



OPEN ACCESS

EDITED BY

Jeff Orchard,
University of Waterloo, Canada

REVIEWED BY

J. Swaroop Guntupalli,
DeepMind Technologies Limited,
United Kingdom

*CORRESPONDENCE

Giorgio Gronchi
✉ giorgio.gronchi@unifi.it

RECEIVED 16 December 2023

ACCEPTED 16 February 2024

PUBLISHED 05 March 2024

CITATION

Gronchi G and Perini A (2024) Dual-process theories of thought as potential architectures for developing neuro-symbolic AI models.

Front. Cognit. 3:1356941.

doi: 10.3389/fcogn.2024.1356941

COPYRIGHT

© 2024 Gronchi and Perini. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Dual-process theories of thought as potential architectures for developing neuro-symbolic AI models

Giorgio Gronchi* and Axel Perini

Section of Psychology, Department of Neuroscience, Psychology, Drug Research and Child's Health, University of Florence, Florence, Italy

KEYWORDS

ChatGPT, large language models, neuro-symbolic AI, dual-process theory of thought, subsymbolic modeling, symbolic modeling

1 Introduction

In recent years, subsymbolic-based artificial intelligence has developed significantly, both from a theoretical and an applied perspective. OpenAI's Chat Generative Pre-trained Transformer (ChatGPT) was launched on November 2022 and became the consumer software application with the quickest growth rate in history (Hu, 2023).

ChatGPT is a large language model (LLM) constructed using either GPT-3.5 or GPT-4, built upon Google's transformer architecture. It is optimized for conversational use through a blend of supervised and reinforcement learning methods (Liu et al., 2023). These models' capabilities include text generation that human evaluators find challenging to differentiate from human-written content (Brown et al., 2020), code computer programs (Chen et al., 2021), and engage in conversations with humans on various subjects (Lin et al., 2020). However, due to the statistical nature of LLMs, they face significant limitations when handling structured tasks that rely on symbolic reasoning (Binz and Schulz, 2023; Chen X. et al., 2023; Hammond and Leake, 2023; Titus, 2023). For example, ChatGPT 4 (with a Wolfram plug-in that allows to solve math problems symbolically) when asked (November 2023) "How many times does the digit 9 appear from 1 to 100?" correctly responds 20 times. Nevertheless, if we say that the answer is wrong and there are 19 digits, the system corrects itself and confirms that there are indeed 19 digits. This simple example testifies the intrinsic difficulties of probabilistic fluency models to deal with mere facts (they can only suggest assertions based on likelihood, and in various instances, they might modify the assertion, see Hammond and Leake, 2023). Although some papers attempted to demonstrate that LLMs alone can solve structured problems without any integration (Noever et al., 2020; Drori et al., 2022), a promising way to address these problems is to integrate systems like chatGPT with symbolic systems (Bengio, 2019; Chaudhuri et al., 2021). A classic problem is how the two distinct systems may interact (Smolensky, 1991). In this opinion paper, we propose that the dual-process theory of thought literature (De Neys, 2018) can provide human cognition-inspired solutions on how two distinct systems, one based on statistic (subsymbolic system) and the other on structured reasoning (symbolic), can interact. In the following, after a brief description of the structure/statistics debate in cognitive science that mirrors the discussion about potentialities and limitations of LLMs (taken as a prototypical example of a subsymbolic model), we propose how different instances of the dual-process theory of thought may serve as potential architectures for hybrid symbolic-subsymbolic models.

2 The structure/statistics dilemma

In the cognitive literature of the 80's, the coexistence of the classic symbolic approach in psychology with the then-emerging subsymbolic approach (based mainly on neural networks) led to an intense debate between these alternative views for modeling the human mind. Smolensky (1991) referred to this dualism as the paradox of cognition. He wrote: "On the one hand, cognition is hard, characterized by the rules of logic [...]. On the other hand, cognition is soft: if you write down the rules, it seems that realizing those rules in automatic formal systems (which AI programs are) gives systems that are just not sufficiently fluid, not robust enough in performance to constitute what we call true intelligence" (p. 282). Of course, Smolensky here was referring to GOFAI¹ and 80's AI technologies such as expert systems. He continued: "In attempting to characterize the laws of cognition, we are pulled in two different directions: when we focus on rules governing high-level cognitive competence, we are pulled toward structured, symbolic representations and processes; when we focus on the variance and complex detail of the real intelligent performance, we are pulled toward statistical, numerical descriptions" (p. 282). Many years have passed since Smolensky wrote these words, and in the meantime, subsymbolic, statistical-based artificial intelligence systems have made incredible progress, which was unthinkable in the early 90's. We think that the limitations of LLMs are again based on this paradox, also called the structure/statistics dilemma. Being probabilistic and statistical-based systems, LLMs are incredibly efficient in handling environmental regularities. However, their weak points are structured, symbolic tasks: for example, LLMs fail in identify conclusions as definitive because they are not based on logic (another case are AI prompt image generators that may produce hands with a different number of fingers than 5, because they do not know symbolically that humans have 5 fingers unless illness or other issues). Smolensky (1991) also summarized the potential solutions of the paradox. Here, we will mention some of them. A first stance of solutions is based on the concept of emergent property, either the soft that emerges from the hard or vice versa. The former means that the essence of cognition is logic and rule-based, in this case soft properties come out when there are many complex rules. The latter means that cognition is intrinsically soft, and rules are an emergent property by means of self-regulation. An alternative way is the so-called *cohabitation* approach (Smolensky, 1991): in this view, cognition is characterized by a soft and hard machine that are one next to the other. This means that our mind includes soft, subsymbolic modules and, at the same time, rule-based modules with some sort of communication mechanism among them. However, this solution opens the door to other questions: what does it mean "one next to other"? How does communication between these modules work? The answers can come from the dual-process theory of thought. This theory, which has undergone significant development over the last 40 years, mirrors the hard/soft distinction (Sloman, 1996; Bengio, 2019). Indeed, the two facets of the cognition paradox are reflected in the two components of the dual-process theory. The first relies mainly on associative principles, encoding and

processing statistical patterns within its surroundings, including frequencies and correlations among different world features (the soft facet of cognition). Instead, the second is rule-based and thus allows to handle symbolic and abstract structures and draw logical inferences.

3 The dual-process theory of thought

A large body of cognitive research (De Neys, 2018) posits that our thought is composed by two processes. The first is commonly characterized as rapid, effortless, automatic, and associative-driven, while the second is deliberate, effortful, controlled, and rule-bound. Based on this distinction, comprehensive theories of mental architecture have been developed. Depending on theoretical distinctions, the two processes are given different names: the former has been called fast thinking, System 1, or the associative process, whereas the latter has been referred to as slow thinking, System 2, or the deliberative system. Among the numerous tasks employed in this literature (which encompasses a substantial portion of research on thinking and decision-making, see Evans, 2000; Evans and Curtis-Holmes, 2005; Ball et al., 2018), a well-established line of research has affirmed that System 2 is responsible for evaluating the logical structure of an argument, independent of its content (Type 2 responses). Meanwhile, System 1 automatically generates a sense of believability for conclusions deemed acceptable (e.g., all crows have wings) and rejects those deemed unbelievable (e.g., all apples are meat products), independently from logical validity (Type 1 responses). When believability conflicts with logical validity, a tension arises between these two forms of thinking. This conflict and broader questions about how these two systems interact form the central focus of the ongoing debate in the dual-process theory literature. Three types of solutions have been proposed: serial models, parallel models and hybrid models.

Serial models, such as the Default-Interventionist model by De Neys and Glumicic (2008) and Evans and Stanovich (2013), assume that System 1 operates as the default mode for generating responses. Subsequently, System 2 may come into play, potentially intervening, provided there are sufficient cognitive resources available. This engagement of System 2 only takes place after System 1 has been activated and is not guaranteed. In this model, individuals are viewed as cognitive misers seeking to minimize cognitive effort (Kahneman, 2011).

Conversely, in parallel models (Denes-Raj and Epstein, 1994; Sloman, 1996) both systems occur simultaneously, with a continuous mutual monitoring. So, System 2-based analytic considerations are taken into account right from the start and detect possible conflicts with the Type 1 processing.

Another perspective is offered by the hybrid models. De Neys and Glumicic (2008) attempted to solve the conflict detection issue between the two forms of thinking dividing System 2 into two distinct processes: an always active, shallow, analytical monitoring process and an optional, deeper, slower process. The former detects potential conflicts with System 1 and activates the latter if necessary (De Neys and Glumicic, 2008; Thompson, 2013; Newell et al., 2015). Some years later, De Neys (2012, 2014) developed the logical intuition model. According to it, responses commonly considered to be computed by System 2 can also be prompted

¹ This acronym stands for Good Old-Fashioned Artificial Intelligence (Haugeland, 1985).

TABLE 1 Dual-process theory of thought models and examples of similar approaches in the neuro-symbolic AI domain (described by Chaudhuri et al., 2021; Manhaeve et al., 2022).

Model	Examples in the corresponding domain	Observations
Dual-process theory of thought models		
Serial (De Neys and Glumicic, 2008; Evans and Stanovich, 2013). Example of empirical evidence: Gillard et al. (2009).	Serial model with symbolic after subsymbolic pipeline (e.g., Valkov et al., 2018) or with subsymbolic after symbolic pipeline (Mao et al., 2019). A specific instance is DeepProbLog (Manhaeve et al., 2021) an integration of neural networks and probabilistic logic programming	Within psychology, the order of the processes in the pipeline is not defined a priori. AI models could be developed where the (serial) intervention of one of the two components depends on the needs.
Parallel (Denes-Raj and Epstein, 1994; Sloman, 1996). Example of empirical evidence: Handley et al. (2011).	Symbolic and subsymbolic models run independently and are aggregated by means of an algebraic operator (Cheng et al., 2019).	See also vector symbolic architectures (Smolensky, 1990; Schlegel et al., 2022) and semantic pointer representations (Eliasmith, 2013).
Hybrid (De Neys and Glumicic, 2008). Example of empirical evidence: Bonner and Newell (2010).	Not found.	Very specific architecture. Interesting for developing new hybrid AI model based on this specific architecture.
Logical intuition (De Neys, 2012, 2014). Example of empirical evidence: Bago and De Neys (2017).	Models that solve structured tasks (maths, chess) using a subsymbolic approach (Noever et al., 2020; Drori et al., 2022).	Open issue: if the subsymbolic process can correctly solve structured tasks, what is the function of the symbolic process? Is the symbolic process necessary?
Three-stage model (Pennycook et al., 2015). Example of empirical evidence: De Neys (2012).	Universal self-consistency for LLM generation (Chen Y. et al., 2023). This approach utilizes multiple reasoning paths sampled from LLMs and then select the most consistent answer among multiple candidates.	Strong similarities between the three-stage model and the AI approach based on self-consistency. Future research should increase the interplay between these two approaches.
Neuro-symbolic AI models (not cited above)		
Neural Module Networks (Andreas et al., 2016)	Cohabitation (Smolensky, 1991).	See also Sejnova et al. (2018) for an example of a biologically inspired implementation.
Statistical Relation models (Milch and Russell, 2006; Marra et al., 2024), i.e., the combination of logic/symbolic relations and probabilities in AI	See Kemp and Tenenbaum (2009) for a psychological model of inductive reasoning based on this approach.	According to Rosenbloom (2010), these models have been explored in the context of individual mechanisms within architectures (Iklé and Goertzel, 2008), but the architectural aspect remains unexplored. See Marra et al. (2024) for a recent review.

AI models that do not have a counterpart in the dual-process theory but are cited by the review papers are listed at the bottom of the table with their respective potential parallels in psychological research.

by System 1. Thus, the latter is posited to generate at least two distinct types of responses: Type 1 responses relying on semantic and other associations and normative Type 1 responses founded on elementary logical and probabilistic principles. In this view, conflict detection can be explained in terms of a conflict between these two different Type 1-responses. More recently, a three-stage dual-process model was proposed (Pennycook et al., 2015). According to this, stimuli may generate several Type 1-based responses, competing in terms of saliency and ease of generation. In the second stage, potential conflicts are detected. In case of no detection, the response is the Type 1 winner output. In case of detection, the response may be the Type 1 output after a rationalization (a *post-hoc* evaluation of System 1 elaboration) or a Type 2 response elaborated analytically.

4 Discussion

Assuming a correspondence between System 1 and subsymbolic processing on one hand, and System 2 and symbolic processing on the other (as stated by Bengio, 2019; see also Booch et al., 2021), we auspicate that different architectures explored within the dual-process theory may help to develop and test the feasibility of specific instances of symbolic/subsymbolic systems integration. AI literature refers to this hybrid approach as neuro-symbolic models (Gulwani et al., 2017). From an architectural point

of view, Chaudhuri et al. (2021) distinguished: (i) serial models either with symbolic after subsymbolic pipeline (e.g., Valkov et al., 2018) or subsymbolic after symbolic pipeline (e.g., Mao et al., 2019); (ii) parallel models where symbolic and subsymbolic programs run independently and then they are aggregated by means of an algebraic operator (e.g., Cheng et al., 2019); (iii) Neural Module Networks (Andreas et al., 2016) where a set of neural modules are components within a programming language. Despite several pioneering works in the AI literature (Garcez et al., 2002; Sun and Alexandre, 2013) and the potentialities of neuro-symbolic programming, its impact has been limited, and current models are still in their infancy. We auspicate that the dual-process architectures described in Section 3 may represent useful ideas to stimulate the development of new models. The dual-process of thought literature emphasizes the problem of the interaction between the two systems, a topic that has been neglected in the neuro-symbolic architectures reviewed by Chaudhuri et al. (2021). The explorations of the feasibility of neuro-symbolic architectures based on the different models proposed in cognitive literature may offer new insights in this field.² Regarding the psychological models described above and the neuro-symbolic models reviewed

² This issue has also been addressed by cognitive architectures such as CLARION (Sun, 2015), ACT-R (Ritter et al., 2019), or SOAR (Laird, 2012).

by Chaudhuri et al. (2021) and Manhaeve et al. (2022), Table 1 outlines the potential parallels between these two fields.

However, models in the psychological literature are designed to effectively describe human mental processes, thus also predicting human errors. Naturally, within the field of AI, it is not desirable to incorporate the limitations of human beings (for example, an increase in Type 1 responses due to time constraints, see also Chen X. et al., 2023). Insights drawn from cognitive literature should be regarded solely as inspiration, considering the goals of a technological system that aims to minimize its errors and achieve optimal performances. The development of these architectures could address issues currently observed in existing LLMs and AI-based image generation software.

Author contributions

GG: Conceptualization, Funding acquisition, Writing – original draft, Writing – review & editing. AP: Conceptualization, Writing – original draft, Writing – review & editing.

References

- Andreas, J., Rohrbach, M., Darrell, T., and Klein, D. (2016). “Neural module networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV), 39–48.
- Bago, B., and De Neys, W. (2017). Fast logic? Examining the time course assumption of dual process theory. *Cognition* 158, 90–109. doi: 10.1016/j.cognition.2016.10.014
- Ball, L. J., Thompson, V. A., and Stuppel, E. J. N. (2018). “Conflict and dual process theory: the case of belief bias,” in *Dual Process Theory*, ed. W. De Neys (London: Routledge/Taylor & Francis Group), 100–120.
- Bengio, Y. (2019). “From system 1 deep learning to system 2 deep learning,” in *2019 Conference on Neural Information Processing Systems*. Vancouver, BC.
- Binz, M., and Schulz, E. (2023). Using cognitive psychology to understand GPT-3. *Proc. Natl. Acad. Sci. U. S. A.* 120:e2218523120. doi: 10.1073/pnas.2218523120
- Bonner, C., and Newell, B. R. (2010). In conflict with ourselves? An investigation of heuristic and analytic processes in decision making. *Mem. Cogn.* 38, 186–196. doi: 10.3758/MC.38.2.186
- Booch, G., Fabiano, F., Horesh, L., Kate, K., Lenchner, J., Linck, N., et al. (2021). “Thinking fast and slow in AI,” in *Proceedings of the AAAI Conference on Artificial Intelligence* (Washington DC).
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. *Adv. Neural Inf. Process.* 33, 1877–1901. Available online at: https://papers.nips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf
- Chaudhuri, S., Ellis, K., Polozov, O., Singh, R., Solar-Lezama, A., and Yue, Y. (2021). Neurosymbolic programming. *Found. Trends Progr. Lang.* 7, 158–243. doi: 10.1561/25000000049
- Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. D. O., Kaplan, J., et al. (2021). Evaluating large language models trained on code. *arXiv [Preprint]*. *arXiv:2107.03374*. doi: 10.48550/arXiv.2107.03374
- Chen, X., Aksitov, R., Alon, U., Ren, J., Xiao, K., Yin, P., et al. (2023). Universal self-consistency for large language model generation. *arXiv [Preprint]*. *arXiv:2311.17311*. doi: 10.48550/arXiv.2311.17311
- Chen, Y., Andiappan, M., Jenkin, T., and Ovchinnikov, A. (2023). A manager and an AI walk into a bar: does ChatGPT make biased decisions like we do? *SSRN. Electron. J.* doi: 10.2139/ssrn.4380365
- Cheng, R., Verma, A., Orosz, G., Chaudhuri, S., Yue, Y., and Burdick, J. W. (2019). “Control regularization for reduced variance reinforcement learning,” in *International Conference on Machine Learning* (Long Beach, CA), 1141–1150.
- De Neys, W. (2012). Bias and conflict: a case for logical intuitions. *Perspect. Psychol. Sci.* 7, 28–38. doi: 10.1177/1745691611429354
- De Neys, W. (2014). Conflict detection, dual processes, and logical intuitions: some clarifications. *Think. Reason.* 20, 169–187. doi: 10.1080/13546783.2013.854725
- De Neys, W. (2018). *Dual Process Theory 2.0*. London: Routledge.
- De Neys, W., and Glumicic, T. (2008). Conflict monitoring in dual process theories of thinking. *Cognition* 106, 1248–1299. doi: 10.1016/j.cognition.2007.06.002
- Denes-Raj, V., and Epstein, S. (1994). Conflict between intuitive and rational processing: when people behave against their better judgment. *J. Pers. Soc. Psychol.* 66, 819–829. doi: 10.1037/0022-3514.66.5.819
- Drori, I., Zhang, S., Shuttleworth, R., Tang, L., Lu, A., Ke, E., et al. (2022). A neural network solves, explains, and generates university math problems by program synthesis and few-shot learning at human level. *Proc. Natl. Acad. Sci. U. S. A.* 119:e2123433119. doi: 10.1073/pnas.2123433119
- Eliasmith, C. (2013). *How to Build a Brain: A Neural Architecture for Biological Cognition*. Oxford: Oxford University Press.
- Evans, J. St. B., and Curtis-Holmes, J. (2005). Rapid responding increases belief bias: evidence for the dual-process theory of reasoning. *Think. Reason.* 11, 382–389. doi: 10.1080/13546780542000005
- Evans, J. St. B., and Stanovich, K. E. (2013). Dual-process theories of higher cognition: advancing the debate. *Perspect. Psychol. Sci.* 8, 223–241. doi: 10.1177/1745691612460685
- Evans, J. St. B. T. (2000). “Thinking and believing,” in *Mental Models in Reasoning*, eds. J. Garcia-Madruga, N. Carriedo, and M. J. Gonzalez-Labra (Madrid: UNED), 41–56.
- Garcez, A. S. A., Broda, K. B., and Gabbay, D. M. (2002). *Neuralsymbolic Learning Systems: Foundations and Applications*. Berlin: Springer Science & Business Media.
- Gillard, E., Schaeken, W., Van Doorden, W., and Verschaffel, L. (2009). “Processing time evidence for a default-interventionist model of probability judgments,” in *Proceedings of the Annual Meeting of the Cognitive Science Society* (Amsterdam), 1792–1797.
- Gulwani, S., Polozov, O., and Singh, R. (2017). Program synthesis. *Found. Trends Program. Lang.* 4, 1–119. doi: 10.1561/9781680832938
- Hammond, K., and Leake, D. (2023). “Large language models need symbolic AI,” in *Proceedings of the 17th International Workshop on Neural-Symbolic Reasoning and Learning*, *CEUR Workshop Proceedings* (Siena), 3–5.
- Handley, S. J., Newstead, S. E., and Trippas, D. (2011). Logic, beliefs, and instruction: a test of the default interventionist account of belief bias. *J. Exp. Psychol. Learn. Mem. Cogn.* 37, 28–43. doi: 10.1037/a0021098
- Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.

Funding

The author(s) declare that no financial support was received for the research, authorship, and/or publication of this article.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Hu, K. (2023). *ChatGPT Sets Record for Fastest-Growing User Base-Analyst Note* (London: Reuters), 12.
- Iklé, M., and Goertzel, B. (2008). *Probabilistic Quantifier Logic for General Intelligence: an Indefinite Probabilities Approach*. Amsterdam: IOS Press.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. London: Macmillan.
- Kemp, C., and Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychol. Rev.* 116, 20–58. doi: 10.1037/a0014282
- Laird, J. E. (2012). *The Soar Cognitive Architecture*. Cambridge, MA: MIT Press.
- Lin, S. C., Yang, J. H., Nogueira, R., Tsai, M. F., Wang, C. J., and Lin, J. (2020). Conversational question reformulation via sequence-to-sequence architectures and pretrained language models. *arXiv [Preprint]*. arXiv:2004.01909. doi: 10.48550/arXiv.2004.01909
- Liu, Y., Han, T., Ma, S., Zhang, J., Yang, Y., Tian, J., et al. (2023). Summary of ChatGPT-related research and perspective towards the future of large language models. *Meta-Radiol.* 2023:100017. doi: 10.1016/j.metrad.2023.100017
- Manhaeve, R., Dumančić, S., Kimmig, A., Demeester, T., and De Raedt, L. (2021). Neural probabilistic logic programming in DeepProbLog. *Artif. Intell.* 298:103504. doi: 10.1016/j.artint.2021.103504
- Manhaeve, R., Marra, G., Demeester, T., Dumančić, S., Kimmig, A., and De Raedt, L. (2022). “Neuro-Symbolic AI = Neural + Logical + Probabilistic AI,” in *Neuro-Symbolic Artificial Intelligence: The State of the Art*, eds. P. Hitzler & K. Sarker (Amsterdam: IOS Press), 173–191.
- Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B., and Wu, J. (2019). The neuro-symbolic concept learner: interpreting scenes, words, and sentences from natural supervision. *arXiv [Preprint]*. arXiv:1904.12584. doi: 10.48550/arXiv.1904.12584
- Marra, G., Dumančić, S., Manhaeve, R., and De Raedt, L. (2024). From statistical relational to neurosymbolic artificial intelligence: a survey. *Artif. Intell.* 2023:104062. doi: 10.1016/j.artint.2023.104062
- Milch, B., and Russell, S. (2006). “First-order probabilistic languages: into the unknown,” in *International Conference on Inductive Logic Programming* (Berlin; Heidelberg: Springer Berlin Heidelberg), 10–24.
- Newell, B. R., Lagnado, D. A., and Shanks, D. R. (2015). *Straight Choices: The Psychology of Decision Making*. New York, NY: Psychology Press.
- Noever, D., Ciolino, M., and Kalin, J. (2020). The chess transformer: mastering play using generative language models. *arXiv*. doi: 10.48550/arXiv.2008.04057
- Pennycook, G., Fugelsang, J. A., and Koehler, D. J. (2015). What makes us think? A three-stage dual-process model of analytic engagement. *Cogn. Psychol.* 80, 34–72. doi: 10.1016/j.cogpsych.2015.05.001
- Ritter, F. E., Tehrani, F., and Oury, J. D. (2019). ACT-R: a cognitive architecture for modeling cognition. *Wiley Interdiscip. Rev. Cogn. Sci.* 10:e1488. doi: 10.1002/wcs.1488
- Rosenbloom, P. S. (2010). “An architectural approach to statistical relational AI,” in *Proceedings of the Association for the Advancement of Artificial Intelligence* (Stanford, CA), 90–91.
- Schlegel, K., Neubert, P., and Protzel, P. (2022). A comparison of vector symbolic architectures. *Artif. Intell. Rev.* 55, 4523–4555. doi: 10.1007/s10462-021-10110-3
- Sejnova, G., Tesar, M., and Vavrecka, M. (2018). Compositional models for VQA: can neural module networks really count? *Procedia Comput. Sci.* 145, 481–487. doi: 10.1016/j.procs.2018.11.110
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychol. Bull.* 119, 3–22. doi: 10.1037/0033-2909.119.1.3
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artif. Intell.* 46, 159–216. doi: 10.1016/0004-3702(90)90007-M
- Smolensky, P. (1991). “The constituent structure of connectionist mental states: a reply to Fodor and Pylyshyn,” in *Connectionism and the Philosophy of Mind*, eds. T. Horgan and J. Tienson (Dordrecht: Springer Netherlands), 281–308.
- Sun, R. (2015). Interpreting psychological notions: a dual-process computational theory. *J. App. Res. Mem. Cogn.* 4, 191–196. doi: 10.1016/j.jarmac.2014.09.001
- Sun, R., and Alexandre, F. (2013). *Connectionist-Symbolic Integration: From Unified to Hybrid Approaches*. London: Psychology Press.
- Thompson, V. A. (2013). Why it matters: the implications of autonomous processes for dual process theories-Commentary on Evans & Stanovich (2013). *Perspect. Psychol. Sci.* 8, 253–256. doi: 10.1177/1745691613483476
- Titus, L. M. (2023). Does ChatGPT have semantic understanding? A problem with the statistics-of-occurrence strategy. *Cogn. Syst. Res.* 2023:101174. doi: 10.1016/j.cogsys.2023.101174
- Valkov, L., Chaudhari, D., Srivastava, A., Sutton, C., and Chaudhuri, S. (2018). HOUDINI: lifelong learning as program synthesis. *Adv. Neural Inf. Process.* 2018, 8701–8712. Available online at: https://proceedings.neurips.cc/paper_files/paper/2018/file/edc27f139c3b4e4bb29d1cdbc45663f9-Paper.pdf