsilon < \frac{1}{k}$, this inequality is satisfied iff $e_{i_\ell} \geq \Delta = 1 - \varepsilon$ for every $\ell \in \{1, \dots, k\}$. The rest of the argument then follows as before.

Note that in each case, we either restrict the setting to embeddings that are approximately binary, or we apply a non-linear transformation to convert the initial embeddings into approximately binary ones. In this way, we replaced the need for discontinuous scoring functions by the assumption that embeddings essentially encode discrete states. It remains unclear whether models that are trained using gradient descent based optimisation can learn to reason under such conditions.

6. Modelling weighted epistemic states

We now consider a setting in which each property p from \mathcal{P} is associated with a certainty level from $\Lambda = \{0, 1, \dots, K\}$. Intuitively, a certainty level of 0 means that we know nothing about p , whereas a level of K means that we are fully certain that p is true. When aggregating evidence from different sources, we assume that the certainty level of p is given by the certainty level of the most confident source, in accordance with possibility theory [18]. This setting has the advantage that we can continue to view pooling operators in terms of accumulating knowledge.

Reduction to the standard setting It is possible to model weighted epistemic states within the standard framework we have considered so far, using scoring functions of the form $\gamma_{\mathcal{Q}}$. We consider the strict semantics here, but an entirely similar argument can be made for the weak semantics. Let us define the set of extended properties as follows:

$$\mathcal{P}^\sharp = \{p\sharp i \mid p \in \mathcal{P}, i \in \Lambda\}$$

For each $p\sharp i \in \mathcal{P}^\sharp$, we let $\gamma_{p\sharp i}$ be a scoring function satisfying (6) for all embeddings $\mathbf{e}, \mathbf{f} \in X$. We can interpret $p\sharp i$ as encoding that the certainty level of p is not equal to i . For $p \in \mathcal{P}$ and $i \in \Lambda \setminus \{0\}$ we define:

$$\mathcal{Q}_{(p,i)} = \{p\sharp j \mid p \in \mathcal{P}, j \in \Lambda, j < i\}$$

Let $\gamma_{\mathcal{Q}_{(p,i)}}$ be a scoring function satisfying (23). Then $\gamma_{\mathcal{Q}_{(p,i)}}(\mathbf{e}) > 0$ means that all certainty levels below i can be excluded for p , i.e. that p is certain at least to degree i .

Weighted epistemic pooling principle The construction from the previous paragraph has an important drawback: the number of properties is increased $(K + 1)$ -fold. From Section 4, we know that this implies that the number of dimensions also has to increase $(K + 1)$ -fold. If for every p and i , we need the ability to model that the certainty of p is not i , this increase in dimensionality is inevitable. However, in practice, we are typically not interested in excluding arbitrary sets of certainty degrees, only in establishing lower bounds on certainty degrees. To study this setting, we introduce a generalisation of the epistemic pooling principle to weighted epistemic states. Let us write $\langle p, i \rangle$ for the fact that the certainty level of p is at least i , where $\langle p, 0 \rangle$ means that nothing is known about p whereas $\langle p, K \rangle$ means that p is known with full certainty. We write $\Lambda_0 = \{1, 2, \dots, K\}$ for the set of non-trivial lower bounds. We furthermore assume that the certainty level of p is determined by the scoring function $\gamma_p : \mathbb{R}^n \rightarrow \mathbb{R}$. In particular, under the strict semantics, we assume that p is known with certainty at least i , for $i \in \Lambda_0$, if $\gamma_p > i - 1$. We then consider the following generalisation of the epistemic pooling principle:

$$\forall p \in \mathcal{P}. \forall i \in \Lambda_0. (\gamma_p(\mathbf{e}) > i - 1) \vee (\gamma_p(\mathbf{f}) > i - 1) \Leftrightarrow \gamma_p(\mathbf{e} \diamond \mathbf{f}) > i - 1 \quad (27)$$

We define Γ_Λ as a set of weighted properties, encoding for each property what is the highest certainty degree with which this property is believed:

$$\Gamma_\Lambda(\mathbf{e}) = \{\langle p, \max(0, \min(K, \lceil \gamma_p(\mathbf{e}) \rceil)) \rangle \mid p \in \mathcal{P}\} \quad (28)$$

Exhaustiveness, i.e. the condition that every weighted epistemic state is modelled by some vector $\mathbf{e} \in X$ can then be formalised as follows:

$$\forall \mu : \mathcal{P} \rightarrow \Lambda. \exists \mathbf{e} \in X. \Gamma_\Lambda(\mathbf{e}) = \{\langle p, \mu(p) \rangle \mid p \in \mathcal{P}\} \quad (29)$$

In other words, for every assignment of certainty degrees to the properties in \mathcal{P} , there exists an embedding $\mathbf{e} \in X$ that encodes the corresponding epistemic state. We also introduce the following notations for $i \in \Lambda_0$:

Table 3

Summary of results about the realizability of the weighted epistemic pooling principles for embeddings of dimensionality $n = |\mathcal{P}|$.

Pooling Operator	Semantics	$n = \mathcal{P} $ possible?
Average	Strict	\times
	Weak	\times
Summation	Strict	\times
	Weak	\times
Max-pooling	Strict	\checkmark
	Weak	\checkmark
Hadamard	Strict	\times
	Weak	\times

$$\text{Pos}_{p,i} = \{\mathbf{e} \in X \mid \gamma_p(\mathbf{e}) > i - 1\}$$

$$\text{Neg}_{p,i} = \{\mathbf{e} \in X \mid \gamma_p(\mathbf{e}) \leq i - 1\} = X \setminus \text{Pos}_{p,i}$$

Similarly, we can also consider a weighted version of the weak epistemic pooling principle:

$$\forall p \in \mathcal{P}. \forall i \in \Lambda_0. (\gamma_p(\mathbf{e}) \geq i - 1) \vee (\gamma_p(\mathbf{f}) \geq i - 1) \Leftrightarrow \gamma_p(\mathbf{e} \diamond \mathbf{f}) \geq i - 1 \quad (30)$$

The weighted epistemic state associated with an embedding $\mathbf{e} \in X$ is now defined as follows:

$$\Gamma'_\Lambda(\mathbf{e}) = \{\langle p, \max(0, \min(K, 1 + \lfloor \gamma_p(\mathbf{e}) \rfloor) \rangle \mid p \in \mathcal{P}\} \quad (31)$$

while the counterpart of exhaustiveness (29) becomes

$$\forall \mu : \mathcal{P} \rightarrow \Lambda. \exists \mathbf{e} \in X. \Gamma'_\Lambda(\mathbf{e}) = \{\langle p, \mu(p) \rangle \mid p \in \mathcal{P}\} \quad (32)$$

We now analyse whether the weighted epistemic pooling principles (27) and (30) can be satisfied in a non-trivial way for embeddings with $|\mathcal{P}|$ dimensions. We find that this is only the case for \diamond_{\max} , as shown in Table 3.

6.1. Realizability of the weighted epistemic pooling principle

Average In entirely the same way as in Section 4.1, we then find that $\text{Pos}_{p,i}$ and $\text{Neg}_{p,i}$ are convex, for every $p \in \mathcal{P}$ and $i \in \Lambda_0$. Similarly, we also find that $\text{Neg}_{p,i} \subseteq \delta(X)$ for every $p \in \mathcal{P}$ and $i \in \Lambda_0$. Since $\text{Neg}_{p,i} \setminus \text{Pos}_{p,K} \subseteq \text{Neg}_{p,K}$, we thus also have that $\text{Pos}_{p,i} \setminus \text{Pos}_{p,K} \subseteq \delta(X)$. We can now show the following result using a similar strategy as in the proof of Proposition 2.

Proposition 15. *Suppose the weighted epistemic pooling principle (27) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{avg}}$ and $X \subseteq \mathbb{R}^n$. Suppose that exhaustiveness (29) is satisfied. It holds that $n \geq |\mathcal{P}| \cdot K$.*

Proof. Let $p_1, \dots, p_{|\mathcal{P}|}$ be an enumeration of the properties in \mathcal{P} . Note that because of (29), we have that $\text{Neg}_{p_1,K} \neq \emptyset$ and $\text{Pos}_{p_1,K} \neq \emptyset$. Since both of these regions are convex, it follows from the hyperplane separation theorem that there exists a hyperplane $H_{1,K}$ which separates $\text{Neg}_{p_1,K}$ and $\text{Pos}_{p_1,K}$. Since $\text{Neg}_{p_1,K} \subseteq \delta(X) = \text{cl}(\text{Pos}_{p_1,K})$, we furthermore know that $\text{Neg}_{p_1,K} \subseteq H_{1,K}$.

Note that $H_{1,K} \cap \text{Neg}_{p_1,K-1}$ and $H_{1,K} \cap \text{Pos}_{p_1,K-1}$ are convex regions. Moreover, since $\text{Neg}_{p_1,K} \subseteq H_{1,K}$, $\text{Neg}_{p_1,K} \cap \text{Neg}_{p_1,K-1} \neq \emptyset$, and $\text{Neg}_{p_1,K} \cap \text{Pos}_{p_1,K-1} \neq \emptyset$, we find that $H_{1,K} \cap \text{Neg}_{p_1,K-1} \neq \emptyset$ and $H_{1,K} \cap \text{Pos}_{p_1,K-1} \neq \emptyset$. It follows from the hyperplane separation theorem that there exists some hyperplane $H_{1,K-1}$ separating $H_{1,K} \cap \text{Neg}_{p_1,K-1}$ and $H_{1,K} \cap \text{Pos}_{p_1,K-1}$. Moreover, it holds that $H_{1,K} \cap \text{Neg}_{p_1,K-1} \subseteq H_{1,K-1}$ (which follows in the same way as $H_1 \cap \text{Neg}_{p_2} \subseteq H_2$ in the proof of Proposition 2). In particular, we then also have that $\text{Neg}_{p_1,K} \cap \text{Neg}_{p_1,K-1} \subseteq H_{1,K-1}$. We can repeat this argument to show that for each $i \in \{0, \dots, K-1\}$ there is a hyperplane $H_{1,K-i}$ separating $H_{1,K} \cap \dots \cap H_{1,K-i+1} \cap \text{Neg}_{p_1,K-i}$ and $H_{1,K} \cap \dots \cap H_{1,K-i+1} \cap \text{Pos}_{p_1,K-i}$, such that $H_{1,K} \cap \dots \cap H_{1,K-i+1} \cap \text{Neg}_{p_1,K-i} \subseteq H_{1,K-i}$. This implies that $H_{1,K} \cap \dots \cap H_{1,K-i+1} \not\subseteq H_{1,K-i}$ and thus $\dim(H_{1,K} \cap \dots \cap H_{1,K-i}) = n - i - 1$. In particular, we have $\dim(H_{1,K} \cap \dots \cap H_{1,1}) = n - K$.

Due to (29), we have that $H_{1,K} \cap \dots \cap H_{1,1} \cap \text{Neg}_{p_2,K} \neq \emptyset$ and $H_{1,K} \cap \dots \cap H_{1,1} \cap \text{Pos}_{p_2,K} \neq \emptyset$. It follows that there exists a hyperplane $H_{2,K}$ separating $H_{1,K} \cap \dots \cap H_{1,1} \cap \text{Neg}_{p_2,K}$ and $H_{1,K} \cap \dots \cap H_{1,1} \cap \text{Pos}_{p_2,K}$, such that $H_{1,K} \cap \dots \cap H_{1,1} \cap \text{Neg}_{p_2,K} \subseteq H_{2,K}$, from which it follows that $\dim(H_{1,K} \cap \dots \cap H_{1,1} \cap H_{2,K}) = n - K - 1$. Continuing as before, we find $\dim(H_{1,K} \cap \dots \cap H_{1,1} \cap H_{2,K} \cap \dots \cap H_{2,1}) = n - 2K$. Repeating this argument for every property in \mathcal{P} , we find:

$$\dim \left(\bigcap_{j=1}^{|\mathcal{P}|} \bigcap_{i=1}^k H_{j,i} \right) = n - |\mathcal{P}| \cdot K.$$

Since this is only possible if $n - |\mathcal{P}| \cdot K \geq 0$ it follows that $n \geq |\mathcal{P}| \cdot K$. \square

In Section 4.1, we found that the epistemic pooling principle cannot be satisfied in a non-trivial way under the weak semantics, if continuous scoring functions are used. This result carries over to the weighted setting.

Summation Entirely similar as in Section 4.2, we find that when the weighted epistemic pooling principle is satisfied for all $\mathbf{e}, \mathbf{f} \in X$ and $\diamond = \diamond_{\text{sum}}$, then it is also satisfied for $\diamond = \diamond_{\text{avg}}$. It thus follows from Proposition 15 that $|\mathcal{P}| \cdot K$ dimensions are needed to satisfy the weighted epistemic pooling principle (27) in a non-trivial way. Moreover, we also have that the weighted epistemic pooling principle under the weak semantics, i.e. the weak weighted epistemic pooling principle (30), cannot be satisfied when continuous scoring functions are used.

Max-pooling For \diamond_{max} , we can satisfy the weighted epistemic pooling principles (27) and (30) for every $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$, while ensuring that every weighted epistemic state is encoded by some vector. Let $\mathcal{P} = \{p_1, \dots, p_n\}$ and let γ_{p_i} be defined as follows:

$$\gamma_{p_i}(e_1, \dots, e_n) = e_i$$

It holds that the strict weighted epistemic pooling principle (27) is satisfied for each $\mathbf{e}, \mathbf{f} \in \mathbb{R}^n$. Indeed, for every $j \in \Lambda_0$ we have:

$$\begin{aligned} & (\gamma_{p_i}(e_1, \dots, e_n) > j - 1) \vee (\gamma_{p_i}(f_1, \dots, f_n) > j - 1) \\ & \Leftrightarrow (e_i > j - 1) \vee (f_i > j - 1) \\ & \Leftrightarrow \max(e_i, f_i) > j - 1 \\ & \Leftrightarrow \gamma_{p_i}((e_1, \dots, e_n) \diamond_{\text{max}} (f_1, \dots, f_n)) > j - 1 \end{aligned}$$

In entirely the same way, we find that the weak weighted epistemic pooling principle (30) is satisfied. We now show that exhaustiveness is satisfied, i.e. (29) and (32). Let $\mu : \mathcal{P} \rightarrow \Lambda$. Then we can define $\mathbf{e} = (e_1, \dots, e_n)$ as follows:

$$e_i = \mu(p_i) - \frac{1}{2}$$

It is trivial to verify that $\Gamma_{\Lambda}(\mathbf{e}) = \Gamma'_{\Lambda}(\mathbf{e}) = \{\langle p, \mu(p) \rangle \mid p \in \mathcal{P}\}$.

Hadamard product As in Section 4.4, we find that $\text{Pos}_{p,\ell}$ is a finite union of regions of the form $X \cap \bigcap_{i \in I} H_i$, with $H_i = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_i = 0\}$. For a given index set $I \subseteq \{1, \dots, n\}$, we define:

$$\mathcal{P}_I^{\Lambda} = \{(p, \ell) \in \mathcal{P} \times \Lambda_0 \mid X \cap \bigcap_{i \in I} H_i \subseteq \text{Pos}_{p,\ell}\}$$

In entirely the same way as in Lemma 9, we find for every $I, J \subseteq \{1, \dots, n\}$ that:

$$\mathcal{P}_{I \cup J}^{\Lambda} = \mathcal{P}_I^{\Lambda} \cup \mathcal{P}_J^{\Lambda} \tag{33}$$

The following lemma is also shown in entirely the same way as Lemma 8.

Lemma 14. *Let $X = \mathbb{R}^n$ or $X = [0, +\infty]^n$. Suppose the weighted epistemic pooling principle (27) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{had}}$. Let $p \in \mathcal{P}$ and $\ell \in \Lambda_0$ such that $\mathbf{e} = (e_1, \dots, e_n) \in \text{Pos}_{p,\ell}$. Let $I = \{i \in \{1, \dots, n\} \mid e_i = 0\}$. It holds that*

$$X \cap \bigcap_{i \in I} H_i \subseteq \text{Pos}_{p,\ell}$$

We can then generalise Proposition 9 as follows.

Proposition 16. *Let $X = \mathbb{R}^n$ or $X = [0, +\infty]^n$. Suppose the weighted epistemic pooling principle (27) is satisfied for all embeddings $\mathbf{e}, \mathbf{f} \in X$, with $\diamond = \diamond_{\text{had}}$. Suppose that exhaustiveness (29) is satisfied. It holds that $n \geq |\mathcal{P}| \cdot K$.*

Proof. Given (29), for each $p \in \mathcal{P}$ and $\ell \in \Lambda_0$, there must exist some $(e_1^{p,\ell}, \dots, e_n^{p,\ell}) \in X$ such that $\Gamma(e_1^{p,\ell}, \dots, e_n^{p,\ell}) = \{\langle p, \ell \rangle\} \cup \{\langle q, 0 \rangle \mid q \in \mathcal{P} \setminus \{p\}\}$. Let us fix such vectors $(e_1^{p,\ell}, \dots, e_n^{p,\ell})$ for each $p \in \mathcal{P}$ and $\ell \in \Lambda_0$. Define $I_{p,\ell} = \{i \in \{1, \dots, n\} \mid e_i^{p,\ell} = 0\}$. Note that by Lemma 14, we have $X \cap \bigcap_{i \in I_{p,\ell}} H_i \subseteq \text{Pos}_{p,\ell}$. Moreover, by construction, we have $X \cap \bigcap_{i \in I_{p,\ell}} H_i \not\subseteq \text{Pos}_{q,\ell'}$ for any $q \neq p$ and $\ell' \in \Lambda_0$, and similarly, we have $X \cap \bigcap_{i \in I_{p,\ell}} H_i \not\subseteq \text{Pos}_{p,\ell'}$ for any $\ell' > \ell$. In other words, we have that all elements in $\mathcal{P}_{I_{p,\ell}}^{\Lambda}$ are of the form (p, ℓ') with $\ell' \leq \ell$. For $p \neq q$ and $\ell, \ell' \in \Lambda_0$, we clearly have $I_{p,\ell} \not\subseteq I_{q,\ell'}$, since $I_{p,\ell} \subseteq I_{q,\ell'}$ would imply $\langle p, \ell \rangle \in \mathcal{P}_{I_{q,\ell'}}$. This means in particular that $I_{p,\ell} \neq \emptyset$ for every $p \in \mathcal{P}$ and $\ell \in \Lambda_0$.

Let $(q_1, \ell_1), \dots, (q_s, \ell_s)$, with $s = |\mathcal{P}| \cdot K$, be an enumeration of the elements from $\mathcal{P} \times \Lambda_0$, such that whenever $\ell < \ell'$, it holds that (p, ℓ) is listed before (p, ℓ') , for any $p \in \mathcal{P}$. Note that we have already established that $I_{q_1, \ell_1} \neq \emptyset$. Now we show that for any given $k \in \{2, \dots, s\}$, it holds that $I_{q_k, \ell_k} \not\subseteq I_{q_1, \ell_1} \cup I_{q_2, \ell_2} \cup \dots \cup I_{q_{k-1}, \ell_{k-1}}$. Indeed $I_{q_k, \ell_k} \subseteq I_{q_1, \ell_1} \cup I_{q_2, \ell_2} \cup \dots \cup I_{q_{k-1}, \ell_{k-1}}$ would imply $(q_k, \ell_k) \in \mathcal{P}_{I_{q_1, \ell_1} \cup \dots \cup I_{q_{k-1}, \ell_{k-1}}}^\Delta$. From (33), we know that $\mathcal{P}_{I_{q_1, \ell_1} \cup \dots \cup I_{q_{k-1}, \ell_{k-1}}}^\Delta = \mathcal{P}_{I_{q_1, \ell_1}}^\Delta \cup \dots \cup \mathcal{P}_{I_{q_{k-1}, \ell_{k-1}}}^\Delta$ and we know that the latter can only contain elements of the form $(q_1, \ell'_1), \dots, (q_{k-1}, \ell'_{k-1})$ with $\ell'_1 \leq \ell_1, \dots, \ell'_{k-1} \leq \ell_{k-1}$. Due to the assumption we made that (p, ℓ) is listed before (p, ℓ') whenever $\ell < \ell'$, we thus find $(q_k, \ell_k) \notin \mathcal{P}_{I_{q_1, \ell_1} \cup \dots \cup I_{q_{k-1}, \ell_{k-1}}}^\Delta$, a contradiction. Hence, we have $I_{q_k, \ell_k} \not\subseteq I_{q_1, \ell_1} \cup I_{q_2, \ell_2} \cup \dots \cup I_{q_{k-1}, \ell_{k-1}}$. This means there is at least one element in $I_{(q_k, \ell_k)}$ which does not occur in $I_{q_1, \ell_1}, \dots, I_{q_{k-1}, \ell_{k-1}}$. Since this needs to hold for every $k \in \{2, \dots, s\}$, there need to be at least s distinct elements in $I_{q_1, \ell_1} \cup \dots \cup I_{q_s, \ell_s}$. In particular, we thus have that $n \geq s = |\mathcal{P}| \cdot K$. \square

The above limitation also applies to the weak semantics. However, if we choose $X = [0, 1]^n$ it is possible to satisfy the weighted epistemic pooling principle (27) for n properties if $K = 2$ (i.e. if we have three certainty degrees). Indeed, let $\mathcal{P} = \{p_1, \dots, p_n\}$. Then we can define:

$$\gamma_{p_i}(e_1, \dots, e_n) = \begin{cases} \frac{3}{2} & \text{if } e_i = 0 \\ -\frac{1}{2} & \text{if } e_i = 1 \\ \frac{1}{2} & \text{otherwise} \end{cases} \quad (34)$$

Then it is straightforward to verify that the weighted epistemic pooling principle (27) is indeed satisfied for every $\mathbf{e}, \mathbf{f} \in [0, 1]^n$ and that (29) also holds. The above construction also provides an example of how the weighted epistemic pooling principle can be satisfied for the weak semantics. Indeed, with the above definition of γ_{p_i} we have that the weak weighted epistemic pooling principle (30) is satisfied for every $\mathbf{e}, \mathbf{f} \in [0, 1]^n$, while exhaustiveness (32) also holds. In Section 4.4, we found that the weak epistemic pooling principle could be satisfied with continuous scoring functions when $\diamond = \diamond_{\text{had}}$. Unfortunately, this strategy does not allow us to obtain a continuous alternative to the scoring functions defined in (34).

7. Discussion and related work

7.1. Logical reasoning with neural networks

The use of neural networks for simulating symbolic reasoning has been extensively studied under the umbrella of neuro-symbolic reasoning [10]. The seminal KBANN method [19], for instance, uses feedforward networks with carefully chosen weights to simulate the process of reasoning with a given rule base. In this case, the neural network simulates a fixed rule base, which is manually specified. More recent work has investigated how a neural network can be trained to simulate the deductive closure of a given logical theory, e.g. a description logic ontology [20]. The idea that rule-based reasoning can be simulated using neural networks also lies at the basis of various strategies for learning rules from data [21–23]. Another recent research line has focused on whether standard neural network architectures, such as LSTMs or transformer based language models [6], can be trained to recognise logical entailment [24,25]. In these works, the input consists of a premise (or a set of premises) and a hypothesis, and the aim is to predict whether the hypothesis can be inferred from the premise(s). The aforementioned works differ from this paper, as our focus is not on whether neural networks can simulate logical reasoning, but on whether vectors can be used for encoding epistemic states. Simulating logical reasoning does not necessarily require that vectors can encode epistemic states, since we can treat reasoning as an abstract symbol manipulation problem. Moreover, the scope of what we address in this paper goes beyond logical reasoning, as the epistemic pooling principle also matters whenever we need to combine evidence from different sources (e.g. features being detected in different regions of an image).

This paper builds on our earlier work [26], where the focus was on the following question: given a set of attributes \mathcal{A} , a pooling function \diamond and a propositional knowledge base K , can we always find an embedding \mathbf{a} and a scoring function γ_a for every $a \in \mathcal{A}$ such that:

$$\gamma_b(\mathbf{a}_1 \diamond \dots \diamond \mathbf{a}_k) \geq 0 \quad \Leftrightarrow \quad K \cup \{a_1, \dots, a_k\} \models b$$

for any $a_1, \dots, a_k, b \in \mathcal{A}$. This analysis fundamentally differs from the results in this paper, because in the case of [26], we only care about the behaviour of the pooling operator and scoring function for a finite set of vectors, i.e. the attribute embeddings. In the notations of this paper, this amounts to limiting X to a finite set of carefully chosen embeddings. As a result, for instance, in [26] it was possible to use average pooling in combination with continuous scoring functions, under the weak semantics, something which we found to be impossible in the more general setting considered in this paper.

One may wonder whether the approach we take in this paper is too strict, e.g. whether it is really necessary to insist that the epistemic pooling principle is satisfied for all embeddings. Note, however, that during training, the embeddings will change after each update step of the (gradient-descent based) optimizer. If we want the model to learn to reason, then it is important that all of these embeddings can be viewed as epistemic states. In other words, if pooling is only meaningful for a particular discrete set of vectors, then we may end up with an architecture that is capable of reasoning in theory, but where those parameters that would lead to meaningful behaviour cannot be learned in practice.

7.2. Reasoning with graph neural networks

In Section 5, we have specifically focused on reasoning in the context of propositional logic. However, our analysis is relevant for reasoning in relational domains as well. A standard approach for relational reasoning with neural networks is to rely on Graph Neural Networks (GNNs) [8,27]. Given a graph $G = (V, E)$, with V a set of nodes and $E \subseteq V \times V$ a set of edges, a GNN aims to learn a vector representation of every node in V . This is achieved by incrementally updating the current representation of each node based on the representations of their neighbours. In particular, let us write $\mathbf{v}^{(i)}$ for the representation of node $v \in V$ in layer i of the GNN. Let $\{u_1, \dots, u_k\}$ be the set of neighbours of v in G . Then the representation of v in layer $i + 1$ is typically computed as follows:

$$\mathbf{v}^{(i+1)} = f_2(\mathbf{v}^{(i)} \diamond_1 f_2(f_3(\mathbf{u}_1^{(i)}) \diamond_2 \dots \diamond_2 f_3(\mathbf{u}_k^{(i)}))) \quad (35)$$

The functions f_1 , f_2 , f_3 , \diamond_1 and \diamond_2 can be defined in various ways. However, regardless of the specifics, we can think of \diamond_1 and \diamond_2 as pooling operators, whereas f_1 , f_2 and f_3 correspond to (possibly non-linear) transformations. Note that we used a non-standard notation in (35) to highlight the connection to this paper. Intuitively, we can think of $f_3(\mathbf{u}_j^{(i)})$ as a vector that captures what we can infer about the entity represented by node v from the fact that it is connected to u_j . In multi-relational settings (e.g. knowledge graphs), where different types of edges occur, f_3 can be replaced by a function that depends on the edge type. The pooling operator \diamond_2 is used to aggregate the evidence coming from the different neighbours of v , whereas \diamond_1 is used to combine the evidence we already have about v with the evidence we can obtain from its neighbours.

The ability of GNNs to simulate logical reasoning has been studied in [11]. The idea is that each node is associated with a set of properties, which are referred to as colours. We can then consider first-order formulas involving unary predicates, referring to these colours, and the binary predicate E , which captures whether two nodes are connected. For instance, consider the following formula:

$$\phi_1(x) \equiv \text{Green}(x) \wedge (\exists y . E(x, y) \wedge \text{Blue}(y)) \wedge \neg(\exists y . E(x, y) \wedge \text{Red}(y))$$

This formula is true for a given node if it is green and it is connected in the graph to a blue node but not to a red node. We can also consider counting quantifiers, as in the following example:

$$\phi_2(x) \equiv (\exists^{\geq 5} y . E(x, y) \wedge \text{Blue}(y))$$

This formula is true for a given node if it is connected to at least 5 blue nodes. The question studied in [11] is which formulas can be recognised by a GNN, i.e. for which class of formulas ϕ can we design a GNN such that we can predict whether ϕ holds for a node n from the final-layer embedding of that node, using some scoring function. In particular, it was shown that the set of formulas that can be recognised (without global read-out) are exactly those that are expressible in graded modal logic [28], which is characterised as follows:

1. for each colour C , the formula $C(x)$ is a graded modal logic formula;
2. if $\phi(x)$ and $\psi(x)$ are graded modal logic formulas and $n \in \mathbb{N}$, then the following formulas are also graded modal logic formulas: $\neg\phi(x)$, $\phi(x) \wedge \psi(x)$ and $\exists^{\geq n} y . E(x, y) \wedge \phi(y)$.

The proof that was provided for the characterisation in [11] is constructive. It relies on the particular choice of \diamond_2 as summation, which appears to be at odds with the limitations that were identified for \diamond_{sum} in this paper. However, the GNN in their construction only uses binary coordinates. As we have seen in Section 5.4, in that case, \diamond_{avg} can be used for propositional reasoning, a result which straightforwardly carries over to \diamond_{sum} . This observation may help to explain the discrepancy between the theoretical ability of GNNs to capture arbitrary formulas from graded logic, and the challenges that have empirically been observed when using GNNs for learning to reason. For instance, GNNs have generally failed to outperform simpler embedding based methods for the task of knowledge graph completion [29], while [30] found that GNNs were limited in their ability to generalise in a systematic way from examples that were more complex than the ones seen during training.

Interestingly, [31] recently proposed a GNN for knowledge graph completion in which \diamond_2 corresponds to max-pooling. Together with a number of other design choices (e.g. avoiding negative weights and using a particular encoding of the knowledge graph), this leads to GNNs that are in some sense equivalent to a set of rules. The suitability of max-pooling, in this context, is not a surprise, given our results from Section 5. Note, however, that our results also suggest that coordinates have to be upper-bounded if we want to identify cases where sets of atomic properties are jointly satisfied. In the approach from [31], this issue is avoided by using high-dimensional embeddings in which each candidate inference corresponds to a separate coordinate, which amounts to treating the formulas of interest as atomic properties in our framework.

7.3. Modelling relations as regions

A popular strategy for making predictions in relational domains consists in learning (i) an embedding $\mathbf{e} \in \mathbb{R}^n$ for each entity of interest e and (ii) scoring function $f_r : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ for each relation r such that $f_r(\mathbf{e}, \mathbf{f})$ indicates the probability

that the fact $r(e, f)$ holds. Despite the fact that such methods intuitively carry out some form of logical inference, and despite their strong empirical performance [29], for most approaches, there is no clear link between the parameters of the model (i.e. the embeddings and the parameters of the scoring functions), on the one hand, and the kinds of inferences that are captured, on the other hand. Region-based methods, however, are a notable exception [32–36]. The central idea of such methods is to represent predicates as regions. For instance, if s is a unary predicate, then the corresponding region R_s is such that $\mathbf{e} \in R_s$ iff $s(e)$ is true. For a binary predicate r , one option is to use a region R_r in \mathbb{R}^{2n} such that $\mathbf{e} \oplus \mathbf{f} \in R_r$ iff $r(e, f)$ is true. In other words, we view the concatenation of \mathbf{e} and \mathbf{f} as the embedding of the tuple (e, f) and model binary predicates as regions over such concatenations. The key advantage of region-based models is that logical relationships can be directly encoded, in terms of spatial relationships between the region-based representations of the predicates involved. As a simple example, the rule $r_1(x, y) \rightarrow r_2(x, y)$ is satisfied if $R_{r_1} \subseteq R_{r_2}$. We refer to [32] for details on how more complex rules can be similarly captured. This correspondence between logical dependencies and spatial relationships can be used to ensure that the predictions of the model are in accordance with a given knowledge base, or to explain the behaviour of a model in terms of the logical rules it captures. Moreover, region based embeddings make it possible to query embeddings of knowledge bases in a principled way [34,35]. Essentially, a given query (e.g. “retrieve all companies whose headquarter is in a European capital city”) is then mapped onto a region, such that the entities that satisfy the query are those whose embedding belongs to the region.

However, it should be noted that the aforementioned region-based embeddings encode a specific possible world, rather than an epistemic state. In other words, they encode which facts are assumed to be true and false, but they cannot encode incomplete knowledge (e.g. that either $r(a, b)$ or $s(a, b)$ holds). In [37] a geometric model based on cones was presented, which has the ability to encode incomplete knowledge to some extent. For instance, these geometric models can capture the fact that it is unknown whether some entity e belongs to some concept A . Essentially, each concept is represented by two cones. Entities whose representation belongs to the first cone are those which are known to instances of the concept; entities whose representation belongs to the second cone are those which are known not to be instances of the concept; and the remaining instances are those whose membership is unknown. Note, however, that not all epistemic states can be captured in this way; e.g. we cannot represent the fact that either e belongs to A or f belongs to B . Along similar lines, the box embeddings proposed in [38,39] can be used to model some epistemic states. For instance, an approach for capturing uncertain knowledge graph embeddings based on box embeddings was proposed in [40]. In this case, the idea is to represent the entities themselves as regions, and as axis-aligned hyperboxes in particular. The problem of pooling such hyperbox representations has not yet been considered, to the best of our knowledge. The most intuitive approach would be to simply take the intersection of the boxes, i.e. if entity e is represented by a hyperbox B_1 , according to one source, and by a hyperbox B_2 , according to another source, we may want to use $B_1 \cap B_2$ as an aggregate representation of entity e , reflecting the information provided by both sources. However, this leads to a number of practical challenges. For instance, if box embeddings are used to parameterise a probabilistic model, as in [40], then it is unclear whether a sound justification can be provided for a pooling operation that relies on intersecting the entity-level box representations. Such probabilistic models also serve a rather different purpose to the framework that we studied in this paper, which is about accumulating knowledge rather than about quantifying uncertainty. Even in settings where the box embeddings are used as a purely qualitative representation, a problem arises when the region $B_1 \cap B_2$ is empty. Intuitively, box embeddings act as constraints on possible worlds, and such constraints can be inconsistent. This is different from the settings we studied in this paper, which were about accumulating knowledge, formalised as sets of properties.

8. Conclusions

Neural networks are often implicitly assumed to perform some kind of reasoning. In this paper, we have particularly focused on the common situation where evidence is obtained from different sources, which then needs to be combined. The core question we addressed is whether it is possible to represent the evidence obtained from each source as a vector, such that pooling these vectors amounts to combining the corresponding evidence, a requirement we refer to as the *epistemic pooling principle*. This question is important for understanding whether, or under which conditions, neural networks that rely on pooling are able to perform reasoning in a principled way. Our analysis shows that standard pooling operators can indeed be used for accumulating evidence, but only under particular conditions. Broadly speaking, the requirement that the epistemic pooling principle is satisfied substantially limits how knowledge can be encoded. For instance, when average pooling is used, we find that embeddings have to be limited to a strict subset X of \mathbb{R}^n , and that vectors which encode that some property p is *not* satisfied have to be located on a bounding hyperplane of X . We also highlighted how such conditions limit the way in which embeddings can be used. For instance, when average pooling is used, it is not possible to use linear scoring functions for checking whether a given propositional formula is satisfied in the epistemic state encoded by a given vector. In general, our results provide valuable insights for the design of neural networks that are required to implement some form of systematic reasoning.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgements

This work was supported by EPSRC grants EP/V025961/1 and EP/W003309/1.

References

- [1] S. Chopra, R. Hadsell, Y. LeCun, Learning a similarity metric discriminatively, with application to face verification, in: Proceedings CVPR, 2005, pp. 539–546.
- [2] A. van den Oord, Y. Li, O. Vinyals, Representation learning with contrastive predictive coding, CoRR, arXiv:1807.03748 [abs], 2018.
- [3] J. Xu, Z. Zhang, T. Friedman, Y. Liang, G. Van den Broeck, A semantic loss function for deep learning with symbolic knowledge, in: Proceedings ICML, 2018, pp. 5498–5507.
- [4] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, L. De Raedt, DeepProbLog: neural probabilistic logic programming, in: Proceedings NeurIPS, 2018, pp. 3753–3763.
- [5] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings CVPR, 2016, pp. 770–778.
- [6] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, in: Proceedings NAACL-HLT, 2019.
- [7] D. Kiela, E. Grave, A. Joulin, T. Mikolov, Efficient large-scale multi-modal classification, in: Proceedings AAAI, 2018, pp. 5198–5204.
- [8] F. Scarselli, M. Gori, A.C. Tsoi, M. Hagenbuchner, G. Monfardini, The graph neural network model, IEEE Trans. Neural Netw. 20 (1) (2008) 61–80.
- [9] M. Schlichtkrull, T.N. Kipf, P. Bloem, R.v.d. Berg, I. Titov, M. Welling, Modeling relational data with graph convolutional networks, in: Proceedings ESWC, 2018, pp. 593–607.
- [10] T.R. Besold, A. d'Ávila Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, K. Kühnberger, L.C. Lamb, P.M.V. Lima, L. de Penning, G. Pinkas, H. Poon, G. Zaverucha, Neural-symbolic learning and reasoning: a survey and interpretation, in: P. Hitzler, M.K. Sarker (Eds.), Neuro-Symbolic Artificial Intelligence: The State of the Art, in: Frontiers in Artificial Intelligence and Applications, vol. 342, IOS Press, 2021, pp. 1–51.
- [11] P. Barceló, E.V. Kostylev, M. Monet, J. Pérez, J.L. Reutter, J.P. Silva, The logical expressiveness of graph neural networks, in: Proceedings ICLR, 2020.
- [12] P. Minervini, S. Riedel, P. Stenetorp, E. Grefenstette, T. Rocktäschel, Learning reasoning strategies in end-to-end differentiable proving, in: Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event, in: Proceedings of Machine Learning Research, PMLR, vol. 119, 2020, pp. 6938–6949, <http://proceedings.mlr.press/v119/minervini20a.html>.
- [13] L. Bergen, T.J. O'Donnell, D. Bahdanau, Systematic generalization with edge transformers, in: M. Ranzato, A. Beygelzimer, Y.N. Dauphin, P. Liang, J.W. Vaughan (Eds.), Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, virtual, NeurIPS 2021, December 6–14, 2021, pp. 1390–1402, <https://proceedings.neurips.cc/paper/2021/hash/0a4dc6dae338c9cb08947c07581f77a2-Abstract.html>.
- [14] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artif. Intell. 55 (1) (1992) 1–60, [https://doi.org/10.1016/0004-3702\(92\)90041-U](https://doi.org/10.1016/0004-3702(92)90041-U).
- [15] J. Pearl, System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: R. Parikh (Ed.), Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning About Knowledge, Pacific Grove, CA, USA, March 199, Morgan Kaufmann, 1990, pp. 121–135.
- [16] S. Benferhat, D. Dubois, H. Prade, Nonmonotonic reasoning, conditional objects and possibility theory, Artif. Intell. 92 (1–2) (1997) 259–276.
- [17] N. Friedman, J.Y. Halpern, Plausibility measures and default reasoning, J. ACM 48 (4) (2001) 648–685, <https://doi.org/10.1145/502090.502092>.
- [18] D. Dubois, H. Prade, Possibility Theory - an Approach to Computerized Processing of Uncertainty, Springer, 1988.
- [19] G.G. Towell, J.W. Shavlik, Knowledge-based artificial neural networks, Artif. Intell. 70 (1–2) (1994) 119–165.
- [20] P. Hohenecker, T. Lukasiewicz, Ontology reasoning with deep neural networks, J. Artif. Intell. Res. 68 (2020) 503–540.
- [21] T. Rocktäschel, S. Riedel, End-to-end differentiable proving, in: Proceedings NIPS, 2017, pp. 3788–3800.
- [22] F. Yang, Z. Yang, W.W. Cohen, Differentiable learning of logical rules for knowledge base reasoning, in: Proceedings NIPS, 2017, pp. 2319–2328.
- [23] A. Sadeghian, M. Armandpour, P. Ding, D.Z. Wang, DRUM: end-to-end differentiable rule mining on knowledge graphs, in: Proceedings NeurIPS, 2019, pp. 15321–15331.
- [24] R. Evans, D. Saxton, D. Amos, P. Kohli, E. Grefenstette, Can neural networks understand logical entailment?, in: Proceedings ICLR, 2018.
- [25] P. Clark, O. Tafjord, K. Richardson, Transformers as soft reasoners over language, in: Proceedings IJCAI, 2020, pp. 3882–3890.
- [26] S. Schockaert, Modelling monotonic and non-monotonic attribute dependencies with embeddings: a theoretical analysis, in: Proceedings of the 3rd Conference on Automated Knowledge Base Construction, 2021.
- [27] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, S.Y. Philip, A comprehensive survey on graph neural networks, IEEE Trans. Neural Netw. Learn. Syst. 32 (1) (2020) 4–24.
- [28] M. de Rijke, A note on graded modal logic, Stud. Log. 64 (2) (2000) 271–283.
- [29] M. Ali, M. Berrendorf, C.T. Hoyt, L. Vermue, M. Galkin, S. Sharifzadeh, A. Fischer, V. Tresp, J. Lehmann, Bringing light into the dark: a large-scale evaluation of knowledge graph embedding models under a unified framework, IEEE Trans. Pattern Anal. Mach. Intell. 44 (1) (2022) 8825–8845, <https://doi.org/10.1109/TPAMI.2021.3124805>.
- [30] K. Sinha, S. Sodhani, J. Dong, J. Pineau, W.L. Hamilton, CLUTRR: a diagnostic benchmark for inductive reasoning from text, in: Proceedings EMNLP-IJCNLP, 2019, pp. 4506–4515.
- [31] D.J.T. Cúcala, B.C. Grau, E.V. Kostylev, B. Motik, Explainable GNN-based models over knowledge graphs, in: Proceedings ICLR, 2021.
- [32] V. Gutiérrez-Basulto, S. Schockaert, From knowledge graph embedding to ontology embedding? An analysis of the compatibility between vector space representations and rules, in: Proceedings KR, 2018, pp. 379–388.
- [33] R. Abboud, İ. Ceylan, T. Lukasiewicz, T. Salvatori, Boxe: a box embedding model for knowledge base completion, in: Proceedings NeurIPS, 2020.
- [34] H. Ren, W. Hu, J. Leskovec, Query2box: reasoning over knowledge graphs in vector space using box embeddings, in: Proceedings ICLR, 2020.
- [35] Z. Zhang, J. Wang, J. Chen, S. Ji, F. Wu, ConE: cone embeddings for multi-hop reasoning over knowledge graphs, in: Proceedings NeurIPS, vol. 34, 2021, pp. 19172–19183.

- [36] A. Pavlovic, E. Sallinger, Expressive: a spatio-functional embedding for knowledge graph completion, CoRR, arXiv:2206.04192 [abs], 2022. <https://doi.org/10.48550/arXiv.2206.04192>.
- [37] M. Leemhuis, Ö.L. Özçep, D. Wolter, Learning with cone-based geometric models and orthologics, *Ann. Math. Artif. Intell.* 90 (11–12) (2022) 1159–1195, <https://doi.org/10.1007/s10472-022-09806-1>.
- [38] L. Vilnis, X. Li, S. Murty, A. McCallum, Probabilistic embedding of knowledge graphs with box lattice measures, in: *Proceedings ACL*, 2018, pp. 263–272.
- [39] X. Li, L. Vilnis, D. Zhang, M. Boratko, A. McCallum, Smoothing the geometry of probabilistic box embeddings, in: *Proceedings ICLR*, 2019.
- [40] X. Chen, M. Boratko, M. Chen, S.S. Dasgupta, X.L. Li, A. McCallum, Probabilistic box embeddings for uncertain knowledge graph reasoning, in: *Proceedings NAACL-HLT*, 2021, pp. 882–893.