

The Tensor Brain: A Unified Theory of Perception, Memory and Semantic Decoding

Volker Tresp

LMU Munich and Siemens Munich, Germany

VOLKER.TRESP@LMU.DE

Sahand Sharifzadeh⁺

LMU Munich, Germany

* SAHAND.SHARIFZADEH@GMAIL.COM

Hang Li⁺

LMU Munich and Siemens Munich, Germany

HANG.LI@CAMPUS.LMU.DE

Dario Konopatzki

LMU Munich, Germany

DK@DKONOPATZKI.DE

Yunpu Ma

LMU Munich and Siemens Munich, Germany

COGNITIVE.YUNPU@GMAIL.COM

Editor: ⁺equal contributions.

Extended Version¹

Abstract

We present a unified computational theory of an agent’s perception and memory. In our model, both perception and memory are realized by different operational modes of the oscillating interactions between a symbolic index layer and a subsymbolic representation layer. The two layers form a bilayer tensor network (BTN). The index layer encodes indices for concepts, predicates, and episodic instances. The representation layer broadcasts information and reflects the cognitive brain state; it is our model of what authors have called the “mental canvas” or the “global workspace.” As a bridge between perceptual input and the index layer, the representation layer enables the grounding of indices by their subsymbolic embeddings, which are implemented as connection weights linking both layers. The brain is a sampling engine: Only activated indices are communicated to the remaining parts of the brain. Although memory appears to be about the past, its main purpose is to support the agent in the present and the future. Recent episodic memory provides the agent with a sense of the here and now. Remote episodic memory retrieves relevant past experiences to provide information about possible future scenarios. This aids the agent in decision-making. “Future” episodic memory, based on expected future events, guides planning and action. Semantic memory retrieves specific information, which is not delivered by current perception, and defines priors for future observations. Our approach explains the great similarity between episodic and semantic memory: Semantic memory models the episodic memory of a future instance. We analyse episodic memory and semantic

1. Sahand Sharifzadeh is now with DeepMind London. Dario Konopatzki is now with ETH Zurich. Extended version of the paper that appeared in *Neural Computation*, 2023. The naming of the models has been changed. The discussion on the *post-observation model* and Dirichlet fusion in Section 3.7 has been added. Section 4.1 describes the BTN as an autoencoder. Section 10.6, which discusses forgetting and modularity, has been added.

memory in the context of meta-learning. We argue that it is important for the agent to encode individual entities, not just classes and attributes. Perception is learning: Episodic memories are constantly being formed, and we demonstrate that a form of self-supervised learning can acquire new concepts and refine existing ones. We test our model on a standard benchmark data set, which we expanded to contain richer representations for attributes, classes, and individuals. Our key hypothesis is that obtaining a better understanding of perception and memory is a crucial prerequisite to comprehending human-level intelligence.

1. Introduction

With an increase in higher animals’ abilities to move and act came a growing demand for high-performing perceptual systems, beyond simple labeling of entities with attributes and classes (Hommel et al., 2001). This might have been a driving force to develop episodic and semantic memory. Episodic memory engrams permit the recall of recent and remote memories and provide guidance for acting right. Semantic memory engrams provide concept grounding and complement perceived information with background knowledge about concepts. The agent does not just acquire new skills but gains and memorizes knowledge: It learns about and remembers things in the world. Memories support the agent in the present and the future: Without memory systems, the brain is literally memoryless.

Episodic memories retrieve previous experiences in an agent’s life. *Recent episodic memory* permits the agent to remember the immediate past since some state information cannot be directly derived from perceptual input. A recall is triggered by nearness in time and relevance. It has evolved such that the agent can remember where it has been before, why it is where it is, and what the general context is. Recent episodic memory is an episodic memory that is almost treated as a current observation. For instance, the agent needs to remember that, even though perception does not give a clue, it is still in the hideout because the bear had been chasing it and might still be lurking outside. *Remote episodic memory* can remind an agent about past situations—similar to the current—and imminent danger and favorable actions associated with that situation. It provides estimates about possible future scenarios and aids the agent in decision-making. Continuing the previous example, the agent might remember previous personal bear encounters and subsequent dangerous situations. Recall in remote episodic memory is triggered by closeness between episodic representation and scene representation. Finally, we define *future episodic memories* as events that are expected to be a memory in the future. Information on future events influences planning and action.

Semantic memory models the time-invariant statistics of statement probabilities and acts as their prior. It is a dictionary view of the agent’s life.² It aggregates entity information, which might have been acquired at different episodes, and this information becomes local in the graph formed by semantic memory. Semantic memory enables multimodal integration. For example, if Sparky is discovered in a scene, semantic memory provides background information, for example, that Sparky is a young dog and is owned by Jack; and although dogs, in general, might be aggressive, Sparky is a friendly dog. Semantic memory support can be essential for survival. An agent simply knows that bears are dangerous, even when a bear looks cozy and sleepy and even if the agent did not yet have an unpleasant encounter

2. Some authors distinguish between a formal dictionary and a grounded cognitive encyclopedia (Evans, 2012). We will not make this distinction.

with a bear. Semantic memory is the default semantic state estimator and is complemented by perception and recent episodic memories, when and where available. Our approach explains the great similarity between episodic and semantic memory: Semantic memory is the expected episodic memory of a future instance. An elementary symbolic statement says that an entity has a particular attribute, belongs to a specific class, or that an entity has a particular relationship with another entity. This motivates our assumption that basic facts are expressed as triple statements of the form (*subject, predicate, object*).³ In most languages, basic facts are expressed in this or a similar format. Thus, triple sentences are arguably of fundamental relevance for encoding and communicating perception, episodic memory, and semantic memory, all of which humans can easily describe by language. Communication is important to make the invisible visible: for example, an agent could inform peers that a bear is lurking outside the hideout, even if it cannot be seen. Triple statements can represent relationships between entities, which enables, for example, rich scene descriptions and reasoning in social and spatial networks. An agent can understand not only that a bear and a deer are in a scene but that, luckily, the bear is not chasing the agent itself but the deer. Symbolic representations and reasoning, classical System-2 properties, depend on the representations of entity-to-entity relationships (Halford et al., 2014).

The starting point of our work is a mathematical model, that is, the bilayer tensor network (BTN). It implements mathematical models for perception and memory that obey some constraints imposed by biology. We then discuss how particular features of our model relate to current discussions in cognition and neuroscience and make predictions about implications of our model to both. The BTN implies a very simple architecture that contains two basic layers: the *index layer* and the *representation layer*. Context is provided by a third layer, the *dynamic context layer*, which interacts with the representation layer, provides context, and stores state information when the brain’s attention moves from one entity to another.

The symbolic index layer contains indices for concepts, predicates, and episodic instances known to the agent. The index layer labels the activation pattern in the representation layer and then feeds back the embedding of that label to the representation layer. The embedding vectors are implemented as connection weights linking both layers. An index is a focal point of activity and competes with other indices, but, since it constantly interacts with the representation layer, it is never active in isolation. Embeddings have an integrative character: the embedding vector for a concept index integrates all that is known about that concept, and the embedding vector for an episodic index represents the world state at that instance. The subsymbolic representation layer is the main communication platform. In cognitive neuroscience, it would correspond to, what authors call, the “mental canvas” or the “global workspace” and reflects the cognitive brain state. In bottom-up mode, scene inputs activate the representation layer, which then activates the index layer. In top-down mode, an index activates the representation layer, which might subsequently activate even earlier processing layers. This last process is called the embodiment of a concept.

Perception and memories first produce subsymbolic representations. They represent the agent’s current and past cognitive state and are subsequently decoded semantically to produce sequences of activated indices that form symbolic triple statements. These

3. In a relational Bayesian network or a Markov logic network, a triple statement is represented as a node; in a graph neural network, it would be a typical output node.

sequences then give feedback to the representation layer and earlier processing layers and thus inform the brain as a whole about what has been decoded. The brain is a sampling engine: Only activated indices are communicated to the remaining parts of the brain.

In our approach, perception and memory produce triple statements by stochastic sampling, which is a central bottleneck, as also discussed by Dehaene (2014). We introduce an attention approximation, which avoids intermediate sampling decisions on visual entities.

One can consider a graph where concepts are represented as nodes, and triples become labeled directed links pointing from subject node to object node. In computer science, this would be called a knowledge graph (KG). Whereas symbolic reasoning would purely act on the graph, embedded reasoning, as realized by the BTN, involves concept embeddings and performs information propagation via the representation layer.

This article is organized as follows. In the next section, we cover related work. Section 3 introduces concepts, triple statements, and probabilistic memory models. In Section 4, we present our proposed bilayer tensor network (BTN) and demonstrate how it realizes perception and the different memory functions. We relate it to hierarchical Bayesian modelling and meta-learning. We show how a posterior model for an instance can be defined and how this posterior model, but also episodic memory and semantic memory, can be viewed as different variants of Dirichlet updates. Section 5 presents and discusses the BTN’s algorithmic implementation, sampling, and attention approximations. The following sections present experimental results and discuss potential relationships to cognition and neuroscience. Section 6 introduces the data set and discusses perception, the representation layer, and the dynamic context layer. Section 7 discusses semantic memory engrams and their semantic decoding. Using social network data, we demonstrate multimodal integration in semantic memory. Section 8 compares embedded reasoning with symbolic reasoning and proposes that *embedded symbolic reasoning* performs symbolic reasoning using embedding vectors. We introduce the one-brain hypothesis, which emphasizes, first, that the brain uses only one representation layer and, second, that perception, episodic memory, semantic memory, and embedded reasoning all rely on the same BTN architecture. We discuss relationships between our approach to cognitive linguistics and consciousness research. We argue that the internal triple-oriented fast speech is transformed into an external, sophisticated slow speech. Section 9 covers episodic memory and the way the agent estimates current and future world states. In Section 10, we discuss that perception involves the storing of new episodic memories and show how a form of self-supervised learning can learn new concepts and refine existing ones. We consider that new perceptual experiences might require representations for new entities; that is, we do not assume domain closure. Section 11 contains our conclusions.

2. Related Work

2.1 Tensor Networks for Modeling Knowledge Graphs

The bilayer tensor network (BTN) is an example of a tensor network. RESCAL was the first tensor-based embedding model for triple prediction in relational data sets and knowledge graphs (Nickel et al., 2011, 2012). Embedding learning for knowledge graphs evolved into a sprawling research area (Bordes et al., 2013; Socher et al., 2013; Yang et al., 2014; Nickel et al., 2015b; Trouillon et al., 2016; Dettmers et al., 2018). (Nickel et al., 2015a) provides

an overview and PyKEEN (Ali et al., 2021) a comprehensive software library. In contrast to previous approaches, the BTN can be implemented as an interaction between an index layer and a representation layer, and thus it is more suitable for brainware implementations.

2.2 Cognitive Tensor Networks and Related Models

Tensors have been used previously as memory models where the main focus was on simple associations (Hintzman, 1984; Kanerva, 1988; Humphreys et al., 1989; Osth and Dennis, 2015) and compositional structures (Smolensky, 1990; Pollack, 1990; Plate, 1997; Halford et al., 1998; Ma et al., 2018a). In the tensor product approach (Smolensky, 1990), encoding or binding is realized by a tensor product (generalized outer product) and compositionality by tensor addition. In the structured tensor analogical reasoning (STAR) model (Halford et al., 1998), a working memory is constructed by a superposition of tensor products. Either standard basis (one-hot) or random vectors are used as representations for concepts and predicates. Some basic symbolic reasoning operations are implemented, such as analogical and transitive inference. The main difference to our work is that we approximate the data tensor by a tensor network, that is, the BTN. This permits generalization and realizes embedded reasoning, as well as the seamless integration of sensory input.

The application of embedding-based tensor models to neurocognitive models started with (Tresp et al., 2015), which introduced tensor networks with index embeddings for perception, as well as semantic and episodic memory. It did not explicitly consider scene bounding boxes and did not contain experimental results. In further work, the connection between temporal and semantic tensor networks was analysed (Tresp and Ma, 2016; Tresp et al., 2017a,b; Ma et al., 2018b). Semantic memory was derived from episodic memory by an integration step performed in latent space. In this article, we avoid explicit integration and extend the approach to include perception.

The neurocognitive tensor network theory evolved in the 1980s and followed the idea of geometrization of biology. It is a theory of brain function, particularly that of the cerebellum. Metric tensors transform sensory space-time coordinates into motor coordinates (Pellionisz and Llinás, 1980). There, tensor networks are used very differently from the work presented here.

2.3 Visual Relationship Detection and Scene Graphs

In 2016, the Stanford Visual Relationship data set was published, which contained images annotated with triple sentences (Lu et al., 2016) and (Krishna et al., 2017). The two articles made their annotated data available, which spawned an explosion of research activity in visual relationship detection (VRD). The background information in (Lu et al., 2016) was extracted from a text corpus. Recent work in this direction is (Luo et al., 2019).

VRD models for knowledge graphs were proposed by (Baier et al., 2017; Zhang et al., 2017; Baier et al., 2018). (Baier et al., 2017) showed how a prior distribution derived from triple occurrences could significantly improve on pure vision-based approaches and on approaches that used prior distributions derived from language models. (Sharifzadeh et al., 2019) showed further improvements by including 3D image information. (Tresp et al., 2019, 2020) describe more recent publications in this tradition. The presented work introduces

more clearly the different operational modes and provides more extensive experimental results.

Triple statements generated from an image form a scene graph (Johnson et al., 2015). Work on scene graphs attempts to find a unique, globally optimal interpretation of an image. State-of-the-art scene graph models are described in (Yang et al., 2018; Zellers et al., 2018; Hudson and Manning, 2019). (Sharifzadeh et al., 2020) captures the interplay between perception and semantic knowledge by introducing schema representations and implementing the classification as an attention layer between image-based representations and the schema.

2.4 Related Modern Technical Models for Memory

(Hochreiter and Schmidhuber, 1997) convincingly demonstrated the importance of memory systems in recurrent neural networks. Important later extensions are the neural Turing machine (NTMs) (Graves et al., 2014) and memory networks (Weston et al., 2014; Sukhbaatar et al., 2015). In those articles, episodic memory acts as an instance buffer. Both use a recurrent neural network in combination with attention mechanisms. Large language models, like OpenAI’s GPT-3, are more recent developments (Brown et al., 2020).

In our BTN, a dynamic context layer is part of a (nonstandard) recurrent neural network. The attention mechanisms in our model, episodic attention and semantic attention, are quite different from the attention mechanisms in those articles. Also, our goal is to derive triple statements, whereas, in those models, the tasks are query answering and improved token embeddings.

2.5 Dual Process Theory and Complementary Learning Systems (CLS)

In psychology, dual process theory concerns the interplay in the mental processing of an *implicit*, automatic, unconscious process (shared with animals) and an *explicit*, controlled, conscious process (uniquely human). See (Evans, 2003) for a review. In our model, the implicit side would be on the level of embeddings and representations, whereas the explicit side is on the level of the concept indices and the extracted triple sentences.

One instance of a dual process theory is Kahneman’s System-1 / System-2 dichotomy (Kahneman, 2011). System-1 is fast, instinctive, and emotional and does not require mental effort. System-2 is slower, more deliberate, and more logical and requires mental effort.

CLARION is a dual-process model of implicit and explicit learning (Sun and Peterson, 1996). It is based on one-shot explicit rule learning (i.e., explicit learning) and gradual implicit tuning (i.e., implicit learning).

A different but related dichotomy can be found in the complementary learning systems (CLS) theory (McClelland et al., 1995; Kumaran et al., 2016), where the formation of time indices and their embeddings would be part of a nonparametric learning system centered on the hippocampus, which allows rapid learning of the specifics of individual items and experiences (Kumaran et al., 2016). Slow training would be part of a parametric learning system, which serves as the basis for the gradual acquisition of structured knowledge about the environment to neocortex (Kumaran et al., 2016). In this article, we introduce self-supervised learning for rapid learning and discuss a consolidation process of learned

knowledge. McClelland et al. (2020) analysed the connection between memory, perception, and language, which is also a focus of this article.

2.6 The Bayesian Brain

Our approach can be related to the tradition of Bayesian approaches to brain modeling (Dayan et al., 1995; Rao and Ballard, 1999; Knill and Pouget, 2004; Körding et al., 2004; Tenenbaum et al., 2006; Griffiths et al., 2008; Friston, 2010). In our approach, all past experiences train the perceptual system and also contribute to perception via memory. In (Baier et al., 2017) an explicit semantic prior distribution was used, describing *a priori* probabilities for triple sentences. For inference, Bayes’ formula is used. The remarkable improvement in performance after integrating the prior information indicates that triple representations might be a powerful abstraction level for formulating prior knowledge, in general. (Sharifzadeh et al., 2020) showed that the probabilistic knowledge graph acts as an inductive bias in perception. It discusses the role of a prior as an integrator of multimodal information and its role in filling in nonperceptual background information.

In the work presented here, we separately model and perform completion on the observed statements, statement priors, and statement posteriors. This is in the context of Bayesian probabilistic inference but not in the sense of Bayesian statistics.

In a Bayesian brain approach, top-down connections are often associated with predictive models; they provide predictions about future state probabilities. Here, we emphasize that, in addition, top-down connections inform the representation layer and earlier processing layers— what the brain has detected.

3. Concepts, Triple Statements and Probabilistic Memory Models

You only see what you know. —Johann Wolfgang von Goethe, an Friedrich von Müller, April 24, 1819

3.1 Triple Statements

To understand what is perceived, an agent needs to have an understanding of the things in the world and their relationships. We assume that the agent’s mind is aware of N_C concepts $\mathcal{C} = \{c_1, \dots, c_{N_C}\}$. A concept can, for example, represent an entity $e \in \mathcal{E} \subset \mathcal{C}$, or a class $k \in \mathcal{K} \subset \mathcal{C}$, which stands for a collection of entities, or an attribute $b \in \mathcal{B} \subset \mathcal{C}$.

In addition to concepts, we also consider a set of predicates $p \in \mathcal{P} = \{p_1, \dots, p_{N_P}\}$, where N_P is the number of predicates the agent is aware of. A triple statement has the form (s, p, o) , where $s \in \mathcal{C}$ assumes the role of the subject, $o \in \mathcal{C}$ assumes the role of the object, and predicate $p \in \mathcal{P}$. Examples of triple statements are: *(Munich, partOf, Bavaria)*, *(Sparky, looksAt, Jack)*, *(AkiraKurosawa, directorOf, SevenSamurai)*, and *(Jack, knows, Mary)*.

Following the notion of the semiotic triangle of (Ogden and Richards, 1923), we can look at triples and their semantics from three different perspectives:

- The agent-independent objective world (world semantics). Some triples have an interpretation in the real world, and they stand for propositions, like *(Munich, partOf,*

Bavaria) and (*Sparky, looksAt, Jack*). This view is mostly taken in a formal analysis of cognition and linguistics (Montague, 1970; Fodor, 1975) and is sometimes referred to as truth-conditional semantics. An agent is only aware of some of the concepts in the world, sometimes called the “projected world” or the “construal world.”

- An agent’s mind and brain (an agent’s neurocognitive semantics): A triple is a statement based on concepts in an agent’s mind. These might have an immediate meaning in the objective world but could also be untestable statements with unclear world semantics such as (*Love, strongerThan, Hate*). A hypothesis of this article is that triple statements correspond to processes in the brain, that is, a sequential firing of index representations. They form the inner fast speech.
- An agent’s utterance (an agent’s linguistic semantics). A triple is a simple language clause involving symbols. Within the approach in this article, we assume that any triple statement in the mind can also be expressed linguistically. But what is spoken, in general, represents external slow speech and is obviously more complex, nuanced, and sophisticated than the inner fast speech, and is modulated, for example, by intent, social context, and cultural background. In general, the relationship between statements in the mind (with concepts) and in language (with symbols) is a matter of an open debate (Evans, 2012).

We are aware that the relationship of the three perspectives touches on fundamental issues in cognition, linguistics, neuroscience, philosophy, and many other scientific fields and academic disciplines. A more detailed discussion would be beyond the scope of this article. This article assumes the position of the agent. We are concerned only about how the agent views the world. We focus on personal Bayesian probabilities, modeling the agent’s neurocognitive semantics, where we restrict ourselves to statements that, in the mind of the agent, can be true or false, for example, (*Sparky, ownedBy, Jack*), and facts that can be related to observations, for example, (*Sparky, hasColor, Black*). This is more of a practical than a principled constraint. On some dimensions thought and reality agree simply to act right and to guarantee survival; on others, there is likely no agreement between agents, no clear relationship to an objective reality, and in determining the truth values on statements, personal emotions and judgments might play a significant role.

3.2 Knowledge Graphs

In a knowledge graph (KG), a triple is represented as a directed labeled link from concept $s \in \mathcal{C}$ to concept $o \in \mathcal{C}$, where the link is labeled by p . Thus, a labeled link represents a triple statement of the form (s, p, o) where s is called *head* or *source node* and o is called *tail* or *target node*. Knowledge graphs currently have a great impact in applications and, as our approach, are entity oriented.

3.3 Unary and Binary Statements

When s and o are entities, (s, p, o) stands for the ground atom $p(s, o)$. (*Munich, partOf, Bavaria*) would stand for $partOf(Munich, Bavaria)$.

We assume a strong default predicate *hasAttribute* (abbreviated as *hA*), which, depending on the subject type and the object type, can stand for certain other predicates from the set $\mathcal{P}^B \subseteq \mathcal{P}$. We write (s, hA, c) , $c \in \mathcal{C}$.

- If both subject and object are entities, the implicit predicate is *sameAs*; thus, $(Jack, hA, John)$ stands for $(Jack, sameAs, John)$.
- If the subject is an entity and the object is an attribute, the implicit predicate is attribute specific but should be obvious; thus $(Jack, hA, Tall)$ stands for $(Jack, height, Tall)$.
- If the subject is an entity and the object is a class, the implicit predicate is *type*, as in $(Sparky, type, Dog)$.
- If both subject and object are classes, the implicit predicate is *subClass*, as in $(Dog, subClass, Mammal)$; note that the transitive *subClass* predicate permits the modeling of deep ontologies.

Triple sentences involving the *hasAttribute* predicate (and its substitutes) we call *unary statements*. The remaining N_B predicates from $\mathcal{P}_B \subseteq \mathcal{P}$ form *binary statements*. We will refer to the object in a unary statement also simply as the unary label of the subject, and to the predicate in a binary statement as the binary label of a concept pair. Binary statements are required when the default unary interpretation is not applicable, as in $(Jane, motherOf, Jack)$ (the default would be $(Jane, sameAs, Jack)$).

Higher-order relations can be reduced to a set of binary statements, for example, by using additional concepts (Noy et al., 2006). For example, $match(Player1, Player2, Location)$ becomes $(matchID, hasPlayer, Player1)$, $(matchID, hasPlayer, Player2)$, and $(matchID, hasLocation, Location)$, where $(matchID)$ is the additional concept.

3.4 Probabilistic Models with Boolean Variables

In the assumed lifetime and the mind of an agent, some statements are always true, as $(Munich, partOf, Bavaria)$, but the truth values of other statements can change in time, as $(Munich, weather, Sunny)$. Thus, with each triple sentence, at each episodic instance t , the agent associates a *Boolean semantic state variable* $Y_{s,p,o,t}$. If the agent is certain that (s, p, o) is true at episodic instance t , then $Y_{s,p,o,t} = 1$ and if (s, p, o) is observed to be false at episodic instance t , then $Y_{s,p,o,t} = 0$. The agent’s *semantic world state* at episodic instance t is defined as the states of all semantic state variables. Here, we assume that the agent is concerned with N_T past episodic instance $\mathcal{T} = \{t_1, \dots, t_{N_T}\}$.

We consider that data, that is, information on the states on triple statements, arrives in chunks, which we call episodes. McClelland et al. (2020) call them situations and in cognition they are called event frames. In knowledge graphs, they could form a namespace. We use the terms events and episodes almost interchangeably, although, in a narrower sense, we reserve the term episode for a sequence of events. (See also the discussion in Section 9.) Each episode provides information about the truth values of a subset of all statements. In this article, we assume that episodic data is either provided by a human annotator (supervised learning) or generated in self-supervised learning (see Section 10).

We can form a data tensor of observed triples as a four-mode triple-observation tensor of the form

$$\underline{\underline{T}} = \sum_{s,p,o,t} y_{s,p,o,t} \mathbf{e}^s \otimes \mathbf{e}^p \otimes \mathbf{e}^o \otimes \mathbf{e}^t.$$

Here, \mathbf{e}^i is our notation for a standard basis vector, that is, a one-hot vector, with the one at position i . Such data tensors have also been the starting point for the tensor memories of Halford et al. (1998) and (Ma et al., 2018a). A tensor entry $y_{s,p,o,t}$ is simply equal to one if the statement (s, p, o) was observed by the agent to be true at t and equal to zero if the corresponding statement was observed or concluded by the agent to be false at t ; otherwise, it is unknown. The temporal KG (tKG) is the graph representation of $\underline{\underline{T}}$, where the labeled links between concepts are time-dependent. We assume that data is acquired from different modalities, like vision or language. We invoke a local closed-world assumption (LCWA): We assume that for entities for which data is acquired in a modality, the truth values of all modality-specific unary statements are available. For entity pairs in a modality, we assume that the truth values of all modality-specific binary statements are available. The issue of which entries in $\underline{\underline{T}}$ are observed might involve some selection or attention process. For example, state changes or singular events might be selected with some priority. We discuss this issue in Section 9.9.

In the next section, we introduce the bilayer tensor network (BTN), which performs data completion. For all $y_{s,p,o,t}$ that are either observed to be true or concluded by the LCWA to be false, the BTN Boolean observation model is

$$\mathbb{E}(Y_{s,p,o,t} | \underline{\underline{T}}) \equiv \mathbb{E}(Y_{s,p,o,t}) \approx y_{s,p,o,t}. \quad (1)$$

The observation model has the characteristics of an autoencoder, where $\underline{\underline{T}}$ represents noisy observations. This relationship is indicated by the approximation symbol (\approx). As indicated, we simplify notations here and in the following by not explicitly conditioning on $\underline{\underline{T}}$.

We also define the quantity $i_{s,p,o,t}$, where $i_{s,p,o,t} = 1$, if $y_{s,p,o,t} = 1$, and $i_{s,p,o,t} = 0$, otherwise. Thus, unknowns in the triple-observation tensor are assigned the value 0. Consider a future instance t' for which no data is yet available. We define the BTN's expected-state model as

$$\mathbb{E}(Y_{s,p,o,\bar{t}} | \underline{\underline{T}}) \equiv \mathbb{E}(Y_{s,p,o,\bar{t}}) \approx \frac{1}{N_{s,p,o,\text{known}}} \sum_t i_{s,p,o,t}. \quad (2)$$

Here, \bar{t} is a new instance not represented in $\underline{\underline{T}}$. $N_{s,p,o,\text{known}} > 0$ counts the times that the truth values of (s, p, o) are known, either because they were observed to be true or as concluded by the LCWA to be false. The expected-state model simply predicts the mean truth value of a future triple statement. The average in the last equation might only consider a limited time horizon to be able to account for state changes (see also Section 10).

If we consider that the observations at an instance describe a knowledge graph, then the expected-state model performs integration by modelling the average truth values of known triples' truth values.

3.5 Observation Model (Episodic Memory)

In this context s, p, o, t are the states of categorical variables. For subject and object, they have N_C states, for the predicate N_P states, and for instances N_T states. The BTN's

observation model is

$$\mathbb{P}(s, p, o|t, \underline{T}) \equiv \mathbb{P}(s, p, o|t) \approx \frac{1}{N_t} i_{s,p,o,t}. \quad (3)$$

It models the distribution of the indices. $N_t = \sum_{s,p,o} i_{s,p,o,t}$ is the total number of triples observed to be true at t .

As we will discuss further down, the observation model mathematically describes the symbolic aspect of episodic memory. Episodic memory is about what was actually observed. It implicitly assumes a closed-world assumption (CWA), in the sense that what was not observed should not be sampled.

From the definitions, it follows that for observed triples and where the LCWA applies

$$\mathbb{P}(s, p, o|t) \approx \mathbb{E}(Y_{s,p,o,t})/N_t.$$

The agent can thus assume that for a given t , a sample generated from the observation model is true. If \underline{T} is a database, $\mathbb{E}(Y_{s,p,o,t})$ describes a probabilistic model of that database, and $\mathbb{P}(s, p, o|t)$ is a database retrieval model.

3.6 Pre-observation model (Semantic Memory)

The BTN’s pre-observation model is

$$\mathbb{P}(s, p, o|\bar{t}, \underline{T}) \equiv \mathbb{P}(s, p, o|\bar{t}) \approx \frac{1}{N_{\text{total}}} \sum_t i_{s,p,o,t}. \quad (4)$$

$N_{\text{total}} = \sum_t N_t$ is the total number of observed true triples. The pre-observation model provides background information on the subject. It represents the predictions of a model with an informative learned prior for some future instance \bar{t} . As we will discuss in the course of the paper, the pre-observation model mathematically describes the symbolic aspect of semantic memory. Samples generated by the pre-observation model correspond to triples often observed to be true. See a more detailed discussion in Appendix B.

If we consider again that the observations at an instance describe a knowledge graph, then the pre-observation model performs integration by summing the truth values of a triple statement at different instances, followed by normalization. Whereas the expected-state model predicts what is true or false, the pre-observation model predicts what should be observed next at a new instance.

3.7 Post-observation Model and Meta-learning

The pre-observation model estimates the semantic world state prior to observations for a new instance t' . It predicts what should be observed next. The post-observation model combines episodic and semantic memory to estimate the state for t' after taking into account observations at t' . Motivated by a hierarchical Bayesian analysis, we consider the pre-observation model to be the base distribution of a Dirichlet distribution with concentration parameter $\gamma > 0$. We obtain the post-observation model by a Dirichlet update of the form

$$\frac{1}{\gamma + \tilde{N}_{t'}} \left(\gamma \mathbb{P}(s, p, o|\bar{t}) + \tilde{N}_{t'} \mathbb{P}(s, p, o|t') \right). \quad (5)$$

The value of γ indicates the trust in the pre-observation model and $\tilde{N}_t = N_t$. We can sample from this distribution by sampling from the pre-observation model (semantic memory) with a frequency proportional to γ and from the observation model (episodic memory) with a frequency proportional to N_t . The latter recalls actual observations. The former adds triples describing the state, prior to the observation. As discussed in Section 9.9, the observations might include samples from recent episodic memories, which are treated as current observations. The post-observation model should not be confused with the updates of the pre-observation model. For the update of the pre-observation model, we would have $\gamma = N_{total}$. In essence, in the post-observation model, the current observations obtain a larger weight. In terms of meta-learning, an instance is a domain, and an instance-specific episodic memory a domain-specific model, but with shared embedding vectors. Semantic memory would be an instance-independent model trained by pooling all data. The post-observation model is the target-domain prediction. Table 1 provides a summary of the different models. Appendix C provides a visualization.

Model	Function	γ	\tilde{N}_t	Embedding
Pre-observation model ($< t'$)	Semantic memory	N_{total}	0	$\bar{\mathbf{a}}$
Observation model (t')	Episodic memory	0	$N_{t'}$	\mathbf{a}_t
Post-observation model (t')	Posterior state at t'	hyperp.	$N_{t'}$	$\approx \mathbf{a}_t + (\gamma/N_{t'})\bar{\mathbf{a}}$
Pre-observation model ($> t'$)	Semantic memory	N_{total}	$N_{t'}$	$\bar{\mathbf{a}}$ (updated)

Table 1: Special cases of the Dirichlet update in Equation 5. The pre-observation model can generate samples from the state prior to observations at t' ($< t'$). The observation model can generate samples from the observations at t' . The post-observation can generate samples from the state posterior to observations at t' , for the state at t' . The pre-observation model ($> t'$) integrates observations at t' to predict future states. Embedding vectors will be introduced in Section 4.

3.8 Modelling Dependencies: Generalized Statements

Can classes or attributes be the subject in a triple statement? Whereas the semantics of (*Sparky*, *hasColor*, *Black*) is clear, the semantics of (*Dog*, *hasColor*, *Black*) is less well defined: For example, it could mean that there is a least one dog that is black, or that some or most dogs are black, or that all dogs are black. In our work, we consider frequencies.

We define a triple (c_1, hA, c_2) with $c_1, c_2 \in \mathcal{C}$ where

$$\begin{aligned} \mathbb{E}(Y_{c_1, hA, c_2, \bar{t}}) &\equiv \mathbb{E}(Y_{\bar{s}, hA, c_2, \bar{t}} | Y_{\bar{s}, hA, c_1, \bar{t}}) \\ &\approx \frac{1}{\sum_s N_{s, hA, c_1, \text{known}}} \sum_{t, s} i_{s, hA, c_1, t} i_{s, hA, c_2, t}. \end{aligned} \tag{6}$$

This is the probability that a new entity \bar{s} at a new time instance \bar{t} that has unary label c_1 will also have unary label c_2 . Generalized statements are useful for predictions at a new instance \bar{t} where only partial measurements are available. With a predicted unary label *Dog*

and using the generalized statement $(Dog, hA, Mammal)$, the unary label $Mammal$ can be predicted.

A generalized statement corresponds to a probabilistic version of the rule $\forall s : (s, hA, Mammal) \leftarrow (s, hA, Dog)$ where s is a variable that stands for an entity. Generalized statements are the basis for embedded symbolic reasoning described in Section 8.3.

Distinguishing between entities, classes, and attributes is important for some of the discussions. But generalized statements permit us to treat all concepts almost identically in the algorithmic implementation in Section 5.

4. A Bilayer Tensor Network (BTN)

4.1 The BTN Observation Model

In this section, we introduce the embedding-based bilayer tensor network (BTN). It assigns an r -dimensional latent embedding vector to each concept c , predicate p , and instance t . The embedding vectors $\{\mathbf{a}_c\}_{c \in \mathcal{C}}$, $\{\mathbf{a}_p\}_{p \in \mathcal{P}_B}$, and $\{\mathbf{a}_t\}_{t \in \mathcal{T}}$ form the columns of the embedding matrix A .

In the BTN, all observed triples are independent, given the embedding matrix A . With some of the children observed, represented as the triple-observation tensor \underline{T} , we get for a triple (s, p, o) at t ,

$$\mathbb{P}(s, p, o | t, \underline{T}) = \int \mathbb{P}(s, p, o | t, \mathbf{a}_s, \mathbf{a}_p, \mathbf{a}_o, \mathbf{a}_t) \mathbb{P}(A | \underline{T}) dA. \quad (7)$$

The last equation has the characteristics of an autoencoder, where \underline{T} represents noisy observations. We use a point estimate approximation as

$$\mathbb{P}(A | \underline{T}) \approx \delta(A - \hat{A}) \quad (8)$$

where $\delta(\cdot)$ is the Dirac delta distribution. The observation model then becomes

$$\mathbb{P}(s, p, o | t, \underline{T}) \approx \mathbb{P}(s, p, o | \hat{\mathbf{a}}_s, \hat{\mathbf{a}}_p, \hat{\mathbf{a}}_o, \hat{\mathbf{a}}_t). \quad (9)$$

We ignore in this notation that due to the softmax transfer function, the conditional probability implicitly depends on the complete embedding matrix.

The point estimates $\hat{\mathbf{a}}_s, \hat{\mathbf{a}}_p, \hat{\mathbf{a}}_o, \hat{\mathbf{a}}_t$ can be sensitive to changes in \underline{T} , implying that there are global dependencies between all triple statements. For example, if a new triple becomes known for the current instance t , this information will become part of \underline{T} and potentially results in major changes to $\hat{\mathbf{a}}_t$, affecting all triple statements at t .

In the following notation we do not use the hat notation; for example, we write \mathbf{a}_t instead of $\hat{\mathbf{a}}_t$.

In the sampling mode, we decompose into conditional probabilities and generate samples for binary statements from

$$\mathbb{P}(s, p, o | t) = \mathbb{P}(s | t) \mathbb{P}(o | s, t) \mathbb{P}(p | s, o, t) \quad (10)$$

where $p \in \mathcal{P}_B$ is a binary label. For unary statements where we use $p = hA$ (*hasAttribute*) and we generate samples from $\mathbb{P}(s | t) \mathbb{P}(c | s, t)$, where c is the unary label.

4.2 The BTN Pre-observation and Post-observation Model

We obtain the pre-observation model (semantic memory) by using the approximation in Equation 9 but by replacing $\hat{\mathbf{a}}_t \leftarrow \bar{\mathbf{a}}$. Technically, $\bar{\mathbf{a}}$ is a hyperparameter vector that was fit on all past data (data pooling).

The post-observation model is achieved simply by applying the Dirichlet update Equation 5 to the learned BTN models. Actually, we would propose that the brain generates samples alternately from the observation model and semantic memory, enabling the mind to keep track of what was observed and what was added by semantic memory, representing background knowledge.

4.3 BTN Architecture

The BTN is implemented by an interaction of two layers, which, as we argue later, might be related to functional brain operations. One layer is the representation layer q with pre-activation vector $\mathbf{q} \in \mathbb{R}^r$; here, $r \in \mathbb{N}$ is the embedding dimension, that is, the rank of the BTN approximation. As we propose later, \mathbf{q} reflects the cognitive state of the brain (see Figure 1). The other one is the index layer n with pre-activation vector $\mathbf{n} \in (0, 1)^d$, with $d = N_C + N_P + N_T$. The index layer contains one dimension or unit for each concept, predicate, and episodic instance. The connection matrix A links both layers, and its columns are the embeddings of the indices. We also introduce a dynamic context layer h with pre-activation vector \mathbf{h} . In the following, we describe the implementations of layers and operations using the unfolded view in Figure 2.

4.4 BTN Algorithm for Memory Recall

In this section, the embedding vectors contain free parameters optimized for good performance. In the following sections, we will ground them in perception, in the spirit of an embodied approach.

We start with an instance t to activate episodic memory. Thus the activation of the index layer is $\mathbf{n}_T = \mathbf{e}^t$. This activates $\mathbf{q}_T = \mathbf{a}_t$. $\mathbb{P}(c_{\text{instance}}|t) = \text{softmax}_c^\beta(\mathbf{n}_C)$ is the probability for a unary label of the instance, for example, a particular vacation day, or a particular location, for example, Paris. Here, $\mathbf{n}_C = \mathbf{a}_c^\top \text{sig}(\mathbf{q}_T)$. In our experiments, we do not use instance labels. We get for the subject

$$\mathbb{P}(s|t) = \text{softmax}_s^\beta(\mathbf{n}_S) \quad (11)$$

with pre-activations $\mathbf{n}_S = \mathbf{a}_s^\top \text{sig}(\mathbf{g}(\mathbf{q}_T))$.

For unary entity labels, we obtain

$$\mathbb{P}(c|s, t) = \text{softmax}_c^\beta(\mathbf{n}_C) \quad (12)$$

with $\mathbf{n}_C = \mathbf{a}_c^\top \text{sig}(\mathbf{q}_S)$. For the object, we obtain

$$\mathbb{P}(o|s, t) = \text{softmax}_o^\beta(\mathbf{n}_O) \quad (13)$$

where $\mathbf{n}_O = \mathbf{a}_o^\top \text{sig}(\mathbf{g}(\mathbf{q}_S))$ and $\mathbf{q}_S = \mathbf{a}_s + \mathbf{g}(\mathbf{q}_T)$. For binary labels, we get

$$\mathbb{P}(p|s, o, t) = \text{softmax}_p^\beta(\mathbf{n}_P) \quad (14)$$




Figure 1: Perception and memory: Our model architecture consists of two main layers, the *representation layer* q and the *index layer* n . \mathbf{q} and \mathbf{n} are the vectors of pre-activations. In perception, the representation layer obtains inputs from the scene. Whereas the computational operations in these layers are practically instantaneous, the *dynamic context layer* h represents internal brain activities and stores information when attention moves from one entity to another. The dotted line represents embodiment, the activation of earlier processing layers via the representation layer.

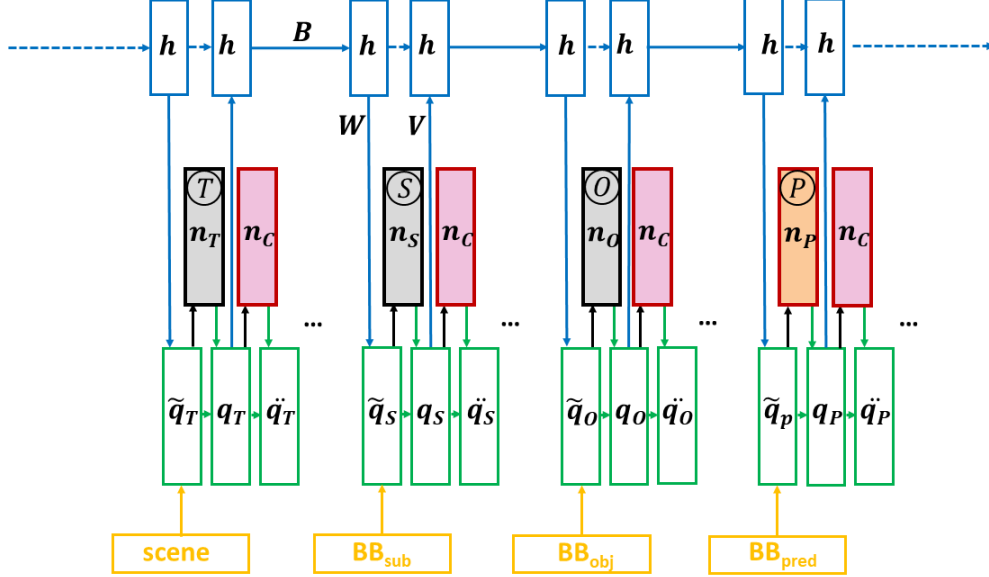


Figure 2: Unfolded representation of episodic and semantic memory. In perception, the structures in yellow color at the bottom are added; they stand for inputs from the visual scene.

with $\mathbf{n}_P = \mathbf{a}_p^\top \text{sig}(\mathbf{g}(\mathbf{q}_O))$ and $\mathbf{q}_O = \mathbf{a}_o + \mathbf{g}(\mathbf{q}_S)$.

For Boolean variables and unary statement, we obtain (cf. Equation 1)

$$\mathbb{E}(Y_{s,h,A,c,t}) = \text{sig}(\mathbf{n}_C) \quad (15)$$

and for binary statements,

$$\mathbb{E}(Y_{s,p,o,t}) = \text{sig}(\mathbf{n}_P). \quad (16)$$

Essentially, we replace the softmax activation function with the sig activation function. This renormalization is discussed in Appendix B.

For the semantic memory model (i.e., the pre-observation model), we use $\mathbf{a}_t \leftarrow \bar{\mathbf{a}}$. Also, we input s .

Here, $\mathbf{g}(\cdot)$ is a nonlinear function. In our approach, we use

$$\mathbf{g}(\mathbf{q}) = W \text{sig}(B \text{sig}(V \text{sig}(\mathbf{q})))$$

which performs computations in the hidden layer h of the recurrent network. W, B, V are learned matrices. Also, $\text{sig}(x) = 1/(1 + \exp(-x))$ is the logistic function and

$$\text{softmax}_i^\beta(\mathbf{x}) = \frac{\exp \beta x_i}{\sum_{i'} \exp \beta x_{i'}}.$$

Then, the post-activation vector $\text{softmax}^\beta(\mathbf{x})$ (without the subscript) is the column vector of all $\{\text{softmax}_i^\beta(\mathbf{x})\}_i$. Here, $\beta \geq 0$ is an inverse-temperature parameter and can be used for making the response more or less focused. \mathbf{x} is the vector of pre-activations.

4.5 Discussion on the BTN

An important property of the BTN is that in the algorithmic implementation and at any iteration step, the representation vectors have the dimension of the embedding vector and, for example, cannot represent the concatenation of two embeddings. We can relate this to brain function. Since the representation layer might occupy a significant portion of the brain, leaving no space for a second concurrent representation, we call this property the “one-brain hypothesis” (discussed in more detail in Section 8.6). Also, due to the nonlinearity of the function $\mathbf{g}(\cdot)$, one breaks the symmetry between subject and object embedding. It is clear which entity is the subject and which one is the object, and the agent can distinguish between *(Sparky, ownedBy, Jack)* and *(Jack, ownedBy, Sparky)*. Thus, proper role labeling is performed. Halford et al. (2014) call this “structural alignment.” We did not select more standard models for several reasons. First, most tensor factorizations (Hackbusch, 2012), such as CPD, the Tucker decomposition, the tensor train, and RESCAL (Nickel et al., 2011), require an excessive multiplication of factors; multiplication is an operation that is not easily implemented in biological hardware. Second, standard tensor networks are functions of several embedding vectors that would need to be presented concurrently; this would violate our one-brain hypothesis. By representing functions of functions, our model is compositional (in the sense of Poggio et al. (2020)), which is a property that is used to explain the superior performance of deep architectures. Other forms of compositionality are discussed in Section 8.

4.6 BTN Algorithm for Perception

We consider the following setting: At a new episodic instance t' , the agent encounters a new *scene $_{t'}$* . Then a bounding box BB_{sub} is segmented, whose content describes a visual entity s' . The visual entity s' might be a known entity, or it might be a novel entity, not yet known to the agent. The agent might also detect a second visual object o' with bounding box BB_{obj} in the scene and might be interested in its relationship to s' . The content of a third bounding box BB_{pred} , typically encompassing the subject and object bounding boxes, describes the predicate. The goal is now to produce statements that are likely true, considering the context of the scene.

We get for the episodic instance

$$\mathbb{P}(t' = t | scene_{t'}) = \text{softmax}_t^\beta(\mathbf{n}_T) \quad (17)$$

where $\mathbf{n}_T = \mathbf{a}_t^\top \text{sig}(\mathbf{f}(scene_{t'}))$.

We obtain for the subject (cf. Equation 11)

$$\mathbb{P}(s' = s | t' = t, BB_{sub}, scene_{t'}) = \text{softmax}_s^\beta(\mathbf{n}_S)$$

with $\mathbf{n}_S = \mathbf{a}_s^\top \text{sig}(\mathbf{f}(BB_{sub}) + \mathbf{g}(\mathbf{q}_T))$ and $\mathbf{q}_T = \mathbf{a}_t + \mathbf{f}(scene_{t'})$.

For unary entity labels, we obtain

$$\mathbb{P}(c' = c | s' = s, t' = t, BB_{sub}, scene_{t'}) = \text{softmax}_c^\beta(\mathbf{n}_C)$$

with $\mathbf{n}_C = \mathbf{a}_c^\top \text{sig}(\mathbf{q}_S)$.

We obtain for the object (cf. Equation 13)

$$\mathbb{P}(o' = o | s' = s, t' = t, BB_{sub}, BB_{obj}, scene_{t'}) = \text{softmax}_o^\beta(\mathbf{q}_O)$$

with $\mathbf{q}_O = \mathbf{a}_o^\top \text{sig}(\mathbf{f}(BB_{obj}) + \mathbf{g}(\mathbf{q}_S))$ and $\mathbf{q}_S = \mathbf{a}_s + \mathbf{f}(BB_{sub}) + \mathbf{g}(\mathbf{q}_T)$.

For binary labels, we get

$$\mathbb{P}(p' = p | o' = o, s' = s, t' = t, BB_{sub}, BB_{obj}, BB_{pred}, scene_{t'}) = \text{softmax}_p^\beta(\mathbf{n}_P)$$

with $\mathbf{n}_P = \mathbf{a}_p^\top \text{sig}(\mathbf{f}(BB_{pred}) + \mathbf{g}(\mathbf{q}_O))$ and $\mathbf{q}_O = \mathbf{a}_o + \mathbf{f}(BB_{obj}) + \mathbf{g}(\mathbf{q}_S)$. For Boolean variables and for unary statements, we obtain

$$\mathbb{E}(Y_{s', hA, c, t'} | s' = s, t' = t, BB_{sub}, scene_{t'}) = \text{sig}(\mathbf{n}_C) \quad (18)$$

and for binary statements

$$\mathbb{E}(Y_{s', p, o', t'} | o' = o, s' = s, t' = t, BB_{sub}, BB_{obj}, BB_{pred}, scene_{t'}) = \text{sig}(\mathbf{n}_P). \quad (19)$$

Here, $\mathbf{f}(\cdot)$ is a representation vector derived from visual inputs realized by a deep convolutional neural network (DCNN, see Section 6).

5. Algorithmic Implementation and the Attention Approximations

5.1 Algorithmic Implementation by Stochastic Sampling

Figure 2 illustrates the unfolded processing steps of the architecture shown in Figure 1. Algorithm 1 describes the processing steps implementing the Equations in the previous sections by generating samples from the conditional probabilities. The main difference between episodic memory and perception is, that, in the latter, visual inputs from the contents of the bounding boxes are integrated. The algorithm outputs t^* , s^* , c^* , o^* , p^* from which we can form triples of the form (s', hA, c^*) and (s', p^*, o') , but also (s^*, hA, c^*) and (s^*, p^*, o^*) . Here, for example, s' is an entity in the scene, and s^* is an entity with a permanent representation in the index layer. In perception, the episodic instance t^* would be the index of a similar episode in the past.

In Figure 2, we indicate that we can also sample a unary label for the complete scene or episode and for the predicate, which we do not actually do in the experiment. We also indicate additional processing steps for calculating embeddings. For example, $\tilde{\mathbf{q}}_S \leftarrow \mathbf{q}_S + \mathbf{a}_{c^*}$ is the embedding of the sentence (s^*, hA, c^*) in the context of the scene. Whereas concepts, predicates, and episodic instances have static embeddings realized as connection weights linking the index layer and the representation layer, the embedding of a sentence is dynamic.

In Algorithm 1, we explicitly introduce the latent vector \mathbf{h} modelling dynamic context, supporting the implementation of the nonlinear function $\mathbf{g}(\cdot)$. Table 2 illustrates the generation of activated indices (“firing indices or neurons”) in different operational modes.

The algorithm for semantic memory (i.e., the pre-observation model) is identical to the one for episodic memory, only that $\mathbf{a}_{t^*} \leftarrow \bar{\mathbf{a}}$. One might also specify $s^* \leftarrow s$ and not sample it.

Sampling is a unique, maybe temporary, commitment to a concept. In a biological interpretation, a single winning unit (neuron) fires. Even a single or a few spikes can be sufficient for communicating the winning concept to later processing. Since sampling commits unique indices t^* and s^* , this allows the association of semantic and episodic memory experiences to the observation.

Algorithm 1: The BTN for Perception and Episodic Memory.

```
1 switch Perception do
  | Input:  $scene, BB_{sub}, BB_{obj}, BB_{pred}; u = 1$ 
2 end
3 switch Episodic Memory do
  | Input:  $t_{in}^*; u = 0$ 
4 end
  Output:  $t^*, s^*, c^*, o^*, p^*$ 
5  $\mathbf{h} = \mathbf{0}$  ▷ Alternatively,  $\mathbf{h}$  is inherited from past decoding
6  $\tilde{\mathbf{q}}_T \leftarrow uf(scene)$  ▷ Representation of overall scene
7 switch Perception do
8   |  $\forall t \in \mathcal{T} : n_{T(t)} \leftarrow \mathbf{a}_t^\top \text{sig}(\tilde{\mathbf{q}}_T)$ 
9   | Sample  $t^* \sim \text{softmax}^\beta(\mathbf{n}_T)$  ▷ Sample a past episodic
10 end
11 switch Episodic Memory do
12   |  $t^* = t_{in}^*$  ▷ Input for episodic memory
13 end
14  $\mathbf{q}_T \leftarrow \tilde{\mathbf{q}}_T + \mathbf{a}_{t^*}$ 
15  $\forall c \in \mathcal{C} : n_{C(c)} \leftarrow \mathbf{a}_c^\top \text{sig}(\mathbf{q}_T)$ 
16  $\mathbf{h} \leftarrow B\text{sig}[\text{sig}(\mathbf{h}) + V\text{sig}(\mathbf{q}_T)]$ 
17  $\tilde{\mathbf{q}}_S \leftarrow uf(BB_{sub}) + W\text{sig}(\mathbf{h})$  ▷  $\mathbf{g}(\cdot) = W\text{sig}(\mathbf{h})$ 
18  $\forall s \in \mathcal{C} : n_S(s) \leftarrow \mathbf{a}_s^\top \text{sig}(\tilde{\mathbf{q}}_S)$ 
19 Sample  $s^* \sim \text{softmax}^\beta(\mathbf{n}_S)$  ▷ Sample a subject entity
20  $\mathbf{q}_S \leftarrow \tilde{\mathbf{q}}_S + \mathbf{a}_{s^*}$ 
21  $\forall c \in \mathcal{C} : n_C(c) \leftarrow \mathbf{a}_c^\top \text{sig}(\mathbf{q}_S)$ 
22 Sample  $c^* \sim \text{softmax}^\beta(\mathbf{n}_C)$  ▷ Sample a unary label for the subject
23  $\mathbf{h} \leftarrow B\text{sig}[\text{sig}(\mathbf{h}) + V\text{sig}(\mathbf{q}_S)]$ 
24  $\tilde{\mathbf{q}}_O \leftarrow uf(BB_{obj}) + W\text{sig}(\mathbf{h})$ 
25  $\forall o \in \mathcal{C} : n_O(o) \leftarrow \mathbf{a}_o^\top \text{sig}(\tilde{\mathbf{q}}_O)$ 
26 Sample  $o^* \sim \text{softmax}^\beta(\mathbf{n}_O)$  ▷ Sample an object entity
27  $\mathbf{q}_O \leftarrow \tilde{\mathbf{q}}_O + \mathbf{a}_{o^*}$ 
28  $\mathbf{h} \leftarrow B\text{sig}[\text{sig}(\mathbf{h}) + V\text{sig}(\mathbf{q}_O)]$ 
29  $\tilde{\mathbf{q}}_P \leftarrow uf(BB_{pred}) + W\text{sig}(\mathbf{h})$ 
30  $\mathbf{q}_P \leftarrow \tilde{\mathbf{q}}_P$ 
31  $\forall p \in \mathcal{P}^B : n_P(p) \leftarrow \mathbf{a}_p^\top \text{sig}(\mathbf{q}_P)$ 
32 Sample  $p^* \sim \text{softmax}^\beta(\mathbf{n}_P)$  ▷ Sample a binary label
33 return  $t^*, s^*, c^*, o^*, p^*$ 
```

5.2 Embedded Symbolic Reasoning by Chaining

As indicated by the four groups of dots in Figure 2, unary decoding can continue during the time of operation of the dynamic context layer. For example, if $\tilde{\mathbf{q}}_S$ represents the triples (*Sparky, hA, Dog*) and the agent has learned the generalized statement (*Dog, hA, Mammal*),

then it is likely that the index for *Mammal* fires next. Thus the sequential index pattern *Sparky, Dog, Mammal* is interpreted by the agent that Sparky is a dog and a mammal. This form of embedded symbolic reasoning, as further discussed in Section 8.3, is based on generalized statements discussed in Section 3.8.

5.3 Sampling and the Attention Approximation

For episodic and semantic memory, we need to generate triples of the form (s^*, hA, c^*) , whereas to label entities in perception, we only need: (s', hA, c^*) , that is, s' , the subject in the scene, does not need to be identified as a stored entity s^* . For example, if the agent is attacked by a bear, it does not care what the name of the bear might be; it might be a bear not yet known. Our approach is motivated by the computational attention approach used in deep learning (Vaswani et al., 2017). It can be related to the concept of internal attention in cognition to distinguish it from external attention, for example, the attention to a particular perceptual entity. With parallel hardware (e.g., brainware) the sampling process can be replaced by a computational attention approximation, which can be executed in parallel and does not commit to a particular s^* .

The *semantic attention* (SA) approximation for the subject replaces line 20 in Algorithm 1 with

$$\mathbf{q}_S \leftarrow \tilde{\mathbf{q}}_S + A \text{softmax}^\beta(A^\top \text{sig}(\tilde{\mathbf{q}}_S)).$$

In matrix A we only consider the columns relating to entities. In terms of standard attention, as defined by (Weston et al., 2014; Sukhbaatar et al., 2015; Vaswani et al., 2017), the argument of the softmax calculates the *score vector*, the softmax output is the *alignment vector*, the multiplication with matrix A calculates the *context vector*, the columns of A are the *key vectors* and the *value vectors*, and $\text{sig}(\tilde{\mathbf{q}}_S)$ is the *query*. The main differences to the standard approaches to attention is that, in our approach, *key vectors* and *value vectors* are *stored* semantic embeddings and the first term on the right is the pre-activation instead of the post-activation; processing several bounding boxes from the same image in parallel, as it is done in visual attention, for example, in (Koner et al., 2020), would violate our one-brain hypothesis from Section 8.6. In the sampling mode of operation, a symbolic unit in the index layer competes with other units in the same layer to be activated. In contrast, in our attention approximation, the index layer acts as a standard neural network layer (without sampling) with softmax activation (i.e., the sum in the equation) and pre-activation skip connections ($\tilde{\mathbf{q}}_S$ in the equation).

In training the model, and also in the attention approximation, we use $\beta = 1$. With $\beta = 1$, predictive samples reflect the predictive uncertainty in prediction. In the actual sampling experiments, we set the inverse-temperature parameter $\beta \rightarrow \infty$, effectively taking the concept with the highest probability. We call this winner-take-all sampling. If there were a dominating dimension or if we would set $\beta \rightarrow \infty$ in the attention approximation, the attention approximation would become identical to winner-take-all sampling.

Similarly, for the object we replace line 27 with

$$\mathbf{q}_O \leftarrow \tilde{\mathbf{q}}_O + A \text{softmax}^\beta(A^\top \text{sig}(\tilde{\mathbf{q}}_O)).$$

In matrix A we only consider the columns relating to entities. Attention can also be applied to episodes in perception. The *episodic attention* (EA) for the episodic index replaces line 14

with

$$\mathbf{q}_T \leftarrow \tilde{\mathbf{q}}_T + A \text{softmax}^\beta(A^\top \text{sig}(\tilde{\mathbf{q}}_T)).$$

In matrix A we only consider the columns relating to episodic instances. EA is the default in all experiments on perception.

With no scene input, $\mathbf{q}_T \leftarrow (1/N_T) \sum_t a_t \approx \bar{a}$. Thus we can consider semantic memory also as a situation where all instance indices are activated equally. Technically, this is a form of an empirical hierarchical Bayesian approach where the embedding average is inherited by a new instance.

Sequential Index Pattern	Equivalent Triple Statements	Mode
<i>Sparky, Dog, Friendly, Black</i>	$(\text{Sparky}, hA, \{\text{Dog}, \text{Friendly}, \text{Black}\})$	SM:U
<i>Sparky, Jack, looksAt, ownedBy</i>	$(\text{Sparky}, \{\text{looksAt}, \text{ownedBy}\}, \text{Jack})$	SM:B
$t, \text{Sparky}, \text{Dog}, \text{Friendly}, \text{Black}$	$t: (\text{Sparky}, hA, \{\text{Dog}, \text{Friendly}, \text{Black}\})$	EM: U
$t, \text{Sparky}, \text{Jack}, \text{looksAt}, \text{ownedBy}$	$t: (\text{Sparky}, \{\text{looksAt}, \text{ownedBy}\}, \text{Jack})$	EM: B
$\text{Dog}, \text{Friendly}, \text{Black}$ $\text{looksAt}, \text{ownedBy}$	$t', s': (s', hA, \{\text{Dog}, \text{Friendly}, \text{Black}\})$ $t', s', o': (s', \{\text{looksAt}, \text{ownedBy}\}, o')$	P-Dir/SA: U P-Dir/SA: B
$\text{Sparky}, \text{Dog}, \text{Friendly}, \text{Black}$	$\text{At } t': (\text{Sparky}, hA, \{\text{Dog}, \text{Friendly}, \text{Black}\})$	P-Samp: U
$\text{Sparky}, \text{Jack}, \text{looksAt}, \text{ownedBy}$	$\text{At } t': (\text{Sparky}, \{\text{looksAt}, \text{ownedBy}\}, \text{Jack})$	P-Samp: B

Table 2: The first column shows the sequence of activated indices (“firing neurons”). The second column shows equivalent triples. The last column shows the operational model: SM (semantic memory), EM (episodic memory), P-Dir (direct perception), P-SA (semantic attention), and P-Samp (Perception with sampling). We indicate that for given episodic instances, several unary labels (U) and binary labels (B) can be generated. If we extend the knowledge graph to also contain nodes for predicates and time instances, then, in sampling, exactly one node is active at a time.

6. Perception, the Representation Layer, and the Dynamic Context Layer

In this and the following sections, we present experimental results. Here, we focus on perception, and in the following sections on engrams, semantic memory, reasoning, language, episodic memory, and self-supervised learning. Intertwined with the experiments, we make the connection to cognition and neuroscience.

There is a long tradition in cognition and neuroscience to distinguish between anatomical brain structure, for example, the actual anatomical structure of the biological neural network, and functional architectures of the cortex (Friston et al., 1995; van den Heuvel and Sporns, 2013; Bassett and Sporns, 2017; Sporns, 2018; Leopold et al., 2019). We discuss the functional architectures primarily but occasionally also consider structure.

6.1 Data Set

We tested our approach experimentally using an augmented version of the VRD data set (Lu et al., 2016). In the past, this data set has been the basis for much research on visual

relationship detection. Each visual entity is labeled as belonging to one out of 100 classes. There are 70 binary labels, with 37,993 binary statements total. We followed other work and assigned 4000 images to the training set and 1000 images to the test set. The training images contain, overall, 26,430 bounding boxes, thus on average, 6.60 per image.

We generated a first derived data set, VRD-E (for VRD-Entity), with additional labels for each visual entity. First, each entity in each image obtains an individual entity index (or name). The 26,430 bounding boxes in the training images describe as many entity indices. Second, we used concept hierarchies from WordNet (Fellbaum (2010); see Figure 3). Each entity is assigned exactly one basis class (or *B-Class*) from VRD (e.g., *Dog*), one parent class (or *P-Class*, for example, *Mammal*), and one grandparent class (or *G-Class*, for example, *LivingBeing*). At any level, we use the default class *Other* for entities that cannot be assigned to a WordNet concept. We perform subclass reasoning in the training data and label an entity with its entity index, its *B-Class*, *P-Class*, and *G-Class*.

In addition, we used pretrained attribute classifiers (Anderson et al., 2018; Wu et al., 2019) to label visual entities using the attribute ontology shown in Figure 3. Each visual entity obtains exactly one color (including the color *Other*) and exactly one activity attribute, for example, a person can be standing or running. We also introduce the unary labels *Young* and *Old*, which are randomly assigned, such that these can only be predicted for test entities that already occurred in training, not for novel entities.

Furthermore, we introduce the nonvisual, or hidden, unary label *Dangerous* to all living things and *Harmless* to all nonliving things. We do not use these labels in perceptual training; we use them instead to demonstrate how semantic memory can supplement nonvisual labels.

In summary, every visual entity receives one entity index and, in addition, seven positive unary labels— for example, *Entity=Sparky*, *B-Class=Dog*, *P-Class=Mammal*, *G-Class=LivingBeing*, *Age=Young*, *Color=Black*, *Activity = Standing*, *Risk=Harmless*.

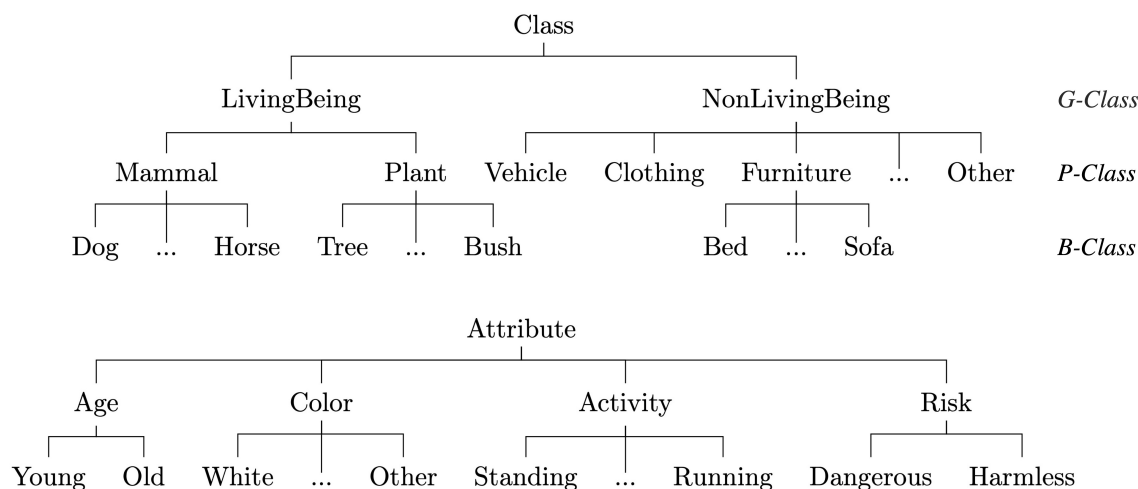


Figure 3: Ontologies. (Top) Class ontology. (Bottom) Attribute ontology.

Based on the VRD-E data set, we generate the VRD-EX data set. Here, we distort each image in the training data set, which generates another 4000 training images⁴. We used an open library to distort the image (Jung et al., 2020). To obtain a distorted image, we apply a sequence of transformations, including translation, rotation, shearing, and horizontal flipping of the original image (Bloice et al., 2017). Each transformation is associated with a probability of actually using it. When the operation has coefficients, such as the displacement of translation, a random value within a reasonable range is generated. Thus VRD-EX has 8000 training images. Then we perform another distortion on each original image and create a set of another 4000 test images.⁵ Note that in these new 4000 test images, every visual entity has already occurred in the training set twice (See Figure 4).

In addition to the visual concepts, we introduce nonvisual entities, which do not occur in any image. In the experiments, we relate visual entities to those hidden entities, for example, by the predicate *ownedBy* or the predicate *lovedBy*. For example, each visual dog is owned by a person who is not in any scene. Table 3 shows the overall statistics.

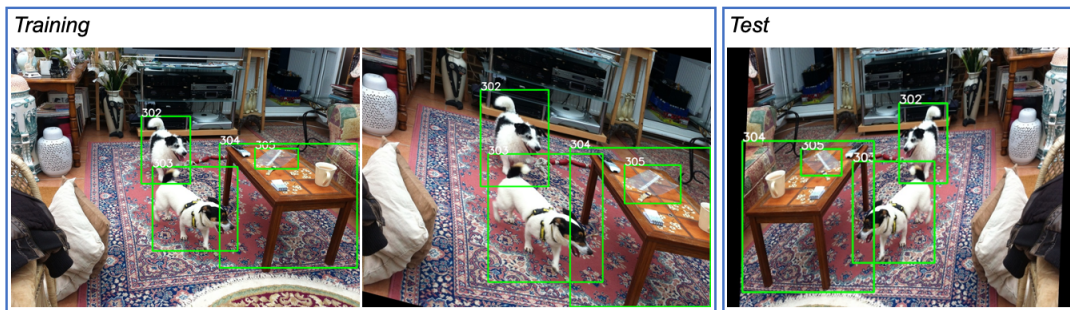


Figure 4: Generation of VRD-EX images. (Left) An original VRD image. (Center) A distorted VRD-EX image that is assigned to the training set. (Right) A distorted VRD-EX image that is assigned to the test set.

data set	Training Images	Test Images	#BB Train	#VisEnt Train	#BinStat Train	# Attr/Ent Train
VRD (Lu et al., 2016)	4000	1000	26430	26430	30355	1
VRD-Entity	4000	1000	26430	26430	30355	8
VRD-EX _{tended}	7737	3753	50910	26430	59095	8

Table 3: Statistics of the data sets VRD-E and VRD-EX. Attr/Ent stands for the average number of unary labels per episodic instance. Overall, 15% of the labels are *Others*.

Training was performed using the Adam optimizer (Kingma and Ba, 2014). For evaluation, we consider top-1 accuracy for unary labels and Hits@k for binary labels and sampling of entities. Hits@k is the fraction of correct entities that appear in the top-k-ranked entities.

5. Due to distortion, some objects and images are discarded, resulting in a reduced number of samples.

6.2 DCNN

As discussed, we use a deep convolutional neural network (DCNN) for the mapping $\mathbf{f}(\cdot)$ from scene and bounding box content to the representation layer. Kriegeskorte and Douglas (2018) and others have discussed how DCNNs might represent functional modules in the brain. This is not the topic of this article. Our approach is object-centric. The specific DCNN architecture we use is the VGG-19 architecture (Simonyan and Zisserman, 2014), pre-trained on ImageNet (Russakovsky et al., 2015) and fine-tuned to our data. The VGG-19 architecture constitutes the backbone layer of the Faster R-CNN (Ren et al., 2015), which we employ to produce a set of bounding boxes as output, where each bounding box contains an object. (For more details, see the Appendix E.) In all of our experiments, we used rank $r = 4096$ as the dimension of the representation layer.

6.3 Direct Perception

In our approach, four bounding boxes are analysed: first, a bounding box describing the complete scene; second, a bounding box describing the subject entity; third, a bounding box describing the object entity; and finally, a bounding box describing the predicate. In *direct perception*, we assume that labels are independently generated by perception, purely based on bounding box content, and entirely in bottom-up mode. There are connections from the representation layer to the index layer but not in the opposite direction. Also, there are no connections between the representation layer and the dynamic context layer.

From a perceptual view, this approach exploits neither dependencies between labels for the same bounding box (that a detected Sparky is known to be a black dog), nor dependencies between labels for the different bounding boxes. Direct perception requires concept indices, together with their associated embedding vectors, which become the connection weights from the representation layer to the index layer. Without connections from the index layer back to the representation layer, no episodic memory or semantic memory can be formed. In direct perception, it is unclear how the overall brain is informed about winning or likely concept labels. This is one of the issues addressed in BTN perception, discussed next.

Table 5 shows results for unary labels on direct perception (rows labeled P-Direct). The results on VRD-E are pretty competitive. The performance is improved in VRD-EX, likely due to some form of memorization by overfitting since different views on the same entities occur in the training set and the test set. The results on the prediction of binary labels (see Table 4) are not very good; as expected, and a well-known result from scene graph analysis, information on the subject and the object bounding boxes is required for binary label prediction.

6.4 BTN Perception

In perception with the BTN, we add connections from the index layer back to the representation layer. We also connect the representation layer with the dynamic context layer. This has several important benefits.

First, in a top-down mode of operation, the representation layer, and thus the whole brain, is informed about which concepts are detected in the scene. Thus, if an entity is

labeled as being *Dangerous* or recognized as being *Jack*, this information is communicated to the whole brain. This pattern completion process is a property typically discussed in the context of associative memories. If this affects earlier processing layers, this process is referred to as embodiment. Some theories on embodiment assume, as the name indicates, that thought and language influence bodily states, or at least the brain regions mapping directly to bodily states (Lakoff et al., 1999).

Second, labels would get biased: If *Sparky* is detected in the scene, it will bias the unary labeling toward the unary label *Black* if this is known from semantic memory to be true.

Third, by connecting the representation layer with the dynamic context layer, we obtain a memory that can transport information from the analysis of one bounding box to another. This change of attention from one bounding box to another might be associated with saccadic eye movement, where the dynamic context layer stores information during the transition phase. The dynamic context layer is crucial for evaluating triples involving binary labels.

Finally, the bidirectional connectivity between index layer and representation layer enables the formation of episodic and semantic memory and can enrich perception with information from either, as will be discussed in the following sections.

One can think of perception as a triple-generating language model in the context of visual inputs. Through the top-down information flow via the representation layer, the whole brain is kept up to date about scene content, about which concepts have been detected in the scene and which statements describe scene content.

The representation layer reflects the context. For example, if *Sparky* is detected in the subject bounding box, then \mathbf{q}_S is an embedding of *Sparky* in the context of the scene. If *Sparky* in the scene is labeled as *Friendly*, then $\mathbf{\ddot{q}}_S$ would be the embedding of the statements (*Sparky*, *hA*, *Friendly*) in the context of the scene. If *Sparky* looks at *Jack* in the scene and this is detected, then \mathbf{q}_P would be the embedding of the statement (*Sparky*, *looksAt*, *Jack*) in the scene’s context.

6.5 Experiments with BTN Perception

In all experiments, we use episodic attention (EA): a sampling of a past episodic instance t^* is used only to evoke a related episodic memory, as discussed in Section 9.

Experiments with semantic attention are labeled P-SA where there is no commitment to a specific s^* or a specific o^* . In semantic attention, one is primarily interested in fast labeling rather than in associating memories.

The results in Table 5 show that on the VRD-E data set with novel entities in the test set, perception with semantic attention (P-SA) shows the best results. However, direct perception (P-Direct) is quite competitive.

On the VRD-EX data set with known entities in testing, perception with sampling, P-Samp, is best, where the algorithm could “remember” past encounters of the same entities. Here, P-Direct is not competitive. For P-Samp, we use winner-take-all sampling.

Table 4 shows results from binary label prediction. For the binary labeling with novel entities on the VRD-E data, P-SA shows the best performance. The inferior performance of P-Direct confirms that the dynamic context layer is important for achieving good performance. We see a large improvement for the VRD-EX data set.

	Model	Binary labels	
		@10	@1
VRD-E	P-Direct	85.45	31.68
	P-Samp	90.39	45.09
	P-SA	91.33	46.84
VRD-EX	P-SA	99.58	80.63

Table 4: Binary label prediction in perception. As the inferior results for P-direct indicate, the dynamic context layer is essential to perform well in binary label prediction (rows 1-3). For known entities (VRD-EX) (fourth row), entity indices permit some memorization and improve performance. Best results in each column are in bold.

In Table 6, we compare our model with other visual relationship detection methods. In the methods from the literature, only binary labels are tested, with one class label for each entity. In our framework, this would be the decoding of generalized statements. The table demonstrates that generalized statements (i.e., probabilistic rules) can be generalized themselves. In a zero-shot situation, where the triple (*subject-class*, *predicate*, *object-class*) never occurred in training, our approach achieved a recall score of 81.61%, which is much better than the result from the BFM (Baier et al., 2017) with 76.05%. In summary, our approach gives competitive results overall and superior results for zero-shot binary labeling.

Figure 5 illustrates perception with unary labels and visual and nonvisual binary statements.

6.6 Representation Layer

In our model, \mathbf{q} , that is, the activation of the representation layer, represents the cognitive brain state. In cognitive neuroscience, the representation layer is referred to as “mental canvas” or the “global workspace,” enabling a global information exchange. It is also known as “theater of the brain,” “communication platform,” “communication bus,” or “blackboard.” It is assumed to have a distributed representation involving large parts of the brain. (Binder and Desai, 2011) states: “The neural systems specialized for storage and retrieval of semantic knowledge are widespread and occupy a large proportion of the cortex in the human brain.” During perception or memory recall, the representation layer integrates information and then makes this information available to the brain as a whole. In particular, the embedding vector of the episodic experience, that is, \mathbf{a}_t , represents a holistic, integrative view of the cognitive state at that instance.

Since the world of an agent mostly changes smoothly, one can develop models to forecast embedding vectors of future instances (Tresp et al., 2015; Han et al., 2020b). The brain’s cognitive state is, to a large degree, determined by perception, which might explain a personal feeling of personal instability and sensitivity to external influences, whereas other individuals, represented by their embedding vectors, appear stable and slowly changing.

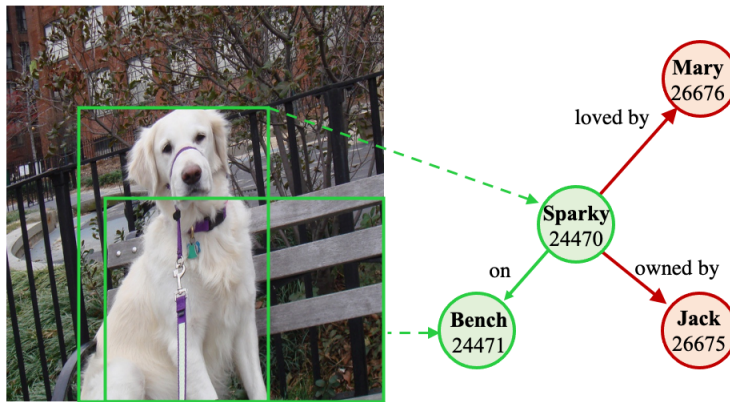


Figure 5: Illustration of perception with known entities (VRD-EX). The left bounding box is identified (sampled) as $s' = s^* = 24470 = \textit{Sparky}$. Highly ranked unary labels are: *Dog*, *Mammal*, *LivingBeing*, *Young*, *White*, *OtherActivity*. Highly ranked unary labels for the second bounding box (24471) are: *Bench*, *Furniture*, *NonLivingBeing*, *Old*, *OtherColor*, *OtherActivity*. Sampled binary statements are: *(Dog, sitsOn, Bench)*, *(LivingBeing, on, Furniture)*, *(Mammal, sitsOn, Old)*, *(White, sitsOn, Bench)*. We also indicate how semantic memory can support perception by adding binary statements to object entities, not in the scene: *(Sparky, ownedBy, Jack(26675))*, *(Sparky, lovedBy, Mary(26676))*, where *Jack* and *Mary* are not on the scene, but in the agent’s semantic memory. The semantic memory experience is further discussed in Section 7.

	Model	Unary labels (accuracy)							
		Entity	B-Class	P-Class	G-Class	Y/O	Color	Activity	Average
VRD-E	P-Direct	-	81.30	88.44	94.67	50.60	69.06	83.73	77.97
	P-Samp	-	80.91	88.85	95.02	47.88	68.02	82.90	77.26
	P-SA	-	81.80	89.34	95.44	49.46	68.93	83.58	78.09
VRD-EX	P-Direct	92.20	90.87	93.54	96.47	76.54	84.49	92.77	89.56
	P-Samp	92.81	96.47	97.49	98.47	94.58	94.08	97.59	95.93
	P-SA	92.65	95.27	97.11	98.20	93.10	92.71	96.83	95.12

Table 5: Unary label prediction in perception. We evaluate (s', hA, c) for the entities s' in the scene. The columns labeled “B-Class, P-Class, G-Class” are class labels, and the columns labeled “Y/O, Color, Activity” are unary labels. “Average” refers to the average over all columns. In P-Direct, there are no links from n to q , and q and h are independent. P-Samp stands for the sampling approach using the index s^* with maximum activation. P-SA uses semantic attention. On the VRD-E data set (rows 1-3), where each entity is novel, P-SA shows the best results. Not surprisingly, P-Direct is also quite competitive. On the VRD-EX data set (rows 4-6), where each entity is known, P-Samp shows the best results since it can recognize specific entities. In particular, the unary label Y/O can only be learned for already known entities. Not surprisingly, P-Direct is significantly worse since it cannot benefit from memory, although overfitting leads to better results than with the VRD-E data. The column labeled “Entity” evaluates if the right entity is recognized. Thus we evaluate $(s', sameAs, s)$ for the entities s' in the scene. Best results in each block are in bold.

Model	ph	z-s-ph	rl	z-s-rl
P-SA	24.50	8.55	93.99	81.61
P-Direct	13.54	5.73	84.64	68.35
BFM (Baier et al., 2017)	25.11	7.96	93.81	76.05
Approach in (Tresp and Ma, 2016)	23.45	10.95	93.32	78.79

Table 6: Binary label prediction in perception. We compare binary labeling with methods published in the literature using the original VRD data set. *Phrase (ph)* shows the recall of binary labels, where also the extracted bounding box contents are evaluated. The *binary (relationship) label (rl)* shows the recall of the binary label given the ground truth class of subject and object. *z-s-ph* and *z-s-rl* denote zero-shot performance for triples that did not occur in the training set. Our proposed model, P-SA (first row), is superior in the last task. The much better performance of P-SA compared to P-Direct (second row) demonstrates the importance of the dynamic context layer. The third and fourth rows show results from approaches published in the literature. Best results in each block are in bold.

6.7 Dynamic Context Layer

To really capture the content of a scene, it is essential to understand the relationships between the scene entities; this requires an additional storage facility since, following our one-brain hypothesis, the brain possesses only exactly one global representation layer (see the discussion in Section 8.6). As illustrated in Figures 1 and 2, we propose that this functionality is represented by the dynamic context layer h , which represents the state of some of the brain modules not explicitly discussed in this article. It receives input from the representation layer, processes that information as a recurrent neural network, and then returns the processed information to the representation layer, which starts processing visual input from a subsequent bounding box. Thus, the dynamic context layer stores and processes information between saccadic eye movements. The dynamic context layer might also involve higher brain regions, like working memory.

Figure 6 shows the Pearson correlation between activities of units in the dynamic context layer. It shows a certain clustering structure, as exhibited in recent studies on functional and structural brain connectivity (e.g., Figure 2 in Pope et al. (2021)). In the brain, such a block structure is a signature of a small-world architecture. Since it is not clear what the dynamic context layer in our model actually represents in the brain, we refrain from a deeper discussion of this apparent similarity.

The dynamic context layer also represents internal mental states that might not be driven directly by perception and memory. In neuroscience, the default network is assumed to be active when a person is not focused on the outside world; instead, the brain is at wakeful rest, such as during daydreaming or mind-wandering (Gazzaniga et al., 2013). Perception, memory recall, the initiation of activity, and other causes can induce the brain to leave the default state. In the default state of our model, the brain’s representation layer interacts with the dynamic context layer, without being driven by perception.

7. Semantic Memory Engrams, Semantic Decoding, and Multimodality

7.1 Background on Semantic Memory

According to (Tulving, 1985), the semantic memory experience is independent of a particular episodic experience. It developed out of perception as an emerging property where semantic enrichment became independent of perceptual input. Thus, in the transition from episodic memory to semantic memory, provenance is lost. It is the longest-lasting and most durable memory since it models the stable statistics in the world. It aggregates information and is a dictionary view of the agent’s life experience. Our approach explains the remarkable similarity between episodic and semantic memory: Semantic memory models the episodic memory of a future instance. Whereas an episodic memory experience requires the activation of the embedding vector of the corresponding instance in the representation layer, a semantic memory experience requires the activation of the semantic memory representation $\bar{\mathbf{a}}$. Semantic memory can also be realized by the activation of all instance indices in the episodic attention approximation (see Section 5.3). Another alternative is represented in (Tresp et al., 2017b) where semantic memory is generated by averaging in embedding space. But that approach is restricted to multilinear models.




Figure 6: Pearson correlation between units in the dynamic context layer, followed by spectral clustering to order columns and rows.

7.2 Neuroscience Perspective: Semantic Memory Engrams

Engrams are memory traces in the brain (Ralph et al., 2017). In our approach, the combination of an index with its embedding vector would be an engram for that concept, although we will typically only refer to the latter part as the engram. An index is a symbol for a concept, whereas embeddings are part of an implicit concept memory and provide symbol grounding (Harnad, 1990; Barsalou et al., 2008).

Here, we are in agreement with several theories on semantic memory engrams from neuroscience. For example, Binder and Desai (2011) state that semantic memory engrams consist of both modal and amodal representations, supported by the “gradual convergence of information throughout large regions of temporal and inferior parietal association cortex.” The amodal representations are described as a high-level convergence zone. In our approach, the embeddings would be the modal distributed representations and the indices the amodal, local, and symbolic representations.

The relationship between the index layer and the representation layer reminds one of the hub-and-spoke model (Ralph et al., 2017). The hub is supposed to be located in the anterior temporal lobes (ATLs), which might be where concept indices are consolidated. The hub is connected to several different areas (e.g., visual cortex, auditory cortex, orbitofrontal cortex), which might be part of the biological realization of the representation layer. Other hubs might be located in the parietal and the temporal lobe (Binder and Desai, 2011) and maybe in the frontal lobe (Tomasello et al., 2017).

Some works propose that the anatomical distinction between the representation layer and the index layer might be blurred in the brain. One should instead assume an “inter-

active continuum of hierarchically ordered neural ensembles, supporting progressively more combinatorial and idealized representations” (Binder and Desai, 2011).

7.3 Neuroscience Perspective: Indices

Indices are explicit concept representations, that is, an index is a symbol for a concept and could be implemented functionally by a single unit. In the sampling mode of operation, a symbolic unit in the index layer competes with other units in the index layer. This is in contrast to the representation layer, where the activations of the ensemble of units contribute to the processing.

We are purposely imprecise about how exactly an index is represented anatomically (i.e., structurally), in brainware. In the one extreme, an index might be a single neuron, realizing localist codes, where neurons respond highly selectively to single entities (“grandmother cells”). The other extreme are densely distributed codes where items are encoded through the activity of many neurons (Kumaran et al., 2016). Most researchers favor a sparse population of cells, realizing a population code. A distributed representation might facilitate the consolidation of new information (see Section 10).

The debate about localized representations in the brain is ongoing. Specific concept cells have been found in the medial lobe (MTL) region of the brain. MTL includes the hippocampus, along with the surrounding hippocampal region consisting of the perirhinal (“what” path), parahippocampal (“where” path), and entorhinal neocortical regions. Researchers have reported on a remarkable subset of MTL neurons that are selectively activated by strikingly different pictures of given individuals, landmarks, or objects and, in some cases, even by letter strings with their names (Quiroga, 2012; Quiroga et al., 2005). Naturally, locality of representation is probably discovered only in well-designed experiments. In our model, an activated concept index activates many units in the representation layer, and a unit in the representation layer in turn activates many indices. Since index activations might change rapidly, the general appearance might be that of a globally activated system, hiding the locality of representation. An index is a focal point of activity, but it is never active in isolation.

7.4 Neuroscience Perspective on Concept Grounding

Our discussion can be related to recent results from neuroscience where, in different brain regions, maps have been discovered that code, for example, for visual appearance, sound, and function. For example, if the concept *Hammer* is activated in the index layer, it might excite brain areas indicating a typical hammer appearance, the sound of hammering, and the required motor movement of hammering, all represented in the biological equivalence of the representation layer (Rueschemeyer et al., 2010). Another example is the concept *Cat*, which includes the information that a cat has four legs, is furry, meows, can move, or can be petted (Kiefer and Pulvermüller, 2012). In terms of (Binder and Desai, 2011), the embeddings are modality specific (here: visual), and the indices represent convergence zones. As (Dehaene, 2014) puts it: “Every cortical site holds its own specialized piece of knowledge.”

Evidence for distributed semantic activation has also been described by (Huth et al., 2016; de Heer et al., 2017). Both papers developed a detailed atlas of semantic categories and




Figure 7: t-SNE visualization of embeddings of 15 classes and randomly selected entities from these classes. Features for the analysis are respective embedding vectors. A dot stands for an entity, for example, a specific dog, *Sparky*. The color of a dot marks the basis class of that entity. A cross stands for a class, for example, *Dog*, which is labeled by the same color as entities belonging to the class. The embedding vector of a class concept lies within the cluster of the classes' entities. Recall that we are not indicating cluster centers but the embeddings that happened to be learned for the class concepts in learning. In cognitive neuroscience, the semantic embedding space is sometimes referred to as conceptual space, where points denote objects, and regions denote concepts (Gärdenfors, 2016).

their topographic organization through extensive fMRI studies, showing the involvement of the lateral temporal cortex, the ventral temporal cortex, the lateral parietal cortex, the medial parietal cortex, the medial prefrontal cortex, the superior prefrontal cortex, and the inferior prefrontal cortex (Huth et al., 2016).

Recently, it has been proposed that embeddings are context-dependent (Popp et al., 2019). Our model suggests that concept embedding is rather stable but that the induced activation of a concept in the representation layer is superimposed with context information, as discussed throughout the article. The representation layer is activated by sensory input and activated concept indices, so the model is informed about a concept in context, even with the connection weights between index and representation layer being fixed.

7.5 Semantic Memory Engram in the BTN

In perception, the embedding vectors determine which indices are activated by the representation layer in a bottom-up operation. However, the operation is bidirectional: An index can activate the representation layer by its embedding, which might even activate earlier perceptual layers. We call this top-down mode the grounding or the embodiment of that concept. Since an embedding vector is optimized for a concept’s role in perception and memory, it implicitly reflects all background that is known about it.

Consider Equation 18, which labels entity s with unary label c . With no visual inputs, $\mathbf{a}_c^\top \text{sig}(\mathbf{a}_s)$ is the contribution from semantic memory; it acts as a prior for the unary statements. The visual input contributes the term $\mathbf{a}_c^\top \text{sig}(\mathbf{f}(BB_{sub}))$, which is the inner product of the embedding vector with high-level features generated from visual inputs. Thus, whereas semantic memory is dependent on the embedding vectors and does not care about the meaning of a dimension, this meaning is provided by perception—the inner product with high-level features. We call this the grounding of embedding vectors.

The embedding vector is a prototypical vector for that concept in a high-dimensional representation space; the assumption is that, in this space, distances are meaningful, and the complex mapping performed by the DCNN performs a normalization, such that, for example, all kinds of dogs with different shapes and appearances form a connected subspace in the semantic embedding space. In cognitive neuroscience, the semantic embedding space is sometimes referred to as conceptual space (Gärdenfors (2016) p. 21ff), where points denote objects, and regions denote concepts.

Figure 7 shows an analysis of the entity embeddings, based on the t-SNE visualization (Van der Maaten and Hinton, 2008). The plot clearly shows a schema-like organization of concept embeddings in this two-dimensional projection, also known as a cognitive map (Tolman, 1948).

7.6 Structure of the Embedding Vector

All memories are implemented as embedding vectors forming the matrix A . In accordance with the standard neurocognitive view, the entries in A correspond to synaptic coupling strengths (Gazzaniga et al., 2013).

In our approach, we implemented symmetric connections between the units in the index layer and the representation layer. Thus, we have connection matrices A and A^\top in Figures 1 and 2. Although backward connections are common in the brain, they are typically not

s^* (ID)	Unary attribute/ Class labels	Unary entity labels (p : <i>sameAs</i>)	Binary statements (s^* , p , o^*)
10830[Person]	Person, 0.96 Mammal, 0.97 LivingBeing, 0.98 Old, 0.98 Other, 0.98 Walking, 0.96	10830[Person]	(10830, wears, Shirt) (10830, wears, 3495[Glasses])
Dog	Dog, 0.99 Mammal, 1.0 LivingBeing, 1.0 Young, 0.52 Brown 0.35 Other 0.99	3537[Dog] 602[Dog] 5976[Dog]	(Dog, on, Grass) (Dog, behind, Person)
Mammal	Dog, 0.38 Mammal 1.0 LivingBeing, 1.0 Young 0.6 Brown 0.31 Other 0.96	3901[Cat] 9100[Horse]	(Mammal, on, Street)
Black	Person 0.24 Other 0.43 NonLivingBeing 0.95 Old 0.59 Black 0.99 Other 0.99	9812[Bag] 3634[Keyboard]	(Black, on, Person) (Black, under, Sky)

Table 7: Semantic memory experience with generalized statements on VRD-EX data. The first column shows the queried s^* (an entity, a class, or an attribute), the input to the algorithm. The second column shows highly rated attribute and class labels describing s^* . The third column shows highly-ranked (sampled) entities for the *sameAs* predicate. The fourth column shows binary statements. We see that person [10830] is a mammal and often wears shirts and glasses. However, we also see that the model “explains” what the class *Dog* stands for, what the class *Mammal* stands for, and what the attribute *Black* stands for: We see that a dog is a mammal with 100% probability, is brown with 35% probability, and is often on the grass or behind a person. We see that if something is black, it is often a person (24%), and black entities are often “on persons” and “under the sky.” See our discussion on generalized statements in Section 3.8.

symmetrical. We did extensive experiments where we removed that constraint. The result was that the performance dropped by about 1% in basically all experiments, so we stayed with symmetric connections in our work.

The representation layer is high-dimensional, although embeddings might be sparse, and a given index only activates a small number of components of the representation layer.

Model	Binary label		Unary labels: attributes/classes					
	@10	@1	B-Class	P-Class	G-Class	Y/O	Color	Act.
SM-givenClass	92.67	49.74						
RESCAL-givenClass	89.95	26.06						
SM-givenEntity	100.0	90.42	100.0	100.0	100.0	100.0	100.0	100.0
RESCAL-givenEntity	100.0	90.12						

Table 8: Semantic memory experience for models trained on the VRD-E data set. In the first two rows, we set the class labels of the subject and the object, e.g, *Dog* or *Person*, and predict binary labels. Our results are better than a state-of-the-art model (RESCAL). (Ruffinelli et al., 2019) showed in a recent study that RESCAL is competitive with state-of-the-art approaches. In the third and fourth rows, we set the entity indices (e.g., *Sparky* or *Jack*). In this memorization, the performance is boosted by the entity representations. In these experiments, we are testing memory retrieval and not generalization, which explains the high scores. Best results in each block are in bold.

Model	Classification Accuracy
	<i>Dangerous</i>
P-SA	52.01
SM	100.0
P-enriched	98.24

Table 9: Semantic memory experience integrated with perception. The task is to predict the unary label *Dangerous* with or without semantic memory. P-SA is the perceptual system, where the label *Dangerous* was not provided in training. It can only be predicted by chance if an entity is dangerous. We trained the label *Dangerous* as part of semantic memory. In SM, the semantic memory experience is activated, which supplements the information from semantic memory if a visual entity is dangerous. P-enriched shows that, when the semantic memory is trained with the *Dangerous* label, this information is also automatically integrated with perception, without an extra activation of a semantic memory experience. Enrichment works well with not-perceptual information, like social network background, which is, in a way, orthogonal to the visual scene input and is an indication that statements become dependent in training by the shared embeddings. The best result is in bold.

Naturally, the index that represents the color *Red* is likely to mostly connect to components of the representation layer that are excited by red images. Another advantage of sparse distributed representations (Rolls, 2016) is that this might lead to increased memory capacity (Ma et al., 2018a). Sparsity in the embedding vectors can be achieved in technical models, for example, by using appropriate regularization terms, like Lasso (Tibshirani, 1996). In our experiments, when we applied Lasso on all parameters, we obtained 70% sparsity.

7.7 Semantic Memory Experience: Semantic Decoding and Embodiment

The engram of a concept is identical to its embedding. However, its treasure is only discovered by the semantic decoding. A complete semantic memory experience is not only the activation of a concept index and the activation of the representation layer with its embedding (i.e., its grounding) but also the subsequent decoding of the embedding into triples. For example, after $s = \textit{Sparky}$ activates $\mathbf{q} \leftarrow \mathbf{a}_s$, the agent can classify, for example, if s is $c = \textit{Friendly}$, producing the triple statement $(\textit{Sparky}, hA, \textit{Friendly})$. This exactly is going on in semantic memory decoding when we generate the sequential index pattern $\textit{Sparky}, \textit{Friendly}$. The decoding produces triples from the pre-observation model which reflects the prior probabilities for triple statements.

In the dual subsymbolic view, the semantic decoding results in an embedding of the decoded sentences: Activating the index for *Sparky*, for example, embodies *Sparky*'s embedding vector in \mathbf{q}_S ; activating the index for *Friendly* then leads to a superposition of both embeddings in $\tilde{\mathbf{q}}_S$ and represents *Sparky* but with an emphasis that *Sparky* is friendly.

In a semantic memory recall, a second entity o can be activated, that is, $o = \textit{Jack}$, and then the binary label *ownedBy* obtains high activation, producing the triple $(\textit{Sparky}, \textit{ownedBy}, \textit{Jack})$. The subsymbolic representation of that statement is activated in \mathbf{q}_P .

Both grounding and semantic decoding are local to the entity and might integrate information on that entity that was collected at different episodic instances and in different modalities. Table 8 shows that in our model, semantic memory can realize a very precise memory recall.

Table 7 illustrates the semantic memory experience, which is triggered by an entity, a class, or an attribute. The latter two correspond to a recall of generalized statements (see Section 3.8). For example, it is shown that semantic memory can recall general information on *Dogs* (that they are mammals with 100%), *Mammals* (that they are dogs with 38% and living beings with 100%), and the color *Black* (which, with 24% probability, is a property of a person).

7.8 Multimodality and Social Networks

An agent can obtain data from different modalities. This could concern subsymbolic sensory data besides vision or symbolic information from conversations, books, movies, and other media. Multimodal data become part of episodic memory and semantic memory and enriches perception in different ways. First, a multimodal episode is represented by an episodic index, and by activating that index, the data from that modality in that episode can be retrieved. Second, semantic memory serves as a site of multimodal integration simply because semantic memory is trained on triples from all modalities. So multimodal data will become part of semantic memory. Third, consider perception. We can now distinguish between visual labels —labels, which were used in the training images like the color *Black*— and multimodal labels, which are the remaining labels, like *Dangerous* or *Rich*. Table 9 shows how multimodality directly enters in perception. We call this enriched perception, P-enriched. The unary label *Dangerous* was not trained in perception but just in semantic memory. The table shows that information from the semantic memory is integrated into perception and episodic memory, even without activating an explicit semantic memory experience. Activating an entity index in perception or episodic memory subsequently

activates the unary labels of that entity in context, but also recalls background of that entity, for example, relating to other modalities. Here it relies on semantic memory (see Section 8.1).

To further examine multimodality, we derived a social network involving all persons in the data set by linking entities representing persons with the predicate *knows*. Each person is linked with five other persons, whom we also denote as friends. In addition, we define a social network’s episodic instance as an instance at which the agent learns about the social contacts of one person. Our social network data set has 4987 person entities (along with their unary labels), 24953 *knows* statements, and 4987 social episodes. We refer to this data set as VRD-S. More details on the generation of the social network data can be found in the Appendix D.

Table 10 shows numerical results on VRD-S. Given a person of interest, the semantic memory can recover friends (object o^*) and predict unary labels. Table 11 illustrates episodic and semantic memory, including social network data recall.

Model	s^*		o^*		Unary labels of s^*					
	@10	@1	@10	@1	B-Class	P-Class	G-Class	Y/O	Color	Activity
Episodic Memory	100	51.47	99.97	65.80	100.0	100.0	100.0	100.0	100.0	100.0
Semantic Memory	-	-	97.39	18.25	100.0	100.0	100.0	100.0	100.0	100.0
RESCAL	-	-	95.76	17.96						

Table 10: Episodic and semantic memory experience for the social network data VRD-S.

The first row shows the performance of an episodic recall, with t^* given. Columns 2 and 3 evaluate the prediction of a subject entity. For columns 3 to 10, the subject s^* is also given. The second row shows the semantic memory experience where the subject s^* is given. The unary labels have perfect recall. Columns 3 and 4 show the performance of predicting a friend. Since each person has five friends, the recall@1 is upper-bounded by 20%. As a comparison, we show the performance of RESCAL in the third row, which shows slightly worse performance.

8. Reasoning and Language

8.1 Embedded Reasoning and Generalization

When the states of new triples are observed, this will be reflected in changes in the $\underline{\underline{T}}$ tensor, which affects $\mathbb{P}(A|\underline{\underline{T}})$ and the embedding matrix A (Equation 8). Thus, although triple statements are independent, given all entries in A , entries in A change with new data. For example, \mathbf{a}_s changes when new truth values of statements involving s become available. Thus, if it becomes known that Sparky is a “Dog”, \mathbf{a}_{Sparky} will be adapted and the unary statement for the label “Mammal” will exhibit a high probability.

Episodic memory, that is, the observation model, generates samples for observed entities in the modality of the instance. Although episodic memory purely models the observations, there is also some degree of generalization by exploiting correlations in observations. Semantic memory, that is, the pre-observation model, can generate samples for all entities



t^*	s^*	Unary labels	o^*	Binary labels
2177	14518 [Person]	Person, Mammal, LivB Young, Other, Sitting	14511 [Bus]	on
	14515 [Building]	Building, Other, NonLivB Old, White, Other	14516 [Roof]	has
	14511 [Bus]	Bus, Vehicle, NonLivB Old, Other, Other	14512 [Road]	nextTo (-)
6662	14518	Person, Mammal, LivB Young, Other, Sitting	10669	knows
	18010	Person, Mammal, LivB Young, Other, Other	14518	knows
	8318	Person, Mammal, LivB Old, Other, Other	14518	knows
	14518	Person, Mammal, LivB Young, Other, Sitting	10669	knows
			25066	knows
			12825 (-)	knows (-)

Table 11: Episodic and semantic memory experience that includes multimodal data (from a social network). The figure shows the original visual scene, which is unavailable during episodic recall. The scene index of the image is $t^* = 2177$. The top segment shows sampled subject entities (s^*), highly ranked unary labels (c^*), a sampled object o^* , and the top predicted binary label (p^*). All labels are correct except for one binary label (indicated by (-)). The second segment shows an episodic recall of a *social episode* with index ($t^* = 6662$). At that episodic instance, social network information (binary labels *knows*) was provided, which is recovered in the episodic recall. Shown are three correct triples that were recovered in the sampling of the episodic memory. The bottom segment shows *knows* statements recovered from semantic memory for $s^* = 14518$ (thus without a recall of a special episodic memory). Two binary statements are correct, and one is incorrect (indicated by (-)).

and modalities and fills in the unobserved. Semantic memory can show a significant degree of generalization. If a particular triple is often sampled in episodic memory but not in semantic memory, this would indicate an anomaly. As an example, consider that Sparky is a friendly dog (semantic memory), but in a particular episode, he was agitated and barking (episodic memory), then that is remarkable.

Perception is based on episodic attention (EA). After perceptual decoding, an instance embedding \mathbf{a}_t is realized, which is fitted to model the extracted triples. Subsequently, \mathbf{a}_t generalizes perception and becomes the basis for episodic memory. See also the discussion in Section 10.

The mind can distinguish between what is currently being observed, what the semantic memory fills in (by activating \bar{t} and $\bar{\mathbf{a}}$), and what an episodic memory could contribute (by activating a past t and \mathbf{a}_t).

8.2 Symbolic Reasoning

In contrast to embedded reasoning, symbolic reasoning solely depends on indices and their relationships, without involving the representation layer. It would be completely amodal, in contrast to the modal embedded reasoning. In particular, it does not involve the adaptation of the embedding matrix A . Examples would be the definitional structure of a bachelor (a male person who is not and has never been married) and reasoning with subclass relationships (any dog is also a mammal). Definitional examples, like the definition of a bachelor, enter the mind of the agent via language and are crisp. They describe how the world should be. Symbolic reasoning can also link singular events with state changes. For example, one might specify that if someone gets married (a singular event), the state of this person changes from single to married. Another example is the triangle rule in social networks: if person A and person B are known to be friends and person A is known to be a friend of person C, then person B is also likely a friend of person C. This example shows that symbolic reasoning is not necessarily deterministic. Descriptive dependencies, like the triangle rule, are learned by means of observations, are often probabilistic, and describe how the world actually is. (Halford et al., 2014) has explored how analogical reasoning can be performed in the data tensor, loaded in working memory. The work assigns symbolic reasoning to System-2. One might call it “inference by contemplation.” (Nickel et al., 2014) is an example where both embodied and symbolic reasoning are combined in an additive model.

8.3 Embedded Symbolic Reasoning

Triple sentences used in symbolic reasoning have an embedding and lead to another form of embedded reasoning. But in contrast to standard embedded reasoning, the embedding matrix A is not modified in embedded symbolic reasoning. A simple symbolic rule would be that a dog is a mammal, that is, $\forall s : (s, hA, Mammal) \leftarrow (s, hA, Dog)$. The corresponding triple statement is $Y_{Dog, hA, Mammal}$ is embedded and $\mathbb{E}(Y_{Dog, hA, Mammal})$ can be learned from data. Section 5.2 describes how the BTN’s embedding approach can produce the sequential index pattern *Sparky, Dog, Mammal*, which is interpreted as *Sparky* being a *Dog* and a *Mammal*. We call this embedded symbolic reasoning. It builds on the same BTN architecture as perception and memory. Embedded symbolic reasoning is integrated

in sampling. Assuming that, based on scene input, the unary label *Mammal* is active with 60% for a given visual entity, so the agent is uncertain about this label. Then, if sampling “decodes” that the visual entity is a *Dog*, embedded symbolic reasoning lifts the probability for the label *Mammal* to 100% (due to the learned rule that every dog is a mammal) and the corresponding unary label is sampled as well. Embedded symbolic reasoning can capture complex dependencies and thus can go beyond embedded reasoning. Embedded symbolic reasoning can be executed fast since it does not depend on a tuning of embeddings for entities or instances.

Table 7 shows generalized statements, which are the basis for embedded symbolic reasoning. We see that $(Dog, hA, Mammal)$ is true with 100% (every dog is a mammal) and $(Mammal, hA, Dog)$ is true with 38% (not every mammal is a dog). Also, Table 4 shows results for generalized binary statements.

8.4 A Foundation for Consciousness?

The dynamic context layer might play a role in working memory. In general, working memory is associated with decision-making and cognitive control (Baddeley, 1992) and is necessary for keeping task-relevant information active, as well as manipulating that information to accomplish behavioral tasks. A modern view is that working memory is distributed across the cortex (Buschman and Miller, 2020). There is an emerging consensus that most working memory tasks recruit a network of the prefrontal cortex (PFC) (front-of-the-brain hypothesis) and/or parietal areas in the prefrontal parietal network (PPN) (back-of-the-brain hypothesis) (Seth and Bayne, 2022).

PPN activity is consistently reported in both attention and consciousness studies (Bor and Seth, 2012). Their publication proposes that the PPN can be viewed as a “core correlate” of consciousness. (Dehaene, 2014) defines consciousness as “global information sharing” where information has entered into a specific storage area that makes it available to the rest of the brain. Koch et al. (2016) argue that the posterior hot zone (PHZ) is the minimal neural substrate essential for conscious perception. The PHZ includes cortical sensory areas in the parietal, temporal, and occipital lobes.

The interactions between the index layer, the representation layer, and the dynamic context layer (which might involve working memory), might be a foundation from which evolution eventually generated human consciousness. The representation layer, associated with the parietal area of the brain, plays a major role in our approach, as well as its interaction with the dynamic context layer, which might include the working memory. The activation of the representation layer encodes the cognitive brain state. It integrates information and then makes this information available to the brain as a whole. In particular, the activation of the representation layer represents a holistic, integrative view of the agent’s cognitive state at that instance. See also our discussion in the remaining part of this section. The role of the memory system in consciousness is discussed in (Budson et al., 2022).

8.5 Serialization of Parallel Computing and the Central Bottleneck

Here are some examples of compositionality and serial processing in the BTN, and possibly in the mind. First, perception analyses a scene that evolves in time, which is a serial process. Second, an episode is composed of several events. Third, an event consists of

several bounding boxes, which are analysed serially. Fourth, many triples are generated serially to decode a scene semantically. Finally, a triple itself is represented as a sequence of indices (see Table 2).

On the other hand, any computation involving the representation layer—in particular, the mappings involving the DCNN—and the mappings between the index layer and the representation layer are highly parallel.

It is this mixture of parallel and sequential processing exhibited in our model that might also make up the brain’s operation. Consider, that in a first step, we map $scene_t$ to $\mathbf{f}(scene_t)$, and this representation can be analysed rapidly and might trigger action with a limited understanding of scene content. Only then is bounding box content analysed in detail, which might be supported by saccadic eye movements, and this is a slower serial process, permitting a transition to language, as discussed later in this section. Maybe not surprisingly, language makes thoughts explicit but at the same time might slow down thinking.

Sequential processing is also a core concept in the theory of a global workspace (Baars, 1997; Bor and Seth, 2012; Dehaene, 2014; Goyal et al., 2021): (Koch, 2014) discusses that Dehaene’s workspace has extremely limited capacity (“the central bottleneck”) and that the mind can be conscious of and pay attention to only one or a few items or events, although these might be quickly varying. In cognition neuroscience, the general understanding is that parallel multitasking of cognitive tasks likely is an illusion. Even working memory is assumed to be able to store only three to four items at a time (Awh and Vogel, 2020). Sequential processing would also contribute to a potential solution of the binding problem (Singer, 2001) since the decoding focuses on concepts in a serial fashion, and associations between activities in the representation layer and the index layer are well defined. The importance of sampling is also recognized in (Dehaene, 2014) in the context of conscious perception. For example, the author states that “consciousness is a slow sampler.”

Another interesting point is that both (Dehaene, 2014) and (Koch et al., 2016) assume mental states, well delineated from all the other states. Switching between different interpretations is also a property of our sampling approach if one interprets a sample as a temporary decision or an interpretation: “It’s a bird, or a plane, or it’s Superman, but not all of them at the same time” (Dehaene, 2014). Dehaene discusses a similar process of “collapsing all unconscious probabilities into a single conscious sample.” His model assumes a “winning neural coalition”; in our approach, the embedding of an episodic instance integrates all information available at that instance, and the winning neural coalition would be the samples and triple sentences associated with that embedding. The optimization of the embedding of an episode, that is, \mathbf{a}_t , is a step that occurs after the decoding of a scene, and thus the conscious experience associated with it might be slightly delayed (see also Section 10.3). A current hypothesis in cognitive neuroscience is that conscious awareness often follows after an observation or a decision is made and serves to explain and justify but not to trigger an action (Gazzaniga et al., 2013; Budson et al., 2022). Only some important decisions are actually made with the help of conscious slow thinking. In the semantic decoding of our model, the representation layer is periodically activated, which might be reflected in neural signals and could be related to some of the neural oscillations found in the brain. A candidate is the beta rhythm (13-35 Hz), considered to be related to consciousness, perception, and motor behavior. Also of relevance might be the gamma

wave (25-140 Hz), which is correlated with large-scale brain network activity and cognitive phenomena such as working memory, attention, and perceptual grouping.

8.6 One-brain Hypothesis

We discuss the one-brain hypothesis in the context of two aspects. First, the representation layer can contain a superposition of embedding vectors and activation vectors, but it cannot represent a concatenation of two vectors. This would, in a way, require two separate representation layers. We propose that this is biologically not plausible, and we adhere to Dehaene’s concept of a central bottleneck. Applied to perception, this means that the mind can only process one bounding box content at a time.

Second, perception, episodic memory, semantic memory, and even embedded symbolic reasoning are realized by the different operational modes of a *single architecture*; thus, we do not assume different modules for these different functions. Instead, the brain repurposes the same architecture for different functions.

Of course, the brain does many things concurrently, and many of these activities are hidden from conscious awareness. This does not contradict our one-brain hypothesis, which is only concerned with the proposed operations of perception and memory, and in particular, when they involve the representation layer, that is, the cognitive state of the brain.

8.7 Cognitive Linguistics

Humans differ from other animals in their ability to express themselves through natural language. Human language is the basis for communication but also a means to argue and reason. Thus, an agent can tell another agent not to leave the hideout since a bear is lurking outside, even when the bear is not visible. Perception, episodic memory, and semantic memory are all declarative (i.e., explicit), and humans can verbally report on either. We propose that the generated triples in our approach are a basis from which the rich human language might have evolved.

Leading approaches in cognition and linguistics are, first, the formal approach, second, the connectionist approach, and third, the embodied approach (Evans, 2012). Our work relates to all three.

Let’s consider first the formal approach. The *language of thought hypothesis* assumes that mental representation has a linguistic structure, as well: thoughts are sentences in the mind. Fodor (1975) describes the nature of thought as possessing “language-like” or compositional structure (sometimes referred to as *mentalese*). In this view, simple concepts combine systematically (akin to the rules of grammar in language) to build thoughts. Also, in our approach, the brain talks to itself by producing triple sentences and their embeddings. We also agree that language offers a window into the operation of the brain. However, in our approach, we do not follow Fodor’s formal logic-based view: Our brain is more of a chatterbox.

Second, we follow a *connectionist approach* (McClelland et al., 2020), since the BTN includes a DCNN and neural processing. However, our notions of discrete symbolic indices, embodiment, and compositionality go beyond a more standard connectionist approach.

Finally, we can relate to the concept of an embodied language. There is general agreement that in a form of bottom-up processing, perception, the state of mind, and the body

influence language. The idea of an embodied language is that in a form of top-down processing, language can influence the perceptual path and the body (Evans, 2012). As discussed, we see our approach as being embodied, at least as far it concerns earlier processing in the brain. One might even consider that language takes over the mind, particularly in memory, with no immediate visual input to be decoded. From a brain, talking to itself, of course, there is a small step to a brain that speaks to others. Most generated triple sentences are not transferred into language; what is actually spoken is obviously more complex, nuanced, and sophisticated, and is modulated, for example, by intent, social context, and cultural background. Thus an internal triple-oriented fast speech is transferred into an external, sophisticated slow speech. When individuals learn a new language, this might involve the mappings between fast speech and slow speech, in addition to the direct mapping from one language to another language. The inner fast speech might already exist in some animals.

If the agent is the recipient of language, say, in a conversation, by reading, or by consuming media, the acquired information could be stored as an episodic memory instance, that is, as an episodic index with its associated embedding vector. This knowledge then can become part of semantic memory. In our simplistic triple-oriented conversation, the agent can learn about facts, for example, (*Sparky, hA, Dog*) and generalized statements, such as Hearst patterns (Hearst, 1992) like (*Dog, hA, Mammal*).

Language compresses information. Significantly different scenes might generate very similar descriptions, demonstrating great invariance in triple descriptions and language in general.

9. Episodic Memory

9.1 Background on Episodic Memory

Episodic memory documents the life of an agent. Tulving (1985) describes episodic memory as a memory that, in contrast to semantic memory, requires a recollection of a prior experience. It is considered to be the result of rapid associative learning in that a single episode and its context become associated and bound together and can be retrieved from memory after a single episode. Episodic memory stores information of general and personal events (Tulving, 1972, 1985, 2002; Gazzaniga et al., 2013) and concerns information we “remember,” including the spatiotemporal context of events (Gluck et al., 2013).

Some theories emphasize the sequential nature of episodic memory and the memory process. Moscovitch et al. (2016) considers an episodic memory experience to be an active process that involves details of the event and its location. Sometimes the reconstruction is regarded as a Bayesian process of reconstructing the past as accurately as possible based on available engram information (Hemmer and Steyvers, 2009).

9.2 Event Memory and Episodic Memory

In this article, we use the terms event memory and episodic memory almost synonymously. In the narrower sense, we consider event memory to be related to a single instance in time and episodic memory to a sequence of events that is, a story. This is in agreement with Mannila et al. (1997), who define an episode as a collection of events that occur relatively close to each other in a given partial order. We propose that the distinction



Figure 8: t-SNE visualization of episodic instances based on their embeddings $\{\mathbf{a}_t\}_{t=1}^{N_T}$. One can see that similar scenes are often in proximity. For example, the circled areas show images of indoor scenes (bottom) and of buses on streets (right).

between event memory and episodic memory is blurred. First, we often remember rather static images of past episodes. Second, even an analysis of a static scene is a sequential process, where each bounding box is recovered in sequence. Third, during a single saccade, the scene might have already changed: Consider, for example, an agent riding a bicycle where the scenery is constantly changing. So the recovery of a single scene might already describe a dynamic event.

9.3 Recall of Episodic Memory Engrams

An activated past episodic instance t restores its embedding \mathbf{a}_t , that is, its engram, in the representation layer. Figure 8 shows that the embedding vectors of episodic memories form meaningful maps and also are organized as a conceptual space.

Model	s^*		Unary labels						Binary labels	
	@10	@1	B-Class	P-Class	G-Class	Y/O	Color	Act.	@10	@1
EM	99.84	37.13	100.0	100.0	100.0	100.0	100.0	100.0	100.0	90.40
P-noI	0.0	0.0	41.19	20.83	51.76	49.03	12.97	71.05	39.34	7.79

Table 12: Top row: episodic memory experience using VRD-E data. For randomly selected past episodic instances t^* as input, we determine the highest-ranked entity indices (first two columns). For the other columns, we set the correct entity indices (s^* and o^*) and predict unary labels and binary labels. The performance is, in many cases, 100%. Thus, an agent might recall seeing a black dog, but it might be less confident that it was *Sparky*. Considering the large number of entities in the data set, the performance on entity prediction (first two columns) is still impressive as well. Bottom row: In P-noI we removed all entity indices, using only class and unary labels. The bad performance demonstrates the relevance of representing entities for memory recall.

9.4 Episodic Memory Experience: Semantic Decoding and Embodiment

The embedding of an episodic index is all there is, but the treasures of an episodic memory are only deciphered by the semantic decoding of that embedding, that is, the symbolic reconstruction of triples involving that episode (see Table 2). The recall of the episodic embedding, but especially the semantic decoding into triples, has the character of a simulation: Episodic memory is a reactivation of a possibly multimodal memory experience (Evans, 2012) and has been described as a reliving of past experience (Shapiro, 2010). Figure 9 illustrates an episodic memory experience. Table 12 provides numerical results.

9.5 Episodic and Semantic Memory in Perception

We propose that in perception, an agent first uses both the episodic and semantic attention approximation, which can be executed fast and in parallel. Only in a second step are specific entities sampled. This permits the integration of specific multimodal background (see Section 7.8), and permits associations with past episodic instances and known entities.

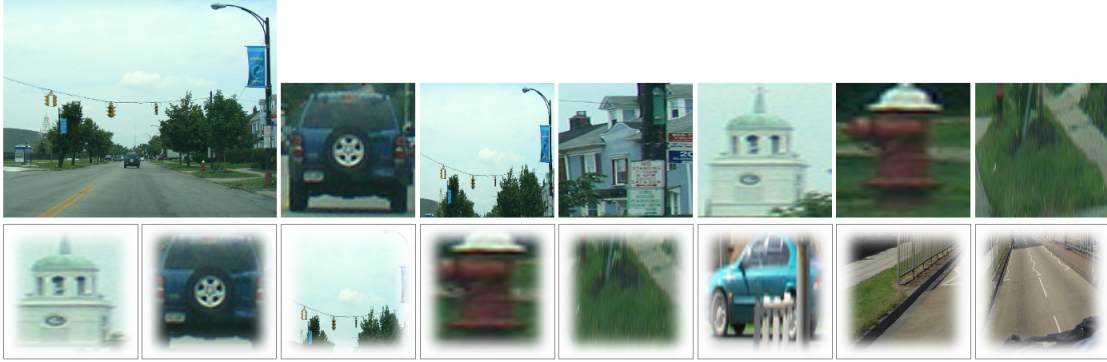


Figure 9: Episodic memory experience using VRD-E data. The first row shows the visual input to perception and then extracted bounding boxes for entities in the scene. In the second row, we show an episodic recall of this scene ($t^* = t'$), which is now a memory. We show the bounding boxes of the scene entities, which were recalled. We see that, in sampling, episodic recall recovers entities with correct bounding box content (e.g., fire hydrant, sky, car). Columns 6 to 8 show incorrect recalls. Note that these incorrect recalls would still be entirely plausible in the scene’s context.


Thus if perception poses the hypothesis that an entity in the image is identical to Sparky, then the semantic memory experience would supplement specific background information on Sparky, which cannot be derived from the visual input. Semantic memory reconstructs what is known about the concept (i.e., the prior information). Similarly, episodic memory would recall relevant past episodes.

9.6 Recent Episodic Memories for Context

Recent episodic memory can provide the agent with information on recent perceptual experiences. A recall is triggered by the recency of the episodic instance and relevance, and it contributes to an agent’s sense of the world state. A recent episodic memory experience is an episodic memory that is almost treated as a current observation. An agent needs to know about the state of the world, even for parts that are not currently being perceived. An example is the “lurking-bear” situation (see Figure 10): “There was a bear strolling around outside the hideout, . . . Remember: it might still be there, although it might not be visible from the hideout.” Figure 11 shows another example of a recent episodic memory recall. This emphasis on episodic closeness can be implemented as “time encoding” (Ma et al., 2018b), in a similar way as “position encoding” is used in the attention literature (Vaswani et al., 2017). Patients who are unable to form new episodic memory show great deficits in personal orientation and context understanding. These deficits are often associated with severe bilateral damage to MTL (Gluck et al., 2013).



Figure 10: Recent episodic memory experience: An illustration of the effect of a recent episodic memory experience using VRD-E data. The left image shows a harmless garden scene, but due to a recall of a recent episodic memory $t^* = 684$, the agent is aware of the lurking bear close by (right scene). Labels for visual entity recovered in the episodic recall (right scene) are *Bear*, *Mammal*, *LivingBeing*, *Old*, *Black*, *OtherActivity*, *Dangerous*. Note that episodic recall is not triggered by closeness in a scene but by recency and relevance.



s^*	Unary labels	o^*	Binary label
346	Person, 1.00, Mammal, 1.00, LivingBeing, 1.00 Young, 1.00, Other, 1.00, Other 1.00	345	on
348	Shirt, 1.00, Clothing, 1.00, NonLivingBeing, 1.00 Young, 1.00, Orange, 1.00, Other, 1.00	347	on
347	Person, 1.00, Mammal, 1.00, LivingBeing, 1.00 Young, 1.00, Other, 1.00, Playing, 1.00	348	wear
345	Grass, 1.00, Plant, 1.00 LivingBeing, 1.00 Old, 1.00, Green, 1.00, Other, 1.00	346	under

Figure 11: Recent episodic memory experience. (Top) The left image shows visual input to perception from VRD-EX data. Then a recent t^* is sampled. The right image shows the image belonging to t^* . (Bottom) The table shows results from this episodic memory experience without the actual image for t^* being available, only based on memory recall. The first column shows samples for s^* . The second column shows unary labels. The third column shows sampled objects o^* . The fourth column shows the most likely binary label.

9.7 Remote Episodic Memory for Decision Support

Remote episodic memory can provide the agent with information on remote perceptual experiences similar to the current perceptual experience, and this can contribute to decision-making. An episodic recall is triggered by $t = t^*$ (a sample in Equation 17) which reflects the similarity between the current scene representation with the episodic embedding. It makes sense that the current event should trigger the same action as in the retrieved episodic memory—if it led to a good outcome—or an alternative action if not. For example, if the agent finds the current situation very similar to a previous one, where it had walked toward a bear and almost got attacked, it would very likely not do this again! Remote episodic memory provides information about possible future scenarios and aids the agent in decision-making.

Thus memory guides behavior. Duncan and Shohamy (2016) describe this process as integration across relational events by imagining possible rewards in the future. The value associated with a memory (e.g., reward, threat) might be an integral aspect of episodic memory. The article also states that there is now extensive empirical data supporting the prevalent use of episodic memory across various decision-making tasks in humans.

Figure 12 shows that a perceptual scene indeed can activate memories of similar scenes in remote episodic memory.



Figure 12: (Top) Remote episodic memory experience using VRD-E data. The two images in the first column show the visual input to perception. Then we sample t^* , representing past episodic memories. The images of the scenes associated with the t^* show that, indeed, recalled past episodic memories are related. (Bottom) The bounding boxes in the first column represent visual entities s' in perception. The right bounding boxes display retrieved entities with high activation.

9.8 Forecasting and Future Episodic Memories

To forecast the state at the next instance t' , the agent might consider recent episodic memories to be current, that is, the instance index for those episodic memories is virtually changed to t' . The post-observation model then combines these memories with current observations and with the semantic memory to estimate the world state.

A future episodic memory is a forecasted event in the future, which at some point in time is predicted to become a regular episodic memory. This helps plan future actions. Events might be predicted using, for example, temporal knowledge graph models (Han et al., 2020a), assuming temporal smoothness and predictability of the embedding of time instances. Here we consider forecasts made entirely on a symbolic level, say, by verbal communication or reading. For example, if the agent has learned by verbal, symbolic communication that there is a football match in the stadium this evening and the weather will be bad, embedded reasoning, based on this assumed future information, predicts that there will be a traffic jam in the city and that driving will be difficult. Technically, the BTN for future episodic memories is identical to regular episodic memories, only that the episodic index embedding is calculated based on assumed events in the future, either without or with imagined and simulated future sensory input.

9.9 Memory Supports the Agent in the Present and the Future

It is of great interest for an agent to estimate the state of the world at the current instance t' , and to predict how it will evolve. We propose that the mind estimates the world’s state using perception and memory: The memory systems provide information that makes the agent act right but is not communicated by current perceptual experience.

For example, to illustrate the relevance of memory in daily life, consider a typical day at the office of agent Mary (see Figure 13). When Mary arrives at the office in the morning, she expects that everything is as usual, as modelled by *semantic memory*. The pre-observation model (semantic memory) will produce statements describing the state of normality. This sets the stage.

Perception will produce triples describing the actual situation, for example, the status of the coffee machine, which is broken. Since the coffee machine has mostly been working, semantic memory by itself would have predicted (incorrectly) that the coffee machine most likely is working. She immediately informs Jack in an office nearby: “Jack, can you believe this?! The coffee machine is broken, again!” (*from triples to language*). *Remote Episodic memory* reminds her that this is not the first time that the coffee machine was broken. *Recent episodic memory* will remind Mary all day that the coffee machine is broken, even when she is not in the same room as the coffee machine and thus does not have immediate perceptual information on its status. As discussed, recent episodic memory is an episodic memory that is almost treated as a current observation. The state change of the coffee machine is only slowly integrated into semantic memory. Being a long-term average, semantic memory is a sluggish state estimator and relies on recent episodic memory to emphasize recent state changes.

Jack might ask Mary if the coffee machine had been working last Tuesday. If she does not have an episodic recall from last Tuesday about the status of the coffee machine, she might consult semantic memory, reminding her that the coffee machine mostly has been

working. But if she has an episodic memory that last Monday, the coffee machine was broken, she might infer that it most likely was also broken on Tuesday. Here, she consults a memory that is recent to the past time instance of interest.

Another *recent episodic memory* might remind Mary that she had met Jane on her way to work and that she had said that she would drop by at the office (*simulation of a recent episodic memory*). Jane is a good friend (*semantic memory*).

Mary recognizes that Sparky is in the office (*perception*), and *semantic memory* adds triples describing background on Sparky, for example, that Sparky is Jack's dog. Mary might also recall that a while ago, Sparky was at the office and behaved well (*remote episodic memory*), although dogs in general can be quite a nuisance (as concluded by a *generalized statement* and by employing embedded symbolic reasoning).

When Mary relaxes, she suddenly recalls that people are repairing the heater today (*episodic memory, not triggered by perception*) and that she should call her roommate to check if they arrived in time.

Then Mary remembers that the football match will be in the stadium this evening and that the weather has been predicted to be bad (*future episodic event*): there will likely be a traffic jam (*generalization of future episodic events*), and, due to expected bad weather conditions, she should better drive carefully (*an association stored in semantic memory*). Future-state prediction uses both episodic memory and generalized future episodic memory.

The story illustrates that background is provided by semantic memory. Semantic memory is mostly relevant as it concerns triple statements, whose states are static or change rather slowly. Recent episodic memory provides information that is not yet or will not be absorbed into semantic memory but contributes to the current state estimate. Episodic memory is mostly relevant in as much as it concerns state changes (broken coffee machine), singular events (meeting Jane on her way to work) or information, not yet known. As our story indicates, narratives often tell causal stories and might also reflect the events and facts the brain pays attention to for decision-making. Future episodic memory supports planning. Remote episodic memory guides behavior: since the last time Sparky behaved well, he might behave well again this time.

The whole story is on a symbolic level. But each index that fires is embedded and leads to embodiment. We would propose that the storyline involves System-1 type associative reasoning, supported possibly by some background knowledge of how the world works. As our story suggests: to deal with daily life, agents might need little effortful System-2 reasoning.

10. To Perceive is to Learn

There is nothing in the mind that was not first in the senses. —John Locke

10.1 Overview

In normal operations, a new perceptual event that becomes an episodic memory requires the establishment of a new episodic index and its associated embedding vector. This roughly corresponds to the MTL-based fast *nonparametric learning system* in the complementary




Figure 13: The horizontal axis stands for time and the vertical for some abstract triple-dimension. We show past, present, and future. The current instance is t' on the right. The light blue background box stands for the predictions of the pre-observation model of semantic memory. The observation model of the episodic memory is represented by the horizontal gray bars; they indicate where, in the past, information was acquired. At episodic instance t' , the agent learns about some statements by perception (orange), and these become part of the episodic memory (recent episodic memory). There might be an association with another portion of the episodic memory, that is, the remote episodic memory (red dotted arrow). If the agent contemplates information unrelated to current perception, it can query on remote episodic memory (blue dotted lines). The right vertical bar indicates a future episodic memory.

learning systems (CLS) theory (McClelland et al., 1995; Kumaran et al., 2016). A new perceptual event might contain novel entities not yet known to the agent. For important ones, the agent needs to establish new indices and their embedding vectors as well. More rarely, new indices for attributes, classes, and predicates need to be established. Being quite stable, those might be represented in neocortex.

The second component of the CLS-theory is the *parametric learning system*, where the neocortex is trained by replay in a slow process from data from the nonparametric learning system. Slow training serves as the basis for the gradual acquisition of structured knowledge about the environment and its transfer to neocortex (Kumaran et al., 2016). An agent might get new information on existing entities, and this information might need to be integrated into their embeddings. In some applications involving online adaptation, interference of new knowledge with old ones can be a problem and might lead to catastrophic forgetting (Kumaran et al., 2016). Catastrophic forgetting does not show up in our preliminary experiments. The large dimensionality and modularity of the embedding vector lead to

robustness and stability; for example, learning first about Jack’s size and hair color might not interfere with the learning of his social network later. And when he later dyes his hair, the changes in his embedding vector are local, that is, it affects only a few dimensions. Learning happens as an online process, so more recent episodes will have a greater effect than more distant ones.

10.2 Neuroscience Perspective on Establishing Episodic Memory Engrams

Episodic memories are first formed in the hippocampus, which is part of medial temporal lobe (MTL). The idea that episodic memory is index based is by now one of several accepted theories (Tonegawa et al., 2018). It goes back to the hippocampal memory indexing theory (Teyler and DiScenna, 1986; Teyler and Rudy, 2007), which was long controversial. The event indices have a relational memory function in that they bind together different pieces of experience. Recent research found evidence for the existence of time cells in the hippocampus (CA1) (Eichenbaum et al., 2012; Eichenbaum, 2014; Kitamura et al., 2015b,a).

There is some evidence that neurogenesis might be involved in forming new episodic memories. Neurogenesis by special stem cells has been discovered in the dentate gyrus (part of the hippocampal formation) and is active throughout adult life; these new neurons may be preferentially recruited in the formation of memories. In fact, it has been observed that the adult macaque monkey forms a few thousand new neurons daily (Gluck et al., 2013; Gould et al., 1999), possibly to encode new information (Becker, 2005).

The establishment of new indices, together with their embedding vectors, are some of the most demanding learning tasks in the brain. Functionally, our model assumes that a new episodic memory engram is quickly stored by establishing a new index and its connections to the representation layer, copying the episodic memory trace. Although there exist several theories, little is known about how exactly new time indices are formed in the brain anatomically and how they quickly set up the bidirectional connection patterns with the representation layer, forming a hippocampal–cortical network (Frankland and Bontempi, 2005). Here, the brain might employ already existing networks, which might also explain how a single index can influence the distributed representation layer with potentially numerous neurons. A structural intra-hub connectivity might facilitate this process.

10.3 Experiments with Self-supervised Learning

Self-supervised learning (SSL) concerns learning without teacher-provided explicit training labels. It is biologically highly relevant. During most of life, an individual has to learn without explicit supervision (lifelong learning). In our approach, self-supervised learning works exactly like supervised training; the difference is that the agent’s predicted winner-take-all unary and binary labels become the training labels. This form of self-supervised learning is a type of bootstrap learning; see the bootstrap Widrow-Hoff rule (Hinton and Nowlan, 1990) and learning with pseudo-labels (Lee et al., 2013). The generated data is then used to train our model using cross-entropy cost-function terms derived from perception, episodic, and semantic memory (cf. Appendix E).

In our SSL experiments (see Table 13), we establish new indices and their embeddings for perceptual episodes (episodic indices) and for new entities in those episodes (concept indices); this would be part of the fast nonparametric learning system. The self-supervised

training for a novel episodic instance and a novel entity⁶ is modular and fast, both technically and most likely also biologically. Thus, we do not assume domain closure and consider that new perceptual experiences might require representations for new entities. The embedding vector of a new episodic instance $\mathbf{a}_{t'}$ is adapted to model the decoded triple statements. Decoded triple statements, which are more certain, will be decoded more often, will be more influential in memory formation, and will then be retrieved more likely in an episodic memory recall. Row 3 in the table shows that the semantic recall on entities is reasonable, but not as good as on entities trained with annotator-labeled data (see row 1). This is understandable since predicted labels are noisier than human-annotator-provided training labels.

Now consider the slow parametric learning system. Row 2 shows that the embeddings of already established entities are not negatively affected by SSL. Rows 4 and 5 show perception performance on unseen entities. We see better performance after SSL has been applied. This shows that embeddings for classes and attributes are improved by SSL.

10.4 Replay for Episodic Memory Consolidation

As self-supervised learning adds information, the required memory capacity grows. In our approach, for each salient episode, a new episodic index with its embedding vector is established. Similarly, for each new entity, a new concept index with its embedding vector is added. The current thinking in neuroscience is that this capacity problem can be solved by a consolidation from MTL, with a limited capacity, to the neocortex, with an essentially unlimited capacity. This is called systems consolidation of memory (SCM).

We now consider the leading two theories about the consolidation of episodic memory. The *standard theory* assumes that at some point, episodic memory becomes independent of the hippocampus and MTL over a period of weeks to years (Squire and Alvarez, 1995; Frankland and Bontempi, 2005). In contrast, the *multiple trace theory* assumes that both the hippocampus and MTL remain involved (Nadel and Moscovitch, 1997; Jonides et al., 2008; Greenberg and Verfaellie, 2010); MTL remains the manager of complex spatial and relational memories (Whittington et al., 2020). In general, it is assumed that consolidation involves both the medial prefrontal cortex (mPFC) and the MTL. The temporal lobe might be where the transferred indices are established (Frankland and Bontempi, 2005; Tonegawa et al., 2018). After consolidation, episodic memories might be organized in temporal order or according to a similarity in representation (see Figure 8). It is assumed that consolidation might be a process executed entirely or partially during sleep (Stickgold, 2005).

We propose consolidation by replay, which, in our model, can be executed as follows. An episodic instance t in MTL is activated, which activates the representation layer with vector \mathbf{a}_t ; this activation is then learned in the connection weights of a newly formed index in the neocortex, for example, by a form of Hebbian learning. Thus, these index duplicates in the neocortex would inherit the connection weights. If the index in the neocortex becomes more distributed, this will lead to greater robustness of memories after consolidation. For a while, both representations exist in parallel, and this facilitates the learning and consolidation of new memories; gradually, the index representation in the neocortex might become

6. The agent might decide that an entity is novel if the activation of all elements in $\text{sig}(\mathbf{n}_S)$ are below a threshold for all known entities.

dominant. Consolidation by replay has the advantage that there is no need for direct interactions of indices in both storage sites, only indirect interactions by a shared activation of the representation layer. Replay might also be one way of how the brain implements large-scale structural changes in the brain, in general, for example, as a consequence of brain damage or as a consequence of a changing world with new statistics. Following the principle “use it or lose it,” replay might also be essential such that relevant information in consolidated memory is not forgotten. From our model’s perspective, the location of the index is irrelevant. It could be in MTL, in neocortex, or both at the same time. In either case, decoding relies on the same machinery (see Figure 2). Cognitive maps might be more pronounced after consolidation in the neocortex (Binder and Desai, 2011). From a cognitive perspective, it makes sense that the brain mainly consolidates unusual episodic instances involving semantic state changes or singular events. It is tempting to assume that recent episodic memory is represented in MTL, whereas remote episodic memory is stored in neocortex.

10.5 Replay for Semantic Memory Consolidation

Training by episodic memory replay might also be important for the gradual transition from episodic to semantic memory, in which episodic memory reduces its sensitivity and association to particular episodes so that the information can be generalized as semantic memory. It might enable the semantic memory to adapt more quickly to state changes, for example, to a friend’s status change from being single to being married. Although this forgetting of the “single” state might require time, the label “married” might quickly show higher activity than “single”; the mind might consult episodic memory to resolve the conflict, discovering that the label “married” was assigned to recent instances. Semantic memory has a built-in forgetting mechanism since Equation 4 shows that if an observation of a triple is not renewed, dividing by N_{total} implies that it will get less significant over time. Some theories speculate that episodic memory may be the gateway to semantic memory (Baddeley and Hitch, 1974; Squire, 1987; Baddeley, 1988; Steyvers et al., 2004; Socher et al., 2009; McClelland et al., 1995; Yee et al., 2014; Kumar et al., 2015).

10.6 Forgetting and Modularity

Our hypothesis is that for a new instance or a new entity, an index is established in MTL. Let’s focus the discussion on a past episodic instance. If an episodic index is never consolidated into neocortex, the instance might eventually be forgotten since the resources in MTL are limited, and the brain might need to reuse them. On the other hand, if an index is consolidated into neocortex, it might be there for the lifetime of the agent. In neuroscience, it is an open question if consolidated memories really ever are forgotten, or if only the access to those memories is lost. The proposed index representation is essential since a particular functional synapse between an index and a dimension in the representation layer only serves this one purpose of connecting the index with the corresponding dimension in the representation layer. It is not reused for any other purpose. Assuming that an index and its embedding vector are stable, there might still be an “aging” of a memory. One reason is that the meaning of a dimension in the representation layer might change over the lifetime of an agent. In other words, the grounding or embedding might later in life be

Training Regime	Test Set	Unary labels (accuracy)						
		B-Class	P-Class	G-Class	Y/O	Color	Activity	Average
SL	SL-D	100.0	100.0	100.0	100.0	100.0	100.0	100.0
+SSL	SL-D	100.0	100.0	100.0	100.0	100.0	99.99	100.0
+SSL	SSL-D	85.51	88.36	93.16	48.32	59.42	79.78	75.76
SL	G	77.43	85.79	92.63	49.08	62.35	81.39	74.78
+SSL	G	77.54	86.15	92.83	48.58	62.62	82.78	75.08

Table 13: Self-supervised learning on the VRD-E data. SL stands for supervised learning, SSL for self-supervised learning. We first trained the model (SL) in a supervised way with human-provided labels (SL-D) with only 50% of the training images (2000 images). Columns 3 to 9 show perfect semantic memory performance for unary labels (first row). Then we continued to train the model with only the other 50% of the training images (SSL-D) in the SSL-mode (+SSL). The second row shows that self-supervised learning does not lead to a deterioration of prediction performance on the entities in SL-D. The third row shows that the performance on SSL-D is quite good, but of course, not as good as on SL-D since predicted labels on this data are noisier than the human-provided training labels. Rows 4 and 5 show performance in perception on new entities (generalization, G). The performance of +SSL is better than SL, which shows that self-supervised learning also improves the detection of classes and attributes, helping to fine-tune their embeddings.

different from the time the memory trace was established. A second issue is that the representations of other indices change over the lifetime of an agent. Thus, if Jack was part of a 10-year-old episode, the embedding vector for the index “Jack” might have changed since then. Reactivation of a past memory in combination with self-supervised learning might stabilize a past episodic memory, associated with the danger of changing that memory. It is well known that the reactivation of a past memory does not just strengthen that memory but also might change it. Note that this stabilization mechanism also only works in the context of index representations.

11. Summary, Conclusions and Future Work

We have shown how perception, memory, reasoning, and foundations for language and consciousness can all be realized by different functional and operational modes of the oscillating interactions between an index layer and a representation layer in the BTN. Whereas the representation layer, as a global workspace, is prominent in the current discussion on consciousness, the introduced index layer is an original contribution. We have proposed that an index representation increases modularity and counteracts forgetting. We have emphasized the role of episodic and semantic memory in perception. We have proposed an associative memory where recency is key to the recall of recent episodic memory and similarity of episodic representations with the present scene representation for remote episodic memory. We argue that episodic memory models the observed data, and semantic memory provides

background knowledge. Our approach explains the great similarity between episodic and semantic memory: semantic memory is the expected episodic memory of a future instance.

We have proposed that perception and memories first produce subsymbolic representations, which are subsequently decoded semantically to produce symbolic triple sentences. Our paper contributed to the discussion (a recent example being (Browning and LeCun, 2022)) of how language, thought, and subsymbolic processing interact.

As part of future work, we will address more systematically subsumption hierarchies and how parts bind to form the whole, which is another form of compositionality addressed, for example, in capsule networks (Sabour et al., 2017). The agent itself is also an important concept: it is part of each event. It might be quite relevant if the angry bear looks at a deer or at the agent itself. Other issues of the “I” are mood, state of mind, current mission, social contexts, and spatial location.

In this article, we did not focus on spatial representations. It is well known that MTL is not only instrumental for forming novel episodic memories but also contains spatial representations, for example, in the form of grid cells and place cells. It is generally assumed that MTL is central for reconstructing spatial memories, permitting complex spatial reasoning. We plan to investigate if instance indices, their embeddings, and their labels might also be the basis for spatial navigation and reasoning.

The brain does many different things simultaneously and addresses a particular problem with several strategies. Our work emphasizes the role of perception and memory, and it might provide some insights into part of the amazing faculties of human intelligence.

Appendix A. Cost Function Terms

$\mathbb{P}(s, p, o, t)$ has four categorical variables. The number of states is $N_C^2 N_P N_T$. So there is one state for each Boolean variable.

OBSERVATION MODEL FOR EPISODIC MEMORY

If $i_{s,p,,o,t} = 1$, then the contribution to the cross-entropy cost function for binary labels is

$$-(\log \mathbb{P}(p|s, o, t) + \log \mathbb{P}(o|s, t) + \log \mathbb{P}(s|t))$$

and for unary labels $-(\log \mathbb{P}(c|s, t) + \log \mathbb{P}(s|t))$.

PRE-OBSERVATION MODEL FOR SEMANTIC MEMORY

If $i_{s,p,,o,t} = 1$, then the contribution to the cross-entropy cost function for binary labels is

$$-(\log \mathbb{P}(p|s, o, \bar{t}) + \log \mathbb{P}(o|s, \bar{t}))$$

and for unary labels $-\log \mathbb{P}(c|s, \bar{t})$.

PERCEPTION

In perception (see Section 4.6)

$$\mathbb{P}(s' = s, p' = p, o' = o|t' = t, BB_{sub}, BB_{obj}, BB_{pred}, scene_{t'}).$$

If $i_{s,p,o,t} = 1$, then the contribution to the cross-entropy cost function is

$$-\log \mathbb{P}(s' = s, p' = p, o' = o | t' = t, BB_{sub}, BB_{obj}, BB_{pred}, scene_{t'}).$$

Appendix B. Boolean versus Categorical Models

As discussed, the interface between the multinomial and the two-state Bernoulli model is via sampling. If the multinomial model generates an (s, p, o) or (s, c) sample, then the random variable $Y_{s,p,o,t}$ is assumed true. Here, we show a close mathematical relationship between the probabilistic scores. The key correspondence is that for entities that were observed in the episodic experience at t ,

$$\mathbb{P}(s, p, o | t) \approx \mathbb{E}(Y_{s,p,o,t}) / N_t.$$

From this, one can derive

$$\mathbb{P}(p | s, o, t) \approx \mathbb{E}(Y_{s,p,o,t}) / \sum_{p'} \mathbb{E}(Y_{s,p',o,t}) \quad (20)$$

$$\mathbb{P}(c | s, t) \approx \mathbb{E}(Y_{s,hA,c,t}) / \sum_{c'} \mathbb{E}(Y_{s,hA,c',t}) \quad (21)$$

For semantic memory, we have

$$\mathbb{P}(s, p, o | \bar{t}) \approx N_{s,p,o,\text{known}} \mathbb{E}(Y_{s,p,o,\bar{t}}) / N_{total}.$$

Thus, the samples from the pre-observation model reflect the expected-state model weighted by the number of measurements on the triple statement. Sometimes it is possible to apply knowledge on mutual exclusiveness. For example, if we define the complementary of c as \bar{c} , then

$$\frac{\mathbb{P}(s, hA, c, \bar{t})}{\mathbb{P}(s, p, hA, c, \bar{t}) + \mathbb{P}(s, hA, \bar{c}, \bar{t})} = \mathbb{E}(Y_{s,hA,c,\bar{t}}).$$

As an example, assume Sparky is observed to be barking in 10 episodic instances and not-barking in 1000 episodic instances. Then, $\mathbb{E}(Y_{Sparky,hA,Barking,\bar{t}}) \approx 1\%$, so Sparky is rarely barking. But since he is sometimes barking, this would be reflected in a firing of “Barking”, following $\mathbb{P}(Sparky, hA, Barking | \bar{t})$. The agent would have problems to understand if Sparky is often barking or not. In contrast, if the agent observes both the firing of “Barking” and “NotBarking”, the fraction of firing rates (see last equation) would tell the agent, that Sparky is mostly quiet.

Furthermore,

$$\mathbb{P}(p | s, o, \bar{t}) \approx \mathbb{E}(Y_{s,p,o,\bar{t}}) / \sum_{p'} \mathbb{E}(Y_{s,p',o,\bar{t}}) \quad (22)$$

$$\mathbb{P}(c | s, \bar{t}) \approx \mathbb{E}(Y_{s,hA,c,\bar{t}}) / \sum_{c'} \mathbb{E}(Y_{s,hA,c',\bar{t}}). \quad (23)$$

In the last two equations, we have exploited that due to the LCWA, $N_{s,p,o,\text{known}} = N_{s,p',o,\text{known}}$ and $N_{s,c,\text{known}} = N_{s,c',\text{known}}$.

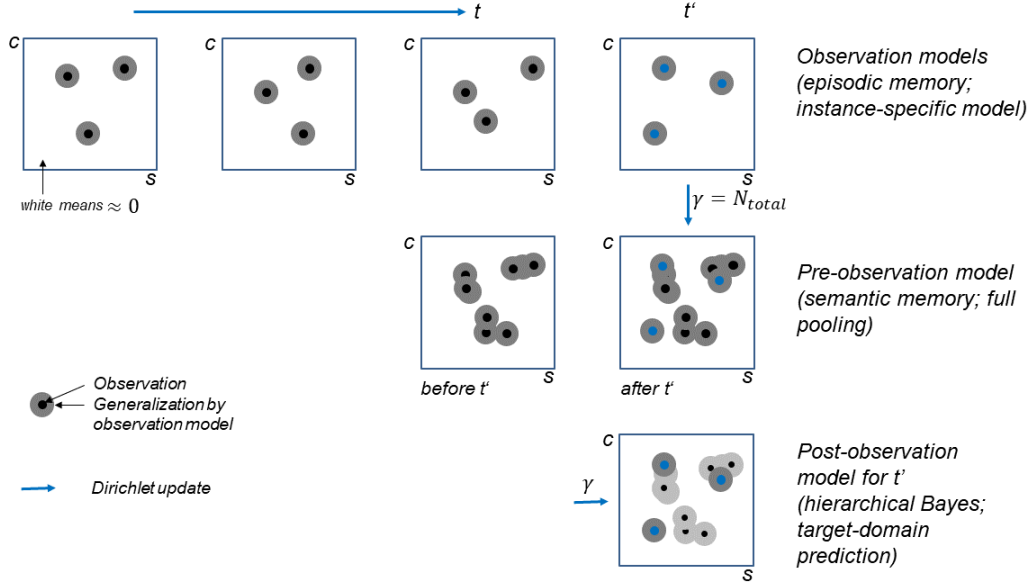


Figure 14: Visualization for categorical variables. Here, we show the situation for unary statements. For binary statements, replace c with (p, o) .

Appendix C. Visualisation

Figure 14 illustrates the situation for categorical variables. The update for the pre-observation and post-observation models are almost identical. The difference is that the observed data have a higher weight in the post-observation model. Figure 15 shows the situation for Boolean variables.

Appendix D. Social Network

We consider all 4987 persons in the data set and link a person s to persons s' , if the score $\mathbf{a}_s^\top \mathbf{a}_{s'}$ is in the top 5 of all scores related to s . Thus links exist between persons with similar embeddings, simulating homophily. We then determine the link direction. Considering two entities s and s' , $\exp \beta \|\mathbf{a}_s\| / (\|\mathbf{a}_s + \mathbf{a}_{s'}\|)$ is proportional to the probability that we determine that (s, \textit{knows}, s') , otherwise, (s', \textit{knows}, s) . At a social network episodic time step t , all links to one person s are added. This defines 4987 episodes for the tKG. The pKG then aggregates the tKG. Overall, we have 24953 *knows* statements. In summary, our social network data set has 4987 person entities (along with their unary labels), 24953 friendship statements, and 4987 episodes of social events.




Figure 15: Same, but for Boolean variables.

Appendix E. Implementation Details

In this section, we focus on the implementation aspects and provide some details about the network architecture and training hyperparameters. Our program is written in Python and utilizes PyTorch.

E.1 Network Architecture

Our model mainly consists of two layers. The representation layer \mathbf{q} has $r = 4096$ neurons. The index layer \mathbf{n} has $N_T + N_C + N_P$ neurons. We use VGG-19 backbone as the deep neural network $\mathbf{f}(\cdot)$ which takes as input a scene or a bounding box and outputs a 4096-dimensional feature vector. The VGG-19 network consists of a sequence of convolutional blocks followed by two fully connected hidden layers. Each convolutional block is a sequence of 2 convolution layers with 3×3 filters, a max-pooling layer, and another two convolution layers with the same parameters. We use the activations before the nonlinear transformation from the last hidden layer and copy them over to \mathbf{q} . Thus the layer \mathbf{q} stores pre-activation values in the range of \mathbb{R} . Next, we apply a nonlinear function (LeakyReLU) to the values in \mathbf{q} and feed them to the index layer via connection weights \mathbf{A}^\top . At different decoding steps, \mathbf{n} covers different sets of indices, namely \mathbf{n}_T has N_T units for episodic instances, \mathbf{n}_S , \mathbf{n}_C , \mathbf{n}_O , all have N_C concepts units (entities, classes, attributes), and \mathbf{n}_P has N_P units for predicates. The index layer in turn activates the representation layer via the same weights \mathbf{A} . To calculate the enhanced representation \mathbf{q}_T , \mathbf{q}_S , and \mathbf{q}_O , we add the pre-activation of layer \mathbf{q} and corresponding activations from \mathbf{n} . To obtain \mathbf{n}_S and \mathbf{n}_O , we apply softmax on the output with an inverse temperature $\beta = 1$. For \mathbf{n}_C we split the concepts into eight sets and apply softmax on each set. Alternatively, we could also use the sigmoid function

as stated in our algorithm. However, softmax fits here as our labels are mutually exclusive, and in practice, softmax leads to faster convergence and slightly better performance. The dynamic context layer \mathbf{h} contains 500 neurons with a self-connection via weight matrix \mathbf{B} . The input to the dynamic context layer is post-activation of \mathbf{q} , that is, values after a nonlinear transformation. There is a direct path for the hidden layer between different decoding steps (the dotted line between \mathbf{h} blocks in figure 2), which stores the state of the working memory and leads to a slight improvement (1%) in relationship prediction. For $\bar{\mathbf{a}}$, we use a learnable embedding vector of length 4096.

E.2 Training Scheme

We train a global set of parameters for our BTN with an objective of minimizing the multi-task loss for perception, episodic, and semantic memory experience. We set the batch size to 128, the learning rate to 0.0001, and dropout $p=0.5$. The BTN is optimized using an Adam optimizer for 60 epochs. During training, we freeze the VGG-19 layers except for its last fully connected layer so that the knowledge in the pre-trained model will not be destroyed. For the last layer of VGG-19, we use a smaller learning rate of 0.00001 to allow it to adapt to the new task. Except for the VGG-19 backbone, we initialize our network using Kaiming uniform initialization proposed in (He et al., 2015). We minimize the summed cross-entropy loss on \mathbf{n}_T , \mathbf{n}_S , \mathbf{n}_C , \mathbf{n}_O , and \mathbf{n}_P for the perception and memory experience. For \mathbf{n}_S and \mathbf{n}_O , we randomly activate an index from the set of entity indices and class/attribute labels for learning generalized statements (cf. 3.8). For \mathbf{n}_C , we apply cross-entropy loss on each subset. For the social network, we train the model for attribute prediction and relationship prediction on the social network data set. Concretely, we minimize the cross-entropy loss for \mathbf{n}_S , \mathbf{n}_C , \mathbf{n}_O , and \mathbf{n}_P . For SSL, we use a batch size of 128, a learning rate of 1e-5, and total training epochs of 20. Except for the embeddings of time and entities, all weights are frozen. Regarding the ablation studies, we train the RESCAL model using the implementation of PyKeen (Ali et al., 2021). The rank of entity embeddings and predicate embeddings are both set to 1000, which gives a comparable number of learnable parameters as our model. All experiments are conducted on an Nvidia GTX 1089 Ti GPU with a 4 core CPU of 16G memory.

Code

The source code related to this article is available at the following link: <https://github.com/hangligit/BTN>

Acknowledgements

We thank the anonymous reviewers for their helpful comments on the manuscript.

References

Mehdi Ali, Max Berrendorf, Charles Tapley Hoyt, Laurent Vermue, Sahand Sharifzadeh, Volker Tresp, and Jens Lehmann. Pykeen 1.0: A python library for training and evalu-

- ating knowledge graph embeddings. *Journal of Machine Learning Research*, 22(82):1–6, 2021.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- E Awh and EK Vogel. Online and off-line memory states in the human brain. *the Cognitive Neurosciences*, 2020.
- Bernard J Baars. *In the theater of consciousness: The workspace of the mind*. Oxford University Press, USA, 1997.
- Alan Baddeley. Cognitive psychology and human memory. *Trends in neurosciences*, 11(4): 176–181, 1988.
- Alan Baddeley. Working memory. *Science*, 255(5044):556–559, 1992.
- Alan D Baddeley and Graham Hitch. Working memory. *The psychology of learning and motivation*, 8:47–89, 1974.
- Stephan Baier, Yunpu Ma, and Volker Tresp. Improving visual relationship detection using semantic modeling of scene descriptions. In *ISWC*. Springer, 2017.
- Stephan Baier, Yunpu Ma, and Volker Tresp. Improving information extraction from images with learned semantic models. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 5214–5218. AAAI Press, 2018.
- Lawrence W Barsalou et al. Grounded cognition. *Annual review of psychology*, 59(1): 617–645, 2008.
- Danielle S Bassett and Olaf Sporns. Network neuroscience. *Nature neuroscience*, 20(3):353, 2017.
- Suzanna Becker. A computational principle for hippocampal learning and neurogenesis. *Hippocampus*, 15(6):722–738, 2005.
- Jeffrey R Binder and Rutvik H Desai. The neurobiology of semantic memory. *Trends in cognitive sciences*, 15(11):527–536, 2011.
- Marcus D Bloice, Christof Stocker, and Andreas Holzinger. Augmentor: an image augmentation library for machine learning. *arXiv preprint arXiv:1708.04680*, 2017.
- Daniel Bor and Anil K Seth. Consciousness and the prefrontal parietal network: insights from attention, working memory, and chunking. *Frontiers in Psychology*, 3:63, 2012.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems 26*, 2013.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Jacob Browning and Yann LeCun. AI and the limits of language. *Noema Magazine*, 2022.
- Andrew E Budson, Kenneth A Richman, and Elizabeth A Kensinger. Consciousness as a memory system. *Cognitive and Behavioral Neurology*, 2022.
- Timothy J Buschman and Earl K Miller. How working memory works. *The Cognitive Neurosciences*, page 357, 2020.
- Peter Dayan, Geoffrey E Hinton, Radford M Neal, and Richard S Zemel. The Helmholtz machine. *Neural computation*, 7(5):889–904, 1995.
- Wendy A de Heer, Alexander G Huth, Thomas L Griffiths, Jack L Gallant, and Frédéric E Theunissen. The hierarchical cortical organization of human speech processing. *Journal of Neuroscience*, 37(27):6539–6557, 2017.
- Stanislas Dehaene. *Consciousness and the brain: Deciphering how the brain codes our thoughts*. Penguin, 2014.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Katherine D Duncan and Daphna Shohamy. Memory states influence value-based decisions. *Journal of Experimental Psychology: General*, 145(11):1420, 2016.
- Howard Eichenbaum. Time cells in the hippocampus. *Nature Reviews Neuroscience*, 15(11), 2014.
- Howard Eichenbaum, Magdalena Sauvage, Norbert Fortin, Robert Komorowski, and Paul Lipton. Towards a functional organization of episodic memory in the medial temporal lobe. *Neuroscience & Biobehavioral Reviews*, 36(7):1597–1608, 2012.
- Jonathan St BT Evans. In two minds: dual-process accounts of reasoning. *Trends in cognitive sciences*, 7(10):454–459, 2003.
- Vyvyan Evans. Cognitive linguistics. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(2):129–141, 2012.
- Christiane Fellbaum. Wordnet. In *Theory and applications of ontology: computer applications*, pages 231–243. Springer, 2010.
- Jerry A Fodor. *The language of thought*, volume 5. Harvard university press, 1975.
- Paul W Frankland and Bruno Bontempi. The organization of recent and remote memories. *Nature Reviews Neuroscience*, 6(2):119–130, 2005.

- Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- Karl J Friston, G Tononi, O Sporns, and GM Edelman. Characterising the complexity of neuronal interactions. *Human Brain Mapping*, 3(4):302–314, 1995.
- Peter Gärdenfors. *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press, 2016.
- Michael S Gazzaniga, Richard B Ivry, and George Ronald Mangun. *Cognitive Neuroscience: The biology of the mind*. New York: WW Norton, fourth edition edition, 2013.
- Mark A Gluck, Eduardo Mercado, and Catherine E Myers. *Learning and memory: From brain to behavior*. Palgrave, 2013.
- Elizabeth Gould, Alison J Reeves, Michael SA Graziano, and Charles G Gross. Neurogenesis in the neocortex of adult primates. *Science*, 286(5439):548–552, 1999.
- Anirudh Goyal, Aniket Didolkar, Alex Lamb, Kartikeya Badola, Nan Rosemary Ke, Nasim Rahaman, Jonathan Binas, Charles Blundell, Michael Mozer, and Yoshua Bengio. Co-ordination among neural modules through a shared global workspace. *arXiv preprint arXiv:2103.01197*, 2021.
- Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Daniel L Greenberg and Mieke Verfaellie. Interdependence of episodic and semantic memory: evidence from neuropsychology. *Journal of the International Neuropsychological society*, 16(05):748–753, 2010.
- Thomas L Griffiths, Charles Kemp, and Joshua B Tenenbaum. Bayesian models of cognition. In *The Cambridge Handbook of Computational Psychology*. Cambridge University Press, 2008.
- Wolfgang Hackbusch. *Tensor spaces and numerical tensor calculus*, volume 42. Springer, 2012.
- Graeme S. Halford, William H. Wilson, and Steven Phillips. Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences*, 21(06):803–831, 1998.
- Graeme S Halford, William H Wilson, Glenda Andrews, and Steven Phillips. *Categorizing cognition: Toward conceptual coherence in the foundations of psychology*. MIT Press, 2014.
- Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. Explainable subgraph reasoning for forecasting on temporal knowledge graphs. In *International Conference on Learning Representations*, 2020a.

- Zhen Han, Yunpu Ma, Yuyi Wang, Stephan Guennemann, and Volker Tresp. Graph hawkes neural network for forecasting on temporal knowledge graphs. In *Automated Knowledge Base Construction*, 2020b.
- Stevan Harnad. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42 (1-3):335–346, 1990.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Marti A Hearst. Automatic acquisition of hyponyms from large text corpora. In *COLING 1992 Volume 2: The 14th International Conference on Computational Linguistics*, 1992.
- Pernille Hemmer and Mark Steyvers. A Bayesian account of reconstructive memory. *Topics in Cognitive Science*, 1(1):189–202, 2009.
- Geoffrey E Hinton and Steven J Nowlan. The bootstrap widrow-hoff rule as a cluster-formation algorithm. *Neural Computation*, 2(3):355–362, 1990.
- Douglas L Hintzman. Minerva 2: A simulation model of human memory. *Behavior Research Methods, Instruments, & Computers*, 16(2):96–101, 1984.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Bernhard Hommel, Jochen Müsseler, Gisa Aschersleben, and Wolfgang Prinz. The theory of event coding (tec): A framework for perception and action planning. *Behavioral and brain sciences*, 24(5):849–878, 2001.
- Drew Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In *Advances in Neural Information Processing Systems*, pages 5901–5914, 2019.
- Michael S Humphreys, John D Bain, and Ray Pike. Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review*, 96 (2):208, 1989.
- Alexander G. Huth, Wendy A. de Heer, Thomas L. Griffiths, Frédéric E. Theunissen, and Jack L. Gallant. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 2016.
- Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015.
- John Jonides, Richard L Lewis, Derek Evan Nee, Cindy A Lustig, Marc G Berman, and Katherine Sledge Moore. The mind and brain of short-term memory. *Annual review of psychology*, 59:193, 2008.

- Alexander B. Jung, Kentaro Wada, Jon Crall, Satoshi Tanaka, Jake Graving, Christoph Reinders, Sarthak Yadav, Joy Banerjee, Gábor Vecsei, Adam Kraft, Zheng Rui, Jirka Borovec, Christian Vallentin, Semen Zhydenko, Kilian Pfeiffer, Ben Cook, Ismael Fernández, François-Michel De Rainville, Chi-Hung Weng, Abner Ayala-Acevedo, Raphael Meudec, and Matias Laporte. Augmentor: an image augmentation library for machine learning. <https://github.com/aleju/imgaug>, 2020. Online; accessed 01-Feb-2020.
- Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.
- Pentti Kanerva. *Sparse distributed memory*. MIT press, 1988.
- Markus Kiefer and Friedemann Pulvermüller. Conceptual representations in mind and brain: theoretical developments, current evidence and future directions. *cortex*, 48(7): 805–825, 2012.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Takashi Kitamura, Christopher J Macdonald, and Susumu Tonegawa. Entorhinal–hippocampal neuronal circuits bridge temporally discontinuous events. *Learning & memory (Cold Spring Harbor, NY)*, 22(9):438–443, 2015a.
- Takashi Kitamura, Chen Sun, Jared Martin, Lacey J Kitch, Mark J Schnitzer, and Susumu Tonegawa. Entorhinal cortical ocean cells encode specific contexts and drive context-specific fear memory. *Neuron*, 87(6):1317–1331, 2015b.
- David C Knill and Alexandre Pouget. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, 2004.
- Christof Koch. Keep it in mind. *Scientific American Mind*, 25(3):26–29, 2014.
- Christof Koch, Marcello Massimini, Melanie Boly, and Giulio Tononi. Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, 17(5):307, 2016.
- Rajat Koner, Suprosanna Shit, and Volker Tresp. Relation transformer network. *arXiv preprint arXiv:2004.06193*, 2020.
- Konrad P Körding, Shih-pi Ku, and Daniel M Wolpert. Bayesian integration in force estimation. *Journal of Neurophysiology*, 92(5):3161–3165, 2004.
- Nikolaus Kriegeskorte and Pamela K Douglas. Cognitive computational neuroscience. *Nature neuroscience*, 21(9):1148–1160, 2018.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, and David A Shamma. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017.

- Ankit Kumar, Ozan Irsoy, Jonathan Su, James Bradbury, Robert English, Brian Pierce, Peter Ondruska, Ishaan Gulrajani, and Richard Socher. Ask me anything: Dynamic memory networks for natural language processing. *arXiv preprint arXiv:1506.07285*, 2015.
- Dharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? *Trends in Cognitive Sciences*, 20(7):512–534, 2016.
- George Lakoff, Mark Johnson, and John F Sowa. Review of philosophy in the flesh: The embodied mind and its challenge to western thought. *Computational Linguistics*, 25(4): 631–634, 1999.
- Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013.
- David A Leopold, Peter L Strick, Danielle S Bassett, Randy M Bruno, Hermann Cuntz, Kristen M Harris, Marcel Oberlaender, and Marcus E Raichle. Functional architecture of the cerebral cortex. In *The Neocortex*, pages 141–164. MIT Press, 2019.
- Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*. Springer, 2016.
- Ruotian Luo, Ning Zhang, Bohyung Han, and Linjie Yang. Context-aware zero-shot recognition. *arXiv preprint arXiv:1904.09320*, 2019.
- Yunpu Ma, Marcel Hildebrandt, Stephan Baier, and Volker Tresp. Holistic representations for memorization and inference. *UAI*, 2018a.
- Yunpu Ma, Volker Tresp, and Erik A Daxberger. Embedding models for episodic knowledge graphs. *Journal of Web Semantics*, 2018b.
- Heikki Mannila, Hannu Toivonen, and A Inkeri Verkamo. Discovery of frequent episodes in event sequences. *Data mining and knowledge discovery*, 1(3):259–289, 1997.
- James L McClelland, Bruce L McNaughton, and Randall C O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419, 1995.
- James L McClelland, Felix Hill, Maja Rudolph, Jason Baldridge, and Hinrich Schütze. Placing language in an integrated understanding system: Next steps toward human-level performance in neural language models. *Proceedings of the National Academy of Sciences*, 117(42):25966–25974, 2020.
- Richard Montague. Universal grammar. *Theoria*, 36(3):373–398, 1970.
- Morris Moscovitch, Roberto Cabeza, Gordon Winocur, and Lynn Nadel. Episodic memory and beyond: the hippocampus and neocortex in transformation. *Annual review of psychology*, 67:105–134, 2016.

- Lynn Nadel and Morris Moscovitch. Memory consolidation, retrograde amnesia and the hippocampal complex. *Current opinion in neurobiology*, 7(2):217–227, 1997.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 809–816, 2011.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. Factorizing YAGO: scalable machine learning for linked data. In *Proceedings of the 21st International Conference on World Wide Web*, WWW '12, pages 271–280, 2012.
- Maximilian Nickel, Xueyan Jiang, and Volker Tresp. Reducing the Rank in Relational Factorization Models by Including Observable Patterns. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 1179–1187. Curran Associates, Inc., 2014.
- Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 2015a.
- Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. Holographic embeddings of knowledge graphs. *arXiv preprint arXiv:1510.04935*, 2015b.
- Natasha Noy, Alan Rector, Pat Hayes, and Chris Welty. Defining n-ary relations on the semantic web. *W3C working group note*, 12(4), 2006.
- Charles Kay Ogden and Ivor Armstrong Richards. *The meaning of meaning: A study of the influence of thought and of the science of symbolism*. Harcourt, Brace & World, Inc., 1923.
- Adam F Osth and Simon Dennis. Sources of interference in item and associative recognition memory. *Psychological review*, 122(2):260, 2015.
- Andras Pellionisz and Rodolfo Llinas. Tensorial approach to the geometry of brain function: Cerebellar coordination via a metric tensor. *Neuroscience*, 5(7):1125–1136, 1980.
- Tony Plate. A common framework for distributed representation schemes for compositional structure. *Connectionist systems for knowledge representation and deduction*, pages 15–34, 1997.
- Tomaso Poggio, Andrzej Banburski, and Qianli Liao. Theoretical issues in deep networks. *Proceedings of the National Academy of Sciences*, 117(48):30039–30045, 2020.
- Jordan B Pollack. Recursive distributed representations. *Artificial Intelligence*, 46(1):77–105, 1990.
- Maria Pope, Makoto Fukushima, Richard F Betzel, and Olaf Sporns. Modular origins of high-amplitude co-fluctuations in fine-scale functional connectivity dynamics. *Proceedings of the National Academy of Sciences*, 118(46):e2109380118, 2021.
- Margot Popp, Natalie M Trumpp, and Markus Kiefer. Processing of action and sound verbs in context: An fmri study. *Translational Neuroscience*, 10(1):200–222, 2019.

- R Quian Quiroga, Leila Reddy, Gabriel Kreiman, Christof Koch, and Itzhak Fried. Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107, 2005.
- Rodrigo Quian Quiroga. Concept cells: the building blocks of declarative memory functions. *Nat Rev Neurosci*, 13(8), 2012.
- Matthew A Lambon Ralph, Elizabeth Jefferies, Karalyn Patterson, and Timothy T Rogers. The neural and computational bases of semantic cognition. *Nature Reviews Neuroscience*, 18(1):42, 2017.
- Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- Edmund T Rolls. *Cerebral cortex: principles of operation*. Oxford University Press, 2016.
- Shirley-Ann Rueschemeyer, Daan van Rooij, Oliver Lindemann, Roel M Willems, and Harold Bekkering. The function of words: Distinct neural correlates for words denoting differently manipulable objects. *Journal of cognitive neuroscience*, 22(8):1844–1851, 2010.
- Daniel Ruffinelli, Samuel Broscheit, and Rainer Gemulla. You can teach an old dog new tricks! on training knowledge graph embeddings. In *International Conference on Learning Representations*, 2019.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, and Michael Bernstein. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. *Advances in neural information processing systems*, 30, 2017.
- Anil K Seth and Tim Bayne. Theories of consciousness. *Nature Reviews Neuroscience*, pages 1–14, 2022.
- Lawrence Shapiro. *Embodied cognition*. Routledge, 2010.
- Sahand Sharifzadeh, Max Berrendorf, and Volker Tresp. Improving visual relation detection using depth maps. *arXiv preprint arXiv:1905.00966*, 2019.
- Sahand Sharifzadeh, Sina Moayed Baharlou, and Volker Tresp. Classification by attention: Scene graph classification with prior knowledge, 2020.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

- Wolf Singer. Consciousness and the binding problem. *Annals of the New York Academy of Sciences*, 929(1):123–146, 2001.
- Paul Smolensky. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, 46(1):159–216, 1990. URL <http://www.sciencedirect.com/science/article/pii/000437029090007M>.
- Richard Socher, Samuel Gershman, Per Sederberg, Kenneth Norman, Adler J Perotte, and David M Blei. A Bayesian analysis of dynamics in free recall. In *Advances in neural information processing systems*, pages 1714–1722, 2009.
- Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. Reasoning With Neural Tensor Networks for Knowledge Base Completion. In *Advances in Neural Information Processing Systems 26*, 2013.
- Olaf Sporns. Graph theory methods: applications in brain networks. *Dialogues in clinical neuroscience*, 20(2):111, 2018.
- Larry R Squire. *Memory and brain*. Oxford UP, 1987.
- Larry R Squire and Pablo Alvarez. Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current opinion in neurobiology*, 5(2):169–177, 1995.
- Mark Steyvers, Richard M Shiffrin, and Douglas L Nelson. Word association spaces for predicting semantic similarity effects in episodic memory. *Experimental cognitive psychology and its applications: Festschrift in honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer*, pages 237–249, 2004.
- Robert Stickgold. Sleep-dependent memory consolidation. *Nature*, 437(7063):1272–1278, 2005.
- Sainbayar Sukhbaatar, Jason Weston, and Rob Fergus. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448, 2015.
- Ron Sun and Todd Peterson. Learning in reactive sequential decision tasks: The clarification model. In *Proceedings of International Conference on Neural Networks (ICNN'96)*, volume 2, pages 1073–1078. IEEE, 1996.
- Joshua B Tenenbaum, Thomas L Griffiths, and Charles Kemp. Theory-based Bayesian models of inductive learning and reasoning. *Trends in cognitive sciences*, 10(7):309–318, 2006.
- Timothy J Teyler and Pascal DiScenna. The hippocampal memory indexing theory. *Behavioral neuroscience*, 1986.
- Timothy J Teyler and Jerry W Rudy. The hippocampal indexing theory and episodic memory. *Hippocampus*, 2007.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

- Edward C Tolman. Cognitive maps in rats and men. *Psychological review*, 55(4):189, 1948.
- Rosario Tomasello, Max Garagnani, Thomas Wennekers, and Friedemann Pulvermüller. Brain connections of words, perceptions and actions: A neurobiological model of spatio-temporal semantic activation in the human cortex. *Neuropsychologia*, 98:111–129, 2017.
- Susumu Tonegawa, Mark D Morrissey, and Takashi Kitamura. The role of engram cells in the systems consolidation of memory. *Nature Reviews Neuroscience*, 19(8):485, 2018.
- Volker Tresp and Yunpu Ma. The tensor memory hypothesis. In *NIPS Workshop on Representation Learning*, 2016.
- Volker Tresp, Cristóbal Esteban, Yinchong Yang, Stephan Baier, and Denis Krompaß. Learning with memory embeddings. *arXiv preprint arXiv:1511.07972*, 2015.
- Volker Tresp, Yunpu Ma, and Stephan Baier. Tensor memories. In *Conference on Cognitive Computational Neuroscience*, 2017a.
- Volker Tresp, Yunpu Ma, Stephan Baier, and Yinchong Yang. Embedding learning for declarative memories. In *European Semantic Web Conference*, pages 202–216. Springer, 2017b.
- Volker Tresp, Sahand Sharifzadeh, and Dario Konopatzki. A model for perception and memory. In *Conference on Cognitive Computational Neuroscience*, 2019.
- Volker Tresp, Sahand Sharifzadeh, Dario Konopatzki, and Yunpu Ma. The tensor brain: Semantic decoding for perception and memory, 2020.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *International Conference on Machine Learning*, pages 2071–2080. PMLR, 2016.
- Endel Tulving. Episodic and semantic memory 1. *Organization of Memory. London: Academic*, 381(e402):4, 1972.
- Endel Tulving. *Elements of episodic memory*. Oxford University Press, 1985.
- Endel Tulving. Episodic memory: from mind to brain. *Annual review of psychology*, 53(1): 1–25, 2002.
- Martijn P van den Heuvel and Olaf Sporns. Network hubs in the human brain. *Trends in cognitive sciences*, 17(12):683–696, 2013.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

- Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014.
- James CR Whittington, Timothy H Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy EJ Behrens. The tolmán-eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell*, 183(5): 1249–1263, 2020.
- Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.
- Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 670–685, 2018.
- Eiling Yee, Evangelia G Chrysikou, and Sharon L Thompson-Schill. *The Cognitive Neuroscience of Semantic Memory*. Oxford Handbook of Cognitive Neuroscience, Oxford University Press, 2014.
- Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018.
- Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5532–5540, 2017.