Semantic Web 0 (0) 1 IOS Press

Is Neuro-Symbolic AI Meeting its Promise in Natural Language Processing? - A Structured Review

Kyle Hamilton^{*}, Aparna Nayak, Bojan Božić and Luca Longo

SFI Centre for Research Training in Machine Learning, School of Computer Science, Technological University Dublin, Republic of Ireland

E-mails: kyle.i.hamilton@mytudublin.ie, aparna.nayak@tudublin.ie, bojan.bozic@tudubulin.ie, luca.longo@tudublin.ie

Abstract. Advocates for Neuro-Symbolic AI (NeSy) assert that combining deep learning with symbolic reasoning will lead to stronger AI than either paradigm on its own. As successful as deep learning has been, it is generally accepted that even our best deep learning systems are not very good at abstract reasoning. And since reasoning is inextricably linked to language, it makes intuitive sense that Natural Language Processing (NLP), would be a particularly well-suited candidate for NeSy. We conduct a structured review of studies implementing NeSy for NLP, the challenges and future directions, and aim to answer the question of whether NeSy is indeed meeting its promises: reasoning, out-of-distribution generalization, interpretability, learning and reasoning from small data, and transferability to new domains. We examine the impact of knowledge representation, such as rules and semantic networks, language structure and relational structure, and whether implicit or explicit reasoning contributes to higher promise scores. We find that knowledge encoded in relational structures and explicit reasoning tend to lead to more NeSy goals being satisfied. We also advocate for a more methodical approach to the application of theories of reasoning, which we hope can reduce some of the friction between the symbolic and sub-symbolic schools of AI.

Keywords: Neuro-Symbolic AI, Natural Language Processing, Knowledge Representation, Structured Review

1. Introduction

1.1. Background

At its core, Neuro-Symbolic AI (NeSy) is "the combination of deep learning and symbolic reasoning" [1]. The goal of NeSy is to address the weaknesses of each of symbolic and sub-symbolic approaches while preserving their strengths (see figure 1). Thus NeSy promises to deliver a best-of-both-worlds approach which embodies the "two most fundamental aspects of intelligent cognitive behavior: the ability to learn from experience, and the ability to reason from what has been learned"[1, 2].

Remarkable progress has been made on the learning side, especially in the area of Natural Language Processing (NLP) and in particular with deep learning architectures such as the transformer [3]. However, these systems display certain intrinsic weaknesses which some researchers argue cannot be addressed by deep learning alone; and that in order to do even the most basic reasoning, we need rich representations which enable precise, human interpretable inference via mathematical logic.

^{*}Corresponding author. E-mail: kyle.i.hamilton@mytudublin.ie.

K. Hamilton et al. / NeSy for NLP: Review



Fig. 1. Symbolic vs Sub-Symbolic strengths and weaknesses. Based on the work of [4]

Historically, rivalry between symbolic and connectionist, or sub-symbolic, AI research has stymied collaboration across these fields. Although it should be acknowledged that not everyone thought the two were incompatible. As early as 1991, Marvin Minsky presciently asked "Why is there so much excitement about Neural Networks today, and how is this related to research on Artificial Intelligence? Much has been said, in the popular press, as though these were conflicting activities. This seems exceedingly strange to me, because both are parts of the very same enterprise" [5].

More recently, a discussion between Gary Marcus and Yoshua Bengio at the 2019 Montreal AI Debate prompted some passionate exchanges in AI circles, with Marcus arguing that "expecting a monolithic architecture to handle abstraction and reasoning is unrealistic", while Bengio defended the stance that "sequential reasoning can be performed while staying in a deep learning framework"[6].

Spurred by this discussion, and almost ironically, by the success of deep learning (and ergo, the clarity into its limitations), research into hybrid solutions has seen a dramatic increase (see figure 2). At the same time, discussion in the AI community has culminated in "violent agreement" [7] that the next phase of AI research will be about "combining neural and symbolic approaches in the sense of NeSy AI [which] is at least a path forward to much stronger AI systems" [8]. Much of this discussion centers around the ability (or inability) of deep learning to reason, and in particular, to reason outside of the training distribution. Indeed, at IJCAI 2021, Yoshua Bengio affirms that "we need a new learning theory to deal with Out-of-Distribution generalization" [9]. Bengio's talk is titled "System 2 Deep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality". Here, System 2 refers to the System 1/System 2 (S1/S2) dual process theory of human reasoning developed by psychologist and Nobel laureate Daniel Kahneman in his 2011 book "Thinking, Fast and Slow" [10]. AI researchers have drawn many parallels between the characteristics of sub-symbolic and symbolic AI systems and human reasoning with S1/S2. Broadly speaking, sub-symbolic (neural, deep-learning) architectures are said to be akin to the fast, intuitive, often biased and/or logically flawed S1. And the more deliberative, slow, sequential S2 can be thought of as symbolic or logical. But this is not the only theory of human reasoning as we'll discuss later in this paper.

1.2. Reasoning & Language

"Language understanding in the broadest sense of the term, including question answering that requires commonsense reasoning, offers probably the most complete application area of neurosymbolic AI"[1]. This makes a lot intuitive sense from a linguistic perspective. If we accept that language is compositional, with rules and structure, then it should be possible to obtain its meaning via logical processing. Compositionality in language was formalized by Richard Montague in the 1970s, in what is now referred to as *Montague grammar*. "The key idea is that compositionality requires the existence of a homomorphism between the expressions of a language and the meanings of those expressions."¹ in other words, there is a direct relationship between syntax and semantics. This is



Fig. 2. Number of Neuro Symbolic articles published since 2010, normalized by the total number of all Computer Science articles published each year. The figure represents the unfiltered results from Scopus given the search keywords described in section 2.2.

in line with Noam Chomsky's *Universal grammar* which states that there is a structure to natural language which is innate and universal to all humans regardless of cultural differences. The challenge lies in representing this structure in a way that both captures the semantics and is computationally efficient.

On the one hand, distributed representations are desirable because they can be efficiently processed by gradient descent (the backbone of deep learning). On the other, the meaning embedded in a distributed representation is difficult if not impossible to decompose. So while a large language model (LLM) may be very good at making certain types of predictions, it is not able to provide an explanation of how it got there. We've also seen that the larger the model - more parameters as well as more training data - the better the predictions. But even as these models get infeasibly large, they still fail on tasks requiring basic commonsense. The example in Figure 3, given by Marcus and Davis in [11] is a case in point.

You are having a small dinner party. You want to serve dinner in the living room. The dining room table is wider than the doorway, so to get it into the living room, you will have to **remove the door. You have a table saw, so you cut the door in half and remove the top half.**

Fig. 3. GPT3 text completion example. The prompt is rendered in regular font, while the GPT3 response is shown in bold [11].

On the other hand, traditional symbolic approaches have also failed to capture the essence of human reasoning. We don't need a scholar to confirm that everyday commonsense reasoning is nothing like the rigorous mathematical logic whose goal is validity. But even when the objective isn't commonsense, but rather tasks which require precise, deterministic answers such as expert reasoning or planning, traditional symbolic reasoners are slow, cumbersome, and computationally intractable at scale. Description Logics (DLs) such as OWL, for example, are used to reason over ontologies and knowledge graphs (KGs) on the Web. However, one must accept a harsh trade-off between expressivity and complexity when choosing a DL flavor. Improving the performance of reasoning over ontologies and knowledge graphs that power search and information retrieval across the Web is particularly relevant to the Semantic Web community. Hitzler et al. report on the current research in this area [12].

1.3. Contributions

45 Several surveys have already been conducted which cover the overall NeSy landscape going as far back as 2005, 46 and as recently as 2021, so we will not attempt to replicate that here [1, 4, 8, 13–25]. In fact, our understanding of 47 the field is guided by the works of these scholars. For a succinct overview we refer the reader to [8]. And for a more 48 in-depth analysis we recommend [4]. Our aim is to synthesize recent work implementing NeSy in the language 49 domain, and to verify if the goals of NeSy are being realized, what the challenges are, and future directions. To our 41 knowledge this is the first attempt at this specific task. In the following sub sections, we briefly describe each of the 42 following goals. These are similar to the benefits described in [8]:

- 1. Out-of-distribution Generalization
- 2. Interpretability
- 3. Reduced size of training data
- 4. Transferability
- 5. Reasoning

1.3.1. Out-of-distribution (OOD) Generalization

OOD generalization refers to the ability of a model to extrapolate to phenomena not previously seen in the training data. The lack of OOD generalization in LLMs is often demonstrated by their inability perform commonsense reasoning, as in the example in figure 3.

1.3.2. Interpretability

As Machine Learning (ML) and AI become increasingly embedded in daily life, the need to hold ML/AI accountable is also growing. This is particularly true in sensitive domains such as healthcare, legal, and some business applications such as lending, where bias mitigation and fairness are critical. These concerns, among others, are the province of Explainable AI (XAI). XAI can be broken down into three main categories:

- 1. *Explainability* the facility for an expert human to see how a model arrived at a prediction or inference often represented as a set of rules derived from the model. It seeks to answer the question: *How does the model work?*
- 2. *Interpretability* the facility for a non-expert human to see how the data used by the model led to a prediction or inference, usually in the form of a cause and effect articulation. It seeks to answer the question: *Why did the model come to this conclusion?*
- 3. *Interactivity* the facility for a human to interrogate the model about counterfactuals. It seeks to answer the question: *What would happen if the data was different?*

In the literature, sometimes the term interpretability is used in place of explainability and vice versa, however, for our purposes, all three categories are subsumed under the general notion of interpretability. XAI is made possible by invoking an explicit reasoning module post hoc, or building interpretability into the system to begin with.

1.3.3. Reduced size of training data

SOTA language models utilize massive amounts of data for training. This can cost in the thousands or even millions of dollars, take a very long time, and is neither environmentally friendly nor accessible to most researchers or businesses. The ability to learn from less data brings obvious benefits. But apart from the practical implications, there is something innately disappointing in LLMs' 'bigger hammer' approach. Science rewards parsimony and elegance, and NeSy promises to deliver results without the need for such massive scale.

1.3.4. Transferability

Transferability is the ability of a model which was trained on one domain, to perform similarly well in a different domain. This can be particularly valuable, when the new domain has very few examples available for training. In such cases we might rely on knowledge transfer similar to the way a human might rely on abstract reasoning when faced with an unfamiliar situation.

1.3.5. Reasoning

According to Encyclopedia Britannica, "To reason is to draw inferences appropriate to the situation" [26]. Reasoning is not only a goal in its own right, but also the means by which the other above mentioned goals can be achieved. Not only is it one of the most difficult problems in AI, it is one of the most contested. In section 4.1 we examine the uses of the term reasoning in more depth.

The remainder of this manuscript is structured as follows. Section 2 describes the research methods employed for searching and analysing relevant studies. In Section 3 we analyze the results of the data extraction, how the studies reviewed fit into Henry Kautz's NeSy taxonomy, and propose a simplified nomenclature for describing Kautz's NeSy categories. Section 4 discusses various existing challenges to implement NeSy. Section 5 presents limitations of the work and future directions of NeSy in NLP followed by conclusion in Section 6.

2. Methods

Our review methodology is guided by the principles described in [27–29].

2.1. Research Questions

- 1. What are the existing studies on neurosymbolic AI (NeSy) in natural language processing (NLP)?
- 2. What are the current applications of NeSy in NLP?
- 3. How are symbolic and sub-symbolic techniques integrated and what are the advantages/disadvantages?
- 4. What are the challenges for NeSy and how might they be addressed (including existing proposals for future work)?
- 5. What areas of NLP might be likely to benefit from the NeSy approach in the future?

2.2. Search Process

We chose Scopus to perform our initial search, as Scopus indexes all the top journals and conferences we were interested in. This obviously precludes some niche publications and it is possible we missed some relevant studies. As our aim is to shed light on the field generally, our assumption is that the top journals are a good representation of the research area as a whole. Since we were looking for studies which combine neural and symbolic approaches, our query consists of combinations of neural and symbolic terms as well as variations of neuro-symbolic terms, listed in table 1:

Table 1 Search Keywords					
Neural Terms	Symbolic Terms	Neuro-Symbolic Terms			
sub-symbolic machine learning deep learning	symbolic reasoning logic	neuro-symbolic neural-symbolic neuro symbolic neural symbolic neurosymbolic			

The initial query was restricted to peer-reviewed English language journal articles from the last 10 years and conference papers from the last 3 years, which produced 2,412 results. The query and additional details can be found in our github repository 2

2.3. Study selection process

We further limit the journal articles to those published by the top 20 publishers as ranked by Scopus's CiteScore, which is based on number of citations normalized by the document count over a 4 year window³, and SJR (SCImago Journal Rank), a measure of prestige inspired by the PageRank algorithm over the citation network⁴, the union of which resulted in 29 publishers, and eliminated 669 articles, for a total of 1,510 journal articles and 232 conference papers for screening. Two researchers independently screened each of the 1,742 studies (articles and conference papers), based on the inclusion/exclusion criteria in Table 2. An overview of the selection process can be seen in Figure 4.

The first round of inclusion/exclusion was performed on the titles and abstracts of the 1,742 studies (1,510 articles and 232 papers) from the above identification step. The inclusion criteria at this stage was intentionally broad, as the process itself was meant to be exploratory, and to inform the researchers of relevant topics within NeSy.

²https://github.com/kyleiwaniec/neuro-symbolic-ai-systematic-review
 ³https://service.elsevier.com/app/answers/detail/a_id/14880/kw/citescore/supporthub/scopus/
 ⁴https://service.elsevier.com/app/answers/detail/a_id/14883/supporthub/scopus/related/1/



To better understand the disagreement and researcher biases, we do an analysis on the term frequency - inverse document frequency (TF-IDF) of one-, bi-, and tri-grams on each of the three areas of the venn diagram: studies included by researcher 1 only, studies included by researcher 2 only, and studies included by both researchers. We calculated the TF-IDF from the abstracts of each of the three groupings, and generated word clouds. At a glance, it appears that researcher 1 chose to include documents where terms related to the symbolic dimension

K. Hamilton et al. / NeSy for NLP: Review

such as "symbolic", "rules", and general terms such as "framework", and "data" appeared more frequently, whereas researcher 2, leaned towards terms along the neural dimension such as "deep learning", "neural", and "networks". Given the academic background of each researcher, we reasoned that the discrepancy was due to bias towards each individual's area of research. In the third grouping, where both researchers agreed on inclusion, a more balanced distribution can be seen with the terms "symbolic", "artificial", and "neural" carrying similar weight. (See Figure 6.)



Fig. 0. IT-IDF word Clouds

We observed that it was not always clear from the abstract alone whether the sub-symbolic and symbolic methods were integrated in a way that meets the inclusion criteria, which may also have led to disagreement.

To facilitate the next round of review, we kept a shared glossary of symbolic and sub-symbolic concepts as they presented themselves in the literature. We each reviewed all 337 studies again, this time skimming the studies themselves. Any disagreement at this stage was discussed in person with respect to the shared glossary. This process led to the elimination of many studies for a final count of 75 studies marked for the next round of review.

2.4. Quality Assessment

The quality of each study was determined through the use of a nine-item questionnaire. Each question was answered with a binary value, and the study's quality was determined by calculating the ratio of positive answers. Studies with a quality score below 50% were excluded.

Quality Questions:

- 1. Is there a clear and measurable research question?
- 2. Is the study put into context of other studies and research, and design decisions justified accordingly? (Number of references in the literature review/ introduction)
- 3. Is it clearly stated in the study which other algorithms the study's algorithm(s) have been compared with?
- 4. Are the performance metrics used in the study explained and justified?
- 5. Is the analysis of the results relevant to the research question?
- 6. Does the test evidence support the findings presented?
- 7. Is the study algorithm sufficiently documented to be reproducible? (independent researchers arriving at the same results using their own data and methods)
 - 8. Is code provided?
 - 9. Are performance metrics provided? (hardware, training time, inference time)

More than 88% of the studies satisfy the requirements listed from Q1 to Q6. However, nearly 82% of the studies fail to provide source code or details related to the computing environment which makes the system difficult to reproduce. This leads to an overall reduction of the average quality score to 76.3% - see Figure 7.

Finally, each of the 34 studies selected for inclusion was evaluated, classified, and data extraction was performed
 for each of the features outlined in Table 3. For acceptable values of individual features see B. The lists of neural and
 symbolic terms referenced in the table constitute the glossary items learned from conducting the selection process.



In this section, we perform quantitative data analysis based on the extracted features in Table 3. Each study was labeled with terms from the aforementioned glossary, and each term in the glossary was classified as either symbolic, or neural. A bi-product of this process are two taxonomies built bottom-up of concepts relevant to the set of studies under review. The two taxonomies are a reflection of the definition of NeSy provided earlier: "the combination of deep learning and symbolic reasoning". Thus on the learning side, we have neural architectures (described in Section 3.2.1), and on the symbolic reasoning side we have knowledge representation (described in Section 3.2.2). These results are rendered in Table 4, with the addition of color representing a simple metric, or promise score, for each study. The promise score is simply the number of goals reported to have been satisfied by the model(s) in the study.

3.1. Exploratory Data Analysis

We plot the relationships between the features extracted from the studies, and the goals from section 1.3 in an effort to identify any correlations between them, and ultimately to identify patterns leading to higher promise scores.

Table 3

	Data extraction features
Feature	Description
Business application	Real world NLP tack of the proposed study
Technical emplication	The of the left starts illustrated in Figure 0
Technical application	Type of model output - illustrated in Figure 9
Type of learning	Indicates learning method (supervised, unsupervised, etc.)
Knowledge representation	One of four categories described in Section 3.2
Type of reasoning	Indicates whether knowledge is represented implicitly (embedded) or explicitly (symbolic)
Language structure	Indicates whether linguistic structure is leveraged to facilitate reasoning
Relational structure	Indicates whether relational structure is leveraged to facilitate reasoning
Symbolic terms	List of symbolic techniques used by the models
Neural terms	List of neural architectures used by the models
Datasets	List of all datasets considered
Model description	Describes model architecture schematically
Evaluation Metrics	Evaluation metrics reported by the authors
Reported score	Model performance reported by the authors
Contribution	Novel contribution reported by the authors
Key-intake	Short description of the study
isNeSv	Indicates whether the authors label their study as Neuro-Symbolic
NeSy goals	For each of the goals listed in Section 1, indicates whether the goal is met as reported by the authors
Kautz category	List of categories from Kautz's taxonomy
NeSy category	List of categories from the proposed nomenclature
Study quality	Percentage of positive answers in the quality assessment questionnaire

3.1.1. Business and Technical Applications

A collection of common NLP tasks is shown in Figure 9. The subset of tasks belonging to Natural Language Understanding (NLU) and Natural Language Generation (NLG) are often regarded as more difficult, and presumed to require reasoning. Given that *reasoning* was one of the keywords used for search, it is not surprising that many studies report reasoning as a characteristic of their model(s). Also unsurprising is the fact that nearly half of the text classification studies (which do not belong to this subset) are not associated with any NeSy goals. The relationship between all tasks, or business applications, and NeSy goals is shown in Figure 10.

The business application largely determines the type of model output, or technical application, as can be seen in the almost one-to-one mapping in Figure 11. The exception being question answering, which has been tackled as both an inference and a classification problem. Question answering is the most frequently occurring task, and is associated mainly with reduced data, and to a much lesser degree, interpretability. On a philosophical level this seems somewhat disappointing, as one would hope that in receiving an answer, one could expect to understand why such an answer was given.

For completeness, the number of studies representing the technical applications and most frequently occurring business application is given in Figure 13, while Figure 12 combines business applications, technical applications, and goals.

K. Hamilton et al. / NeSy for NLP: Review

		No Inmhair a C	urai & Syn	idolic Combina	ho 5 desert	din Castier 1.2		
	1 2 3 4 M	Number of I Note: s	vesy goais : ome studies	sausned out of t use multiple ter	ne 5 describe	u in Section 1.3.		
		11010. 5	ome snares	use manple let	innques.			
				Knowledge	Representatio	on		
			Fromas	Logia	Pulas	Semantic		
			Frames	Logic	Kules	network		
					[311 [32]			
	Linear Models	SVM	[30]		[33]	[34]		
					[]			
	Early	MLP	[35]	[36, 37]	[38] [39]			
	Generations				[40]			
		CNN	[41]	[42] [37]	[43]	[44, 45]		
	Graphical	DBN		[42]				
	Models	CNN	5461		F 4771	[40]		
		GNN	[46]		[47]	[48]		
			[41 /0]	[51] [52]	[54] [55]	[44] [59]		
	Sequence-	RNN	[50]	[42] [53]	[56] [57]	[45] [60]		
	to-Sequence		[50]		[58]			
		RcNN		[61]	[62]			
		I TO I	5603					
	Neuro-	LIN	[63]					
	Symbolic	RNKN		[61]				
						[64]		
	Neuroevolution							
61 A 11						oturol I	ndorotondin-	
Natural La	(NLP)	Toyt Col	orizotion/		N	atural Language U (NLU)	nderstanding	
	· · · · · · · · · · · · · · · · · · ·	Classifi	cation			Natural La	nguage Generation	
	Named E	ntity Recogni	tion /	Relation Extra	ction		(NLG)	
		(NER)	/ Ques	tion Answering (Q	۹)	$\langle \rangle$		
		ech Tagging		Para	aphrase & Natura	al		
	Part-of-Spe	051	Contim	nt Analysis				
	Part-of-Spe (P	OS)	Sentimer	it Analysis Lar	nguage Inference	e		
	Part-of-Spe (P Syntactio	OS)	Sentimer	antic Parsing	nguage Inference Dialog Agents	e 		
	Part-of-Spe (P Syntaction Coreferent	OS) Parsing ence Resolutio	Sentimer	antic Parsing	nguage Inference Dialog Agents	9 /		
	Part-of-Spe (P Syntactic Corefere	OS) c Parsing ence Resolution Machine Tra	Sentimer	antic Parsing [Summarizat	nguage Inference	a /		
	Part-of-Spe (P Syntactio Corefere	COS) c Parsing ence Resolution Machine Tra	Sentimer	antic Parsing [Summarizat	nguage Inference	a /		
	Part-of-Spe (P Syntactio Corefere	COS) C Parsing ence Resolution Machine Tr	Sentimer Sem on anslation	nt Analysis Lar Lantic Parsing [Summarizat	nguage Inference	a /		
	Part-of-Spe (P Syntaction Corefere	OS) c Parsing ance Resolution Machine Tr	Sentimer Sem on anslation	nt Analysis Lar Lantic Parsing Summarizat	nguage Inference	a /		
	Part-of-Spe (P Syntactio Corefere	OS) c Parsing ence Resolution Machine Tr	Fig. 9. Con	nt Analysis Lar Lartic Parsing Summarizat	nguage Inference	a /		
	Part-of-Spe (P Syntactio Corefere	OS) c Parsing ence Resolution Machine Tr	sentimer Sem anslation Fig. 9. Con	nt Analysis Lar Summarizat	nguage Inference	a /		
3.1.2. Type of la	Part-of-Spe (P Syntactio Corefere	OS) c Parsing ence Resolution Machine Tr	Fig. 9. Con	nt Analysis Lar Lar Summarizat	nguage Inference	a /		
3.1.2. Type of le Machine lear	Part-of-Spe (P Syntactic Corefere earning rning algorithms ar	e classif	Fig. 9. Con	nmon NLP task	nguage Inference Dialog Agents ion	ed, semi-supe	rvised, curriculum or	

14 demonstrates that the supervised method outweighs all other approaches.





3.1.3. Implicit vs Explicit Reasoning

How reasoning is performed often depends on the underlying representation and what it facilitates. Sometimes the representations are obtained via explicit rules or logic, but are subsequently transformed into non-decomposable embeddings for learning. As such, we can say that any reasoning during the learning process is done implicitly. Studies utilizing Graph Neural Networks (GNNs) [46-48] would also be considered to be doing reasoning implicitly. The majority of the studies doing implicit reasoning leverage linguistic and/or relational structure to generate those internal representations. These studies meet 21 out of a possible 102 NeSy goals, where 102 = #goals * #studies, or 20.6%. For reasoning to be considered explicit, rules or logic must be applied during or after training. Studies which implement explicit reasoning perform slightly better, meeting 21 out of 72 goals, or 29.2% and generally require less training data. Additionally, 4 studies implement both implicit and explicit reasoning, at a NeSy promise rate of 40%. Of particular interest in this grouping is Bianchi et al. [63]'s implementation of Logic Tensor Networks (LTNs), originally proposed by Serafini and Garcez in [68]. "LTNs can be be used to do after-training reasoning over combinations of axioms which it was not trained on. Since LTNs are based on Neural Networks, they reach similar results while also achieving high explainability due to the fact that they ground first-order logic.[63]" Also in this grouping, Jiang et al. [61] propose a model where embeddings are learned by following the logic expressions encoded in huffman trees to represent deep first-order logic knowledge. Each node of the tree is a logic expression, thus hidden layers are interpretable.

Figure 15 shows the relationship between implicit & explict reasoning and goals, while the relationship between
 knowledge representation, type of reasoning, and goals is shown in Figure 16.



extracted features and their interactions - see Figure 17. Perhaps the most telling chart is the mapping between structures and goals, where nearly half the studies leveraging linguistic structure don't meet any of the goals. This runs counter to the intuition that language is a natural fit for NeSy. However, to properly test this intuition, future work on more targeted experiments would need to be performed.

3.1.5. Datasets and Benchmarks

Each study in our survey is based on a unique dataset, and a variety of metrics. Given that there are nearly as many business applications, or tasks, as there are studies, this is not surprising. As such it is not possible to compare the performance of the models reviewed. However, this brings up an interesting question, and that is how one might design a benchmark for NeSy in the first place. A discussion about benchmarks at the IBM Neuro-Symbolic AI Workshop 2022⁵ resulted in general agreement that the most important characteristic of a good benchmark



for NeSy is in the diversity of tasks tackled. Gary Marcus pointed out that current benchmarks can be solved extensionally, meaning they can be "gamed". With enough attempts, a model can become very good at a specific task without solving the fundamental reasoning challenge. Ostensibly, this leads to models which are not able to generalize out of distribution. In contrast, to solve a task intensionally is to demonstrate "understanding" which is transferable to different tasks. This view is controversial with advocates of purely connectionist approaches arguing that "understanding" is not only ill defined, but also a moving target. So instead of worrying about the semantics of "understanding", the panelists agreed that to make the benchmarks robust to gaming is to build in enormous variance in the types of tasks they tackle. Taking this a step further, Luis Lamb proposed that instead of designing benchmarks for testing models, we should be designing challenges which encourage people to work on important real world problems.

3.2. Taxonomies: Neural, Symbolic, & NeSy

We now describe the taxonomies introduced in Section 3 as well as NeSy categories.

3.2.1. Neural Architectures

In the main, the extracted neural terms refer to the neural architecture implemented in a given study. We group these into higher level categories such as Linear models, Early generation (which includes CNNs), Graphical models, and Sequence-to-Sequence. We also include here Neuro-Symbolic architectures such as Logic Tensor Networks (LTN) and Recursive Neural Knowledge Networks (RNKN) because they are suitable to optimization via gradient descent.

We include one study [64] which doesn't strictly meet the neural criteria in the sense that it does not implement gradient descent, but rather Neuroevolution (NE). Neuroevolution involves genetic algorithms for modifying neural network weights, topologies, or ensembles of networks by taking inspiration from biological nervous systems [69, 70]. Neuroevolution is often employed in the service of Reinforcement Learning (RL).

Studies which do not mention any specific models are categorised as Multilayer Perceptron (MLP).

3.2.2. Symbolic Knowledge Representation

The definition we adopted states that NeSy is *the combination of deep learning and symbolic reasoning*. Our neural taxonomy described above reflects the *deep learning* component. Since no reasoning can be done without knowledge, we use Knowledge Representation (KR) as a means of categorizing the *symbolic reasoning* component. Common KR categories include: production rules, logical representation, frames, and semantic networks [71–74].

- Production rules The production rules are a set of condition-action pairs that represent knowledge [75].
- Logical representation Logical representation is the study of entailment relations—languages, truth conditions, and rules of inference [73].
- Frames Frames are objects which hold entities, their properties and methods.
- *Semantic networks* Semantic networks (frame networks) are graphs where nodes are frames, and edges are the relationships between nodes which can represent information.

Table 4 shows which studies combine which of the above neural (3.2.1) and symbolic (3.2.2) categories.

3.2.3. NeSy Categories

NeSy systems can be categorized according to the nature of the combination of neural and symbolic techniques.
At AAAI-20, Henry Kautz presented a 6 level taxonomy of Neuro-Symbolic architectures with a brief example of
each [7]. While Kautz has not provided any additional information beyond his talk at AAAI-20, several researchers
have formed their own interpretations [1, 8, 21]. We have categorized all the reviewed studies according to Kautz's
taxonomy only to discover that nearly half (N=15) of the studies belong to Level 1, which arguably is not neurosymbolic, but rather "standard-operating-procedure" as Kautz himself put it.

Level 1 *symbolic Neuro symbolic* is a special case where symbolic knowledge (such as words) is transformed into continuous vector space and thus encoded in the feature embeddings of an otherwise "standard" ML model. We opted to include these studies if the derived input features belong to the set of symbolic knowledge representations described in Section 3.2. One could still argue that this is simply a case of good old fashioned feature engineering, into continuous vector space and thus encoded in the feature specific transformed the set of symbolic knowledge representations into continuous vector space and thus encoded in the feature specific transformed the set of symbolic knowledge representations into continuous vector space and thus encoded in the feature specific transformed to the set of symbolic knowledge representations is described in Section 3.2. One could still argue that this is simply a case of good old fashioned feature engineering,

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

and not particularly special, but we want to explore the idea that deep learning can perform reasoning, albeit implicitly, if provided with a rich knowledge representation in the pre-processing phase. We classify these studies as *Sequential*. Evaluating these studies as a group was particularly challenging as they have very little in common including different datasets, benchmarks and business applications. Half of the studies don't mention reasoning at all, and the ones that do are mainly executing rules on candidate solutions output by the neural models post hoc. In aggregate, only 16 out of a total of 75 (15 studies * 5 goals), or 21.3%, possible NeSy goals were met.

Level 2 *Symbolic[Neuro]* is what we describe as a *Nested* architecture, where a symbolic reasoning system is
 the primary system with neural components driving certain internal decisions. AlphaGo is the example provided by
 Kautz, where the symbolic system is a Monte Carlo Tree Search with neural state estimators nominating next states.
 We found three studies that fit this architecture [32, 40, 42].

Level 3 Neuro; Symbolic appears to hold more promise in terms of NeSy goals met. Three of the five studies with 11 a Promise score of 3 or more belong to this category (the remaining two belong to levels 4 and 5, which we will 12 discuss in the next subsection). There are six studies in this category, all but one of which leverage relational structure 13 in some manner. A common theme is the use of graph representations and/or GNNs which aligns with recent 14 research directions proposed by Garcez et al. in [21], as well as Zhang et al. in [24]. We call Level 3 Cooperative, 15 16 as it is conceptually similar to Reinforcement Learning (RL). Here, a neural network focuses on one task (e.g. object detection) and interacts via input/output with a symbolic reasoner specializing in a complementary task (e.g. 17 query answering). Unstructured input is converted into symbolic representations which can be solved by a symbolic 18 reasoner, which in turn informs the neural component which learns from the errors of the symbolic component. The 19 Neuro-symbolic Concept Learner [57] is an example of Level 3, meeting 4 out of the 5 NeSy goals, where reasoning 20 21 is performed explicitly in a "symbolic and deterministic manner." Its ability to perform well with reduced data is particularly impressive: "Using only 10% of the training images, our model is able to achieve comparable results 22 with the baselines trained on the full dataset." Similarly, [39] report perfect performance on small datasets which 23 they also attribute to the use of explicit and precise reasoning. Both studies display similar limitations, the use of 24 synthetic datasets, and the need for handcrafted logic, a DSL (Domain Specific Language) in the case of [57], and 25 26 Image Schemas in [39].

²⁷ Levels 4 and 5, *Neuro: Symbolic* \rightarrow *Neuro* and *Neuro_Symbolic* respectively, were originally presented by Kautz ²⁸ under one heading. After his presentation, Kautz modified the slide deck separating these two levels into systems ²⁹ where knowledge is compiled into the weights, and where knowledge is compiled into the loss function.

³⁰ Deep Learning For Mathematics [76] is an example of Level 4 where the input and output to the model are ³¹ mathematical expressions. The model performs symbolic differentiation or integration, for example, given x^2 as ³² input, the model outputs 2x. The model exploits the tree structure of mathematical expressions, which are fed into a ³³ sequence-to-sequence architecture. This seems like a particularly fitting paradigm for natural language applications ³⁴ on the basis that structures such as parse trees can be similarly leveraged to output other meaningful structures such ³⁵ as for example: cause and effect relationships, argument schemes, or rhetorical devices, to name a few.

Level 5 comprises Tensor Product Representations (TPRs)[77], Logic Tensor Networks (LTNs)[78], Neural Tensor Networks (NTN)[79] and more broadly is referred to as tensorization, where logic acts as a constraint. In Levels 4 and 5, reasoning can be performed both implicitly and explicitly, in that it is calculated via gradient descent, but can also be performed post hoc. We have grouped studies belonging to these two categories under the moniker of *Compiled* systems, of which there are ten.

An example of a *Compiled* system in our set of studies is proposed in [61] which we mentioned in Section 3.1.3. Here, knowledge is encoded in the form of huffman trees made of triples, and logic expressions, in order to jointly learn embeddings and model weights. The first layer of the tree consists of entities, the second layer consists of relations $(x \rightarrow y)$. Higher layers compute logic rules. The root node is the final embedding representing a document (in this case a single health record). Back propagation is used for optimization with softmax for calculating class probabilities. The model is intended for medical diagnosis decision support, where a much desirable characteristic is interpretability, and this model meets that goal.

Figure 18 illustrates our mapping from Kautz's levels to our proposed nomenclature. There are no studies in Level 6, nor in the ensemble major category, but which we include for completeness. Figure 19 shows the number of studies per category, and Figure 20 illustrates the relationship between categories and goals. Table 5 shows the number of studies in each category per goal.



	. 1.	Table 5	. • • •				
Number of s	studies per goa	dies per goal reported as met, in each category					
	Compiled	Cooperative	Nested	Sequential			
Reasoning	7		2				
OOD	3	2	1	1			
Interpretability	3	3	0	2			
Reduced data	2	3	1	3			
Transferability	2	1	1	0			
TOTAL	17	13	5	15			

4. Discussion

All studies report performance either on par or above benchmarks, but we can't compare studies based on performance as nearly every study uses a different dataset and benchmark as discussed Section 3.1.5. Our focus is instead on whether the goals of NeSy are being met. Our *Promise Score* metric is not necessarily what the studies' authors were optimizing for or even reporting, especially studies which have not labeled themselves as NeSy per se. So we want to make it very clear that our analysis is not a judgement of the success of any particular study, but rather we seek to understand if the hypotheses about NeSy are materializing (namely, that the combination of symbolic and sub-symbolic techniques will fulfill the goals described in Section 1.3. And the short answer is we're not there yet, as can be seen in Figure 21. For a detailed breakdown of each goal and study see Table 6.



Fig. 21. Proportion of studies which have met one or more of the 5 goals

In Section 1.3.5 we put forward the hypothesis that reasoning is the means by which the other goals can be achieved. This is not evidenced in the studies we reviewed. Some possible explanations for this finding are: 1) The kind of reasoning required to fulfill the other goals is not the kind being implemented; 2) The approaches are theoretically promising, but the technical solutions need further development. Next we look at each of these possibilities.

4.1. Reasoning Challenges

Twenty two out of the thirty four studies included for review mention reasoning as a characteristic of their solution. But there is a lot of variation in how reasoning is described and implemented. Given the overwhelming evidence of the fallibility of human reasoning, to understand language, AI researchers have sought guidance from disciplines such as psychology, cognitive linguistics, neuroscience, and philosophy. The challenge is that there are multiple competing theories of human reasoning and logic both across and within these disciplines. What we've discovered in our review, is a blurring of the lines between various types of logic, human reasoning, and

Table 6 NeSy Promises OOD Reduced Interpretability Transferability Reasoning isNeSy Generalization Data

mathematical reasoning, as well as counter-productive assumptions about which theory to adopt. For example, drawing inspiration from "how people think", accepting that how people think is flawed, and subsequently attempting to build a model with a logical component, which by definition, is rooted in validity seems counter productive to us. This logical component doesn't have to be binary or even monotonic - the sheer fact that it is implemented in code necessitates a 'valid' outcome, however that may be defined by the particular logic theory. Additionally, the justification of "because that's how people think" is inconsistent. Some examples from the studies we reviewed include:

- [63] describe human reasoning in terms of a dual process of "subsymbolic commonsense" (strongly correlated with associative learning), and "axiomatic" knowledge (predicates and logic formulas) for structured inference.

Score

Ref.

[54]

[49]

[51]

[63]

[44]

[46]

[35]

[31]

[64]

[52]

[30]

[61]

[32]

[34]

[59]

[38]

[39]

[42]

[36]

[80]

[45]

[37]

[55]

[56]

[43]

[48]

[60]

[40]

[53]

[57]

[33]

[62]

[41]

[47]

[50]

[58]

 In [34] humans reason by way of analogy, and commonsense knowledge is represented in ConceptNet, a graphical representation of common concepts and their relationships.

- For [39] human reasoning can be modeled by Image Schemas (IS). Schemas are made up of logical rules on (Entity1,Relation,Entity2) tuples, such as transitivity, or inversion.
 - [37] explain their choice of fuzzy logic for "its resemblance to human reasoning and natural language." This is a probabilistic approach which attempts to deal with uncertainty.
- [43] propose that human thought constructs can be modelled as cause-effect pairs. Commonsense is often
 described as the ability to draw causal conclusions from basic knowledge, for example: If I drop the glass, it
 will break.
- And [50] state that "when people perform explicit reasoning, they can typically describe the way to the conclusion step by step via relational descriptions."

But the most plausible hypothesis in our view is that of Schon et al.[52]: in order to emulate human reasoning, systems need to be flexible, be able to deal with contradicting evidence, evolving evidence, have access to enormous amounts of background knowledge, and include a combination of different techniques and logics. Most notably, no particular theory of reasoning is given. The argument put forward by Leslie Kaelbling at IBM Neuro-Symbolic AI Workshop 2022 is similarly appealing. She points to the over-reliance on the System1/System2 analogy, and advocates for a much more diverse and dynamic approach. We posit that the type of reasoning employed shouldn't be based solely on how we think people think, but on the attendant objective. This is in line with the "goal oriented" theory from neuroscience, in that reasoning involves many sub-systems: perception, information retrieval, decision making, planning, controlling, and executing, utilizing working memory, calculation, and pragmatics. But here the irony is not lost on us, and we acknowledge that by resorting to neuroscience for inspiration, we have just committed the same mischief for which we have been decrying our peers! But if we must resort to analogies with human reasoning then it is imperative to be as rigorous as possible. In their recent book, A FORMAL THEORY OF COMMONSENSE PSYCHOLOGY How People Think People Think [81], Gordon and Hobbs present a "large-scale logical formalization of commonsense psychology in support of humanlike artificial intelligence" to act as a baseline for researchers building intelligent AI systems. Santos et al [82] take this a step in the direction we are advocating, by testing whether there is human annotator agreement when categorizing texts into Gordon and Hobbs' theories. "Our end-goal is to advocate for better design of commonsense benchmarks [and to] support the development of a formal logic for commonsense reasoning". It is difficult to imagine a single formal logic which would afford all of Gordon and Hobbs' 48 categories of reasoning tasks. Besold at al. [19] dedicate several pages to this topic under the heading of Neural-Symbolic Integration in and for Cognitive Science: Building Mental Models. In short, computational modelling of cognitive tasks and especially language processing is still considered a hard challenge.

Others contend that to understand language, one should approach it through the lens of argumentation; that language is a communication tool, and as such should be understood as the interaction of two or more interlocutors engaged in the exercise of persuasion. Or at a minimum there is a messenger and a receiver, each with a responsibility of encoding and decoding the information exchange. While deduction and induction have traditionally been the realm of symbolic and sub-symbolic systems respectively, abduction and non-monotonic logic are the tools of computational argumentation systems. Certain of Gordon and Hobbs' theories lend themselves well to argumentation, such as Causality, Envisioning, or Explanation. None of the studies we reviewed undertook this approach. We believe this is a gap worth exploring in NeSy for NLP.

4.2. Technical challenges

There is strong agreement that a successful NeSy system will be characterized by compositionality. Compositionality allows for the construction of new meaning from learned building blocks thus enabling extrapolation beyond the training distribution. To paraphrase Garcez et al., one should be able to query the trained network using a rich description language at an adequate level of abstraction [1]. The challenge is to come up with dense/compact differentialble representations while preserving the ability to decompose, or unbind, the learned representations for downstream reasoning tasks.

⁵⁰ Bianchi et al. [63] propose LTN_{EE} , an extention of Logic Tensor Networks (LTNs), in which pre-trained ⁵¹ embeddings are fed into the LTN. They show promising results on small datasets which have the important

characteristic of being capable of after-training logical inferences. However, LTN_{EE} is limited by heavy computational requirements as the logic becomes more expressive, for example by the use of quantifiers.

Several other studies [39, 57] introduce logical inference within their solutions, but all require manually designed rules, and are limited by the domain expertise of the designer. Learning rules from data, or structure learning [83] is an ongoing research topic as pointed out by [22]. In [42] Chaturvedi et al. use fuzzy logic for emotion classification where explicit membership functions are learned. However, as stated by the authors, the classifier becomes very slow with the number of functions.

Other (compiled) approaches involve translating logic into differentialble functions, which are either directly included as network nodes as in [61], or added as a constraint to the loss function, as in [30]. To achieve this, First Order Logic (FOL) can be operationalized using t-norms for example. However, even if we can't agree on how people reason, we can probably agree how they do not, and that's with FOL. To address the many types of reasoning as discussed in the previous section, we need to be able incorporate other types of logic, such as temporal logic, modal logics, epistemic logic, non-monotonic logics and more, but which have no obvious differentiable form.

In summary, formulating logic, or more broadly reasoning, in a differentiable fashion remains challenging.

5. Limitations & Future Work

We organized our analysis according to the characteristics extracted from the studies to test whether there were any patterns leading to NeSy goals. Another approach would be to reverse this perspective, and look at each goal separately to understand the characteristics leading to its fulfillment. However, each goal is really an entire field of study in and of itself, and we don't think we could have done justice to any of them by taking this approach. We spent a lot of time looking for signal in a very noisy environment where the studies we reviewed had very little in common. More can be said about what we did not find, than what we did. Another approach might be to narrow the criteria for the type of NLP task, while expanding the technical domain. In particular, a subset of tasks from the NLU domain could be a good starting point, as these tasks are often said to require reasoning.

We tried to be comprehensive in respect to the selected studies which led to the trade-off of less space dedicated to technical details or additional context from the neuro-symbolic discussion. There are a lot of ideas and concepts which we did not cover, such as, and in no particular order, Relational Statistical Learning (RSL), Inductive Logic Programming (ILP), DeepProbLog [84], Connectionist Modal Logics (CML), Extreme Learning Machines (ELM), Genetic Programming, grounding and proposinalization, Case Based Reasoning (CBR), Abstract Meaning Representation (AMR), to name but a few, some of which are covered in detail in other surveys [4, 19].

Furthermore, we argued that we need differentiable forms of different types of logic, but we did not discuss how they might be implemented. A comprehensive point of reference such as this would be a very valuable contribution to the NeSy community, especially if the implementations were anchored in cognitive science and linguistics as discussed in 4.1.

Finally, the need for common datasets and benchmarks cannot be overstated.

6. Conclusion

We analyzed recent studies implementing NeSy for NLP in order to test whether the promises of NeSy are materializing in NLP. We attempted to find a pattern in a small and widely variable set of studies, and ultimately we do not believe there are enough results to draw definitive conclusions. Only 34 studies met the criteria for our review, and many of them (in the Sequential category) we would not consider truly integrated NeSy systems. The one thing studies which show promise [39, 47, 57, 63, 64] in meeting the specified goals have in common is that they all belong to the tightly integrated set of NeSy categories, *Cooperative* and *Compiled* which is good news for NeSy. Two out of these five report lower computational cost than baselines, and performance on par or slightly above baselines, though we must reiterate that performance comparisons are not possible as discussed in Section 3.1.5. We've seen that studies implementing both implicit and explicit reasoning meet the most NeSy goals, but suffer from high computational cost. Furthermore, explicit reasoning still often requires hand crafted domain specific rules and

K. Hamilton et al. / NeSy for NLP: Review

logic which makes them difficult to scale or generalize to other applications. Indeed, of the five goals, transferability to new domains was the least frequently satisfied.

Our view is that the lack of consensus around theories of reasoning and appropriate benchmarks is hindering our ability to evaluate progress. Hence we advocate for the development of robust reasoning theories and formal logics as well as the development of challenging benchmarks which not only measure the performance of specific implementations, but have the potential to address real world problems. Systems capable of capturing the nuances of natural language (ie., ones that "understand" human reasoning) while returning sound conclusions (ie., perform logical reasoning) could help combat some of the most consequential issues of our times such as mis- and disinformation, corporate propaganda such as in climate change denialism, divisive political speech, and other harmful rhetoric in the social discourse.

Acknowledgements

This publication has emanated from research supported in part by a grant from Science Foundation Ireland under Grant number 18/CRT/6183. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Appendix A. Acronyms

Table 7

	Acronyms and Abbreviations	
	reconjuis and reoreviations	
AI	Artificial Intelligence	
Attn	Attention Network	
CL	Curriculum Learning	
CNN	Convolutional Neural Network	
DBN	Deep Belief Network	
DL	Deep Learning	
GNN	Graph Neural Network	
LLM	Large Language Models	
LTN	Logic Tenson Network	
ML	Machine Learning	
NE	Neuroevolution	
NeSy	Neuro-Symbolic AI	
NLP	Natural Language Processing	
NLU	Natural Language Understanding	
NN	Neural Network	
OOD	Out-of-distribution	
OOP	Object-oriented programming(paradigm)	
RcNN	Recursive Neural Network	
RL	Reinforcement Learning	
RNKN	Recursive Neural Knowledge Network	
RNN	Recurrent Neural Network	
SVM	Support Vector Machine	
TF-IDF	Term frequency - Inverse document frequency	
TPR	Tensor Product Representation	
XAI	Explainable AI	

Appendix B. Allowed Values

		Allowed values
I	Feature	Allowed values
ŀ	Business application	Annotation, Constituency Parsing, Cause-effect Identification, Dialog system, Emotion detection, Fact verification, Image captioning, Information extraction, POS tagging, Question answering, Translation, KG Completion / link prediction, Language comprehension, Recommendation model, Relation extraction, Semantic Search, Semantic Similarity, Sentiment analysis, Text classification, Text games, Decision making, Topic modeling / categorization, Textual reasoning, Decision making, Text summarization
]	Technical application	Clustering, Generative, Regression, Inference, Classification, Dimentionality reduction, Recommendation, Information extraction
1	Type of learning	Supervised, Unsupervised, Semi-supervised, Reinforcement, Curriculum
1	Type of reasoning	Implicit, Explicit, Both
Ι	Language structure	Yes, No
F	Relational structure	Yes, No
١	NeSy goals	Reasoning, OOD Generalization, Interpretability, Reduced data, Transferability
ŀ	Kautz category	 symbolic Neuro symbolic, 2. Symbolic[Neuro], 3. Neuro; Symbolic, Neuro: Symbolic → Neuro, 5. Neuro_Symbolic, 6. Neuro[Symbolic]
r	NeSv category	Sequential Cooperative Compiled Nested
Refe	erences	
[1] _	erences A.d. Garcez and L.C. I	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020).
(1] [2] [3] [3]	e rences A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazee	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). oblems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i>
[1] [2] [3] [4]	e rences A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i>	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). oblems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i> processing systems, 2017, pp. 5998–6008.
[1] [2] [3] [4]	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). oblems in Computer Science, J. ACM 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance processing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine L carping and Reasoning. <i>arXiv:1905.06088 [cs]</i> (2019)
(1) (2) (2) (2) (2) (2) (2) (2) (2) (2) (2	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical	 Lamb, Neurosymbolic AI: The 3rd Wave, arXiv:2012.05876 [cs] (2020). Joblems in Computer Science, J. ACM 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance processing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, arXiv:1905.06088 [cs] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, AI magazine 12(2) (1991), 34–34.
(1) (2) (3) (4) (4) (5) (6)	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus	 Lamb, Neurosymbolic AI: The 3rd Wave, arXiv:2012.05876 [cs] (2020). bblems in Computer Science, J. ACM 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance processing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, arXiv:1905.06088 [cs] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, AI magazine 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence
[1] 1 [2] 1 [3] 1 [4] 1 [5] 1 [6] 1 [7] 1	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/.	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). oblems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i> <i>processing systems</i> , 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088 [cs]</i> (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12 (2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificial.letelligang
(1) [2] 1 [3] 2 [4] 2 [5] 1 [6] 3 (7] 1 [7] 1	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeet <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York NY https://	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876</i> [<i>cs</i>] (2020). oblems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i> <i>processing systems</i> , 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088</i> [<i>cs</i>] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12 (2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/
[1] 1 [2] 1 [3] 1 [4] 1 [5] 1 [6] 1 [7] 1 [8] 1	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou,	 Lamb, Neurosymbolic AI: The 3rd Wave, arXiv:2012.05876 [cs] (2020). bblems in Computer Science, J. ACM 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance processing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology fc of Machine Learning and Reasoning, arXiv:1905.06088 [cs] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, AI magazine 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligencc Usummer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, arXiv:2105.05330 [cs] (2021).
[1] 1 [2] 1 [3] 1 [4] 1 [5] 1 [6] 1 [7] 1 [8] 1	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 D	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). bblems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i> <i>processing systems</i> , 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088 [cs]</i> (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12 (2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, <i>arXiv:2105.05330 [cs]</i> (2021) ep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internation
(1) (2) (1) (2) (2) (2) (2) (2) (2) (2) (2) (2) (2	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Du Joint Conference on Ar	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). oblems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i> <i>processing systems</i> , 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088 [cs]</i> (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12 (2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence 'www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, <i>arXiv:2105.05330 [cs]</i> (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internationartificial Intelligence. <i>trificial Intelligence. https://ijcai-21.org/invited-talks/.</i> <i>a fast and slow</i> . Farrar, Straus and Giroux, New York, 2011, ISBN 9780374275631 0374275637
[1] 1 [2] 1 [3] 1 [4] 1 [5] 1 [6] 1 [7] 1 [8] 1 [9] 1 [1] 0	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeet <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Du Joint Conference on Al Dint Conference on Al D. Kahneman, <i>Thinkin</i> , G. Marcus and E. Dav	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). bolems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i> <i>processing systems</i> , 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088 [cs]</i> (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12 (2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, <i>arXiv:2105.05330 [cs]</i> (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internation rificial Intelligence. https://jcai-21.org/invited-talks/. <i>g. fast and slow</i> , Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about MIT Technology Review
(1) 2 (2) 1 (3) 2 (3) 2 (4) 2 (5) 1 (5) 1 (5) 1 (6) 2 (7) 1 (7) 1 (7) 1 (8) 1 (7) 1	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Do Joint Conference on Ar D. Kahneman, <i>Thinkin</i> G. Marcus and E. Dav https://www.technolog	Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). bblems in Computer Science, <i>J. ACM</i> 50 (1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance</i> <i>processing systems</i> , 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088 [cs]</i> (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12 (2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, <i>arXiv:2105.05330 [cs]</i> (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internation rtificial Intelligence. https://jcai-21.org/invited-talks/. <i>g, fast and slow</i> , Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about MIT Technology Review yreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/.
[1] [2]] [3] [4] [5]] [6] [7]] [8]] [9] [1] [1] [2]	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 D Joint Conference on A: D. Kahneman, <i>Thinkin</i> , G. Marcus and E. Dav https://www.technolog, P. Hitzler, F. Bianchi,	 Jamb, Neurosymbolic AI: The 3rd Wave, arXiv:2012.05876 [cs] (2020). Joblems in Computer Science, J. ACM 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance mocessing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, arXiv:1905.06088 [cs] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, AI magazine 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence Versus. Cochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, arXiv:2105.05330 [cs] (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internation trificial Intelligence. https://jcai-21.org/invited-talks/. g, fast and slow, Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about MIT Technology Review yreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/. M. Eberahimi and M.K. Sarker, Neural-symbolic integration and the Semantic Web, Semantic Web 11(1) (2020)
[1] 1 [2] 1 [3] 1 [4] 1 [5] 1 [6] 1 [7] 1 [8] 1 [9] 1 [1] 0 [1] 1 [2] 1	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Do Joint Conference on Ai D. Kahneman, <i>Thinkin</i> , G. Marcus and E. Dav https://www.technolog P. Hitzler, F. Bianchi, 3–11–. doi:10.3233/SV	 Lamb, Neurosymbolic AI: The 3rd Wave, arXiv:2012.05876 [cs] (2020). Soblems in Computer Science, J. ACM 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance mocessing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology fc of Machine Learning and Reasoning, arXiv:1905.06088 [cs] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, AI magazine 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, arXiv:2105.05330 [cs] (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internationartificial Intelligence. https://j.gi-21.org/invited-talks/. g, fast and slow, Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about MIT Technology Review yreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/. M. Ebrahimi and M.K. Sarker, Neural-symbolic integration and the Semantic Web, Semantic Web 11(1) (2020 V-190368.
[1] 2 [2] 1 [2] 1 [3] 2 [4] 2 [5] 1 [6] 3 [7] 1 [8] 1 [9] 2 [11] 0 [12] 1 [3] 2	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeer <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Do Joint Conference on Ar D. Kahneman, <i>Thinkin</i> , G. Marcus and E. Dav https://www.technolog P. Hitzler, F. Bianchi, 3–11–. doi:10.3233/SW	 Lamb, Neurosymbolic AI: The 3rd Wave, arXiv:2012.05876 [cs] (2020). bblems in Computer Science, J. ACM 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance processing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, arXiv:1905.06088 [cs] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, AI magazine 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, arXiv:2105.05330 [cs] (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internationartificial Intelligence. https://jcai-21.org/invited-talks/. g. fast and slow, Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about 1 MIT Technology Review yreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/. M. Ebrahimi and M.K. Sarker, Neural-symbolic integration and the Semantic Web, Semantic Web 11(1) (2020 V-190368. <i>p. Dimensions of Neural-symbolic Integration – A Structured Survey, in: We Will Show Them! Essays in Honour One, Sne, S.N.</i>, College Publications, 2005.
[1] 2 [2] 1 [3] 2 [4] 2 [6] 1 [6] 1 [7] 1 [8] 1 [9] 1 [2] 1 [3] 2 [4] 1	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeet <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Du Joint Conference on Ai Doint Conference on Ai Doint Conference on Ai Doint Conference on Ai D. Kahneman, <i>Thinkin</i> G. Marcus and E. Dav https://www.technolog P. Hitzler, F. Bianchi, 3–11–. doi:10.3233/SW S. Bader and P. Hitzler Dov Gabbay, Volume C B. Hammer and P. Hitz	 Lamb, Neurosymbolic AI: The 3rd Wave, arXiv:2012.05876 [cs] (2020). bblems in Computer Science, J. ACM 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: Advance processing systems, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, arXiv:1905.06088 [cs] (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, AI magazine 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, arXiv:2105.05330 [cs] (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internationartificial Intelligence. https://ijcai-21.org/invited-talks/. g. fast and slow, Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about MIT Technology Review yreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/. M. Ebrahimi and M.K. Sarker, Neural-symbolic integration and the Semantic Web, Semantic Web 11(1) (2020 V-190368. <i>t</i>, Dimensions of Neural-symbolic Integration – A Structured Survey, in: We Will Show Them! Essays in Honour of Dae, S.N. College Publications, 2005. ther (eds), Perspectives of Neural-Symbolic Integration, Vol. 77, Springer, 2007. ISBN 978-3-540-73953-1.
fe 1 <t< td=""><td>A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeet <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Du Joint Conference on Al Dint Conference</td><td> Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). biblems in Computer Science, <i>J. ACM</i> 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance maccessing systems</i>, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088 [cs]</i> (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, <i>arXiv:2105.05330 [cs]</i> (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internationartificial Intelligence. https://ijcai-21.org/invited-talks/. <i>g, fast and slow</i>, Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about 1 MIT Technology Review yreview.com/2020/08/221/007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/. M. Ebrahimi and M.K. Sarker, Neural-symbolic integration and the Semantic Web, <i>Semantic Web</i> 11(1) (2020 V-190368. <i>c, Dimensions of Neural-symbolic Integration – A Structured Survey, in: We Will Show Them! Essays in Honour of Dine, S.N.</i>, College Publications, 2005. Let (eds), <i>Perspectives of Neural-Symbolic Integration</i>, Vol. 77, Springer, 2007. ISBN 978-3-540-7395-1. th and D.M. Gabbay, <i>Neural-Symbolic Integration</i>, Vol. 77, Springer, 2007. ISBN 978-3</td></t<>	A.d. Garcez and L.C. I L.G. Valiant, Three Pro A. Vaswani, N. Shazeet <i>in neural information p</i> A.d. Garcez, M. Gori, Principled Integration of M.L. Minsky, Logical Y. Bengio, G. Marcus com/aidebate/. H. Kautz, The Third Al New York, NY. https:// M.K. Sarker, L. Zhou, Y. Bengio, System 2 Du Joint Conference on Al Dint Conference	 Lamb, Neurosymbolic AI: The 3rd Wave, <i>arXiv:2012.05876 [cs]</i> (2020). biblems in Computer Science, <i>J. ACM</i> 50(1) (2003), 96–99–. doi:10.1145/602382.602410. r, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, in: <i>Advance maccessing systems</i>, 2017, pp. 5998–6008. L.C. Lamb, L. Serafini, M. Spranger and S.N. Tran, Neural-Symbolic Computing: An Effective Methodology for of Machine Learning and Reasoning, <i>arXiv:1905.06088 [cs]</i> (2019). versus analogical or symbolic versus connectionist or neat versus scruffy, <i>AI magazine</i> 12(2) (1991), 34–34. and V. Boucher, AI DEBATE! Yoshua Bengio vs Gary Marcus, Montreal.AI. https://montrealartificialintelligence I Summer, AAAI Robert S. Engelmore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence www.cs.rochester.edu/u/kautz/talks/. A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, <i>arXiv:2105.05330 [cs]</i> (2021) eep Learning: Higher-Level Cognition, Agency, Out-of-Distribution Generalization and Causality, 30th Internationartificial Intelligence. https://ijcai-21.org/invited-talks/. <i>g, fast and slow</i>, Farrar, Straus and Giroux, New York, 2011. ISBN 9780374275631 0374275637. is, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about 1 MIT Technology Review yreview.com/2020/08/221/007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/. M. Ebrahimi and M.K. Sarker, Neural-symbolic integration and the Semantic Web, <i>Semantic Web</i> 11(1) (2020 V-190368. <i>c, Dimensions of Neural-symbolic Integration – A Structured Survey, in: We Will Show Them! Essays in Honour of Dine, S.N.</i>, College Publications, 2005. Let (eds), <i>Perspectives of Neural-Symbolic Integration</i>, Vol. 77, Springer, 2007. ISBN 978-3-540-7395-1. th and D.M. Gabbay, <i>Neural-Symbolic Integration</i>, Vol. 77, Springer, 2007. ISBN 978-3

1	[16]	T.R. Besold and KU. Kühnberger, Towards integrated neural-symbolic systems for human-level AI: Two research programs helping to	1
2	[17]	bridge the gaps, <i>Biologically Inspired Cognitive Architectures</i> 14 (2015), 97–110–. doi:10.1016/j.bica.2015.09.003.	2
3	[1/]	L. Gabriovich, K. Guna, A. McCanum and K. Murphy, Knowledge Representation and Reasoning: Integrating Symbolic and Neural	3
4	[18]	Ad Garcez TR Besold L De Raedt P Földiak P Hitzler T Icard K -U Kühnberger L C Lamb R Miikkulainen and D L Silver	4
5	[10]	Neural-symbolic learning and reasoning: contributions and challenges, in: 2015 AAAI Spring Symposium Series, 2015.	5
6	[19]	T.R. Besold, A.d. Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, KU. Kühnberger, L.C. Lamb, D. Lowd, P.M.V. Lima et al.,	6
7		Neural-symbolic learning and reasoning: A survey and interpretation, arXiv preprint arXiv:1711.03902 (2017).	7
8	[20]	V. Belle, Symbolic Logic Meets Machine Learning: A Brief Survey in Infinite Domains, in: Scalable Uncertainty Management, Springer	8
9		International Publishing, 2020, pp. 3–16–. ISBN 978-3-030-58449-8.	9
10	[21]	L.C. Lamb, A.d. Garcez, M. Gori, M.O.R. Prates, P.H.C. Avelar and M.Y. Vardi, Graph Neural Networks Meet Neural-	1.0
10		Symbolic Computing: A Survey and Perspective, in: Proceedings of the Twenty-Ninth International Joint Conference on Artificial	10
11		Intelligence, international joint Conferences on Artificial Intelligence Organization, 2020, pp. 4877–4884–. ISBN 978-0-9992411-0-5.	11
12	[22]	L von Rueden S Mayer K Beckh B Georgiev S Giesselbach R Heese B Kirsch M Walczak I Pfrommer A Pick R Ramamurthy	12
13	[22]	J. Garcke, C. Bauckhage and J. Schuecker. Informed Machine Learning - A Taxonomy and Survey of Integrating Prior Knowledge into	13
14		Learning Systems, <i>IEEE Transactions on Knowledge and Data Engineering</i> (2021), 1–1–. doi:10.1109/TKDE.2021.3079836.	14
15	[23]	D. Yu, B. Yang, D. Liu and H. Wang, A Survey on Neural-symbolic Systems, arXiv:2111.08164 [cs] (2021).	15
16	[24]	J. Zhang, B. Chen, L. Zhang, X. Ke and H. Ding, Neural, symbolic and neural-symbolic reasoning on knowledge graphs, AI Open 2 (2021),	16
17		14–35–. doi:10.1016/j.aiopen.2021.03.001.	17
18	[25]	E. Tsamoura, T. Hospedales and L. Michael, Neural-Symbolic Integration: A Compositional Perspective, Proceedings of the AAAI	18
19		Conference on Artificial Intelligence 35 (66) (2021), 5051–5060–.	19
20	[26]	Reasoning, Encyclopædia Britannica, inc. https://www.britannica.com/technology/artificial-intelligence/Reasoning.	20
21	[27]	B. Kitchennam and S. Charters, Guidelines for performing Systematic Literature Reviews in Software Engineering, 2007.	20
21	[20]	<i>& Management</i> 52 (2) (2015) 183–199– doi:https://doi.org/10.1016/i im 2014.08.008	21
22	[29]	M I Page JE McKenzie PM Bossuvt I Boutron TC Hoffmann CD Mulrow L Shamseer JM Tetzlaff E A Akl SE Brennan	22
23	[=>]	R. Chou, J. Glanville, J.M. Grimshaw, A. Hróbjartsson, M.M. Lalu, T. Li, E.W. Loder, E. Mayo-Wilson, S. McDonald, L.A. McGuinness,	23
24		L.A. Stewart, J. Thomas, A.C. Tricco, V.A. Welch, P. Whiting and D. Moher, The PRISMA 2020 statement: an updated guideline for	24
25		reporting systematic reviews, Systematic Reviews 10(1) (2021), 89. doi:10.1186/s13643-021-01626-4.	25
26	[30]	M. Diligenti, M. Gori and C. Saccà, Semantic-based regularization for learning and inference, Artificial Intelligence 244 (2017), 143–165.	26
27		doi:10.1016/j.artint.2015.08.011.	27
28	[31]	J. Silva, L. Coheur, A.C. Mendes and A. Wichert, From symbolic to sub-symbolic information in question classification, <i>Artificial</i>	28
29	[20]	Intelligence Review 35 (2) (2011), 137–154. doi:10.1007/s10462-010-9188-4.	29
30	[32]	(2016) 43 56 doi:10.1016/j.eswo.2016.01.007	30
31	[33]	I D'Souza I O Mulang and S. Auer. Team SVMrank: Leveraging feature-rich support vector machines for ranking explanations to	.31
32	[00]	elementary science questions, 2019, pp. 90–100.	32
22	[34]	A. Hussain and E. Cambria, Semi-supervised learning for big social data analysis, <i>Neurocomputing</i> 275 (2018), 1662–1673.	22
23		doi:10.1016/j.neucom.2017.10.010.	33
34	[35]	Q. Cui, Y. Zhou and M. Zheng, Sememes-Based Framework for Knowledge Graph Embedding with Comprehensive-Information, Lecture	34
35		Notes in Computer Science 12816 LNAI (2021), 419–426. doi:10.1007/978-3-030-82147-0_34.	35
36	[36]	M. Bounabi, K. Elmoutaouakil and K. Satori, A new neutrosophic TF-IDF term weighting for text mining tasks: text classification use case,	36
37	[27]	International Journal of Web Information Systems 17(3) (2021), 229–249. doi:10.1108/IJW1S-11-2020-0067.	37
38	[37]	F. ES-Sabery, A. Hair, J. Qadir, B. Sainz-De-Adajo, B. Garcia-Zapirain and I. Torre-Diez, Sentence-Level Classification Using Parallel Eurzy Deep Learning Classifier <i>IEEE Access</i> 9 (2021), 17043–17085. doi:10.1100/ACCESS.2021.3053017	38
39	[38]	A Mehler A Lücking and P Menke Assessing cognitive alignment in different types of dialog by means of a network model <i>Neural</i>	39
40	[50]	Networks 32 (2012), 159–164. doi:10.1016/i.neunet.2012.02.013.	40
41	[39]	Y. Yao, J. Xu, J. Shi and B. Xu, Learning to activate logic rules for textual reasoning, <i>Neural Networks</i> 106 (2018), 42–49.	41
42		doi:10.1016/j.neunet.2018.06.012.	42
4.3	[40]	Q. Chen, A. Lamoreaux, X. Wang, G. Durrett, O. Bastani and I. Dillig, Web question answering with neurosymbolic program synthesis,	43
4.4		2021, pp. 328–343. doi:10.1145/3453483.3454047.	44
45	[41]	A. Tato, R. Nkambou and A. Dufresne, Hybrid deep neural networks to predict socio-moral reasoning skills, 2019, pp. 623–626.	11
10	[42]	I. Unaturvedi, K. Satapathy, S. Cavallari and E. Cambria, Fuzzy commonsense reasoning for multimodal sentiment analysis, <i>Pattern</i>	40
40	[43]	Recognition Letters 123 (2019), 204–270. doi:10.1010/j.patrec.2019.04.024.	46
4 /	[+3]	doi:10.1109/INDICON47234.2019.9028985	4.7
48	[44]	M. Ezzat, A Framework for a Comprehensive Conceptualization of Urban Constructs, 2020, pp. 111–120.	48
49	[45]	J. Gong, H. Ma, Z. Teng, Q. Teng, H. Zhang, L. Du, S. Chen, M.Z.A. Bhuiyan, J. Li and M. Liu, Hierarchical Graph	49
50	-	Transformer-Based Deep Learning Model for Large-Scale Multi-Label Text Classification, IEEE Access 8 (2020), 30885-30896.	50
51		doi:10.1109/ACCESS.2020.2972751.	51

[46]	H. Lemos, P. Avelar, M. Prates, A. Garcez and L. Lamb, Neural-Symbolic Relational Reasoning on Graph Models: Effective Link Inference and Computation from Knowledge Bases, <i>Lecture Notes in Computer Science</i> 12396 LNCS (2020), 647–659. doi:10.1007/978-3-030-61600.0.51	1 2
F 4 77 1		3
[47]	S. Huo, I. Ma, J. Chen, M. Chang, L. Wu and M. Wildowck, Graph ennanced cross-domain text-to-SQL generation, 2019, pp. 159–165.	4
[48]	M. Zhou, D. J and F. Li, Kelation Extraction in Dialogues: A Deep Learning Model Based on the Generatily and Specialty of Dialogue The IEEE/ACM Transactions on Automatic Special and Learning Deep Learning Model Based on the Generatily and Specialty of Dialogue	5
F401	Text, IEEE/ACM Transactions on Auato Speech and Language Processing 29 (2021), 2015–2020. doi:10.1109/1ASLP.2021.3082295.	6
[49]	A.M.P. Braşoveanu and K. Andome, Semanuc rake News Detection: A Machine Learning Perspective, <i>Lecture Notes in Computer Science</i>	-
[50]	HS00 LINCS (2019), 030-007. doi:10.1007/978-3-030-20321-20321-204.	1
[30]	K. Chen, Q. Huang, H. Parangi, F. Sinotensky, K.D. Polous and J. Gao, Mapping natural-ranguage problems to formal-ranguage solutions	8
[51]	using structured neural representations, 2020, pp. 1544–1555.	9
[31]	<i>Science</i> 11730 LNCS (2019), 185–197. doi:10.1007/978-3-030-30490-4_16.	10
[52]	C. Schon, S. Stebert and F. Stolzenburg, The CoRg Project: Cognitive Reasoning, KI - Kunstliche Intelligenz 33(3) (2019), 293–299. doi:10.1007/s13218-019-00601-5.	12
[53]	L.B. Fazlic, A. Hallawa, A. Schmeink, A. Peine, L. Martin and G. Dartmann, A novel NLP-fuzzy system prototype for information	13
	extraction from medical guidelines, 2019, pp. 1025–1030. doi:10.23919/MIPRO.2019.8756929.	14
[54]	N. Ramrakhiyani, G. Palshikar and V. Varma, A simple neural approach to spatial role labelling, <i>Lecture Notes in Computer Science</i> 11438 LNCS (2019), 102–108. doi:10.1007/978-3-030-15719-7_13.	15
[55]	K. Amin, S. Kapetanakis, KD. Althoff, A. Dengel and M. Petridis, Cases without borders: Automating knowledge acquisition approach	16
	using deep autoencoders and siamese networks in case-based reasoning, 2019, pp. 133-140. doi:10.1109/ICTAI.2019.00027.	17
56]	A. Sutherland, S. Magg and S. Wermter, Leveraging Recursive Processing for Neural-Symbolic Affect-Target Associations, 2019.	18
	doi:10.1109/IJCNN.2019.8851875.	19
[57]	J. Mao, C. Gan, P. Kohli, J.B. Tenenbaum and J. Wu, The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from	2.0
	natural supervision, 2019.	21
58]	D. Demeter and D. Downey, Just Add Functions: A Neural-Symbolic Language Model, 2020, pp. 7634–7642.	21
59]	W. Liu, J. Tang, X. Liang and Q. Cai, Heterogeneous graph reasoning for knowledge-grounded medical dialogue system, <i>Neurocomputing</i>	22
	442 (2021), 260–268. doi:10.1016/j.neucom.2021.02.021.	23
60]	P. Manda, S. Sayedahmed and S.D. Mohanty, Automated ontology-based annotation of scientific literature using deep learning, 2020.	24
	doi:10.1145/3391274.3393636.	25
51]	J. Jiang, H. Wang, J. Xie, X. Guo, Y. Guan and Q. Yu, Medical knowledge embedding based on recursive neural network for multi-disease	26
	diagnosis, Artificial Intelligence in Medicine 103 (2020). doi:10.1016/j.artmed.2019.101772.	27
52]	Q. Huang, L. Deng, D. Wu, C. Liu and X. He, Attentive tensor product learning, 2019, pp. 1344–1351.	27
53]	F. Bianchi, M. Palmonari, P. Hitzler and L. Serafini, Complementing logical reasoning with sub-symbolic commonsense, <i>Lecture Notes in Computer Science</i> 11784 LNCS (2019), 161–170. doi:10.1007/978-3-030-31095-0_11.	28 29
64]	B. Škrlj, M. Martinc, N. Lavrač and S. Pollak, autoBOT: evolving neuro-symbolic representations for explainable low resource text	30
	classification, Machine Learning 110(5) (2021), 989–1028. doi:10.1007/s10994-021-05968-x.	31
65]	M. Kang and N.J. Jameson, Machine Learning: Fundamentals, Prognostics and Health Management of Electronics: Fundamentals,	30
	Machine Learning, and the Internet of Things (2018), 85–109.	52
56]	G. Bonaccorso, Machine learning algorithms, Packt Publishing Ltd, 2017.	33
67]	Y. Bengio, J. Louradour, R. Collobert and J. Weston, Curriculum learning, in: Proceedings of the 26th annual international conference on	34
	machine learning, 2009, pp. 41–48.	35
68]	L. Seratini and A.d. Garcez, Logic Tensor Networks: Deep Learning and Logical Reasoning from Data and Knowledge, <i>arXiv:1606.04422</i>	36
<i></i>	[<i>cs</i>] (2016).	37
69]	R. Milkkulainen, Neuroevolution, in: Encyclopedia of Machine Learning, Springer, New York, 2010.	38
70]	J. Lehman and R. Miikkulainen, Neuroevolution, <i>Scholarpedia</i> 8 (6) (2013), 30977. doi:10.4249/scholarpedia.30977.	
71]	R. Davis, H. Shrobe and P. Szolovits, What is a knowledge representation?, AI magazine 14(1) (1993), 17–17.	39
72]	T.J. Bench-Capon, <i>Knowledge representation: An approach to artificial intelligence</i> , Vol. 32, Elsevier, 2014.	40
73]	H.J. Levesque, Knowledge representation and reasoning, <i>Annual review of computer science</i> 1 (1) (1986), 255–287.	41
/4]	I.L. Travis, Knowledge Representation in Artificial Intelligence, <i>Clinic on Library Applications of Data Processing (27th: 1990)</i> (1990).	42
75]	E. Shortliffe, <i>Computer-based medical consultations: MYCIN</i> , Vol. 2, Elsevier, 2012.	43
/6]	G. Lample and F. Charton, Deep Learning for Symbolic Mathematics, 2019.	44
//]	46 (1) (1990), 159–216–. doi:https://doi.org/10.1016/0004-3702(90)90007-M.	45
[78]	L. Serafini and A.S. d'Avila Garcez, Learning and Reasoning with Logic Tensor Networks, in: AI*IA 2016 Advances in Artificial	46
	Intelligence, G. Adorni, S. Cagnoni, M. Gori and M. Maratea, eds, Lecture Notes in Computer Science, Springer International Publishing,	47
	2016, pp. 334–348–. ISBN 978-3-319-49130-1. doi:10.1007/978-3-319-49130-1_25.	48
[79]	R. Socher, D. Chen, C.D. Manning and A. Ng, Reasoning With Neural Tensor Networks for Knowledge Base Completion, in: <i>Advances in</i>	49
0.01	Neural Information Processing Systems, Vol. 26, Curran Associates, Inc., 2013.	50
80]	H. Honda and M. Hagiwara, Question Answering Systems with Deep Learning-Based Symbolic Processing, <i>IEEE Access</i> 7 (2019),	50
	152508-152578. doi:10.1109/ACCESS.2019.2948081.	51

K. Hamilton et al. / NeSy for NLP: Review

1	[81]	A.S. Gordon and J.R. Hobbs, A Formal Theory of Commonsense Psychology: How People Think People Think, Cambridge University Press,	1
2	1021	2017. doi:10.1017/9781316584705.	2
3	[02]	categorization agreement on commonsense sentences. <i>Experimental Results</i> 2 (2021), e19. doi:10.1017/exp.2021.9.	3
4	[83]	V. Embar, D. Sridhar, G. Farnadi and L. Getoor, Scalable Structure Learning for Probabilistic Soft Logic, <i>arXiv:1807.00973 [cs, stat]</i>	4
5		(2018).	5
6	[84]	R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester and L. De Raedt, Deepproblog: Neural probabilistic logic programming, <i>Advances</i>	6
7		in Neural Information Processing Systems 31 (2018).	7
8			8
9			9
10			10
11			11
12			12
13			13
14			14
15			15
10			10
10			10
10			10
20			20
21			20
22			22
23			23
24			24
25			25
26			26
27			27
28			28
29			29
30			30
31			31
32			32
33			33
34			34
35			35
36			36
37			37
38			38
39			39
40			40
41			41
42			42
43			43
44			44
45			45
46			46
47			47
48			48
49			49
50			50
JL			51