Trends in Cognitive Sciences

CellPress

25th Anniversary Series: Looking Forward

Forum

Episodes of experience and generative intelligence

Linda B. Smith [1,*] and Hadar Karmazyn-Raz[1]

How do humans, including toddlers, take knowledge from past experiences and apply this knowledge in new ways? Current approaches to human and artificial intelligence (AI) fail to offer satisfactory explanations. We suggest the explanation will be found in the coherence statistics of the individual time-extended episodes of human experience and the cognitive processes those statistics engage.

Innovative generalization is an unrealized goal of machine learning [1], but child's play for humans. Consider this observed example: a 1-year-old toddler is given a stuffed toy, an elephant, with four cup-shaped shoes made of felt on the four legs. Despite never having seen this configuration before, the toddler immediately knows the cup-like things are shoes and directly compares them to his own shoes. Later, the child takes cups out of the cupboard and fits them on his own feet as shoes [2]. How did the toddler think to do this?

The traditional assumption in human and machine learning is that far generalization requires experience with many variable training examples. This assumption is challenged by observations such as the toddler's charming use of cups as shoes. Humans, even young ones, appear to discover regularities from a few specific experiences and readily generate novel instances [2]. The field has limited insight into the nature of the internal representations that could accomplish innovative generalization and the learning machinery that could give rise to those representations [1]. The missing ingredient for understanding both human cognitive development and reverse-engineering human intelligence may lie in the precision study of individual episodes of real-world experience, the kinds of specific experience, for example, that toddlers have with shoes and cups. By way of analogy: classic approaches have tried to optimize generalization by capturing massive forests of training data; however, the solution may lie in the details of the individual trees of lived experience.

Observations of the lives of 1-year-old toddlers suggest that an individual 'shoes go on feet' experience might go something like this: the parent sits the toddler on the floor with a toy giraffe while the parent retrieves socks from a drawer. The toddler finds a cheerio on the floor and puts the giraffe down. The parent puts one sock on the child's foot, taking away the old cheerio just as the toddler is about to put it in their mouth. The parent re-gives the child the giraffe before putting the second sock on the other foot. The parent reaches for shoes nearby but, at the same time, the dog tries to take one. Meanwhile, the child has pulled one sock off. The parent puts the sock back on. The dog approaches the shoe again; there is laughter and talk about the dog, but the parent holds onto the shoe and gives it to the toddler while putting the other shoe on the toddler's foot. The child finds another old cheerio, it is taken away, and the child is given the giraffe and continues to handle the giraffe as the parent puts the second shoe on the child. This example illustrates how day-in and day-out experiences are extended-in-time episodes characterized by regularities (shoes go on feet) but loaded with idiosyncrasies. Multiple factors (yesterday's dropped cheerios, the child's momentary goals, the dog's behavior, the parent's persistent objective of getting the shoes on the feet) converge to create a never-to-be-exactly-repeated episode of 'shoes go on feet'.

Given these idiosyncrasies, how could a learner discover the regularity 'shoes go on feet?' Current approaches to AI gather up a forest of 'shoes on feet' experiences and feed them to a learning algorithm designed to find optimal features that recognize all and only 'shoes on feet' and reject negative cases (cups on feet?). This approach has multiple weaknesses, including that the learner must pre-know the relevance of the experience, that an episode is about 'shoes on feet' and not, say, 'do not eat old cheerios off the floor'. Given presorted (labeled) experiences, the algorithm might be able to recognize all varieties of shoes, but what about 'cups on feet?' Could such an algorithm, in a catastrophe such as a flood or fire, find a functional substitute for shoes among available non-shoes? Humans regularly make these functional leaps. We believe that the precision study of the natural statistics of individual episodes of experience, and the memory processes that make and learn from these statistics, will provide the solution.

The everyday episodes that drive human cognitive development are, similar to the toddler's 'shoes on feet' episode, created through human behavior: those of the young learner and the learner's social partners. It is well known that human behavior generates time-series of interconnected events characterized by a suite of statistics, including Zipfian-like frequency distributions [3], burstiness [4], and co-occurrence and transitional patterns that are well described by small-world networks [5]. These statistics also characterize episodes of toddlers' everyday experiences [6,7] and result in experiences that are structured like a story: these experiences have beginnings, middles, ends, main themes ('shoes on feet'), subthemes ('dogs and shoes'), and patterns of repetition at shorter and longer lags. Graph theoretic analyses reveal that toddlers' everyday experiences, similar to strings of human-generated events more generally, exhibit coherent network structure: a few main theme elements [sock–shoe–foot] connect all the events, including the idiosyncrasies [dogs, giraffes, cheerios], into a coherent story-like whole.

The natural statistics of human-generated events are generally understood as emerging from a complex system (the human cognitive system) that has a 'memory', such that each generated event influences the likelihood of future events [4]. The relevant memories are the transient activations that emerge as long-range recurrent feedback loops across multiple cortical networks and are commonly known as working memories; these transient memories maintain recent events in the episode as well as predictable future events in both active and ready-to-be-reactivated states for the duration of a context-bound episode [8,9]. These transient memories are known to guide attention and action in the moment and, thus, are the likely causes of the moment-to-moment behavior acts within the episode. Evidence from human and animal research shows that persistent transient memories form when supported by the coherence statistics described previously and persistent transient memories create durable and integrated memories of single episodes [8–10].

These transient memory processes are fundamental to learning. We propose they are also fundamental to innovation. Transient memory activations driven by sensory-motor input reactivate more permanent memories. Thus, the sight of a four-legged elephant with black cups at the bottom of each leg may remind (and activate) memories of the relation 'shoes on feet' on other days, strengthening that concept, as well as expanding the notion of what can serve as a shoe. Later, the sight of cups in the cupboard and the toddler's own bare feet may activate the earlier experience of the elephant's cup-like shoe and suggest the behavior of putting the cups on the feet. This can occur only if durable memories retain enough detail of specific experiences. Notice also, if this is the case, the memory of a single episode may contribute to the learning of multiple concepts, such as 'shoes go on feet', 'parent will try to distract you with a toy', as well as 'don't eat old cheerios off the floor'.

These ideas differ from current approaches in cognitive science (but see [11]). The relevant input is story-like episodes of experiences: individual trees not forests. The input is created by humans in real-time as the lesson unfolds. The learner has a direct role in creating the temporal and co-occurrence patterns. The internal mechanisms that create the stream of data for learning are also involved in learning. These ideas are empirically and computationally underdetermined. Nevertheless, they have a grounding in human, and toddler, everyday behavior and the underlying memory processes that create and learn from lived experience.

These ideas also suggest a path forward: cognitive scientists should study the time-extended coherence statistics of lived experience. Time-extended story-like episodes of behavior-generated experience are the context in which human intelligence evolved and develops. This is also the context in which autonomous intelligence systems must function. The missing secret ingredient to understanding (and reversing engineering) human intelligence may be the 'stories' of everyday life.

Acknowledgments

We thank Erin Anderson and the entire Cognitive Developmental Laboratory at Indiana University for influencing our thinking on these issues.

Declaration of interests

None declared by authors.

[1]Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN, USA

*Correspondence:
smith4@indiana.edu (L.B. Smith).

https://doi.org/10.1016/j.tics.2022.09.012

References

1. Bengio, Y. *et al.* (2021) Deep learning for AI. *Commun. ACM* 64, 58–65
2. Shore, C. *et al.* (1984) First sentences in language and symbolic play. *Dev. Psychol.* 20, 872
3. Zipf, G.K. (1949) *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*, Addison-Wesley Press
4. Barabási, A.L. (2010) *Bursts: The Hidden Patterns Behind Everything We Do, From Your E-Mail to Bloody Crusades*, Penguin
5. Watts, D.J. and Strogatz, S.H. (1998) Collective dynamics of 'small-world' networks. *Nature* 393, 440–442
6. Karmazyn-Raz, H. and Smith, L.B. (2022) Discourse with few words: coherence statistics, parent-infant actions on objects, and object names. *Lang. Acquis.* Published online July 4, 2022. https://doi.org/10.1080/10489223.2022.2054342
7. Karmazyn-Raz, H. and Smith, L.B. Sampling statistics are like story creation: a network analysis of parent-toddler exploratory play. *Philos. Trans. R. Soc. B: Biol. Sci.* (in press)
8. Olivers, C.N. and Roelfsema, P.R. (2020) Attention for action in visual working memory. *Cortex* 131, 179–194
9. Manohar, S.G. *et al.* (2019) Neural mechanisms of attending to items in working memory. *Neurosci. Biobehav. Rev.* 101, 1–12
10. Hebscher, M. *et al.* (2019) Rapid cortical plasticity supports long-term memory formation. *Trends Cogn. Sci.* 23, 989–1002
11. Chan, S.C. *et al.* (2022) Data distributional properties drive emergent in-context learning in transformers. *arXiv* 2022, arXiv:2205.05055