Check for updates

The Challenge of Modeling the Acquisition of Mathematical Concepts

Alberto Testolin [1,2] *

[1]Department of General Psychology, University of Padova, Padova, Italy, [2]Department of Information Engineering, University of Padova, Padova, Italy

As a full-blown research topic, numerical cognition is investigated by a variety of disciplines including cognitive science, developmental and educational psychology, linguistics, anthropology and, more recently, biology and neuroscience. However, despite the great progress achieved by such a broad and diversified scientific inquiry, we are still lacking a comprehensive theory that could explain how numerical concepts are learned by the human brain. In this perspective, I argue that computer simulation should have a primary role in filling this gap because it allows identifying the finer-grained computational mechanisms underlying complex behavior and cognition. Modeling efforts will be most effective if carried out at cross-disciplinary intersections, as attested by the recent success in simulating human cognition using techniques developed in the fields of artificial intelligence and machine learning. In this respect, deep learning models have provided valuable insights into our most basic quantification abilities, showing how numerosity perception could emerge in multi-layered neural networks that learn the statistical structure of their visual environment. Nevertheless, this modeling approach has not yet scaled to more sophisticated cognitive skills that are foundational to higher-level mathematical thinking, such as those involving the use of symbolic numbers and arithmetic principles. I will discuss promising directions to push deep learning into this uncharted territory. If successful, such endeavor would allow simulating the acquisition of numerical concepts in its full complexity, guiding empirical investigation on the richest soil and possibly offering far-reaching implications for educational practice.

Keywords: computational modeling, artificial neural networks, deep learning, number sense, symbol grounding, mathematical learning, embodied cognition, material culture

INTRODUCTION

Despite the importance of mathematics in modern societies, the cognitive foundations of mathematical learning are still mysterious and hotly debated. At the one end of the bridge, the idealistic view conceives mathematical concepts as purely abstract entities that humans discover using logical reasoning; at the other end, empiricists argue that mathematics is the product of our sensory experiences, and therefore it is essentially an activity of construction (Brown, 2012). A somehow intermediate position is taken by modern neurocognitive theories, which identify a set of "core" brain systems specifically evolved to support basic intuitions about quantity (Butterworth, 1999; Feigenson et al., 2004; Piazza, 2010; Dehaene, 2011) but also acknowledge that higher-level numerical knowledge has materialized only recently, via cultural practices supported by language and symbolic reference (Núñez, 2017).

In recent years, the finding that measures of basic quantification skills correlate to later mathematical achievement (e.g., Halberda et al., 2008; Libertus et al., 2011; Starr et al., 2013) has led to the hypothesis that our "number sense" might indeed constitute the starting point to learn more complex mathematical concepts. However, the relationship between numerosity perception and symbolic math remains controversial (Negen and Sarnecka, 2015; Schneider et al., 2017; Wilkey and Ansari, 2019), calling for a deeper theoretical investigation that should be carried out with the support of formal models.

Here I will argue that the quest for artificial intelligence provides an extremely rich soil for the development of a computational theory of mathematical learning. Indeed, although computers largely outperform humans on numerical tasks requiring the mere application of syntactic manipulations (e.g., performing algebraic operations on large numbers, or iteratively computing the value of a function), they are completely blind about the *meaning* of such operations because they lack a conceptual semantics of number. Grounding abstract symbols into some form of intrinsic meaning is a longstanding issue in artificial intelligence (Searle, 1980; Harnad, 1990), and mathematics likely constitutes the most challenging domain for investigating how high-level knowledge could be linked to bottom-up, sensorimotor primitives (Leibovich and Ansari, 2016).

By framing a theory in computational terms, scientists are forced to adopt a precise, formal language, because all the details of the theory should be explicitly stated to simulate it on a computer. Modeling also requires to carefully think about the tasks that are being simulated and the possible ways in which a computational device can (or cannot) solve them. In this perspective article, I will focus in particular on *connectionist* models, where cognition is conceived as an emergent property of networks of units that self-organize according to physical principles (Rumelhart and McClelland, 1986; Elman et al., 1996; McClelland et al., 2010). According to this view, knowledge is implicitly stored in the connections among neurons, and learning processes adaptively change the strength of these connections according to experience. Notably, the recent breakthroughs in *deep learning* (LeCun et al., 2015) have revealed the true potential of this approach, by showing how machines endowed with domain-general learning mechanisms can simulate a variety of high-level cognitive skills, ranging from visual object recognition (He et al., 2016) to natural language understanding (Devlin et al., 2018) and strategic planning (Silver et al., 2017).

Computational Models of Basic Quantification Skills

According to the "number sense" view, numerical cognition is grounded in basic quantification skills, such as the ability to rapidly estimate the number of items in a visual display (Dehaene, 2011). Numerosity is thus conceived as a primary perceptual attribute (Anobile et al., 2016) processed by a specialized (and possibly innate) system yielding an approximate representation of numerical quantity (Feigenson et al., 2004). The seminal neural network model by Dehaene and Changeux (1993) incorporated these principles: numerosity perception

was hardwired in the model, reflecting the assumption that this ability is present at birth. Successive models revisited this nativist stance, by showing that numerosity representations can emerge as a result of learning and sensory experience (Verguts and Fias, 2004). In particular, recent work based on unsupervised deep learning has demonstrated that human-like numerosity perception can emerge in multi-layer neural networks that learn a hierarchical generative model of the sensory data (Stoianov and Zorzi, 2012; Zorzi and Testolin, 2018; see **Figure 1A**).
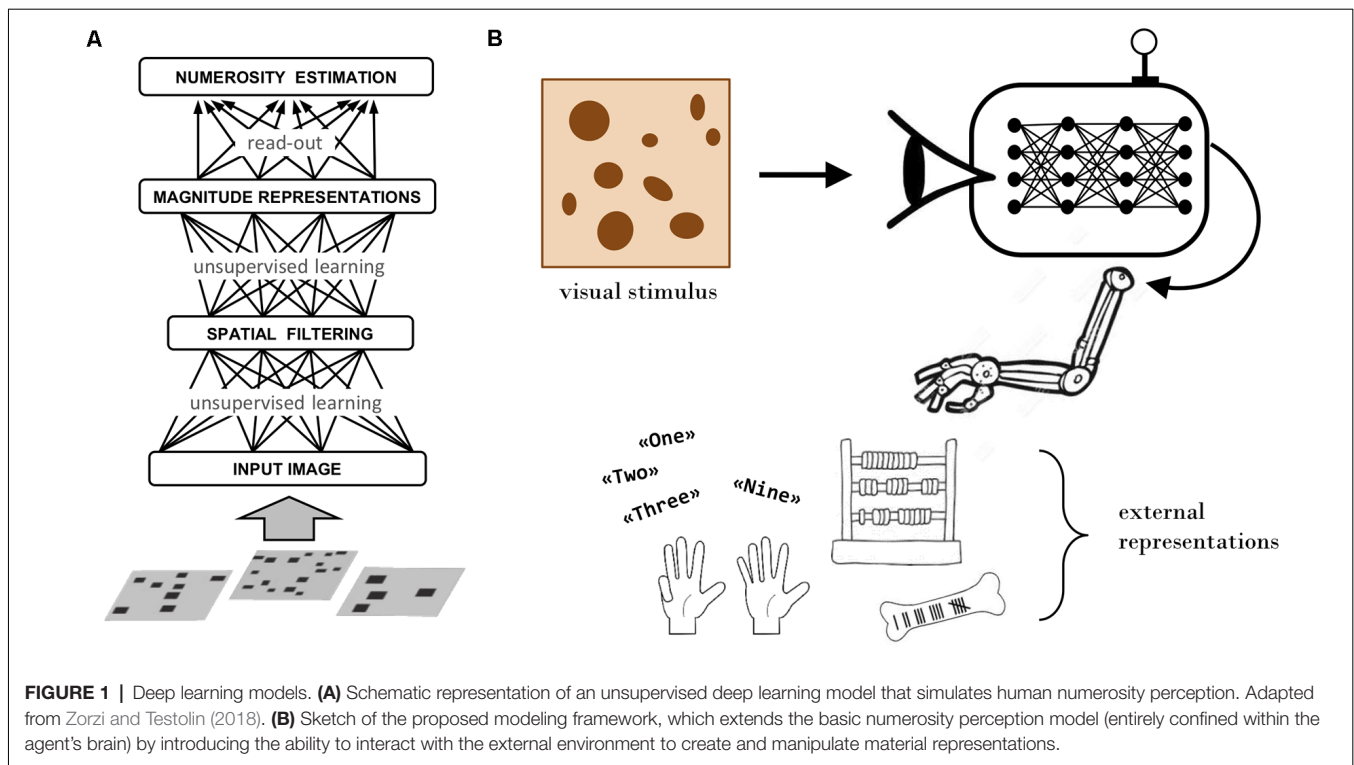
Deep learning models account for a wide range of empirical phenomena in the number sense literature. They can accurately simulate Weber-like responses in numerosity comparison tasks (Stoianov and Zorzi, 2012), also accounting for congruency effects (Zorzi and Testolin, 2018) and for the fine-grained contribution of non-numerical magnitudes in biasing behavioral responses (Testolin et al., 2019). Notably, the number acuity of randomly initialized deep networks rivals that of newborns, and its gradual development follows trajectories similar to those observed in human longitudinal studies (Testolin et al., 2020). Deep networks have also been successfully tested in subitizing (Wever and Runia, 2019) and numerosity estimation tasks (Chen et al., 2018). Last, but not least, artificial neurons often reproduce neurophysiological properties observed in single-cell recording studies, for example by exhibiting number-sensitive tuning functions (Zorzi and Testolin, 2018; Nasr et al., 2019).

Several questions remain under investigation: Is it possible to fully disentangle numerosity from continuous magnitudes by only relying on unsupervised learning (Zanetti et al., 2019)? Can generative models generalize to unseen numerosities (Zhao et al., 2018)? Are there computational limitations in tracking multiple objects in dynamic scenes (Cenzato et al., 2019)? What is the contribution of explicit feedback and multi-sensory integration in shaping numerosity representations? How do deep learning models map into the cortical processing hierarchy? Nevertheless, despite these open questions, we can safely argue that deep learning has paved the way toward a computational theory about the origin of our number sense, confirming the appeal of deep networks as models of human sensory processing (Testolin and Zorzi, 2016; Yamins and DiCarlo, 2016; Testolin et al., 2017). Unfortunately, simulating the transition from approximate to symbolic numbers turns out to be much more challenging, as we discuss in the next section.

Modeling the Acquisition of Higher-Level Mathematical Concepts

One of the most ambitious questions to be addressed is whether deep learning models could develop even more sophisticated numerical abilities, such as those involving arithmetic and symbolic math. Symbolic reasoning is notoriously difficult for connectionist models (Marcus, 2003), and despite recent progress, deep neural networks still struggle with tasks requiring procedural and compositional knowledge (Garnelo and Shanahan, 2019).

Only a few modeling studies have investigated how arithmetic could be learned by artificial neural networks. Since early

FIGURE 1 | Deep learning models. (A) Schematic representation of an unsupervised deep learning model that simulates human numerosity perception. Adapted from Zorzi and Testolin (2018). (B) Sketch of the proposed modeling framework, which extends the basic numerosity perception model (entirely confined within the agent's brain) by introducing the ability to interact with the external environment to create and manipulate material representations.

attempts, associative memories have been used to simulate mental calculation as a process of storage and retrieval of arithmetic facts (McCloskey and Lindemann, 1992): during the learning phase, the two arguments and the result of a simple operation (e.g., single-digit multiplication) are given as input to an associative memory, whose learning goal is to accurately store them as a global, stable state. During the testing phase, only the operands are given, and the network must recover the missing information (i.e., the result) by gradually settling into the correct configuration. Building on this approach, successive simulations have shown that numerosity-based ("semantic") representations can facilitate the learning of arithmetic facts (Zorzi et al., 2005) and equivalence problems (Mickey and McClelland, 2014). Others have shown that multi-digit addition and subtraction (but not multiplication) can be acquired through end-to-end supervised learning from pixel-level images (Hoshen and Peleg, 2015). One critical limitation of these approaches, however, is that they conceive arithmetic learning as a mere process of storing and recall, which gradually develops through the massive reiteration of all possible arithmetic facts that need to be learned. Besides being psychologically implausible and computationally unfeasible, this approach does not guarantee that the system will be able to generalize the acquired knowledge to unseen numbers and, even less, to exploit the acquired knowledge to more effectively learn new mathematical concepts.

The challenge of developing learning models that can exhibit algebraic generalization with the robustness and flexibility exhibited by humans is so fundamental that major players in deep learning research are intensively investigating these issues. For example, Google's DeepMind company has recently evaluated several deep learning models on a set of benchmark problems taken from UK national school mathematics curriculums, covering arithmetic, algebra, elementary calculus, et cetera (Saxton et al., 2019). DeepMind's best model correctly solved only 14 out of 40 problems, which would be equivalent to an "E" grade. Although such difficulties have led some researchers to argue that neural networks are incapable of exhibiting compositional abilities (Marcus, 2018), others argue for the opposite (Baroni, 2020; Martin and Baggio, 2020).

Even the acquisition of the concept of *exact number* is still out of reach for deep networks, which often cannot generalize outside of the range of numerical values encountered during training (Trask et al., 2018). Integer numbers are one of the pillars of arithmetic, so they constitute the perfect testbed for developing and testing computational models of mathematical learning. Developmental studies show that integers are gradually acquired by children during formal education through the acquisition of number words and counting skills: Indeed, although sequential (item-by-item) enumeration skills are present in animal species (Platt and Johnson, 1971; Beran and Beran, 2004; Dacke and Srinivasan, 2008), even in humans counting is not culturally universal (Gordon, 2004) and there is evidence that young children and people from cultures lacking number words have an incomplete understanding of what it means for two sets of items to have exactly the same number of items (Izard et al., 2008, 2014).

Some authors have sought to characterize the acquisition of exact numbers as the semantic induction of a "cardinality principle" (Sarnecka and Carey, 2008). This hypothesis has

been exemplified in a computational model based on Bayesian inference, which simulated the stage-like development of counting abilities by relying on a pre-determined set of "core" cognitive operations (Piantadosi et al., 2012). The repertoire of innate abilities included the capacity to exactly identify cardinalities up to 3, perform basic operations on sets (e.g., difference, union, intersection), retrieve the next or previous word from an ordered counting list, and to operate these functions recursively. Although such modeling approach offers a rational interpretation of the process that might underlie the acquisition of an abstract cardinality principle, it assumes a certain amount of *a priori* symbolic knowledge and procedural skills, which is in contrast to empirical data suggesting, for example, that a complete understanding of the successor principle arises only after considerable interaction with the teaching environment (Davidson et al., 2012).

TOWARD A COMPREHENSIVE NEUROCOMPUTATIONAL FRAMEWORK

The Downplayed Role of External Representations

A central tenet of connectionist models is that semantics intrinsically emerges in a system interacting with its surrounding environment. However, this idea is usually superficially implemented in deep learning models, because the interaction is often limited to *passive* observation of statistical properties of the world (Zorzi et al., 2013). Taking inspiration from constructivist theories in developmental psychology, here I argue that a step forward will require to build computational models that learn by *actively* manipulating the environment, that is, by causally interacting with objects in their perceptual space. Crucially, the notion of "environment" should include embodiment (Lakoff and Núñez, 2000) and—most importantly—the social, cultural and educational environment (Vygotsky, 1980; Clark, 2011). Indeed, according to the Vygotskyan perspective, students actively construct abstract knowledge through interactions with teachers and peers, gradually moving their dependency on explicit forms of mediation to more implicit (internalized) forms (Walshaw, 2017).

The possibility to manipulate the environment greatly increases the complexity of the learning agent but also enables the functional use of external entities to create powerful representational systems, which can be manipulated in simple ways to get answers to difficult problems. The underlying assumption is that cultural evolution and history are foundational forces for the emergence of superior cognitive functions and that great intellectual achievements (such as the invention of mathematics) have been triggered by our ability to create artifacts serving as physical representations of abstract concepts. Some investigators have recently emphasized the role of material culture in numerical cognition (Menary, 2015; Overmann, 2016, 2018), for example by highlighting that our mental organization of numbers into an ordered "number line" might be related to the linearity of the material forms used to represent and manipulate them (Núñez, 2011).

Primitive devices used for representing numbers date back to notched bones in the Paleolithic period (d'Errico et al., 2018) and clay tokens in the Neolithic period (Schmandt-Besserat, 1992), which predated the subsequent diffusion of abaci, positional systems and increasingly more sophisticated numerical notations (Menninger, 1992). However, despite the concept of external representations was foreseen in early connectionist theories[1], it has been seldomly explored in practice.

Learning to Create and Manipulate Symbolic Representations

We can now sketch a concrete proposal for building more realistic simulations of mathematical learning. The computational framework should incorporate the following key components, summarized in **Figure 1B**.

- **Perceptual system**. This is where computational modeling has been mostly focused (and successful) up to now (see Section "Computational Models of Basic Quantification Skills"). The challenge will be to scale-up the existing models to more realistic sensory input (e.g., naturalistic visual scenes) and to incorporate a larger repertoire of pattern recognition abilities, which should not only allow to approximately represent visual quantities but also to recognize structured configurations of object arrays (e.g., sequences of tally marks, geometric displacements of items, patterns encoded in an abacus, etc.) and symbolic notations (e.g., written digits and operands).
- **Embodiment**. Of particular interest to the development of exact numbers is finger counting (Butterworth, 1999; Andres et al., 2007; Domahs et al., 2012), which not only helps children to keep track and coordinate the production of number words (Alibali and DiRusso, 1999) but may also allow to organize numbers spatially (Fischer, 2008). Hand-based representations are ubiquitous across cultures (Bender and Beller, 2012) and play a key role in the subsequent acquisition of number words (Gunderson et al., 2015; Gibson et al., 2019), possibly influencing symbolic number processing even in adulthood (Domahs et al., 2010). It has been recently shown that neural networks can learn to count the number of items in visual displays and that the ability to sequentially point to individual objects helps in speeding up counting acquisition (Fang et al., 2018). A further step is taken by cognitive developmental robotics, which explores the instantiation of these principles in physically embodied agents (Di Nuovo and Jay, 2019). Interestingly, pointing gestures significantly improved counting accuracy in a humanoid robot, and learning was more effective when both fingers and words were provided as input (Rucinski et al., 2012; De La Cruz et al., 2014).
- **Material representations**. The ability to manipulate external objects might be the key missing piece for simulating the acquisition of exact numbers. Indeed, although hand gestures might serve as placeholders to

---

[1]See for example, the section "External Representations and Formal Reasoning" in Rumelhart et al. (1986).

learn more efficient arithmetic strategies (Siegler and Jenkins, 1989; for a computational account see Hansen et al., 2014), material representations allow for a much more precise encoding of numerical information. For example, the agent can learn to establish the cardinality of a set by organizing items in regular configurations that promote "groupitizing" (Starkey and McCandliss, 2014), or to exactly compare the cardinality of two sets by disposing of items in one-to-one correspondence. More sophisticated devices such as abaci and Cuisenaire rods further extend our ability to represent exact numbers, for example by exploiting inter-exponential relations to precisely (but compactly) encode large numbers, or to explicitly represent compositionality to promote generalization (Overmann, 2018).

- **Diversified learning signals**. In addition to unsupervised learning, the agent should exploit *reinforcement learning* (Sutton and Barto, 1998) to predict the outcome of its actions. This learning modality would also play a key role in simulating curiosity-driven behavior and active engagement with material representations. Notably, deep reinforcement learning has recently achieved impressive performance in difficult cognitive tasks, for example by discovering complex strategies in board games (Silver et al., 2017). However, learning through reinforcement can be challenging in the presence of very large action spaces (i.e., the correct action has to be chosen from a wide range of possible actions) and sparse rewards (i.e., feedback is given only once the whole task has been carried out). Taking inspiration from the notions of *transfer learning* and *curriculum learning* used in machine learning (Bengio et al., 2009) and from *shaping* procedures used in animal conditioning (Skinner, 1953), these issues can be mitigated by decomposing the task into simpler sub-tasks. For example, rather than rewarding only the trials where the agent has correctly counted all items in a display, rewards can be initially given every time the agent touches an object, to first promote the acquisition of sequential pointing skills. Similarly, the agent could first be rewarded simply for being able to accurately reproduce the abacus configuration corresponding to a specific number, rather than for being able to correctly manipulate the abacus to solve an addition problem. This idea of "gradually walking the agent through the word" also implies the exploitation of *supervised learning*, because explicit teaching signals must be used to stimulate learning by imitation and adult guidance.
- **Linguistic input**. Despite language might not be crucial for the acquisition of elementary numerical concepts (Gelman and Butterworth, 2005; Butterworth et al., 2008), it provides useful cues during the development of basic algebraic notions: for example, morphological cues allow single/plural distinction, number words can act as stable placeholders during counting acquisition, and learning natural language quantifiers seems a key step for mastering the ordering principle (Le Corre, 2014). A recent deep learning model has shown that learning quantifiers allows to more easily carry out approximate numerosity judgments (Pezzelle et al.,

2018); however, the role of linguistic input for simulating the acquisition of exact numbers has yet to be explored. Furthermore, later in development language becomes the primary medium to acquire higher-level mathematical knowledge, hence it will need to be taken into account to design computational models approaching that level of complexity.

DISCUSSION

Symbolic numbers are a hallmark of human intelligence, but we are still lacking a comprehensive theory explaining how the brain learns to master them. Here I argued that computational modeling should have a primary role in this enterprise. Taking the acquisition of natural numbers as a case study, I emphasized the role of material representations in supporting the transition from approximate to symbolic numerical concepts. According to this view, exact numbers do not emerge from the mere association between number words and perceptual magnitudes: such mapping is strongly mediated by the acquisition of procedural skills (e.g., finger counting) and the ability to effectively manipulate representational devices (Leibovich and Ansari, 2016; Overmann, 2018; Carey and Barner, 2019).

In line with the idea that improved problem representation is a key mechanism for the joint development of conceptual and procedural knowledge (Rittle-Johnson et al., 2001), cognitive development in artificial agents must thus be supported by an adequate learning environment, which should provide feedback, teaching signals, and representational media commensurate with the current level of development. Notably, once a procedural skill has been mastered it might become internalized: the agent can simply "imagine" carrying out operations on the material device, without the need to physically operate over it. Some representations might thus serve just as intermediate steps for the acquisition of more abstract and efficient notations: as finger counting allows us to gradually grasp the meaning of number words, manipulating an abacus allows to ground numerical symbols into concrete visuospatial representations. A historical case that illustrates this perspective is the famous dispute between "abacists" and "algorists", which was undoubtedly won by the latter, who demonstrated the superiority of symbolic notation for carrying out arithmetic operations (see **Figure 2**). However, one might wonder whether Boethius could have mastered arithmetic algorithms without first grounding his numerical concepts into a set of more concrete representations.

In addition to providing a useful framework to interpret empirical findings, the proposed approach can raise important questions that would stimulate further theoretical and experimental work. For example, a critical aspect of our school system is to teach how to effectively discover useful strategies and representational schemes for solving difficult problems. In computational simulations, the necessity for appropriate teacher guidance stems from the fact that it is very difficult to invent new representations for problems we might wish to solve: it may even be that the process of

FIGURE 2 | Allegory of Arithmetic. Engraving from the encyclopedic book *Margarita Philosophica* by Gregor Reisch (1503) depicting the "abacists vs. algorists" debate. Arithmetica (female figure) is supervising a calculation contest between Pythagoras (right), represented as using a counting board, and Boethius (left), who embraces algorithmic calculation with Arabic numbers. The struggle of Pythagoras suggests who is going to be the winner. Reproduced from Wikipedia.

inventing such representations is one of our highest intellectual abilities (Rumelhart et al., 1986). Computational frameworks that allow simulating a more complex interaction between artificial agents and their learning environment might thus eventually provide insights also about the teaching practices that could be most effective to guide numerical development in our children.

AUTHOR CONTRIBUTIONS

AT is fully responsible for the content of this article.

FUNDING

ACKNOWLEDGMENTS

REFERENCES

Alibali, M. W., and DiRusso, A. A. (1999). The function of gesture in learning to count: more than keeping track. *Cogn. Dev.* 56, 37–56. doi: 10.1016/s0885-2014(99)80017-3

Andres, M., Seron, X., and Olivier, E. (2007). Contribution of hand motor circuits to counting. *J. Cogn. Neurosci.* 19, 563–576. doi: 10.1162/jocn.2007.19.4.563

Anobile, G., Cicchini, G. M., and Burr, D. C. (2016). Number as a primary perceptual attribute: a review. *Perception* 45, 5–31. doi: 10.1177/0301006615602599

Baroni, M. (2020). Linguistic generalization and compositionality in modern artificial neural networks. *Philos. Trans. R. Soc. B Biol. Sci.* 375:20190307. doi: 10.1098/rstb.2019.0307

Bender, A., and Beller, S. (2012). Nature and culture of finger counting: diversity and representational effects of an embodied cognitive tool. *Cognition* 124, 156–182. doi: 10.1016/j.cognition.2012.05.005

Bengio, Y., Louradour, J., Collobert, R., and Weston, J. (2009). "Curriculum learning," in *Proceedings of the 26th Annual International Conference on Machine Learning* (New York, NY, USA: ACM), 1–8.

Beran, M. J., and Beran, M. M. (2004). Chimpanzees remember the results of one-by-one addition of food items to sets over extended time periods. *Psychol. Sci.* 15, 94–99. doi: 10.1111/j.0963-7214.2004.01502004.x

Brown, J. R. (2012). *Platonism, Naturalism and Mathematical Knowledge.* New York, NY: Routledge.

Butterworth, B. (1999). *The Mathematical Brain.* London, UK: MacMillan.

Butterworth, B., Reeve, R., Reynolds, F., and Lloyd, D. (2008). Numerical thought with and without words: evidence from indigenous Australian children. *Proc. Natl. Acad. Sci. U S A* 105, 13179–13184. doi: 10.1073/pnas.0806045105

Carey, S., and Barner, D. (2019). Ontogenetic origins of human integer representations. *Trends Cogn. Sci.* 23, 823–835. doi: 10.1016/j.tics.2019.07.004

Cenzato, A., Testolin, A., and Zorzi, M. (2019). On the difficulty of learning and predicting the long-term dynamics of bouncing objects. *arXiv:1907.13494* [Preprint].

Chen, S. Y., Zhou, Z., Fang, M., and McClelland, J. L. (2018). "Can generic neural networks estimate numerosity like humans?," in *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, eds T. T. Rogers, M. Rau, X. Zhu and C. W. Kalish (Austin, TX: Cognitive Science Society), 202–207.

Clark, A. (2011). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension.* New York, NY: Oxford University Press.

d'Errico, F., Doyon, L., Colagé, I., Queffelec, A., Le Vraux, E., Giacobini, G., et al. (2018). From number sense to number symbols. An archaeological perspective. *Philos. Trans. R. Soc. B Biol. Sci.* 373:20160518. doi: 10.1098/rstb.2016.0518

Dacke, M., and Srinivasan, M. V. (2008). Evidence for counting in insects. *Anim. Cogn.* 11, 683–689. doi: 10.1007/s10071-008-0159-y

Davidson, K., Eng, K., and Barner, D. (2012). Does learning to count involve a semantic induction? *Cognition* 123, 162–173. doi: 10.1016/j.cognition.2011.12.013

De La Cruz, V. M., Di Nuovo, A., Di Nuovo, S., and Cangelosi, A. (2014). Making fingers and words count in a cognitive robot. *Front. Behav. Neurosci.* 8:13. doi: 10.3389/fnbeh.2014.00013

Dehaene, S. (2011). *The Number Sense: How the Mind Creates Mathematics.* New York, NY: Oxford University Press.

Dehaene, S., and Changeux, J.-P. J. (1993). Development of elementary numerical abilities: a neuronal model. *J. Cogn. Neurosci.* 5, 390–407. doi: 10.1162/jocn.1993.5.4.390

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). BERT: pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805* [Preprint].

Di Nuovo, A., and Jay, T. (2019). Development of numerical cognition in children and artificial systems: a review of the current knowledge and proposals for multi-disciplinary research. *Cogn. Comput. Syst.* 1, 2–11. doi: 10.1049/ccs.2018.0004

Domahs, F., Kaufmann, L., and Fischer, M. H. (2012). *Handy Numbers: Finger Counting and Numerical Cognition.* Lausanne: Special Research Topic in Frontiers Media SA.

Domahs, F., Moeller, K., Huber, S., Willmes, K., and Nuerk, H.-C. (2010). Embodied numerosity: implicit hand-based representations influence symbolic number processing across cultures. *Cognition* 116, 251–266. doi: 10.1016/j.cognition.2010.05.007

Elman, J. L., Bates, E., Johnson, M., Karmiloff-smith, A., Parisi, D., and Plunkett, K. (1996). *Rethinking Innateness: A Connectionist Perspective on Development.* Cambridge, MA: MIT Press.

Fang, M., Zhou, Z., Chen, S. Y., and McClelland, J. L. (2018). "Can a recurrent neural network learn to count things?" in *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, eds T. T. Rogers, M. Rau, X. Zhu, and C. W. Kalish (Austin, TX: Cognitive Science Society), 360–365.

Feigenson, L., Dehaene, S., and Spelke, E. S. (2004). Core systems of number. *Trends Cogn. Sci.* 8, 307–314. doi: 10.1016/j.tics.2004.05.002

Fischer, M. H. (2008). Finger counting habits modulate spatial-numerical associations. *Cortex* 44, 386–392. doi: 10.1016/j.cortex.2007.08.004

Garnelo, M., and Shanahan, M. (2019). Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Curr. Opin. Behav. Sci.* 29, 17–23. doi: 10.1016/j.cobeha.2018.12.010

Gelman, R., and Butterworth, B. (2005). Number and language: how are they related? *Trends Cogn. Sci.* 9, 6–10. doi: 10.1016/j.tics.2004.11.004

Gibson, D. J., Gunderson, E. A., Spaepen, E., Levine, S. C., and Goldin-Meadow, S. (2019). Number gestures predict learning of number words. *Dev. Sci.* 22:e12791. doi: 10.1111/desc.12791

Gordon, P. (2004). Numerical cognition without words: evidence from amazonia. *Science* 306, 496–499. doi: 10.1126/science.1094492

Gunderson, E. A., Spaepen, E., Gibson, D., Goldin-Meadow, S., and Levine, S. C. (2015). Gesture as a window onto children's number knowledge. *Cognition* 144, 14–28. doi: 10.1016/j.cognition.2015.07.008

Halberda, J., Mazzocco, M. M., and Feigenson, L. (2008). Individual differences in non-verbal number acuity correlate with maths achievement. *Nature* 455, 665–668. doi: 10.1038/nature07246

Hansen, S. S., McKenzie, C., and McClelland, J. L. (2014). "Two plus three is five: discovering efficient addition strategies without metacognition," in *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (Montreal, QC: Cognitive Science Society), 583–588.

Harnad, S. (1990). The symbol grounding problem. *Phys. D Nonlinar Phenom.* 42, 335–346. doi: 10.1016/0167-2789(90)90087-6

He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition," in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (Las Vegas, NV: IEEE), 770–778.

Hoshen, Y., and Peleg, S. (2015). "Visual learning of arithmetic operations," in *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, (Phoenix, AZ: AAAI), 3733–3739.

Izard, V., Pica, P., Spelke, E. S., and Dehaene, S. (2008). Exact equality and successor function: two key concepts on the path towards understanding exact numbers. *Philos. Psychol.* 21, 491–505. doi: 10.1080/09515080802285354

Izard, V., Streri, A., and Spelke, E. S. (2014). Toward exact number: young children use one-to-one correspondence to measure set identity but not numerical equality. *Cogn. Psychol.* 72, 27–53. doi: 10.1016/j.cogpsych.2014.01.004

Lakoff, G., and Núñez, R. E. (2000). *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics Into Being.* New York, NY: Basic Books.

Le Corre, M. (2014). Children acquire the later-greater principle after the cardinal principle. *Br. J. Dev. Psychol.* 32, 163–177. doi: 10.1111/bjdp.12029

LeCun, Y., Bengio, Y., and Hinton, G. E. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539

Leibovich, T., and Ansari, D. (2016). The symbol-grounding problem in numerical cognition: a review of theory, evidence, and outstanding questions. *Can. J. Exp. Psychol.* 70, 12–23. doi: 10.1037/cep0000070

Libertus, M. E., Feigenson, L., and Halberda, J. (2011). Preschool acuity of the approximate number system correlates with school math ability. *Dev. Sci.* 14, 1292–1300. doi: 10.1111/j.1467-7687.2011.01080.x

Marcus, G. F. (2003). *The Algebraic Mind: Integrating Connectionism and Cognitive Science.* Cambridge, MA: MIT Press.

Marcus, G. F. (2018). Deep learning: a critical appraisal. *arXiv:1801.00631* [Preprint].

Martin, A. E., and Baggio, G. (2020). Modelling meaning composition from formalism to mechanism. *Philos. Trans. R. Soc. B Biol. Sci.* 375:20190298. doi: 10.1098/rstb.2019.0298

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., et al. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends Cogn. Sci.* 14, 348–356. doi: 10.1016/j.tics.2010.06.002

McCloskey, M., and Lindemann, A. M. (1992). Mathnet: preliminary results from a distributed model of arithmetic fact retrieval. *Adv. Psychol.* 91, 365–409. doi: 10.1016/s0166-4115(08)60892-4

Menary, R. (2015). Mathematical cognition—a case of enculturation. *Open Mind* 25, 1–20. doi: 10.15502/9783958570818

Menninger, K. (1992). *Number Words and Number Symbols: A Cultural History of Numbers.* New York: Dover Pubns.

Mickey, K. W., and McClelland, J. L. (2014). "A neural network model of learning mathematical equivalence," in *Proceedings of the 36th Annual Conference of the Cognitive Science Society,* (QC, Canada: Cognitive Science Society), 1012–1017.

Nasr, K., Viswanathan, P., and Nieder, A. (2019). Number detectors spontaneously emerge in a deep neural network designed for visual object recognition. *Sci. Adv.* 5:eaav7903. doi: 10.1126/sciadv.aav7903

Negen, J., and Sarnecka, B. W. (2015). Is there really a link between exact-number knowledge and approximate number system acuity in young children? *Br. J. Dev. Psychol.* 33, 92–105. doi: 10.1111/bjdp.12071

Núñez, R. E. (2011). No innate number line in the human brain. *J. Cross. Cult. Psychol.* 42, 651–668. doi: 10.1177/0022022111406097

Núñez, R. E. (2017). Is there really an evolved capacity for number? *Trends Cogn. Sci.* 21, 409–424. doi: 10.1016/j.tics.2017.03.005

Overmann, K. A. (2016). The role of materiality in numerical cognition. *Quat. Int.* 405, 42–51. doi: 10.1016/j.quaint.2015.05.026

Overmann, K. A. (2018). Constructing a concept of number. *J. Numer. Cogn.* 4, 464–493. doi: 10.5964/jnc.v4i2.161

Pezzelle, S., Sorodoc, I.-T., and Bernardi, R. (2018). "Comparatives, quantifiers, proportions: a multi-task model for the learning of quantities from vision," in *Proceedings of the 2018 Conference of the Association for Computational Linguistics,* (New Orleans, LA: Association for Computational Linguistics), 419–430.

Piantadosi, S. T., Tenenbaum, J. B., and Goodman, N. D. (2012). Bootstrapping in a language of thought: a formal model of numerical concept learning. *Cognition* 123, 199–217. doi: 10.1016/j.cognition.2011.11.005

Piazza, M. (2010). Neurocognitive start-up tools for symbolic number representations. *Trends Cogn. Sci.* 14, 542–551. doi: 10.1016/j.tics.2010.09.008

Platt, J. R., and Johnson, D. M. (1971). Localization of position within a homogeneous behavior chain: effects of error contingencies. *Learn. Motiv.* 2, 386–414. doi: 10.1016/0023-9690(71)90020-8

Rittle-Johnson, B., Siegler, R. S., and Alibali, M. W. (2001). Developing conceptual understanding and procedural skill in mathematics: an iterative process. *J. Educ. Psychol.* 93, 346–362. doi: 10.1037/0022-0663.93.2.346

Rucinski, M., Cangelosi, A., and Belpaeme, T. (2012). "Robotic model of the contribution of gesture to learning to count," in *IEEE International Conference on Development and Learning and Epigenetic Robotics* (San Diego, CA, USA: IEEE), 1–6. doi: 10.1109/DevLrn.2012.6400579

Rumelhart, D. E., and McClelland, J. L. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations.* Cambridge, MA: MIT Press.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., and Hinton, G. E. (1986). "Schemata and sequential thought processes in PDP models," in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition Volume 2: Psychological and Biological Models* eds J. L. McClelland and D. E. Rumelhart (Cambridge, MA: MIT Press), 7–57.

Sarnecka, B. W., and Carey, S. (2008). How counting represents number: what children must learn and when they learn it. *Cognition* 108, 662–674. doi: 10.1016/j.cognition.2008.05.007

Saxton, D., Grefenstette, E., Hill, F., and Kohli, P. (2019). "Analysing mathematical reasoning abilities of neural models," in *International Conference on Learning Representations,* (New Orleans, LA: ICLR), 1–17.

Schmandt-Besserat, D. (1992). *Before Writing: From Counting to Cuneiform.* Austin, TX: University of Texas Press.

Schneider, M., Beeres, K., Coban, L., Merz, S., Susan Schmidt, S., Stricker, J., et al. (2017). Associations of non-symbolic and symbolic numerical magnitude processing with mathematical competence: a meta-analysis. *Dev. Sci.* 20:e12372. doi: 10.1111/desc.12372

Searle, J. R. (1980). Minds, brains, and programs. *Behav. Brain Sci.* 3, 417–424. doi: 10.1017/S0140525X00005756

Siegler, R. S., and Jenkins, E. (1989). *How Children Discover New Strategies.* New York, NY: Psychology Press.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., et al. (2017). Mastering the game of Go without human knowledge. *Nature* 550, 354–359. doi: 10.1038/nature24270

Skinner, B. (1953). *Science and Human Behavior.* New York, NY: MacMillan.

Starkey, G. S., and McCandliss, B. D. (2014). The emergence of "groupitizing" in children's numerical cognition. *J. Exp. Child Psychol.* 126, 120–137. doi: 10.1016/j.jecp.2014.03.006

Starr, A., Libertus, M. E., and Brannon, E. M. (2013). Number sense in infancy predicts mathematical abilities in childhood. *Proc. Natl. Acad. Sci. U S A* 110, 18116–18120. doi: 10.1073/pnas.1302751110

Stoianov, I., and Zorzi, M. (2012). Emergence of a "visual number sense" in hierarchical generative models. *Nat. Neurosci.* 15, 194–196. doi: 10.1038/nn.2996

Sutton, R. S., and Barto, A. G. (1998). *Reinforcement Learning.* Cambridge, MA: MIT Press.

Testolin, A., Dolfi, S., Rochus, M., and Zorzi, M. (2019). Perception of visual numerosity in humans and machines *arXiv:1907.06996* [Preprint].

Testolin, A., Stoianov, I., and Zorzi, M. (2017). Letter perception emerges from unsupervised deep learning and recycling of natural image features. *Nat. Hum. Behav.* 1, 657–664. doi: 10.1038/s41562-017-0186-2

Testolin, A., and Zorzi, M. (2016). Probabilistic models and generative neural networks: towards an unified framework for modeling normal and impaired neurocognitive functions. *Front. Comput. Neurosci.* 10:73. doi: 10.3389/fncom.2016.00073

Testolin, A., Zou, W. Y., and McClelland, J. L. (2020). Numerosity discrimination in deep neural networks: initial competence, developmental refinement and experience statistics. *Dev. Sci.* 12:e12940 [Epub ahead of print].

Trask, A., Hill, F., Reed, S., Rae, J., Dyer, C., and Blunsom, P. (2018). Neural arithmetic logic units. *arXiv:1808.00508* [Preprint].

Verguts, T., and Fias, W. (2004). Representation of number in animals and humans: a neural model. *J. Cogn. Neurosci.* 16, 1493–1504. doi: 10.1162/0898929042568497

Vygotsky, L. S. (1980). *Mind in Society: The Development of Higher Psychological Processes.* Cambridge, MA: Harvard University Press.

Walshaw, M. (2017). Understanding mathematical development through Vygotsky. *Res. Math. Educ.* 19, 293–309. doi: 10.1080/14794802.2017.1379728

Wever, R., and Runia, T. F. H. (2019). Subitizing with variational autoencoders. *arXiv:1808.00257* [Preprint].

Wilkey, E. D., and Ansari, D. (2019). Challenging the neurobiological link between number sense and symbolic numerical abilities. *Ann. N Y Acad. Sci.* doi: 10.1111/nyas.14225 [Epub ahead of print].

Yamins, D. L. K., and DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* 19, 356–365. doi: 10.1038/nn.4244

Zanetti, A., Testolin, A., Zorzi, M., and Wawrzynski, P. (2019). Numerosity representation in InfoGAN: an empirical study *Advances in Computational Intelligence. IWANN,* eds I. Rojas, G. Joya and A. Catala (Cham: Springer International Publishing), 49–60.

Zhao, S., Ren, H., Yuan, A., Song, J., Goodman, N., and Ermon, S. (2018). "Bias and generalization in deep generative models: an empirical study," in *32nd*

International Conference on Neural Information Processing Systems (Montréal, QC: NeurIPS), 10792–10801.

Zorzi, M., Stoianov, I., and Umiltà, C. (2005). "Computational modeling of numerical cognition," in *Handbook of Mathematical Cognition*, ed. J. Campbell (New York, NY: Psychology Press), 67–84.

Zorzi, M., and Testolin, A. (2018). An emergentist perspective on the origin of number sense. *Philos. Trans. R. Soc. B Biol. Sci.* 373:20170043. doi: 10.1098/rstb.2017.0043

Zorzi, M., Testolin, A., and Stoianov, I. P. I. P. (2013). Modeling language and cognition with deep unsupervised learning: a tutorial overview. *Front. Psychol.* 4:515. doi: 10.3389/fpsyg.2013.00515