

THE ROLE OF GENERAL INTELLIGENCE IN MATHEMATICAL REASONING

Aviv Keren *

ABSTRACT

Objects are a centerpiece of the mathematical realm and our interaction with and reasoning about it, just as they are of the physical one (if not more). And humans' mathematical reasoning must ultimately be grounded in our general intelligence. Yet in contemporary cognitive science and A.I., the physical and mathematical domains are customarily explored separately, which allows for baking in assumptions for what objects *are* for the system - and missing potential connections.

In this paper, I put the issue into its philosophical and cognitive context. I then describe an abstract theoretical framework for learning object representations, that makes room for mathematical objects *on par* with non-mathematical ones. Finally, I describe a case study that builds on that view to show how our general ability for integrating different aspects of objects effects our conception of the natural numbers.

1 BACKGROUND – MATHEMATICAL REASONING

Getting A.I. to achieve some skill often forces us to gain a better, more concrete and detailed understanding of that skill and what it requires – and the skill of mathematical reasoning is no different. But what is the current state of our understanding of it?

There is no understanding of a form of reasoning without a decent understanding of the domain which that reasoning is about, but the case of mathematics makes it particularly challenging on the former. The common conception of mathematics, in philosophy of mathematics as well as amongst working mathematicians, is of a non-physical, non-psychological, timeless realm. This gives rise to the principled Benacerraf challenge: “[T]o provide an account of the mechanisms that explain how our beliefs about these remote entities can so well reflect the facts about them.” (Benacerraf, 1973; Field, 1989)

A central component here would supposedly be mathematical language, its logic, to which us humans relate – however that language in turn relates to the mathematical realm. Yet limitations of (first-order) logic (both Gödelean incompleteness and the existence of non-standard models) prohibit it from grounding even a fundamental and seemingly simple structure such as \mathbb{N} . This nevertheless still standard (if orthodox)¹ logic-centered view also puts aside the question of how mathematics is applied, used.²

This situation, however, does not deter mathematical cognition, the scientific field in charge of investigating and accounting for math-related reasoning, from exploring its topic. This field has been growing immensely in the past few decades.³ It reveals aspects of what mathematical objects are to us, that we may not be aware of. One such example is the distance effect, whereby the discrimination of numbers turns out to depend (logarithmically) on their magnitudes rather than

*aviv.keren@gmail.com
<https://avivkeren.wixsite.com/website>

¹A new wave of Martin-Löf type-theoretic approaches to the foundations of mathematics (esp. homotopy type theory) is on the rise (Univalent Foundations Program, 2013; Grayson, 2018). While aiming to capture what mathematicians consciously intend, there is no presumption there to go *beyond* what they can perceive, to reflect rich *non-conscious* processing à la cognitive science and A.I.

²To put these delicate matters extremely crudely. For general introductions, see the classic (Benacerraf & Putnam, 1984) or the modern (Shapiro, 2005)

³For a modern intro, see e.g. (Gillmore et al., 2018)

their representation in digits. (e.g., it is in a sense harder to tell that $69 < 71$ than that $61 < 69$, in spite of the leading digit). More generally, we seem to possess an innate Approximate Number System, which is considered to underlie our so-called “number sense” and even our conception of number. This is in contrast to the aforementioned logical point of view, by which approximations seem irrelevant to the *constitution* of numbers.

The contrast with a logic-based approach runs deep, for a few important reasons. With a substantial attention to not only human-general mathematical abilities, but those of infants and even many animals, the place of language seems more of a late addendum than any sort of substantial grounding of the mathematics involved. And the evolutionary pragmatic view downplays the absolute, precise nature of mathematics anyhow (the main ingredient for which logic seems to be the natural, even necessary, candidate). Instead, it accentuates mathematics’ “applied”, real life manifestation. And even for us adult humans, language is after all ordinarily only a referential expression of an independently existing, independently perceived domain. In particular here, we don’t seem to come to be acquainted with the natural numbers via some Peano Arithmetic like axiomatization (e.g. Two numbers of which the successors are equal are themselves equal), qua fact registration. In fact, we come to see such axioms *as* true (and there is no obvious more-fundamental, psychologically-viable candidate axiomatization). Yet that we do grasp that number structure, seems like an undisputable starting point. As Dehaene (2011) puts it:

Ironically, any 5-year-old has an intimate understanding of those very numbers that the brightest logicians struggle to define. No need for a formal definition: We know intuitively what integers are. Among the infinite number of models that satisfy Peano’s axioms, we can immediately distinguish genuine integers from other meaningless and artificial fantasies. Hence our brain does not rely on axioms.

Dehaene attributes our basic numerical abilities and concepts to an *innate* “number sense”. Against this common position, others take a more *learned* view. Carey (2009), for example, proposes a theory of how we come to acquire concepts that ascend beyond a current conceptual capacity (numbers being a prime example), through a process of “Quinean bootstrapping” (in which symbols are initially only partly interpreted, while we nevertheless learn relations between them). Other positions in-between complete innateness and cultural acquisition abound (see e.g. (Spelke, 2017)), and the scientific issue is not settled (Núñez, 2017).

Computational models (including neural ones) for the distance effect and other findings can and are being given (e.g. (Zorzi & Butterworth, 1999)). But the current methodological limitations of cognitive science as an empirical field do not allow for approaching a sufficiently-definite account how of the related cognitive machinery actually functions, computationally. And so, the tension, between a mechanistically-grounded conception of numbers that somehow gives rise to their truths too, and the logical impossibility results (which apply quite generally), can be kept repressed, for now. What in particular is still missing is how numbers, and mathematical objects in general, are represented (if not merely as the logic that’s about them, their theories).

2 LEARNING GENERAL OBJECT REPRESENTATION

Objects seem to take a central place in the structure of our mental realm, at least at its highest, conscious levels. Accordingly, how our representations of objects are *learned*, or what aspects of it are, rather, supported innately, is an important if not central topic and general theme in cognitive science. It is less eminent in contemporary A.I., although the field of developmental robotics, for example, “aims to show how a robot can start with the ‘blooming, buzzing confusion’ of low-level sensorimotor interaction, and can learn higher-level symbolic structures”(Modayil & Kuipers, 2008); and although there is interest in locating the issue within contemporary settings (NeurIPS, 2020).

Mathematical objects are different from other sorts of objects in fundamental respects. But we do consider them, and seem to come to master their grasp and manipulation, as objects. (Some even talk in *literal* terms of perception (Gödel, 1947).) Ruling out a specific evolutionary adaptation suggests that this is done very much through pre-existing, general cognitive mechanisms. This poses, for both cognitive science and A.I., the central challenge that I wish to put forth here:

To account for the learning of object representations in sufficiently abstract theoretical terms, so as to pertain over both mathematical and non-mathematical objects (before getting into the differences, within such a system).

Such abstraction may be possible, I suggest, from an approach such as (Modayil & Kuipers, 2008)’s:

[i]nstead of considering an object to be a physical entity, we consider an object to be an explanation for some subset of an agent’s experience [Or: “a hypothesized entity that accounts for a spatiotemporally coherent cluster of sensory experience.”]. With this approach, the semantics of an object are intrinsically defined from the agent’s sensorimotor experience”

While the interest in robotics is in the physical domain, of sensory experience and motoric influence within the spatiotemporal, the underlying notions and principles at play need not be limited to that. Experience, in particular, also includes experience with the internal workings and patterns within the cognitive system itself (to some extent); and coherence is certainly an abstract notion to begin with. In this respect, consider our related abilities with physical objects:

The core system of object representation centers on a set of principles governing object motion: cohesion (objects move as connected and bounded wholes), continuity (objects move on connected, unobstructed paths), and contact (objects influence each others’ motion when and only when they touch). These principles allow infants.. to perceive the boundaries and shapes of objects that are visible or partly out of view, and to predict when objects will move and where they will come to rest... research with older infants suggests that a single system underlies infants’ object representations. (Kinzler & Spelke, 2007)

Even if such principles develop innately, they need not be coded in in such explicitly-physical terms. The theoretical challenge, then, is to subsume such principles under a general computational or information-theoretic notion.

In (Keren, 2019, Chapter 6), I develop a preliminary theoretical account of such an abstract framework (restricted to a specific aspect), describing the following:

1. The fundamentals regarding such a system’s treatment of types of objects (e.g. apples).
2. How such types can (in a technical sense) be extended, e.g. oranges as a kind of fruit that can be squeezed. Or combined, “amalgamated” with other types, e.g. bats as a kind of winged mammals; or smartphones as phones + computers + cameras + g.p.s...; or a 3-d structure perceived through two 2-d channels (eyes).
3. How these issues depend on the environment, on the regularities it presents. e.g., cast shadows that tend to appear under balls, being integrated into their very perception, that of their location in 3d-space.
4. What principles guide actual development – learning, mastering, and maintaining this sort of representation.

The immense scope of a full such theory aside, my own restricted focus sets to substantiate the following picture: sufficiently regular interactions between represented, coherent object-types (e.g. a person’s looks and their voice) allow for amalgamating them into a new coherent object-type, assimilating the interaction between the original object-types, that are re-conceived a *aspects* of the new, “higher level” object-type. Such higher level object-types are more valuable, involving a dimensionality reduction and offering in particular an enlarged space of affordances. Thus, the system is directed towards pursuing such amalgamations whenever it detects empirical regularities that allow for it, structuring the perceived ontology accordingly.

On a methodological note: Beyond having to stand on its own merits as a reasonable ‘story’ of how such a system would operate, the full challenge for such a theoretical account is two-fold. On the one hand, it should be able to serve as the engineering starting-point for the construction of a working A.I. system (also showing how related topics such as binding (Greff et al., 2020) and disentanglement (Higgins et al., 2018) fit in). On the other, it should square empirically, as a generalization, with the many related issues from cognitive science (the core system of object representation, multi-modal integration, various forms of binding, categorization, and more); and hopefully produce new predictions as well.

A structurally related theory in cognitive science is that of *conceptual blending* (Fauconnier & Turner, 2002). It accounts for concepts such as *houseboat* (as well as more temporary or contextual blends) as the blended "mental space" that shares some counterparts of the input spaces of e.g. house & boat, aligned across a more abstract generic space that they both share. It offers optimality principles that govern the blending process. Blending has in particular been used to account for concepts in mathematics as well (Lakoff & Núñez, 2000; Guhe et al., 2011). Various computational implementations have followed the cognitive theory, formalizing its notions and principles (Eppe et al., 2018); and even a theoretical Category Theoretic formulation that unifies these various A.I. concretizations (Schorlemmer & Plaza, 2021).

The phenomenon explored here, however, is treated as a very high-level (and rather peripheral) conceptual one, concerning (perhaps human-only) creativity. And the computational implementations are grounded in logic, accordingly. The different focus I suggest here is on such combination as a fundamental component of the cognitive construction of object-representations at large. In particular, how the regularities that govern the possibility and value of such combinations are handled, should be understood as part of the more general working of the system, picking up on and integrating *statistical* patterns into the construction of its ontology. In terms of A.I., this reflects the on-going quest to find the proper place for the symbolic within modern architectures, which are chiefly neural.

Given the suggested picture, of how an object-centered system might in general come to attribute objecthood or break it apart, we can approach the development of our conception of numbers and account for a hidden, sub-symbolic intricacy in that conception.

3 CASE STUDY: ORDINALS AND CARDINALS IN THE NATURAL NUMBERS

The common mathematical narrative is that only in going to infinity, numbers had to be broken up into ordinals and cardinals (as representatives of order-type and of quantity). In the finite, these are customarily considered as merely different aspects of the very same *natural* numbers, in how they are used. This narrative can be argued to be conceptually problematic, however, given how distinct ordinals and cardinals turned out to be, not only conceptually and in terms of their usage, but ontologically and arithmetically.

Steiner (2005) substantiates this doubt with an analysis of the intricate interplay between ordinals and cardinals already *in the finite realm*. We learn to transition between these "aspects" from the very beginning, where counting (through some particular ordering) is to produce the (order-independent) quantity (the so-called "Cardinality Principle" (Gelman & Gallistel, 1978; Sarnecka & Carey, 2008)). But the child may still wonder, "Why is it always the case that adding x , y times, gives the same result as adding y , x times?" (that is, in terms of orderings, with addition as concatenation, not quantities and disjoint union). And she would be "right" to wonder, since the commutativity of multiplication is indeed **not** a property of ordinal arithmetic, only cardinal arithmetic (if evident only in the infinite). Steiner gives a multi-stage analysis of how we transition from ordinals to cardinals, commute there through the (cardinality) isomorphism between the Cartesian products $X \times Y$ and $Y \times X$, and return to ordinals, in order to see that fact.

Steiner's analysis reveals finer workings within a seemingly extremely familiar and basic realm. Though his account does not concern with cognition or A.I., we can take it to show that we were consciously missing this complexity that our brains have nonetheless learned to manage successfully and systematically. In (Keren, 2019, Chapter 7), I account for this in terms of the framework described in the previous section, as follows.

If ordinals and cardinals be amalgamated ("combined") into numbers, then by the time we acquire the latter, the former object-types must conceivably already be grasped by the system (at least in a preliminary, *implicit* form). The first stage thus substantiates them as separate object-types, basic enough to play a role in daily life and be learnt. Ordinals are relevant for example wherever ordered operations are involved. Shifting the grasp (representation) of a song between its lines and its multi-line verses actually involves ordinal multiplication. Likewise for cardinality, a child may be able to match forks to plates when setting a table, without paying attention - representing - any particular order by which to do so. Moreover, she may understand that approaching the table from a different side doesn't change anything about the task. Such are feats that the cognitive system achieves.

Given these acquired object-types, we can consider the interaction between them. Unlike the general case, in the *restricted context* of the finite domain, their interaction presents a substantial regularity: We can move back and forth between an ordinal and its matching cardinal. And we can switch between ordinal arithmetic and cardinal arithmetic – these are coordinated accordingly. Falling right in line with the framework’s abstract principles (omitted here), these object-types can be merged coherently into a higher-level (dimension-reduced) object-type – that of numbers. This structure, which seamlessly integrates both as aspects, is precisely where the non-conscious, automated interplay that allows for “seeing” the commutativity of their multiplication resides. This is an instance of the value of the higher-level representation, towards which the system is driven.

With this ontological structure clarified, we can consider how our developmental track revolves around it (beginning with the system’s drive to amalgamate object-types when it can). For example, why would we (consciously) forget about ordinals and cardinals, while we do not forget of mammals nor winged animals upon learning of bats? How do teachers fit into this track? And more generally, how does this view square with the vast empirical research on the development of these abilities?

4 CONCLUSION

For mathematics as currently conceived, the ordinals-cardinals case study is a cautionary tale. The finite is the implicit context in which these mathematical objects happen to first be experienced and discovered or taught. The tight connection between them is in this sense a contingent, empirical regularity rather than an absolute mathematical necessity. Yet the cognitive system, by its general, generally-useful object-learning nature, comes to integrate it implicitly into our conception. Even logical codification is susceptible to the partiality of such explicit conceptions (as in coding the natural numbers through Peano axioms, which does not code the way they are used).

The more general lesson, however, I take to be more optimistic. Yes, mathematics is unique, not least because fundamental aspects of mathematical reasoning are unique. But mathematical reasoning as done by us humans, within the context of our full cognitive system, makes room for *general* intelligence to enable and support it in various ways. Perceptual patterns, in particular, may run deep, integrated into mathematical reasoning pre-symbolically. This exploration holds the promise for advancing our understanding of humans’ understanding of mathematics, and through that, indeed, of (the realm of) mathematics itself (as I have set to demonstrate with a very particular example). But for that, we must approach it all in sufficiently abstract terms. After all, our ability to reason about this unique, newly-discovered realm, is a testimony to the extreme generality of our general intelligence. (Sloman, 2008)

ACKNOWLEDGMENTS

The author thanks Amit Miller, Gilad Landau, and three anonymous referees for comments on this paper. And most of all, Carl Posy and Aaron Sloman, for supervising the PhD on which this paper is based.

REFERENCES

- Paul Benacerraf. Mathematical truth. *The Journal of Philosophy*, 70:661–679, 1973.
- Paul Benacerraf and Hilary Putnam. *Philosophy of Mathematics: Selected Readings*. Cambridge University Press, 2nd edition, 1984.
- Susan Carey. *The origin of concepts*. Oxford University Press, 2009.
- Stanislas Dehaene. *The Number Sense: How the Mind Creates Mathematics, Revised and Updated Edition*. Oxford University Press, 2011.
- Manfred Eppe, Ewen Maclean, Roberto Confalonieri, Oliver Kutz, Marco Schorlemmer, Enric Plaza, and Kai-Uwe Kühnberger. A computational framework for conceptual blending. *Artificial Intelligence*, 256:105–129, 2018. ISSN 0004-3702.
- Gilles Fauconnier and Mark Turner. *The way we think: Conceptual blending and the mind’s hidden complexities*. Basic Books, 2002.

- Hartry Field. *Realism, Mathematics, and Modality*. Blackwell, 1989.
- Rochel Gelman and C. R. Gallistel. *The child's understanding of number*. Harvard University Press, 1978.
- Camilla Gilmore, Silke M. Göbel, and Matthew Inglis. *An Introduction to Mathematical Cognition*. International Texts in Developmental Psychology. Routledge, 2018. ISBN 9781138923959.
- Daniel R. Grayson. An introduction to univalent foundations for mathematicians. *Bulletin of the American Mathematical Society*, 55(4):427–450, Mar 2018. ISSN 1088-9485.
- Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. On the Binding Problem in Artificial Neural Networks, 2020.
- Markus Guhe, Alison Pease, Alan Smaill, Maricarmen Martinez, Martin Schmidt, Helmar Gust, Kai-Uwe Kühnberger, and Ulf Krumnack. A computational account of conceptual blending in basic mathematics. *Cognitive Systems Research*, 12(3):249–265, 2011. ISSN 1389-0417. doi: <https://doi.org/10.1016/j.cogsys.2011.01.004>. URL <https://www.sciencedirect.com/science/article/pii/S1389041711000155>. Special Issue on Complex Cognition.
- Kurt Gödel. What is Cantor’s Continuum Problem? *The American Mathematical Monthly*, 54(9): 515–525, 1947.
- Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, and Alexander Lerchner. Towards a Definition of Disentangled Representations, 2018.
- Aviv Keren. *Towards a Cognitive Foundation of Mathematics*. PhD dissertation, The Hebrew University of Jerusalem, 2019.
- Katherine D. Kinzler and Elizabeth S. Spelke. Core systems in human cognition. *Progress in Brain Research*, 164:257–264, 2007.
- George Lakoff and Rafael Núñez. *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*. Basic Books, 2000.
- Joseph Modayil and Benjamin Kuipers. The initial development of object knowledge by a learning robot. *Robotics and Autonomous Systems*, 56(11), 2008.
- NeurIPS. Workshop: Object Representations for Learning and Reasoning, 2020. Thirty-fourth Conference on Neural Information Processing Systems.
- Rafael E. Núñez. Is There Really an Evolved Capacity for Number? *Trends in Cognitive Sciences*, 21(6):409–424, 2017. ISSN 1364-6613.
- Barbara W. Sarnecka and Susan Carey. How counting represents number: What children must learn and when they learn it. *Cognition*, 108(3):662–674, 2008. ISSN 0010-0277.
- Marco Schorlemmer and Enric Plaza. A uniform model of computational conceptual blending. *Cognitive Systems Research*, 65:118–137, 2021. ISSN 1389-0417.
- Stewart Shapiro. *The Oxford handbook of philosophy of mathematics and logic*. Oxford University Press, United Kingdom, 2005. ISBN 0-19-514877-0.
- Aaron Sloman. The well-designed young mathematician. *Artificial Intelligence*, 172(18):2015–2034, 2008. ISSN 0004-3702. Special Review Issue.
- Elizabeth S. Spelke. Core Knowledge, Language, and Number. *Language Learning and Development*, 13(2):147–170, 2017.
- Mark Steiner. Mathematics — Application and Applicability. In *The Oxford Handbook of Philosophy of Mathematics and Logic*. Oxford University Press, 2005.
- The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. <https://homotopytypetheory.org/book>, Institute for Advanced Study, 2013.
- Marco Zorzi and Brian Butterworth. A computational model of number comparison. In *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society*, 1999.