social to in-fact social statements. They are intended to illustrate that simple cases where associative (i.e., social learning) and propositional logic are easily distinguishable, and self-referential projection is no confound are difficult to find in actual social reasoning.

Second, the high artificiality of the experimental context and task in both of the authors' research examples should be taken into account when interpreting their results as evidence for a specific reasoning hypothesis. The experimental context itself is expected to increase participants' cognitive alertness and motivation for accuracy (e.g., Orne, 1962; Zizzo, 2010). The artificiality of most experimental reasoning tests, including the authors' examples, is further likely to encourage participants' deliberate instead of intuitive thinking regarding reasoning statements (see, as process explanation; Evans, 2008; Kahneman, 2011; recently, De Neys, 2022) as stimulus materials. In this respect, a strong interpretation of the discussed experiments might commit the same fallacy as early interpretations of human bias (e.g., Kahneman, Slovic, Slovic, & Tversky, 1982) that were later challenged to contain experimental artefacts (e.g., Gigerenzer, 1996; Hertwig, Leuker, Pachur, Spiliopoulos, & Pleskac, 2022; but also see, Vranas, 2000). I hasten to note that this is largely an inherent problem of the experimental context created by conversational norms and the idiosyncrasy of the experimental design (e.g., Schwarz, 1994, 1999), not a shortcoming by the authors. Experimental exploration is, by all means, meaningful. However, at the same time, it is just a first step to investigate a psychological phenomenon, even more so when considering social cognition phenomena like social reasoning. The experimental artificiality can be fled by also using observational and field study designs, exchanging some internal for ecological and external validity. Before the experiments, that Quilty-Dunn et al. call upon to argue for LoT in the social psychological space, have been extended to more ecologically valid contexts, generalizable claims of any sort, including the LoT hypothesis, should be modest.

Concluding, I welcome Quilty-Dunn et al.'s attempt for an exhaustive integration of the LoT hypothesis in psychological theory and empirics. Relevantly, with my commentary I do not attempt to rebut or support the LoT hypothesis. I seek to make the authors and readers aware of the fact that for a robust, that is, a persuasive, test of the LoT hypothesis in the social context, researchers cannot exclusively revert to simple experimental imitations of social reasoning. Instead, existing findings from realistic social inference-making scenarios have to be considered by the authors and observational and field experimental approaches need to be focused on in the future. Cross-cultural exploration, as an advanced extension of social psychology, would provide an additional opportunity to test the generalizability of the LoT hypothesis.

Financial support. This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

Competing interest. None.

References

Colman, A. M. (2003). Depth of strategic reasoning in games. *Trends in Cognitive Sciences*, *7*(1), 2–4. https://doi.org/10.1016/S1364-6613(02)00006-2
De Neys, W. (2022). Advancing theorizing about fast-and-slow thinking. *Behavioral and Brain Sciences*, 1–68. https://doi.org/10.1017/S0140525X2200142X
Evans, J. St. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, *59*, 255–278. https://doi.org/10.1146/annurev.psych.59.103006.093629
Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky. *Psychological Review*, *103*(3), 592–596. https://doi.org/10.1037/0033-295X.103.3.592
Grüning, D. J., & Krueger, J. I. (2021). Strategic thinking: A random walk into the rabbit hole. *Collabra: Psychology*, *7*(1), 24921. https://doi.org/10.1525/collabra.24921
Grüning, D. J., & Krueger, J. I. (2022). Strategic thinking in the shadow of self-enhancement: Benefits and costs. *PsyArXiv*. https://doi.org/10.31234/osf.io/gtc2m
Hedden, T., & Zhang, J. (2002). What do you think I think you think?: Strategic reasoning in matrix games. *Cognition*, *85*(1), 1–36. https://doi.org/10.1016/S0010-0277(02)00054-9
Hertwig, R., Leuker, C., Pachur, T., Spiliopoulos, L., & Pleskac, T. J. (2022). Studies in ecological rationality. *Topics in Cognitive Science*, *14*(3), 467–491. https://doi.org/10.1111/tops.12567
Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
Kahneman, D., Slovic, S. P., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.
Krueger, J. I. (2008). From social projection to social behaviour. *European Review of Social Psychology*, *18*(1), 1–35. https://doi.org/10.1080/10463280701284645
Krueger, J. I. (2013). Social projection as a source of cooperation. *Current Directions in Psychological Science*, *22*(4), 289–294. https://doi.org/10.1177/0963721413481352
Krueger, J. I., & Grüning, D. J. (2021). Psychological perversities and populism. In J. P. Forgas, W. D. Crano, & K. Fiedler (Eds.), *The psychology of populism* (pp. 125–142). Routledge.
Krueger, J. I., Grüning, D. J., & Heck, P. R. (2023). Inductive reasoning model. *PsyArXiv*. https://doi.org/10.31234/osf.io/3yasf
Kurdi, B., & Dunham, Y. (2021). Sensitivity of implicit evaluations to accurate and erroneous propositional inferences. *Cognition*, *214*, 104792. https://doi.org/10.1016/j.cognition.2021.104792
Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, *17*(11), 776–783. https://doi.org/10.1037/h0043424
Schwarz, N. (1994). Judgment in a social context: Biases, shortcomings, and the logic of conversation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 26, pp. 123–162). Academic Press. https://doi.org/10.1016/S0065-2601(08)60153-7
Schwarz, N. (1999). Self-reports: How the questions shape the answers. *American Psychologist*, *54*(2), 93–105. https://doi.org/10.1037/0003-066X.54.2.93
Vranas, P. B. (2000). Gigerenzer's normative critique of Kahneman and Tversky. *Cognition*, *76*(3), 179–193. https://doi.org/10.1016/S0010-0277(99)00084-0
Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics*, *13*, 75–98. https://doi.org/10.1007/s10683-009-9230-z

Compositionality in visual perception

Alon Hafri[a], E. J. Green[b] and Chaz Firestone[c]

[a]Department of Linguistics and Cognitive Science, University of Delaware, Newark, DE, USA; [b]Department of Linguistics and Philosophy, Massachusetts Institute of Technology, Cambridge, MA, USA and [c]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD, USA
alon@udel.edu; https://pal.lingcogsci.udel.edu/
ejgr@mit.edu; https://sites.google.com/site/greenedwinj/
chaz@jhu.edu; https://perception.jhu.edu/

doi:10.1017/S0140525X23001838, e277

Abstract

Quilty-Dunn et al.'s wide-ranging defense of the Language of Thought Hypothesis (LoTH) argues that vision traffics in abstract, structured representational formats. We agree: Vision, like language, is *compositional* – just as words compose into phrases, many visual representations contain discrete constituents that combine in systematic ways. Here, we amass evidence extending this proposal, and explore its implications for how vision interfaces with the rest of the mind.

The world we see is populated by colors, textures, edges, and countless other visual features. Yet we see more than a collection of features: We also see whole objects, and relations within and between those objects. How are these entities represented? Here, we advance the case for LoT-like representation in perception. We argue that at least two types of visual representations are compositional, and we explore their connections with the rest of the mind.

Consider the hands in Figure 1A. Although they differ in various superficial features, they appear to share something: their *structure* – specifically, their *skeletal structure*. The same parts are connected in the same ways, just in different poses. Similarly, the middle shape in Figure 1B shares its structure with the left shape but not the right shape, even though the middle and right shapes share other features. Skeletal representations describe shapes via their parts' intrinsic axes and connections, often in a hierarchical *tree* format, wherein certain parts "descend" or "offshoot" from others (Feldman & Singh, 2006). Copious evidence suggests that skeletal representations are psychologically real, implicated in detection (Kovács & Julesz, 1994; Wilder, Feldman, & Singh, 2016), discrimination (Lowet, Firestone, & Scholl, 2018), categorization (Wilder, Feldman, & Singh, 2011), aesthetics (Van Tonder, Lyons, & Ejima, 2002), and more (Firestone & Scholl, 2014; Psotka, 1978).

We contend that skeletal representations exhibit several of Quilty-Dunn et al.'s LoT properties: *Discrete constituents*, *role-filler independence*, and *abstract content*. First, skeletal representations contain discrete constituents that represent axis structure independently of surrounding boundaries, composing with boundary representations to describe overall shape. This may explain why infants (Ayzenberg & Lourenco, 2022) and adults (Wilder et al., 2011) categorize novel shapes by skeletal structure despite differences in surface properties. Second, representations of individual parts exhibit role-filler independence, retaining identity over changes in position within the overall skeletal representation. Such *transportability* (Fodor, 1987) explains why we can easily determine when distinct shapes share the same parts, and why such shapes prime one another (Cacciamani, Ayars, & Peterson, 2014). Third, skeletal representations are abstract, expressing aspects of shape that appear stable despite part articulations (Fig. 1A), changes in surface properties (Fig. 1B; Green, 2019), and sense modality (Green, 2022). Moreover, visual brain areas encode skeletal structure across surface changes (Ayzenberg, Kamps, Dilks, & Lourenco, 2022; Hung, Carlson, & Connor, 2012; Lescroart & Biederman, 2013). Skeletal representations may also encode nonmetric, categorical properties – for example, *straight/curved* and *symmetric/asymmetric* (Amir, Biederman, & Hayworth, 2012; Green, 2017; Hafri, Gleitman, Landau, & Trueswell, 2023).

We suggest that these LoT properties make skeletal representations compositional: Discrete constituents encoding different geometrical elements and properties combine to form representations of global shape.

Compositionality in vision extends to relations *between* objects. Consider the object pairs in Figure 1C. They appear to share something: the relation *containment*. Visual processing respects this commonality – it represents relations between objects, beyond the objects themselves (Hafri & Firestone, 2021). Such representations also exhibit several LoT properties. First, visual processing represents relations abstractly and categorically: Observers are more sensitive to metric changes across relational category boundaries (e.g., from containing to merely touching) than within (e.g., from one instance of containment to another; Lovett & Franconeri, 2017), and even "confuse" instances of the same relation for one another (Hafri, Bonner, Landau, & Firestone, 2020). Furthermore, visual brain areas encode eventive relations abstractly, generalizing across event participants (Hafri, Trueswell, & Epstein, 2017; Wurm & Lingnau, 2015).

Second, such representations contain discrete constituents and exhibit role-filler independence, in ways that augment Quilty-Dunn et al.'s discussion. Consider Figure 1D. Both images

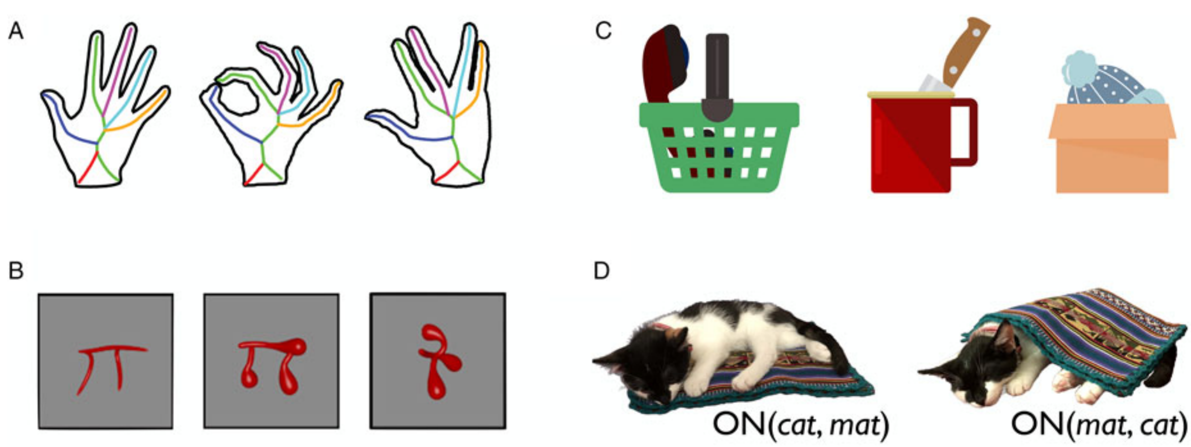

Figure 1 (Hafri et al.). Demonstrations of compositionality in visual perception. (A) The three hands shown here differ in global shape, the locations of their boundaries, and other surface features; however, they appear to share something: Their *structure* – specifically, their *skeletal structure* (indicated by the inset colored lines). The same parts have taken on different poses. Skeletal shape representations describe objects in terms of the axes of their parts, including how those parts are arranged with respect to one another, in ways that instantiate several core LoT properties. (Adapted from Lowet et al., 2018.) (B) Skeletal shape representations explain why infants and adults can see that the middle shape shares something with the leftmost shape that it does not share with the rightmost shape, even though the middle and rightmost shapes share other features. (Adapted from Ayzenberg & Lourenco, 2019.) (C) The three object pairs shown here differ in a variety of visual features, and even involve different objects – but each seems to instantiate the same relation: *containment*. Recent evidence suggests that the mind rapidly and automatically encodes such relations, representing the relation itself separately from the objects participating in it. (Adapted from Hafri et al., 2020.) (D) These two images depict the same objects (cat and mat) and the same relation (*support*), but differ in their *structure* – a cat on a mat is a very different scene from a mat on a cat. Put differently, "argument order" matters: $R(x,y)$ may be quite different than $R(y,x)$, and there is evidence that visual processing is sensitive to this difference in compositional structure. (Adapted from Hafri & Firestone, 2021.)

https://doi.org/10.1017/S0140525X22002849 Published online by Cambridge University Press

involve the same objects (cat and mat) and relation (*support*), but cat-on-mat differs from mat-on-cat in compositional structure. Thus, "argument order" matters – the "fillers" map to different roles. Recent work shows that vision is sensitive to this difference. When observers repeatedly reported the location of a target individual (e.g., blue-shirted man) in a stream of action photographs (e.g., blue-kicking-red, red-pushing-blue), a "switching cost" emerged: Slower responses when the target individual's role (*Agent*/*Patient*) switched (e.g., pusher on trial $n-1$ but kickee on trial n), suggesting that observers encoded relational structure automatically (Hafri, Trueswell, & Strickland, 2018).

These properties make representations of categorical between-object relations compositional: Discrete constituents encoding entities and relations combine to form representations of structured situations.

The prospect of LoT-like, compositional visual representations impacts broader debates about perception's *format*. Many claim that perceptual representations are constitutively iconic, analog, or "picture-like" (Burge, 2022; Carey, 2009; Dretske, 1981; Kosslyn, Thompson, & Ganis, 2006). However, although LoT-like formats clearly suffice to encode categorical, nondegreed relations (e.g., containment), many iconic formats may not – particularly accounts requiring perceptual icons to mirror graded degrees of difference in perceptible properties (e.g., orientation or brightness; Block, 2023).

This perspective also raises exciting questions and research directions. For example, it may partially explain how information from perception is "readily consumed" by cognitive and linguistic systems (because of the similar formats of some perceptual and higher-level representations; Cavanagh, 2021; Quilty-Dunn, 2020). Recent work explores these connections explicitly: Skeletal shape representations impact aesthetic preferences and linguistic descriptions of shapes (Sun & Firestone, 2022a, 2022b), and representations of symmetry and roles may be shared across perception and language (Hafri et al., 2018, 2023; Rissman & Majid, 2019; Strickland, 2017). One could also investigate the "psychophysics" of compositional processes – the timing and ordering of how relational representations are built from their parts.

Nevertheless, LoT-like perceptual representations may not be *fully* language-like. Although perception plausibly predicates properties of individuals (Quilty-Dunn & Green, 2023), it may lack the full *expressive freedom* of first-order logic (Camp, 2018), especially logical connectives needed for truth-functional completeness (Mandelbaum et al., 2022). Perception may be able to represent that an object is red but not that it is *not* red. Moreover, certain perceptual formats may impose constraints on which properties are attributable to which individuals – constraints absent from higher-level cognition. Perhaps perception cannot explicitly represent relations between nonadjacent object parts, or eventive relations of long durations (e.g., a jack slowly lifting a car).

Because perception and thought confront multifarious tasks with different computational demands, we contend that they comprise a multiplicity of formats (Marr, 1982; Yousif, 2022), each optimized for different computations, and some more LoT-like than others. Thus, any theory positing a single-privileged format for perception or thought should be met with suspicion. Instead, researchers should heed Quilty-Dunn et al.'s advice to "let a thousand representational formats bloom" (target article, sect. 2, para. 2).

Acknowledgments. For comments on an earlier draft, the authors acknowledge members of the JHU Perception and Mind Laboratory.

Financial support. This study was supported by NSF BCS no. 2021053 awarded to C. F.

Competing interest. None.

References

Amir, O., Biederman, I., & Hayworth, K. J. (2012). Sensitivity to nonaccidental properties across various shape dimensions. *Vision Research*, *62*, 35–43.
Ayzenberg, V., Kamps, F. S., Dilks, D. D., & Lourenco, S. F. (2022). Skeletal representations of shape in the human visual cortex. *Neuropsychologia*, *164*, 108092.
Ayzenberg, V., & Lourenco, S. F. (2019). Skeletal descriptions of shape provide unique perceptual information for object recognition. *Scientific Reports*, *9*, 1–13.
Ayzenberg, V., & Lourenco, S. F. (2022). Perception of an object's global shape is best described by a model of skeletal structure in human infants. *eLife*, *11*, e74943.
Block, N. (2023). *The border between seeing and thinking*. Oxford University Press.
Burge, T. (2022). *Perception: First form of mind*. Oxford University Press.
Cacciamani, L., Ayars, A. A., & Peterson, M. A. (2014). Spatially rearranged object parts can facilitate perception of intact whole objects. *Frontiers in Psychology*, *5*, 482.
Camp, E. (2018). Why maps are not propositional. In A. Grzankowski & M. Montague (Eds.), *Non-propositional intentionality* (pp. 19–45). Oxford University Press.
Carey, S. (2009). *The origin of concepts*. Oxford University Press.
Cavanagh, P. (2021). The language of vision. *Perception*, *50*, 195–215.
Dretske, F. I. (1981). *Knowledge and the flow of information*. MIT Press.
Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences of the United States of America*, *103*, 18014–18019.
Firestone, C., & Scholl, B. J. (2014). "Please tap the shape, anywhere you like": Shape skeletons in human vision revealed by an exceedingly simple measure. *Psychological Science*, *25*, 377–386.
Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. MIT Press.
Green, E. J. (2017). A layered view of shape perception. *The British Journal for the Philosophy of Science*, *68*, 355–387.
Green, E. J. (2019). On the perception of structure. *Noûs*, *53*, 564–592.
Green, E. J. (2022). The puzzle of cross-modal shape experience. *Noûs*, *56*, 867–896.
Hafri, A., Bonner, M. F., Landau, B., & Firestone, C. (2020). A phone in a basket looks like a knife in a cup: Role-filler independence in visual processing. *PsyArXiv*. https://psyarxiv.com/jx4yg
Hafri, A., & Firestone, C. (2021). The perception of relations. *Trends in Cognitive Sciences*, *25*, 475–492.
Hafri, A., Gleitman, L. R., Landau, B., & Trueswell, J. C. (2023). Where word and world meet: Language and vision share an abstract representation of symmetry. *Journal of Experimental Psychology: General*, *152*, 509–527.
Hafri, A., Trueswell, J. C., & Epstein, R. A. (2017). Neural representations of observed actions generalize across static and dynamic visual input. *The Journal of Neuroscience*, *37*, 3056–3071.
Hafri, A., Trueswell, J. C., & Strickland, B. (2018). Encoding of event roles from visual scenes is rapid, spontaneous, and interacts with higher-level visual processing. *Cognition*, *175*, 36–52.
Hung, C. C., Carlson, E. T., & Connor, C. E. (2012). Medial axis shape coding in macaque inferotemporal cortex. *Neuron*, *74*, 1099–1113.
Kosslyn, S. M., Thompson, W. L., & Ganis, G. (2006). *The case for mental imagery*. Oxford University Press.
Kovács, I., & Julesz, B. (1994). Perceptual sensitivity maps within globally defined visual shapes. *Nature*, *370*, 644–646.
Lescroart, M. D., & Biederman, I. (2013). Cortical representation of medial axis structure. *Cerebral Cortex*, *23*, 629–637.
Lovett, A., & Franconeri, S. L. (2017). Topological relations between objects are categorically coded. *Psychological Science*, *28*, 1408–1418.
Lowet, A. S., Firestone, C., & Scholl, B. J. (2018). Seeing structure: Shape skeletons modulate perceived similarity. *Attention, Perception, & Psychophysics*, *80*, 1278–1289.
Mandelbaum, E., Dunham, Y., Feiman, R., Firestone, C., Green, E. J., Harris, D., … Quilty-Dunn, J. (2022). Problems and mysteries of the many languages of thought. *Cognitive Science*, *46*, e13225.
Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. W.H. Freeman.
Psotka, J. (1978). Perceptual processes that may create stick figures and balance. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 101–111.
Quilty-Dunn, J. (2020). Concepts and predication from perception to cognition. *Philosophical Issues*, *30*, 273–292.
Quilty-Dunn, J., & Green, E. J. (2023). Perceptual attribution and perceptual reference. *Philosophy and Phenomenological Research*, *106*, 273–298. doi: 10.1111/phpr.12847.
Rissman, L., & Majid, A. (2019). Thematic roles: Core knowledge or linguistic construct? *Psychonomic Bulletin & Review*, *26*, 1850–1869.
Strickland, B. (2017). Language reflects "core" cognition: A new theory about the origin of cross-linguistic regularities. *Cognitive Science*, *41*, 70–101.

Sun, Z., & Firestone, C. (2022a). Beautiful on the inside: Aesthetic preferences and the skeletal complexity of shapes. *Perception*, *51*, 904–918.
Sun, Z., & Firestone, C. (2022b). Seeing and speaking: How verbal "description length" encodes visual complexity. *Journal of Experimental Psychology: General*, *151*, 82–96.
Van Tonder, G. J., Lyons, M. J., & Ejima, Y. (2002). Visual structure of a Japanese Zen garden. *Nature*, *419*, 359–360.
Wilder, J., Feldman, J., & Singh, M. (2011). Superordinate shape classification using natural shape statistics. *Cognition*, *119*, 325–340.
Wilder, J., Feldman, J., & Singh, M. (2016). The role of shape complexity in the detection of closed contours. *Vision Research*, *126*, 220–231.
Wurm, M. F., & Lingnau, A. (2015). Decoding actions at different levels of abstraction. *Journal of Neuroscience*, *35*, 7727–7735.
Yousif, S. R. (2022). Redundancy and reducibility in the formats of spatial representations. *Perspectives on Psychological Science*, *17*, 1778–1793.

Incomplete language-of-thought in infancy

Jean-Rémy Hochmann[a,b]

[a]CNRS UMR5229 – Institut des Sciences Cognitives Marc Jeannerod, Bron, France and [b]Université Lyon 1 Claude Bernard, Lyon, France
hochmann@isc.cnrs.fr
https://sites.google.com/site/jrhochmann/

doi:10.1017/S0140525X23001826, e278

Abstract

The view that infants possess a full-fledged propositional language-of-thought (LoT) is appealing, providing a unifying account for infants' precocious reasoning skills in many domains. However, careful appraisal of empirical evidence suggests that there is still no convincing evidence that infants possess discrete representations of abstract relations, suggesting that infants' LoT remains incomplete. Parallel arguments hold for perception.

The view that infants possess a propositional language-of-thought (LoT) appeals as a unifying account for precocious physical (Stahl & Feigenson, 2015), logical (Cesana-Arlotti et al., 2018), probabilistic (Denison & Xu, 2010; Téglás, Girotto, Gonzalez, & Bonatti, 2007), and social reasoning (Baillargeon, Scott, & He, 2010; Hamlin, Wynn, & Bloom, 2007; Powell & Spelke, 2013). It suggests continuity along development in the format of human thought. But arguing for such continuity also raises questions. Most, if not all, of the cognitive skills of young infants are also documented in nonhuman species (Engelmann et al., 2022; Krupenye, Kano, Hirata, Call, & Tomasello, 2016), suggesting continuity along evolution. We should thus attribute the same type of thoughts to nonhuman animals and human infants, to animals and human adults. How, then, do we account for animals' failure to acquire human natural languages and develop unique human cognitive skills? Careful appraisal of the available data and careful experimental designs may instead highlight important discontinuities in the format of thought along both developmental and evolutionary scales, suggesting that a full-fledged LoT, involving all six properties identified by Quilty-Dunn et al. is not yet available to young infants (nor to animals).

I applaud the project of Quilty-Dunn et al. to list specific properties of a propositional LoT and evaluate the presence of these properties in various subdomains of cognitive science. The strength of the evidence for each property in all domains is however unequal. In particular, before concluding that infants possess a full-fledged LoT, we need to provide evidence for each property, individually, and also investigate the limits of each property. I will focus on the first property, "discrete constituents." It is the most important, as it is presupposed by most other properties: Roles are attributed to discrete constituents; predication combines discrete constituents; logical operators are conceived as discrete constituents. Contrary to Quilty-Dunn et al., I will argue that, although both perception and infant cognition certainly possess discrete representations of objects and possibly of features, there is no evidence for discrete representations of relations in perception nor in prelexical infants.

Although experimental evidence suggests that perceptual representations of relational events and scenes are generalizable to a certain extent (e.g., Goupil, Papeo, & Hochmann, 2022; Papeo, 2020; see Kominsky & Scholl, 2020, for the limits of those generalizations), there is no evidence that those representations are discrete, dissociated from the object representations. Rather, relations may well be represented by perceptual schema composed of discrete *object* representations. The generalizability can be obtained through the underspecification of object representations, a process we previously called "abstraction by impoverishment" (Hochmann & Papeo, 2021). For instance, in perception, a schematic social interaction would consist in two schematic bodies facing each other (Papeo, 2020), a schematic relation of support would consist in an empty object file on top of another empty object file, and so on. Similar representations, with object files possibly enriched with thematic roles, may account for the representation of many relational events in infancy (Leslie & Keeble, 1987; Rochat, Striano, & Morgan, 2004; Tatone, Geraci, & Csibra, 2015).

We recently provided direct evidence supporting the proposal that prelexical infants lack discrete representations for abstract relations (Hochmann, 2022). We showed that infants can represent the relation same in a format that is abstract, as it can generalize to novel instances of the relation. However those representations are limited to four same individuals, suggesting that the format of infants' representations is not something like $S(A,B)$, where A and B would be object representations and S the representation of the relation between those objects, but rather $(X\ X)$, where X is a variable for an object (see Hochmann, 2022, for the full argumentation). The repetition of the variable carries the relational content *same*, but only symbols for objects are explicitly represented. This view is reinforced by the systematic failure of young children and other animal species in the relational match-to-sample task, where they need to match pairs of the same or different images (e.g., matching square–square to circle–circle and square-star to moon-triangle). If infants and young children possessed discrete symbols S and D for the relations same and different, they should activate S for both square–square and circle–circle, and D for both square-star and moon-triangle, and easily match S to S or D to D. Instead children fail until the age of 4, and only succeed when actively using the words "same" and "different" (Hochmann et al., 2017). Likewise, chimpanzees (and other animal species) fail the relational match-to-sample task, unless they previously acquired external unitary symbols that refer to the relations same and different (Premack, 1983; Thompson, Oden, & Boysen, 1997). These observations highlight a discontinuity along human development. They put forward the hypothesis that relations are initially represented in mental models, and that discrete representations of relations are related to the acquisition of words for