

Ontology as manifold: towards symbolic and numerical artificial embedding

Chloé Mercier, Frédéric Alexandre, Thierry Viéville

▶ To cite this version:

Chloé Mercier, Frédéric Alexandre, Thierry Viéville. Ontology as manifold: towards symbolic and numerical artificial embedding. Workshop on VSA and hyperdimensional computing - Winter 2022 session, Luleå University of Technology, Sweden and UC Berkeley, USA, Jan 2022, Virtual, Sweden. hal-03550354

HAL Id: hal-03550354 https://inria.hal.science/hal-03550354

Submitted on 1 Feb2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Ontology as manifold: towards symbolic and numeric artificial embedding

Chloé Mercier, Frédéric Alexandre, Thierry Viéville



How to model human cognition with regard to...

Psychologically inspired symbolic reasoning

Biologically plausible spiking neural networks

We propose a geometric mapping of ontologies onto neuronal manifolds

RDF(S): Resource Description Framework (Schema) OWL: Web Ontology Language VSA: Vector Symbolic Architecture SPA: Semantic Pointer Architecture NEF: Neural Engineering Framework



- Each symbol encoded as a vector sampled from the unit hypersphere
- High-dimensional vector space => random vectors approximately orthogonal





The neural manifold hypothesis (Gallego et al, 2017)

- Each symbol encoded as a vector sampled from the unit hypersphere
- High-dimensional vector space => random vectors approximately orthogonal
- Dot product as a similarity measure:

compares the semantic meaning of 2 vectors.



- Each symbol encoded as a vector sampled from the unit hypersphere
- High-dimensional vector space => random vectors approximately orthogonal
- Dot product as a similarity measure:

compares the semantic meaning of 2 vectors.

Closed-world reasoning	Open-world reasoning
Anything that cannot be stated true is false	Anything might be true, unless proven false
$\tau \in \{0,1\}$	$\tau \in \{-1, 0, 1\}$



- Each symbol encoded as a vector sampled from the unit hypersphere
- High-dimensional vector space => random vectors approximately orthogonal
- Dot product as a similarity measure:

compares the semantic meaning of 2 vectors.

Closed-world reasoning	Open-world reasoning	Possibility theory
Anything that cannot be stated true is false	Anything might be true, unless proven false	Anything has a certain degree of possibility and necessity
$\tau \in \{0,1\}$	$\tau \in \{-1, 0, 1\}$	$\tau \in [-1,1]$





Structuring the universe of discourse

• symbols: *resources*

- classes (concepts)
- individuals (instances)
- properties (roles)
- facts: (weighted) triples

(subject, predicate, object), τ



Structuring the universe of discourse

• symbols: *resources*

- classes (concepts)
- individuals (instances)
- properties (roles)
- facts: (weighted) triples





RDF/RDFS: Resource Description Framework (Schema)



- **RDF: Resource Description Framework**
 - → syntax: (subject, predicate, object) triples
- RDFS: RDF Schema
 - → metamodel of RDF, written in RDF
 - → allows to write (lightweight) ontologies and draw inferences

A triplestore using associative memories





key	value
S	т <i>В</i> (о,р)



(subject₁, predicate₁, object₁), T₁

key	value
s ₁	→ T ₁ B(0 ₁ ,p ₁)



(subject₁, predicate₁, object₁), τ_1

(subject₂, predicate₂, object₂), T₂

key	value
s ₁ _	⊾ т ₁ <i>В</i> (о ₁ ,р ₁)
s ₂ -	► т ₂ <i>В</i> (о ₂ ,р ₂)



(subject₁, predicate₁, object₁), τ_1

(subject₂, predicate₂, object₂), T₂

(subject₂, predicate₃, object₃), T_3

key	value
s ₁ _	→ T ₁ B(o ₁ ,p ₁)
s ₂ -	$\Rightarrow T_{2} \mathcal{B}(o_{2},p_{2}) + T_{3} \mathcal{B}(o_{3},p_{3})$



```
(subject<sub>1</sub>, predicate<sub>1</sub>, object<sub>1</sub>), T_1
(subject<sub>2</sub>, predicate<sub>2</sub>, object<sub>2</sub>), T_2
(subject<sub>2</sub>, predicate<sub>3</sub>, object<sub>3</sub>), T_3
```

Need to introduce:

Mnemosyne

- scalar multiplication a = та
- vector superposition c = a + b
- vector binding $c = \mathcal{B}(a,b)$

⇒ Semantic Pointer Architecture (Eliasmith et al, 2013)

key	value
s ₁ _	→ T ₁ B(o ₁ ,p ₁)
s ₂ -	$\Rightarrow T_{2} \mathcal{B}(o_{2},p_{2}) + T_{3} \mathcal{B}(o_{3},p_{3})$

(subject₁, predicate₁, object₁), T₁

(subject₂, predicate₂, object₂), T₂

(subject₂, predicate₃, object₃), τ₃

Need to introduce:

Mnemosyne

- scalar multiplication a = τ a
- vector superposition c = a + b
- vector binding $c = \mathcal{B}(a,b)$

⇒ Semantic Pointer Architecture (Eliasmith et al, 2013)

key	value
s ₁ –	⊸т ₁ ℬ(о ₁ ,р ₁)
s ₂ -	$+ T_{2} \mathcal{B}(o_{2},p_{2}) + T_{3} \mathcal{B}(o_{3},p_{3})$

Vector-derived Transformation Bindi	ng			
(Gosmann and Eliasmith, 2019)				
 Square dimensionality d with d'² = d 				
 Binding operation defined as: 	_			٦
$\mathbf{x}, \mathbf{y} = \mathbf{B}, \mathbf{x} = \begin{bmatrix} \mathbf{B}'_{\mathbf{y}} & 0 & 0 \\ 0 & \mathbf{B}'_{\mathbf{x}} & 0 \end{bmatrix} \mathbf{x}$ where $\mathbf{B}'_{\mathbf{x}} = d^{\frac{1}{4}}$	$egin{array}{l} y_1 \ y_{d'+1} \end{array}$	$egin{array}{l} y_2 \ y_{d'+2} \end{array}$	· · · · · · ·	$egin{array}{c} y_{d'} \ y_{2d'} \end{array}$
$\begin{bmatrix} 0 & \mathbf{B}_{\mathbf{y}} & 0 \\ 0 & 0 & \mathbf{B}_{\mathbf{y}} \end{bmatrix}^{\mathbf{H}} \begin{bmatrix} 0 & \mathbf{D}_{\mathbf{y}} & 0 \\ 0 & 0 & \mathbf{B}_{\mathbf{y}} \end{bmatrix}$	$\vdots y_{d-d'+1}$	$dots y_{d-d'+2}$	••. 	$\left[\begin{array}{c} \vdots \\ y_d \end{array} \right]$

- Non commutative, non associative
- Distributive and bilinear

 $\mathcal{B}($

Admits right identity and right inverse

(subject₁, predicate₁, object₁), τ₁ (subject₂, predicate₂, object₂), τ₂ (subject₂, predicate₃, object₃), τ₃

key	value
s ₁ –	→ T ₁ B(o ₁ ,p ₁)
s ₂ -	$+ T_{2} \mathcal{B}(o_{2},p_{2}) + T_{3} \mathcal{B}(o_{3},p_{3})$

Need to introduce:

Mnemosyne

- scalar multiplication a = τ a
- vector superposition c = a + b
- vector binding $c = \mathcal{B}(a,b)$

 ⇒ Semantic Pointer Architecture (Eliasmith et al, 2013)



NEF Associative Memory (Voelker et al, 2014)

yellow circles: passthrough nodes blue circles: neuron ensembles orange connections: inhibitory green connections: modulatory

```
(subject<sub>1</sub>, predicate<sub>1</sub>, object<sub>1</sub>)
(subject<sub>2</sub>, predicate<sub>2</sub>, object<sub>2</sub>)
(subject<sub>2</sub>, predicate<sub>3</sub>, object<sub>3</sub>)
```

key	value
s ₁ _	→ B(o ₁ ,p ₁)
s ₂ -	$+ \mathcal{B}(o_2,p_2) + \mathcal{B}(o_3,p_3)$

Need to introduce:

Mnemosyne

- scalar multiplication a = та
- vector superposition c = a + b
- vector binding $c = \mathcal{B}(a,b)$

 ⇒ Semantic Pointer Architecture (Eliasmith et al, 2013)



NEF Associative Memory (Voelker et al, 2014)

yellow circles: passthrough nodes blue circles: neuron ensembles orange connections: inhibitory green connections: modulatory

inspired by https://protege.stanford.edu/ ontologies/pizza/pizza.owl

Example of inference

Inheritance inference rule:

X rdf:type C & C rdfs:subClassOf C' ⇒ X rdf:type C'





inspired by https://protege.stanford.edu/ ontologies/pizza/pizza.owl

Example of inference

Inheritance inference rule:

X rdf:type C & C rdfs:subClassOf C' ⇒ X rdf:type C' Sufficient condition: **B**type **B**subclass \rightarrow **B**type "inference rule" stored into a 2nd AM: type \oslash subclass \rightarrow type

where $x \oslash y$ is such that $B_{x \oslash y} = B_x B_y$

Mnemosyne



CUE 1: subject of interest





VSA ONLINE - Winter 2022





VSA ONLINE - Winter 2022





Simulation on Nengo



Plotted here are the input cue signals: accounting respectively for the subject, the predicate and the rule to apply.



Simulation on Nengo



Plotted here are the outputs given by the model: the object associated to the given subject and predicate, the simplified predicate obtained from the rule (subClass * type -> type, from the 2nd AM) and the final statement inferred.



Underlying mathematics



Vector-derived Transformation Binding (VTB)

(Gosmann and Eliasmith, 2019)

- Square dimensionality d with d'² = d
- Binding operation defined as:

$$\mathcal{B}(\mathbf{x},\mathbf{y}) = \mathbf{B}_{\mathbf{y}}\mathbf{x} = egin{bmatrix} \mathbf{B}'_{\mathbf{y}} & 0 & 0 \ 0 & \mathbf{B}'_{\mathbf{y}} & 0 \ 0 & 0 & \mathbf{B}'_{\mathbf{y}} \end{bmatrix} \mathbf{x} \quad ext{ where } \mathbf{B}'_{\mathbf{y}} = d^{rac{1}{4}} egin{pmatrix} y_1 & y_2 & \dots & y_{d'} \ y_{d'+1} & y_{d'+2} & \dots & y_{2d'} \ dots & dots & \ddots & dots \ y_{d-d'+1} & y_{d-d'+2} & \dots & y_{d} \end{bmatrix}$$

Γ.

- Non commutative, non associative
- Distributive and bilinear



VT unbinding

(Gosmann and Eliasmith, 2019)

- VTB admits right approximate inverses:
 - the right approximate inverse **y~** for **y** is such that:

 $\forall \mathbf{x}, \mathcal{B}(\mathcal{B}(\mathbf{x}, \mathbf{y}), \mathbf{y}^{\sim})) = \mathbf{B}_{\mathbf{y}^{\sim}} \mathbf{B}_{\mathbf{y}} \mathbf{x} = \mathbf{x}$

• By almost orthogonal: $\mathbf{B}_{\mathbf{y}^{\sim}} = \mathbf{B}_{\mathbf{y}}^{\top}$

=> just permute the elements of y to get y~

• Right identity vector is such that $orall x, \mathcal{B}(\mathbf{x}, \mathbf{i_B}) = \mathbf{x}$ ie $\mathbf{B_{i_B}} = \mathbf{I}$



VT unbinding

Check if a property is a predicate for a given resource and, if so, retrieve the corresponding object:

$$\mathbf{x}
ightarrow \mathbf{z} = \mathcal{B}(\mathbf{y},\mathbf{p})$$
 ,

- Unbind p: $\mathcal{B}(\mathbf{z},\mathbf{p}^{\sim})=\mathcal{B}(\mathcal{B}(\mathbf{y},\mathbf{p}),\mathbf{p}^{\sim})=\mathbf{y}$
- Check if the result belongs to the vocabulary: $\exists y' \in voc, y \cdot y' = 1$? (if similarity is $-1 < \tau < 1$, the statement is possible to some extent)



VT unbinding

What if want to check if two resources are linked to each other by a predicate and, if so, retrieve the corresponding property (left unbinding)?

We can flip the order of the operands using a well-chosen matrix $\mathbf{B}_{\leftrightarrow}$: $\mathcal{B}(\mathbf{y},\mathbf{p}) = \mathbf{B}_{\leftrightarrow}\mathcal{B}(\mathbf{p},\mathbf{y})$

=> Check if two resources are linked to each other by a predicate and, if so, retrieve the corresponding property:

- Flip the operands (multiply by $\mathbf{B}_{\leftrightarrow}$)
- Unbind y (multiply by the transpose of By)
- Check if the result is semantically close to a word from the vocabulary (compute the dot product)



About the operator ⊘

We introduce a vector composition operator making explicit the composition of two VTB operations, namely:

$$\mathbf{B}_{\mathbf{v}} = \mathbf{B}_{\mathbf{y}} \mathbf{B}_{\mathbf{x}} \Leftrightarrow \mathbf{v} \stackrel{\text{def}}{=} \mathbf{y} \oslash \mathbf{x} \quad \text{computable in } O\left(d^{\frac{3}{2}}\right)$$
$$[\mathbf{v}]_{i} = \sqrt{d'} \sum_{k=1}^{k=d'} [\mathbf{y}]_{k+(i-1)} \operatorname{div}_{d'} [\mathbf{x}]_{1+d'(k-1)+(i-1)} \operatorname{mod}_{d'}$$

It commutes with the approximate inverse:(vérifier formulation)

$$(\mathbf{y} \oslash \mathbf{x})^\sim = \mathbf{x}^\sim \oslash \mathbf{y}^\sim$$
 where $\mathbf{x}^\sim \oslash \mathbf{x} \simeq \mathbf{i_B}$



Towards more reasoning power



Beyond the type inheritance inference

- Covering all RDFS entailments reduced to the rule schema developed here, e.g.:
 - property inference (e.g., if Bob is the son of Ali then Bob is of the same family as Ali)
 - range inference (e.g., if Bob is the son of Ali, then Bob is a male)
 - \circ $\,$ domain inference (e.g., if Bob is the son of Ali, then Ali is a human) etc...
- To what extent could it be applicable to description logics and OWL entailment rules?
- SWRL rules (i.e., inference rules expressed directly as rewriting rules on variables) stored in associative memories?



All RDFS entailments

- Everything is a resource
 - IFxpyTHEN
 - x rdf:type rdfs:Resource
 - y rdf:type rdfs:Resource
- Any class is a subclass of rdfs:Resource
 - IF c rdf:type rdfs:Class THEN
 - c rdfs:subClassOf rdfs:Resource
- Type propagation
 - IF c2 rdfs:subClassOf c1
 - **AND** x rdf:type c2
 - THEN x rdf:type c1
- Subsumption reflexivity
 - IF c rdf:type rdfs:Class
 - THEN c rdfs:subClassOf c
- Subsumption transitivity

Mnemosvne

- IF c2 rdfs:subClassOf c1
- AND c3 rdfs:subClassOf c3
 - THEN c3 rdfs:subClassOf c1

- Property propagation
 - IF p2 rdfs:subPropertyOf p1
 - AND x p2 y
 - THEN x pl y
- Property subsumption reflexivity
 - IF p rdf:type rdf:Property
 - THEN p rdfs:subPropertyOf p
- Property subsumption transitivity
 - IF p2 rdfs:subPropertyOf p1
 - AND p3 rdfs:subPropertyOf p2
 - THEN p3 rdfs:subPropertyOf p1
- Type inference (domain)
 - IF p rdfs:domain d AND x p y
 - THEN x rdf:type d
- Type inference (range)
 - $\circ \quad \mathsf{IF} \texttt{p} \texttt{ rdfs:range r AND x p y}$
 - THEN y rdf:type r

inspired by https://protege.stanford.edu/ ontologies/pizza/pizza.owl

Type inheritance



Sufficient condition:

 $B \text{type } B \text{subclass} \to B \text{type}$

"inference rule" stored into a 2nd AM:

type \circ subclass \rightarrow type

Mnemosyne

```
where x \oslash y is such that B_{x \oslash y} = B_x B_y
```





rdfs:range

rdfs:subClassOf

rdfs:subClassOf



Domain and range inferences p rdfs:domain C & хру \Rightarrow X rdf:type C **p**rdfs:range **C** & хру \Rightarrow y rdf:type C Mozzarella We introduce an UNKNOWN vector Υ (accounting for OWA) such that: rdf:type $\mathbf{B}_{\mathbf{B}_{\mathrm{domain}}c} \Upsilon \Rightarrow \mathbf{B}_{\mathrm{type}}c$ thisMozzarella

hasTopping rdfs:range rdf:type Food rdf:type rdfs:subClassOf Pizza rdf:type rdfs:subClassOf thisPizza MargheritaPizza hasTopping rdf:type eats (same with range inference) Luigi Person rdf:type

AM1
$$p
ightarrow \mathbf{B_{domain}} c$$
 $x
ightarrow \mathbf{B_p} \Upsilon$ AM2 $x
ightarrow \mathbf{B_{type}} c$

Mnemosyne

rdfs:subPropertyOf hasTopping hasIngredient Property propagation **p** rdfs:subPropertyOf **q** & хру \Rightarrow xqy Food rdf:type rdfs:subClassOf Using our previous architecture, this kind of inference is only possible when Pizza Mozzarella rdf:type $subPropertyOf = i_B$ rdfs:subClassOf Is this consistent? Why not do the same for all "is a" thisMozzarella thisPizza MargheritaPizza hasTopping relationships (subClassOf, type) ?... rdf:type => ongoing work hasIngredient eats Luigi Person rdf:type Mnemosyne

Gabriel's internship: subproperty inference

$${}^{po}_{s}\mathcal{B} \;\; {}^{po}_{s}\mathcal{B}\mathbf{subPropertyOf} = {}^{po}_{s}\mathcal{B}$$

would mean $\mathbf{subPropertyOf} = i_B$

Is this consistent? Why not do the same for all "is_a" relationships (subClassOf, type) ?... => ongoing work (requires a slightly different architecture)



Property composition

Rules such as

eats \circ hasTopping \rightarrow eats

to infer relational compositions

(not part of RDFS entailments, but similar to OWL *property chains*)





Mnemosyne

OWL: another representation for knowledge



OWL + SWRL







Alternate implementation: a compressed triplestore vector



An alternate implementation of entailment closure

All triplets are memorized in a store that is merely a sum of vectors:

$$\mathbf{s} = \sum_i \mathbf{B}_{\$ \mathbf{p}_i} \, \mathbf{B}_{\$ \mathbf{s}_i} \, \$ \mathbf{o}_i = \sum_i \mathbf{B}_{\$ \mathbf{p}_i \oslash \$ \mathbf{s}_i} \, \$ \mathbf{o}_i = \sum_i \mathbf{B}_{\$ \mathbf{p}_i} \, \mathbf{B}_{\leftrightarrow} \, \mathbf{B}_{\$ \mathbf{o}_i} \, \$ \mathbf{s}_i$$

allowing to enumerate all (subject, object) of a property, i.e., select by predicate, also by predicate and subject or by predicate and object.

Given an entailment rule of the form:

$$(\$s_1 \$p_1 \$o_1.)$$
 and $(\$s_2 \$p_2 \$o_2.) \Rightarrow (\$s_0 \lambda \$p_0 \$o_0.)$

where $\lambda \in [0, 1]$ allows to consider approximate inference.



Note that for an efficient retrieval, when storing triples, we can index them either by subject, object or property:

$$egin{aligned} &{}_{s}^{po}\mathcal{B}(p,y) = \mathcal{B}(p,y) = \mathcal{B}_{y}p\ &{}_{s}^{so}\mathcal{B}(x,y) = \mathcal{B}_{y}^{T}x = \mathcal{B}_{y^{\sim}}x\ &{}_{s}^{sp}\mathcal{B}(x,p) = \mathcal{B}_{p}^{T}x^{\sim} = \mathcal{B}_{p^{\sim}}x^{\sim}\ &{}_{s}^{op}\mathcal{B}(y,p) = \mathcal{B}(p,y) = \mathcal{B}_{y}p\ &{}_{o}^{os}\mathcal{B}(y,x) = \mathcal{B}_{y}^{T}x = \mathcal{B}_{y^{\sim}}x\ &{}_{p}^{s}\mathcal{B}(p,x) = \mathcal{B}_{p}^{T}x^{\sim} = \mathcal{B}_{p^{\sim}}x^{\sim} \end{aligned}$$

(indexing triples by subject)

(indexing triples by property)

(indexing triples by object)

B is a bilinear form.



The class inheritance example:

```
(\$s \text{ rdf:type } \$c_1.) \text{ and } (\$c_1 \text{ rdfs:subClassOf } \$c_2.) \Rightarrow (\$s \text{ rdf:type } \$c_2.)
```



The class inheritance example:

- The rule corresponds to a hardwiring of a 2 input 1 output operator
- Approximate equality (i.e. similarity) is intrinsic to the calculus
- Approximately true property is intrinsic to the calculus
- The '&' operator can be implemented as a product (but not only).





The class inheritance example:

- The rule corresponds to a hardwiring of a 2 input 1 output operator
- Approximate equality (i.e. similarity) is intrinsic to the calculus
- Approximately true property is intrinsic to the calculus
- The '&' operator can be implemented as a product (but not only).



- For "negative" inferences (i.e. $\lambda < 0$) a different (symmetric) rule is required :

 $(\$s \neg rdf:type \$c_1.)$ and $(\$c_2 rdfs:subClassOf \$c_2.) \Rightarrow (\$s \neg rdf:type \$c_2.)$



Accounting for uncertainty: possibility, necessity



Linking with possibility π / necessity v theory $0 \le \nu(E) \le \pi(E) \le \pi(E) \le 1$

$$\begin{aligned} \tau &= \pi + \nu - 1 \in [-1, 1] \\ \rho &= \nu - \pi + 1 \in [0, 1] \end{aligned}$$

A new interpretation of **necessity** and **possibility** as probability bounds



Linking with possibility π / necessity v theory

Given a set of a-priori knowledge \mathcal{K} and observation O, (π ,v) is properly defined as:

 $\nu(E) = \min_{\kappa \in \mathcal{K}} P_{\kappa}(E|O)$ $\pi(E) = \max_{\kappa \in \mathcal{K}} P_{\kappa}(E|O)$

Given a possibility/necessity distribution (π, v) we can define

- a probability distribution of maximal entropy
- for "minimally corrected" compatible distributions $\tilde{\nu} \leq \nu, \pi \leq \tilde{\pi}$



Linking with possibility π / necessity v theory

Compatibility with Boolean and three-value logic calculus can be derived

Partially obtained by distribution monotonicity

 $E \subseteq E' \Rightarrow \nu(E) \leq \nu(E') \text{ and } \pi(E) \leq \pi(E')$

Fully compatible in the "deterministic case" (vanilla theory) Requires additional constraints in the 2D case for 0 or \pm 1 values

In an open world (i.e. if we do not consider an exhaustive universe), logical derivation related to the the complement of an event are not possible.

The correct model underneath is modal logic.



Linking with possibility π / necessity v theory

In VSA (π , v) could be implemented either as a **pair of real numbers**, or as a **complex number**.

- Similarity can be defined as the dot product $Re(xy^*)$ in the hermitian space \mathbb{C}^d with the same interpretation
- VTB algebra and related operators seem generalizable

However there is no real "gain" to consider complex number apart from using the complex product instead of 2d matrix product.



Conclusion/Perspectives

- Distributive, tightly coupled approach, drawing inspiration from biology and psychology
- 2 associative memories:
 - AM1 to store the ontological "facts"
 - AM2 to store the inference rules => used to disclose implicit relationships/memberships
- Alternate implementation: compressed triplestore (computing the transitive closure)
- Achievable inferences:
 - RDFS entailments:
 - concept inclusion an inheritance
 - property domains and ranges
 - extension to OWL ontologies: class complement, intersection, etc
- + Scalability, completeness, soundness on larger databases?
- + How to manage uncertainty? (possibility-necessity approach?)
- + How to update the associative memory to add the inferred relationships?
- + How to learn the inference rules from patterns discovered in AM1 ? Nengo associative memory is read-only (though some work is aiming to implement learning within AMs, cf Voelker et al 2014)

Thank you!



Special thanks to:

Hugo Chateau-Laurent Théophane Valleys Gabriel Doriath Döhler Terrence Stewart

References: Alexandre 2019; Allemang, Hendler, and Gandon 2020; Ayadi et al. 2019; Bekolay et al. 2014; Choo 2018; Chui and Mhaskar 2018; Cohen, Yang, and Mazaitis 2017; Crawford, Gingerich, and Eliasmith 2016; Denœux, Dubois, and Prade 2020; Eidoon, Yazdani, and Oroumchian 2008; Eliasmith 2013; Eliasmith and Anderson 2002; Eliasmith et al. 2012; Gallego et al, 2017, Garcez and Lamb 2020; Gayler 2003; Gosmann 2018; Gosmann and Eliasmith 2019; Graves, Wayne, and Danihelka 2014; Grosof et al. 2003; Guo et al. 2016; Hohenecker and Lukasiewicz 2020; Horridge 2011; Horst 2005; Jiménez, Elizalde, and Raj 2018; Komer et al. 2019; Lallement, Hilario, and Alexandre 1995; Levesque 1986; Levy and Gayler 2008; Mandler 2011; McClelland and Rogers 2003; Mercier et al. 2021; Nallapu 2019; Petrucci, Ghidini, and Rospocher 2016; Phan et al. 2017; Plate 1995; Pulvermüller 2013; Riegel et al. 2020; Romero, David, and Lille 2019;Rusawuk 2018; Sajjad, Docherty, and Tyshetskiy 2019; Sauerwald and Zanetti 2019; Schlegel, Neubert, and Protzel 2020; Shi et al. 2020; Simpkin et al. 2018; Smith 1994; Stewart, Choo, and Eliasmith 2010; Stewart, Tang, and Eliasmith 2011; Sun and Alexandre 2013; Tettamanzi, Zucker, and Gandon 2017; Tous and Delgado 2006; Voelker, Crawford, and Eliasmith 2014; Wang, Qiu, and Wang 2021; Xiao, Huang, and Zhu 2015; Zhu et al. 2017

