



Placing language in an integrated understanding system: Next steps toward human-level performance in neural language models

James L. McClelland^{a,b,1}, Felix Hill^{b,1}, Maja Rudolph^{c,1}, Jason Baldridge^{d,1,2}, and Hinrich Schütze^{e,1,2}

Edited by Miguel P. Eckstein, University of California, Santa Barbara, CA, and accepted by Editorial Board Member Michael S. Gazzaniga August 14, 2020 (received for review December 6, 2019)

Language is crucial for human intelligence, but what exactly is its role? We take language to be a part of a system for understanding and communicating about situations. In humans, these abilities emerge gradually from experience and depend on domain-general principles of biological neural networks: connection-based learning, distributed representation, and context-sensitive, mutual constraint satisfaction-based processing. Current artificial language processing systems rely on the same domain general principles, embodied in artificial neural networks. Indeed, recent progress in this field depends on query-based attention, which extends the ability of these systems to exploit context and has contributed to remarkable breakthroughs. Nevertheless, most current models focus exclusively on language-internal tasks, limiting their ability to perform tasks that depend on understanding situations. These systems also lack memory for the contents of prior situations outside of a fixed contextual span. We describe the organization of the brain's distributed understanding system, which includes a fast learning system that addresses the memory problem. We sketch a framework for future models of understanding drawing equally on cognitive neuroscience and artificial intelligence and exploiting query-based attention. We highlight relevant current directions and consider further developments needed to fully capture human-level language understanding in a computational system.

natural language understanding | deep learning | situation models | cognitive neuroscience | artificial intelligence

Striking recent advances in machine intelligence have appeared in language tasks. Machines better transcribe speech and respond in ever more natural sounding voices. Widely available applications allow one to say something in one language and hear its translation in another. Humans still perform better than machines in most language tasks, but these systems work well enough to be used by billions of people every day.

What underlies these successes? What limitations do they face? We argue that progress has come from exploiting principles of neural computation employed by the human brain, while a key limitation is that these systems treat language as if it can stand alone. We propose that language works in concert with other inputs to understand and communicate about situations. We describe key aspects of human understanding and key

components of the brain's understanding system. We then propose next steps toward a model informed by both cognitive neuroscience and artificial intelligence and point to extensions addressing understanding of abstract situations.

Principles of Neural Computation

The principles of neural computation are domain general, inspired by the human brain and human abilities. They were first articulated in the 1950s (1) and further developed in the 1980s in the parallel distributed processing (PDP) framework for modeling cognition (2). A central principle in this work is the idea that cognition depends on mutual constraint satisfaction (3). For example, interpreting a sentence requires resolving both syntactic and semantic ambiguity. If we hear "a boy hit

^aDepartment of Psychology, Stanford University, Stanford, CA 94305; ^bDeepMind, London N1C 4AG, United Kingdom; ^cBosch Center for Artificial Intelligence, Renningen 71272, Germany; ^dGoogle Research, Austin, TX 78701; and ^eCenter for Information and Language Processing, Ludwig Maximilian University of Munich, Munich 80538, Germany

Author contributions: J.L.M., F.H., M.R., J.B., and H.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. M.P.E. is a guest editor invited by the Editorial Board.

Published under the PNAS license.

¹To whom correspondence may be addressed. Email: jlmcc@stanford.edu, felixhill@google.com, majarita.rudolph@de.bosch.com, jasonbaldridge@google.com, or inquiries@cislmu.org.

²J.B. and H.S. contributed equally to this work.

First published September 28, 2020.

Jack and Jill went up the hill.

The pole vault was the last event.

Fig. 1. Context influences the identification of letters in written text: the visual input we read as “went” in the first sentence and “event” in the second is the same bit of Rumelhart’s handwriting, cut and pasted into each context. Adapted with permission from figure 3 of ref. 3.

a man with a bat,” we tend to treat “with a bat” as the instrument used to hit (semantics) and therefore as part of the verb phrase of the sentence (syntax). However, if “beard” replaces “bat,” then “with a beard” is an attribute of the man and is treated as part of a noun phrase headed by “the man” (4). Even segmenting language into elementary units depends on meaning and context (Fig. 1). Rumelhart (3) envisioned a model in which estimates of the probabilities of all aspects of an input constrain estimates of the probabilities of all others, motivating a model of context effects in perception (5) that launched the PDP approach.

This work also introduced the idea that structure in cognition and language is emergent: it is captured in learned connection weights supporting the construction of context-sensitive representations whose characteristics reflect a gradual, input statistics-dependent, learning process (6). In classical linguistic theory and most past work in computational linguistics, discrete symbols and explicit rules are used to characterize language structure and relationships. In neural networks, these symbols are replaced by continuous, multivariate pattern vectors called distributed representations or embeddings, and the rules are replaced by continuous, multivalued arrays of connection weights that map patterns to other patterns.

Since its introduction (6), debate has raged about this approach to language processing (7). Protagonists argue it supports nuanced, context- and similarity-sensitive processing that is reflected in the quasiregular relationships between phrases and their sounds, spellings, and meanings (8, 9). These models also capture subtle aspects of human performance in language tasks (10). However, critics note that neural networks often fail to generalize beyond their training data, blaming these failures on the absence of explicit rules (11–13).

Neural Language Modeling

Initial Steps. Elman (14) introduced a simple recurrent neural network (RNN) (Fig. 2A) that captured key characteristics of language structure through learning, a feat once considered impossible (15). It was trained to predict the next word in a sequence $[w(t+1)]$ based on the current word $[w(t)]$ and its own hidden (that is, learned internal) distributed representation from the previous time step $[h(t-1)]$. Each of these inputs is multiplied by a matrix of connection weights (arrows labeled W_{hi} and W_{hh} in Fig. 2A), and the results are added to produce the input to the hidden units. The elements of this vector pass through a function limiting the range of their values, producing the hidden representation. This in turn is multiplied with weights to the output layer from the hidden layer (W_{oh}) to generate a vector used to predict the probability of each of the possible successor words. Learning is based on the discrepancy between the network’s output and the actual next word; the values of the connection weights are adjusted by a small amount to reduce the discrepancy. The network

is recurrent because the same connection weights (denoted by arrows in the figure) are used to process each successive word.

Elman showed two things. First, after training his network to predict the next word in sentences like “man eats bread,” “dog chases cat,” and “girl sleeps,” the network’s representations captured the syntactic distinction between nouns and verbs as well as interpretable subcategories among nouns and verbs (14), as shown for the nouns in Fig. 2B. This illustrates a key feature of learned representations: they capture specific as well as general or abstract information. By using a different learned representation for each word, its specific predictive consequences can be exploited. Because representations for words that make similar predictions are similar and because neural networks exploit similarity, the network can share knowledge about predictions among related words.

Second, Elman (16) used both simple sentences like “boy chases dogs” and more complex ones like “boy who sees girls chases dogs.” In the latter, the verb “chases” must agree with the first noun (“boy”), not the closest noun (“girls”), since the sentence contains a main clause (“boy chases dogs”) interrupted by a reduced relative clause (“[who] sees girls”). The model learned to predict the verb form correctly despite the intervening clause, showing that it acquired sensitivity to the syntactic structure of language, not just local co-occurrence statistics.

Scaling up to Natural Text. Elman’s task of predicting words based on context has been central to neural language modeling. However, Elman trained his networks with tiny, toy languages. For many years, it seemed they would not scale up, and language modeling was dominated by simple n -gram models and systems designed to assign explicit structural descriptions to sentences, aided by advances in probabilistic computations (17). Over the past 10 years, breakthroughs have allowed networks to predict and fill in words in huge natural language corpora.

One challenge is the large size of a natural language’s vocabulary. A key step was the introduction of methods for learning word representations (now called embeddings) from co-occurrence relationships in large text corpora (18, 19). These embeddings exploit both general and specific predictive relationships of all of the words in the corpus, improving generalization: task-focused neural models trained on small datasets better generalize to infrequent words (e.g., “settee”) based on frequent words (e.g., “couch”) with similar embeddings.

A second challenge is the indefinite length of the context that might be relevant for prediction. Consider this passage:

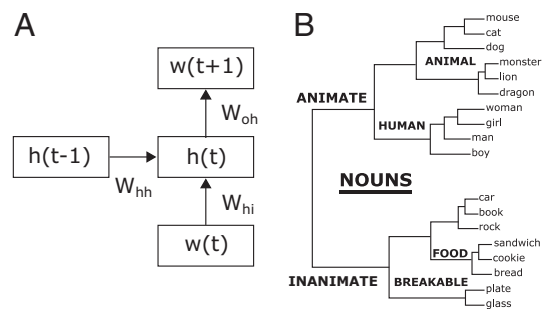


Fig. 2. (A) Elman’s (14) simple recurrent network and (B) his hierarchical clustering of the representations of the nouns in the corpus. Adapted with permission from figure 7 of ref. 14.

“John put some beer in a cooler and went out with his friends to play volleyball. Soon after he left, someone took the beer out of the cooler. John and his friends were thirsty after the game, and went back to his place for some beers. When John opened the cooler, he discovered that the beer was ____.”

Here, a reader expects the missing word to be “gone.” Yet, if we replace “took the beer” with “took the ice,” the expected word is “warm.” Any amount of additional text between “beer” and “gone” does not change the predictive relationship, challenging RNNs like Elman’s. An innovation called long short-term memory (LSTM) (20) partially addressed this problem by augmenting the recurrent network architecture with learned connection weights that gate information into and out of a network’s internal state. However, LSTMs did not fully alleviate the context bottleneck problem (21): a network’s internal state was still a fixed-length vector, limiting its ability to capture contextual information.

Query-Based Attention. Recent breakthroughs depend on an innovation we call query-based attention (QBA) (21). It was used in the Google Neural Machine Translation system (22), a system that attained a sudden leap in performance and attracted widespread public interest (23).

We illustrate QBA in Fig. 3 with the sentence “John hit the ball with the bat.” Context is required to determine whether “bat” refers to an animal or a baseball bat. QBA addresses this by issuing queries for relevant information. A query might ask, “is there an action and relation in the context that would indicate which kind of bat fits best?” The embeddings of words that match the query then receive high weightings in the weighted attention vector. In our example, the query matches the embedding of “hit” closely and of “with” to some extent; the returned attention vector captures the content needed to determine that a baseball bat fits in this context.

There are many variants of QBA; the figure is intended to capture the essence of the concept (21). In one important QBA model called BERT, which stands for bidirectional encoder representations from transformers (25), the network is trained to correct missing or replaced words in blocks of text, typically spanning two sentences. BERT relies on multiple attention heads, each employing QBA (26), to query all of the words in the text

block including itself, concatenating the returned weighted vectors to form a composite attention vector. The process is iterated across several stages, so that contextually constrained representations of each word computed at intermediate stages in turn constrain the representations of every word in later stages. In this way, BERT employs mutual constraint satisfaction, as Rumelhart (3) envisioned. The embeddings that result from QBA capture gradations within the set of meanings a language maps to a given word, aiding translation and other downstream tasks. For example, in English, we use “ball” for many types of balls, whereas in French, some are “balles” and others “ballons.” Subtly different embeddings for “ball” in different English sentences aid selecting the correct French translation.

In QBA architectures, the vectors all depend on learned connection weights, and analysis of BERT’s representations shows that they capture syntactic structure well (24). Different attention heads capture different linguistic relationships, and the similarity relationships among the contextually shaded word representations can be used to reconstruct a sentence’s syntactic structural description. These representations capture this structure without building it in, supporting the emergence principle. That said, careful analysis (27) indicates that the deep network’s sensitivity to grammatical structure is still imperfect and only partially understood.

Some attention-based models (28) proceed sequentially, predicting each word using QBA over prior context, while BERT operates in parallel, using mutual QBA simultaneously on all of the words in an input text block. Humans appear to exploit past context and a limited window of subsequent context (29), suggesting a hybrid strategy. Some machine models adopt this approach (30), and below we adopt a hybrid approach as well.

Attention-based models have produced remarkable improvements on a wide range of language tasks. The models can be pretrained on massive text corpora, providing useful representations for subsequent fine tuning to perform other tasks. A recent model called GPT-3, which stands for generative pretrained transformer, achieves impressive gains on several benchmarks without requiring fine tuning (28). However, this model still falls short of human performance on tasks that depend on what the authors call “common sense physics” and on carefully crafted tests of their ability to determine if a sentence follows from a preceding text (31). Further, the text corpora these models rely on are far

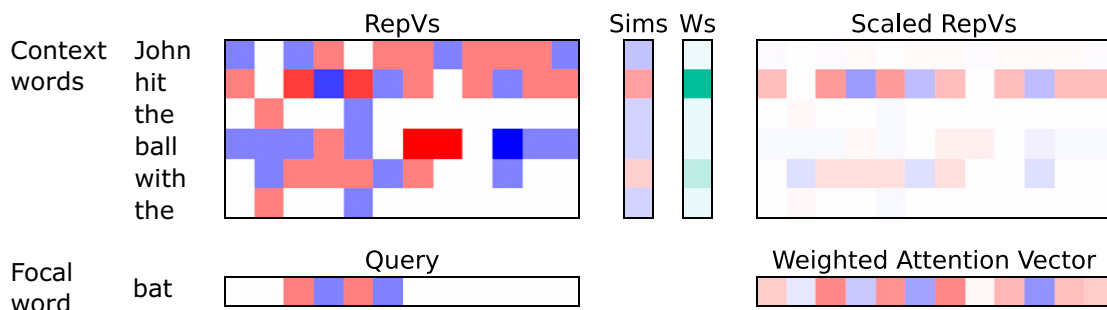


Fig. 3. Query-based attention (QBA). To constrain the interpretation of the word “bat” in the context “John hit the ball with the ____,” a query generated from “bat” can be used to construct a weighted attention vector, which shapes the word’s interpretation. The query is compared with each of the learned representation vectors (RepVs) of the context words; this creates a set of similarity scores (Sims), which in turn, produce a set of weightings (Ws; a set of positive numbers summing to one). The Ws are used to scale the RepVs of the context words, creating Scaled RepVs. The weighted attention vector is the element-wise sum of the Scaled RepVs. The Query, RepVs, Sims, Scaled RepVs, and weighted attention vector use red color intensity for positive magnitudes and blue for negative magnitudes. Ws are shown as green color intensity. White = 0 throughout. The Query and RepVs were made up for illustration, inspired by ref. 24. Mathematical details: for query q and representation vector v_j for context word j , the similarity score s_j is $\cos(q, v_j)$. The s_j are converted into weightings w_j by the softmax function, $w_j = e^{(gs_j)} / (\sum_j e^{(gs_j)})$, where the sum in the denominator runs over all words in the context span, and g is a scale factor.

larger than a human learner could process in a lifetime. Gains from further increases may be diminishing, and human learners appear to require far less training experience. The authors of GPT-3 note these limitations and express the view that further improvements may require more fundamental changes.

Language in an Integrated Understanding System

Where should we look for further progress addressing the limitations of current language models? In concert with others (32), we argue that part of the solution will come from treating language as part of a larger system for understanding and communicating.

Situations. We adopt the perspective that the targets of understanding are situations. Situations are collections of entities, their properties and relations, and patterns of change in them. A situation can be static (e.g., a cat on a mat). Situations include events (e.g., a boy hitting a ball). Situations can embed within each other; the cat may be on a mat inside a house on a particular street in a particular town, and the same applies to events like baseball games. A situation can be conceptual, social, or legal, such as one where a court invalidates a law. A situation may even be imaginary. The entities participating in a situation or event may be real or fictitious physical objects or locations; animals, persons, groups, or organizations; beliefs or other states of mind; sets of objects (e.g., all dogs); symbolic objects such as symbols, tokens, or words; or even contracts, laws, or theories. Situations can even involve changes in beliefs about relationships among classes of objects (e.g., biologists' beliefs about the genus a species of trees belongs in).

What it means to understand a situation is to construct a representation of it that captures aspects of the participating objects, their properties, relationships and interactions, and resulting outcomes. This understanding should be thought of as a construal or interpretation that may be incomplete or inaccurate; it will depend on the culture and context of the agent and the agent's purpose. When other agents are the source of the input, the target agent's construal of the knowledge and purpose of these other agents also play important roles. As such, the construal process must be considered to be completely open ended and to potentially involve interaction between the construer and the situation, including exploration of the world and discourse between the agent and participating interlocutors.

Within this construal of understanding, we emphasize that language should be seen as a component of an understanding system. This idea is not new, but historically, it was not universally accepted. Chomsky (33), Fodor and Pylyshyn (11), and Fodor (34) argued that grammatical knowledge sits in a distinct, encapsulated subsystem. Our proposal to focus on language as part of a system representing situations builds on a long tradition in linguistics (35), human cognitive psychology (36), psycholinguistics (37), philosophy (38), and artificial intelligence (39). The approach was adopted in an early neural network model (40) and aligns with other current perspectives in cognitive science (41), cognitive neuroscience (42), and artificial intelligence (32).

People Construct Situation Representations. A person processing language constructs a representation of the described situation in real time, using both the stream of words and other available information. Words and their sequencing serve as clues to meaning (43) that jointly constrain the understanding of the situation (40, 44). Consider this passage: "John spread jam on some bread. The knife had been dipped in poison." We make many inferences here: the jam was spread with the poisoned

knife, and poison has been transferred to the bread. If John eats it, he may die! Note the entities are objects, not words, and the situation could be conveyed by a silent movie.

Evidence that humans construct situation representations from language comes from classic work by Bransford and colleagues (36, 45). This work demonstrates that 1) we understand and remember texts better when we can relate the text to a familiar situation; 2) relevant information can come from a picture accompanying the text; 3) what we remember from a text depends on the framing context; 4) we represent implied objects in memory; and 5) after hearing a sentence describing spatial or conceptual relationships, we remember these relationships, not the language itself. Given "Three turtles rested beside a floating log and a fish swam under it," the situation changes if "it" is replaced by "them." After hearing the original sentence, people reject the variant with "it" in it as the sentence they heard before, but if the initial sentence said "Three turtles rested on a floating log," the situation is unchanged by replacing "it" with "them," and people accept this variant.

Evidence from eye movements shows that people use linguistic and nonlinguistic input jointly and immediately (46). Just after hearing "The man will drink," participants look at a full wine glass rather than an empty beer glass (47). After hearing "The man drank," they look at the empty beer glass. Understanding thus involves constructing, in real time, a representation conveyed jointly by vision and language.

The Compositionality of Situations. An important debate centers on compositionality. Fodor and Pylyshyn (11) argued that our cognitive systems must be compositional by design to allow language to express arbitrary relationships and noted that early neural network models failed tests of compositionality. Such failures are still reported (48), leading some to propose building compositionality in (13); yet, as we have seen, the most successful language models avoid doing so. We suggest that a focus on situations may enhance compositionality because situations are themselves compositional. Suppose a person picks an apple and gives it to someone. A small number of objects and persons are focally involved, and a sentence like "John picked an apple and gave it to Mary" could describe this situation, capturing the most relevant participants and their relationships. We emphasize that compositionality is not universal or absolute (another apple could fall, and Mary's boyfriend might be jealous), so it is best to allow for matters of degree. Letting situation representations emerge through experience will help our models to achieve greater systematicity while leaving them the flexibility to capture nuance that has led to their successes to date.

Language Informs Us about Situations. Situations ground the representations we construct from language; equally importantly, language informs us about situations. Language tells us about situations we have not witnessed and describes aspects that we cannot observe. Language also communicates folk or scientific construals that shape listener's construals, such as the idea that an all-knowing being took 6 days to create the world or the idea that natural processes gave rise to our world and, ultimately, ourselves over billions of years. Language can be used to communicate information about properties that only arise in a social setting, such as ownership, or that have been identified by a culture as important, such as exact number. Language thus enriches and extends the information we have about situations and provides the primary medium conveying properties of many kinds of objects and many kinds of relationships.

Toward a Brain- and Artificial Intelligence-Inspired Model of Understanding

Capturing the full range of situations is clearly a long-term challenge. We return to this later, focusing first on concrete situations involving animate beings and physical objects. We seek to integrate insights from cognitive neuroscience and artificial intelligence (AI) toward the goal of building an integrated understanding model. We start with our construal of the understanding system in the human brain and then sketch aspects of what an artificial implementation might look like.

The Understanding System in the Brain. Our construal of the human integrated understanding system builds on the principles of mutual constraint satisfaction and emergence and the idea that understanding centers on the construction of situation representations. It is consistent with a wide range of evidence, some of which we review, and is broadly consistent with recent characterizations in cognitive neuroscience (42, 49). However, researchers hold diverse views about the details of these systems and how they work together.

We focus first on the part of the system located primarily in the neocortex of the brain, as schematized in the large blue box of Fig. 4. Together with input and output systems, this system allows a person to combine linguistic and visual input to understand what we will call a microsituation such as the one involving a batter hitting a ball with a bat as in Fig. 4A. It is important to note that the neocortex is very richly structured, with on the order of 100 well-defined anatomical subdivisions. However, it is common and useful to group these divisions into subsystems. The ones we focus on here are each indicated by blue ovals in the figure. One subsystem subserves the formation of a visual representation of the given situation, and another subserves the formation of an auditory representation capturing the spatiotemporal structure of the co-occurring spoken language. The three ovals above these provide representations of more integrative/abstract types of information (see below).

Within each subsystem, and between each connected pair of subsystems, the neurons are reciprocally interconnected via learning-dependent pathways, allowing mutual constraint satisfaction among all of the elements of each of the representation types, as indicated by the looping blue arrows from each oval to itself and by the bidirectional blue arrows between these ovals. Brain regions for representing visual and auditory inputs are well established, and the evidence for their involvement in a mutual constraint satisfaction process with more integrative brain areas has been reviewed elsewhere (50, 51). Here, we consider the three more integrative subsystems.

Object representations. A brain area near the front of the temporal lobe houses neurons whose activity provides an embedding capturing the properties of objects (52). Damage to this area impairs the ability to name objects, to grasp them correctly for their intended use, to match them with their names or the sounds they make, and to pair objects that go together, either from their names or from pictures. This brain area is itself an intermodal area, receiving visual, language, and other information about objects such as the sounds they make and how they feel to the touch. Models capturing these findings (53) treat this area as the hidden layer of an interactive, recurrent network with bidirectional connections to other layers representing different types of object properties, including the object's name. In these models, an input to any of these other layers activates the corresponding pattern in the hidden layer, which in turn activates the corresponding

patterns in the other layers. This supports, for example, the ability to produce the name of an object from visual input. Damage (simulated by removing neurons in the hidden layer) degrades the model's representations, capturing the patterns of errors made by patients with the corresponding brain damage.

Representation of context. There is a network of areas in the brain that capture the spatiotemporal context, here called the scene, within which microsituations are experienced. These context representations arise in a set of interconnected areas primarily within the parietal lobes (42, 49). In recent work, brain imaging data are used to analyze the patterns of neural activity in this network while humans process a temporally extended narrative. The brain activity patterns that represent scenes extending over tens of seconds (e.g., a detective searching a suspect's apartment for evidence) are largely the same, whether the information comes from watching a movie, hearing or reading a narrative description, or recalling the movie after having seen it (54, 55).

Activity in different brain areas tracks information on different timescales. Activity in modality-specific areas associated with speech and visual processing follows the moment-by-moment time course of spoken and/or visual information, while activity in the network associated with context representations fluctuates on a much longer timescale. During processing of narrative information, activations in these regions tend to be relatively stable within a




Fig. 4. Sketch of the brain's understanding system. Ovals in the blue box stand for neocortical brain areas representing different kinds of information. Arrows in the neocortex stand for connections, allowing representations to constrain each other. The MTL (medial temporal lobe, red box) stores an integrated representation of the neocortical system state arising from a situation. The red arrow represents fast-learning connections that store this pattern for later reactivation and use. Green arrows stand for gradually learned connections supporting bidirectional influence between the MTL and neocortex. A and B are two example inputs discussed in the text. Hitter and numbat images credit: Mary Reusch (artist).

scene, punctuated with larger changes at boundaries between these scenes, and these patterns lose their coherence when the narrative structure is scrambled (42, 55). Larger-scale spatial transitions (e.g., transitions between rooms) also create large changes in neural activity (49) in these context representations.

The role of language. Where in the brain should we look for representations of the relations among the objects participating in a microsituation? The effects of brain damage suggest that representations of relations may be integrated, at least in part, with the representation of language itself. Injuries affecting the lateral surface of the frontal and parietal lobes produce profound deficits in the production of fluent language but can leave the ability to read and understand concrete nouns largely intact. Such lesions produce intermediate degrees of impairment to abstract nouns, verbs, and modifiers and profound impairment to words like “if” and “by” that capture grammatical and spatial relations (56). This striking pattern is consistent with the view that language itself is intimately tied to the representation of relations and changes in relations (information conveyed by verbs, prepositions, and grammatical markers). Indeed, the frontal and parietal lobes are associated with representation of space and action (which causes change in relations), and patients with lesions to the frontal and parietal language-related areas have profound deficits in relational reasoning tasks (57). We therefore tentatively suggest that understanding of microsituations depends jointly on the object and language systems and that language is intimately linked to representation of spatial relationships and actions.

Complementary learning systems. The brain systems described above support understanding of situations that draw on general knowledge as well as oft-repeated specific information, but they do not support the rapid formation of memories for arbitrary new information encountered only one or a few times and that can be accessed and used at arbitrary future times. This ability depends on structures that include the hippocampus in the medial temporal lobes (MTLs) (Fig. 4, red box). While these areas are critical for new learning, damage to them does not affect knowledge of general and frequently encountered specific information, acquired skills, or the ability to process language and other forms of input to understand a situation, except when this depends on remote information experienced briefly outside the immediate current context (58). These findings are captured in the neural-network based complementary learning systems (CLSs) theory (59–61), which holds that connections within the neocortex gradually acquire the knowledge that allows a human to understand objects and their properties, to link words and objects, and to understand and communicate about generic and familiar situations as these are conveyed through language and other forms of experience. The MTL provides a complementary fast-learning system supporting the formation of new arbitrary associations, linking the elements of an experience together, including the objects and language encountered in a situation and the co-occurring spatiotemporal context, as might arise in the situation depicted in Fig. 4B, where a person encounters a novel animal called a numbat for the first time, from both visual and language input.

It is generally accepted that knowledge that depends initially on the MTL can be integrated into the neocortex through a consolidation process (58). In CLS (60), the neocortex learns arbitrary new information gradually through interleaved presentations of new and familiar items, weaving it into the fabric of knowledge and avoiding interference with existing knowledge (62). The details are subjects of current debate and ongoing investigation (63).

As in our example of the beer John left in the cooler, understanding often depends on remote information. People exploit such information during language processing (64), and patients with MTL damage have difficulty understanding or producing extended narratives (65); they are also profoundly impaired in learning new words for later use (66). Neural language models, including those using QBA, face similar challenges. In BERT and related models, bidirectional attention operates within a span of a couple of sentences at a time. Other models (28) employ QBA over longer spans of prior context, but there is still a finite limit. These models learn gradually like the human neocortical system, allowing them to acquire knowledge that guides their processing of current inputs. For example, GPT-3 (28) is impressive in its ability to use a word encountered for the first time within its contextual span appropriately in a subsequent sentence, but this information is lost forever when the context is reinitialized, as it would be in humans without their MTLs. Including an MTL-like system in future AI understanding models will address this limitation.

The brain’s CLSs may provide a means to address the challenge of learning to use a word encountered in a single context appropriately across a wide range of contexts. Deep neural networks that learn gradually through many repetitions of an item that occurs in a single context, interleaved with presentations of other items occurring in a diversity of contexts, do not show this ability (48). We attribute this failure to the fact that the distribution of training experiences they receive conveys the information that the target item is in fact restricted to its single context. Further research should explore whether augmenting a model like GPT-3 with an MTL-like memory would enable more human-like extension of a novel word encountered just once to a wider range of contexts.

Next Steps toward a Brain- and Artificial Intelligence-Inspired Model.

Given the construal we have described of the human understanding system, we now ask: what might an implementation of a model consistent with it be like? This is a long-term research question—addressing it will benefit from a greater convergence of cognitive neuroscience and AI. Toward this goal, we sketch a proposal for a brain- and AI-inspired model. We rely on the principles of mutual constraint satisfaction and emergence, the QBA architecture from AI and deep learning, and the components and their interconnections in the understanding system in the brain, as illustrated in Fig. 4.

For simplicity, we treat the system as one that receives sequences of microsituations each consisting of a picture–description (PD) pair, grouped into scenes that in turn form story-length narratives, with the language conveyed by text rather than speech. Our sketch leaves important issues open and will require substantial development. Steps toward addressing some of these issues are already being taken: mutual QBA is already being used (e.g., in ref. 67, to exploit audio and visual information from movies).

In our proposed model, each PD pair is processed by interacting object and language subsystems, receiving visual and text input at the same time. Each subsystem must learn to restore missing or distorted elements (words in the language subsystem, objects in the object subsystem) by using mutual QBA as in BERT, allowing each element in each subsystem to be constrained by all of the elements in both subsystems. Additionally, these systems will query the context and memory subsystems, as described below.

The context subsystem encodes a sequence of compressed representations of the previous PD pairs within the current scene.

Processing in this subsystem would be sequential over pairs, allowing the network constructing the current compressed representation to query the representations of past pairs within the scene. Within the processing of a pair, the context system would engage in mutual QBA with the object and language subsystems, allowing the language and object subsystems to indirectly exploit information from past pairs within the scene.

Our system also includes an MTL-like memory to allow it to use remote information beyond the current scene. A neural network with learned connection weights constructs an invertible reduced description of the states of the object, language, and context subsystems along with the states of the vision and text subsystems after processing each PD pair. This compressed vector is then stored in an MTL-like memory. While a superpositional, distributed memory (68) will ultimately be the best model of this system, an initial implementation could store each such vector in a separate slot (69, 70). Use of these vectors later would rely on the flexible querying scheme of ref. 70, such that any subset of the object, language, or context representations of an input currently being processed could contribute to the query. The resulting weighted memory vector would then be decompressed and the relevant portions made available to the cortical subsystems. Thus, after the initial experience with the numbat shown in Fig. 4B, seeing a second numbat at an arbitrary later time would query the representation formed during the earlier experience, allowing access to the animal's name and the fact that it eats termites to inform processing of the second numbat. Developing efficient implementations of this system will be important for cognitive science and AI.

Our model will benefit from an additional subsystem that guides processing in all of the subsystems we have described. The brain has such a system in its frontal lobes; frontal damage leads to impairments in guiding behavior and cognition according to current task demands, an ability current AI systems lack (71). Incorporating such a subsystem into understanding system models is therefore an important future step.

Enhancing Understanding by Incorporating Interaction with the Physical and Social World. A complete model of the human understanding system will require integration of many additional information sources. These include sounds, touch and force sensing, and information about one's own actions. Every source provides opportunities to predict information of each type, relying on every other type. Information salient in one source can bootstrap learning and inference in the other, and all are likely to contribute to enhancing compositionality and addressing the data inefficiency of learning from language alone. This affords the human learner an important opportunity to experience the compositional structure of the environment through its own actions. Ultimately, an ability to link one's actions to their consequences as one behaves in the world should contribute to the emergence of, and appreciation for, the compositional structure of events and provide a basis for acquiring notions of cause and effect, of agency, and of object permanence (72).

These considerations motivate recent work on agent-based language learning in simulated interactive three-dimensional environments (73–76). In ref. 77, an agent was trained to identify, lift, carry, and place objects relative to other objects in a virtual room, as specified by simplified language instructions. At each time step, the agent received a first-person visual observation and processed the pixels to obtain a representation of the scene. This was concatenated to the final state of an LSTM that processed the instruction and then passed to an integrative LSTM whose output was used to select a motor action. The agent gradually learned to follow

instructions of the form “find a pencil,” “lift up a basketball,” and “put the teddy bear on the bed,” encompassing 50 objects and requiring up to 70 action steps. Such instructions require constructing representations based on language stimuli that support identification of objects and relations across space and time and the integration of this information to inform motor behaviors.

Importantly, without building in explicit object representations, the learned system was able to interpret novel instructions. For instance, an agent trained to lift each of 20 objects, but only trained to put 10 of those in a specific location, could place the remaining objects in the same location on command with over 90% accuracy, demonstrating a degree of compositionality in its behavior. Notably, the agent's egocentric, multimodal, and temporally extended experience contributed to this outcome; both an alternative agent with a fixed perspective on a two-dimensional grid world and a static neural network classifier that received only individual still images exhibited significantly worse generalization. This underscores how affording neural networks access to rich, multimodal interactive environments can stimulate the development of capacities that are essential for language learning and contribute toward emergent compositionality.

Beyond Concrete Situations. How might our approach be extended beyond concrete situations to those involving relationships among objects like laws, belief systems, and scientific theories? Basic word embeddings themselves capture some abstract relations via vector similarity (e.g., encoding that “justice” is closer to “law” than “peanut”). Words are uttered in real-world contexts, and there is a continuum between grounding and language-based linking for different words and different uses of words. For example, “career” is not only linked to other abstract words like “work and specialization” but also, to more concrete ones like “path” and its extended metaphorical use as the means to achieve goals (78). Embodied, simulation-based approaches to meaning (79, 80) build on this observation to bridge from concrete to abstract situations via metaphor. They posit that understanding the word “grasp” is linked to neural representations of the action of grabbing and that this circuitry is recruited for understanding contexts such as grasping an idea. We consider situated agents as a critical catalyst for learning about how to represent and compose concepts pertaining to spatial, physical, and other perceptually immediate phenomena—thereby providing a grounded edifice that can connect both to brain circuitry for motor action and to representations derived primarily from language.

Conclusion

Language does not stand alone. The understanding system in the brain connects language to representations of objects and situations and enhances language understanding by exploiting the full range of our multisensory experience of the world, our representations of our motor actions, and our memory of previous situations. We believe next generation language understanding systems should emulate this system, and we have sketched an approach that incorporates recent machine learning breakthroughs to build a jointly brain- and AI-inspired understanding system. We emphasize understanding of concrete situations and argue that understanding abstract language should build upon this foundation, pointing toward the possibility of someday building artificial systems that understand abstract situations far beyond concrete, here-and-now situations. In sum, combining insights from neuroscience and AI will take us closer to human-level language understanding.

Acknowledgments

This article grew out of a workshop organized by H.S. at *Meaning in Context* 3, Stanford University, September 2017. We thank Janice Chen, Maryellen

MacDonald, Karalyn Patterson, Chris Potts, and especially Mark Seidenberg for discussion. H.S. was supported by European Research Council Advanced Grant 740516.

- 1 F. Rosenblatt, *Principles of Neurodynamics. Perceptions and the Theory of Brain Mechanisms* (Spartan Books, Cornell University, Ithaca, NY, 1961).
- 2 D. E. Rumelhart, J. L. McClelland, The PDP research group, *Parallel Distributed Processing Explorations in the Microstructure of Cognition. Volume 1: Foundations* (MIT Press, Cambridge, MA, 1986).
- 3 D. E. Rumelhart, "Toward an interactive model of reading" in *Attention & Performance VI*, S. Dornic, Ed. (LEA, Hillsdale, NJ, 1977), pp. 573–603.
- 4 R. Taraban, J. L. McClelland, Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *J. Mem. Lang.* **27**, 597–632 (1988).
- 5 J. L. McClelland, D. E. Rumelhart, An interactive activation model of context effects in letter perception. I. An account of basic findings. *Psychol. Rev.* **88**, 375–407 (1981).
- 6 D. E. Rumelhart, J. L. McClelland, "On learning the past tenses of English verbs" in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 2: Psychological and Biological Models*, J. L. McClelland, D. E. Rumelhart, Eds. (MIT Press, Cambridge, MA, 1986), vol. 2, pp. 216–271.
- 7 S. Pinker, J. Mehler, *Connections and Symbols* (MIT Press, Cambridge, MA, 1988).
- 8 B. MacWhinney, J. Leinbach, Implementations are not conceptualizations: Revising the verb learning model. *Cognition* **40**, 121–157 (1991).
- 9 J. Bybee, J. L. McClelland, Alternatives to the combinatorial paradigm of linguistic theory based on domain general principles of human cognition. *Linguist. Rev.* **22**, 381–410 (2005).
- 10 M. S. Seidenberg, D. C. Plaut, Quasiregularity and its discontents: The legacy of the past tense debate. *Cogn. Sci.* **38**, 1190–1228 (2014).
- 11 J. A. Fodor, Z. W. Pylyshyn, Connectionism and cognitive architecture: A critical analysis. *Cognition* **28**, 3–71 (1988).
- 12 G. Marcus, *The Algebraic Mind* (MIT Press, Cambridge, MA, 2001).
- 13 B. M. Lake, T. D. Ullman, J. B. Tenenbaum, S. J. Gershman, Building machines that learn and think like people. *Behav. Brain Sci.* **40**, e253 (2017).
- 14 J. L. Elman, Finding structure in time. *Cogn. Sci.* **14**, 179–211 (1990).
- 15 E. M. Gold, Language identification in the limit. *Inf. Contr.* **10**, 447–474 (1967).
- 16 J. L. Elman, Distributed representations, simple recurrent networks, and grammatical structure. *Mach. Learn.* **7**, 195–225 (1991).
- 17 C. D. Manning, H. Schütze, *Foundations of Statistical Language Processing* (MIT Press, Cambridge, MA, 1999).
- 18 R. Collobert et al., Natural language processing (almost) from scratch. *J. Mach. Learn. Res.* **12**, 2493–2537 (2011).
- 19 T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, "Distributed representations of words and phrases and their compositionality" in *Proceedings of Neural Information Processing Systems*, C. J. C. Burges, L. Bottou, Z. Ghahramani, K. Q. Weinberger, Eds. (Curran Associates, Red Hook, NY, 2013), pp. 3111–3119.
- 20 S. Hochreiter, J. Schmidhuber, Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
- 21 D. Bahdanau, K. Cho, Y. Bengio, "Neural machine translation by jointly learning to align and translate" in *3rd International Conference on Learning Representations*. arXiv:1409.0473v7 (19 May 2016).
- 22 Y. Wu et al., Google's neural machine translation system: Bridging the gap between human and machine translation. arXiv:1609.08144 (26 September 2016).
- 23 G. Lewis-Kraus, The great AI awakening. *NY Times Magazine*, 14 December 2016. <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>. Accessed 23 May 2020.
- 24 C. D. Manning, K. Clark, J. Hewitt, U. Khandelwal, O. Levy, Emergent linguistic structure in artificial neural networks trained by self-supervision. *Proc. Natl. Acad. Sci. U.S.A.*, 10.1073/pnas.1907367117 (2020).
- 25 J. Devlin, M. W. Chang, K. Lee, K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding" in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*, J. Burstein, C. Doran, T. Solorio, Eds. (Association for Computational Linguistics, Stroudsburg, PA, 2019), pp. 4171–4186.
- 26 A. Vaswani et al., "Attention is all you need" in *Advances in Neural Information Processing Systems*, I. Guyon, et al., Eds. (Curran Associates, Inc., Red Hook, NY, 2017), vol. 30, pp. 5998–6008.
- 27 T. Linzen, M. Baroni, Syntactic structure from deep learning. *Annu. Rev. Linguist.*, 10.1146/annurev-linguistics-032020-051035 (2020).
- 28 T. B. Brown et al., Language models are few-shot learners. arXiv:2005.14165 (28 May 2020).
- 29 R. M. Warren, Perceptual restoration of missing speech sounds. *Science* **167**, 392–393 (1970).
- 30 Z. Yang et al., "Xlnet: Generalized autoregressive pretraining for language understanding" in *Advances in Neural Information Processing Systems*, H. Wallach et al., Eds. (Curran Associates, Inc., Red Hook, NY, 2019), vol. 32, pp. 5753–5763.
- 31 Y. Nie et al., Adversarial NLI: A new benchmark for natural language understanding. arXiv:1910.14599 (31 October 2019).
- 32 Y. Bisk et al., Experience grounds language. arXiv:2004.10151 (21 April 2020).
- 33 N. Chomsky, "Deep structure, surface structure, and semantic interpretation" in *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics, and Psychology*, D. Steinberg, L. A. Jakobovits, Eds. (Cambridge University Press, Cambridge, UK, 1971), pp. 183–216.
- 34 J. A. Fodor, *The Modularity of Mind* (MIT Press, 1983).
- 35 G. Lakoff, *Women, Fire, and Dangerous Things* (The University of Chicago Press, 1987).
- 36 J. D. Bransford, M. K. Johnson, Contextual prerequisites for understanding: Some investigations of comprehension and recall. *J. Verb. Learn. Verb. Behav.* **11**, 717–726 (1972).
- 37 S. Crain, M. Steedman, "Context and the psychological syntax processor" in *Natural Language Parsing*, D. R. Dowty, L. Karttunen, A. M. Zwicky, Eds. (Cambridge University Press, Cambridge, UK, 1985), pp. 320–348.
- 38 R. Montague, "The proper treatment of quantification in ordinary English" in *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, J. Hintikka, J. Moravcsik, P. Suppes, Eds. (Riedel, Dordrecht, Netherlands, 1973), pp. 221–242.
- 39 R. C. Schank, *Dynamic Memory: A Theory of Reminding and Learning in Computers and People* (Cambridge University Press, 1983).
- 40 M. F. St. John, J. L. McClelland, Learning and applying contextual constraints in sentence comprehension. *Artif. Intell.* **46**, 217–257 (1990).
- 41 S. Gershman, J. B. Tenenbaum, "Phrase similarity in humans and machines" in *Proceedings of the 37th Annual Meeting of the Cognitive Science Society* (Cognitive Sciences [COGSCI], 2015).
- 42 U. Hasson, G. Egidio, M. Marelli, R. M. Willems, Grounding the neurobiology of language in first principles: The necessity of non-language-centric explanations for language comprehension. *Cognition* **180**, 135–157 (2018).
- 43 D. E. Rumelhart, "Some problems with the notion that words have literal meanings" in *Metaphor and Thought*, A. Ortony, Ed. (Cambridge University Press, Cambridge, UK, 1979), pp. 71–82.
- 44 M. Rabovsky, S. S. Hansen, J. L. McClelland, Modeling the n400 brain potential as change in a probabilistic representation of meaning. *Nat. Hum. Behav.* **2**, 693–705 (2018).
- 45 J. D. Bransford, J. R. Barclay, J. J. Franks, Sentence memory: A constructive versus interpretive approach. *Cogn. Psychol.* **3**, 193–209 (1972).
- 46 M. Tanenhaus, M. Spivey-Knowlton, K. Eberhard, J. Sedivy, Integration of visual and linguistic information in spoken language comprehension. *Science* **268**, 1632–1634 (1995).

- 47 G. T. Altmann, Y. Kamide, The real-time mediation of visual attention by language and world knowledge: Linking anticipatory (and other) eye movements to linguistic processing. *J. Mem. Lang.* **57**, 502–518 (2007).
- 48 B. Lake, M. Baroni, “Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks” in *International Conference on Machine Learning*, J. G. Dy, A. Krause, Eds. (Proceedings of Machine Learning Research, 2018), pp. 2873–2882.
- 49 C. Ranganath, M. Ritchey, Two cortical systems for memory-guided behavior. *Nat. Rev. Neurosci.* **13**, 713–726 (2012).
- 50 J. L. McClelland, D. Mirman, D. J. Bolger, P. Khaitan, Interactive activation and mutual constraint satisfaction in perception and cognition. *Cogn. Sci.* **38**, 1139–1189 (2014).
- 51 M. Heilbron, D. Richter, M. Ekman, P. Hagoort, F. P. de Lange, Word contexts enhance the neural representation of individual letters in early visual cortex. *Nat. Commun.* **11**, 321 (2020).
- 52 K. Patterson, P. J. Nestor, T. T. Rogers, Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat. Rev. Neurosci.* **8**, 976–987 (2007).
- 53 T. T. Rogers et al., Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychol. Rev.* **111**, 205–235 (2004).
- 54 A. Zadbood, J. Chen, Y. Leong, K. Norman, U. Hasson, How we transmit memories to other brains: Constructing shared neural representations via communication. *Cereb. Cortex* **27**, 4988–5000 (2017).
- 55 C. Baldassano et al., Discovering event structure in continuous narrative perception and memory. *Neuron* **95**, 709–721 (2017).
- 56 J. Morton, K. Patterson, “Little words–no!” in *Deep Dyslexia*, M. Coltheart, K. Patterson, J. C. Marshall, Eds. (Routledge and Kegan Paul, London, UK, 1980), pp. 270–285.
- 57 J. V. Baldo, S. A. Bunge, S. M. Wilson, N. F. Dronkers, Is relational reasoning dependent on language? A voxel-based lesion symptom mapping study. *Brain Lang.* **113**, 59–64 (2010).
- 58 B. Milner, S. Corkin, H. L. Teuber, Further analysis of the hippocampal amnesic syndrome: 14-Year follow-up study of H.M. *Neuropsychologia* **6**, 215–234 (1968).
- 59 D. Marr, Simple memory: A theory for archicortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **262**, 23–81 (1971).
- 60 J. L. McClelland, B. L. McNaughton, R. C. O’Reilly, Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**, 419–457 (1995).
- 61 D. Kumaran, D. Hassabis, J. L. McClelland, What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends Cogn. Sci.* **20**, 512–534 (2016).
- 62 J. L. McClelland, B. L. McNaughton, A. K. Lampinen, Integration of new information in memory: New insights from a complementary learning systems perspective. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **375**, 20190637 (2020).
- 63 A. Yonelinas, C. Ranganath, A. Ekstrom, B. Wiltgen, A contextual binding theory of episodic memory: Systems consolidation reconsidered. *Nat. Rev. Neurosci.* **20**, 364–375 (2019).
- 64 L. Menenti, K. M. Petersson, R. Scheeringa, P. Hagoort, When elephants fly: Differential sensitivity of right and left inferior frontal gyri to discourse and world knowledge. *J. Cogn. Neurosci.* **21**, 2358–2368 (2009).
- 65 X. Zuo et al., Temporal integration of narrative information in a hippocampal amnesic patient. *NeuroImage* **213**, 116658 (2020).
- 66 J. D. Gabrieli, N. J. Cohen, S. Corkin, The impaired learning of semantic knowledge following bilateral medial temporal-lobe resection. *Brain Cogn.* **7**, 157–177 (1988).
- 67 C. Sun, F. Baradel, K. Murphy, C. Schmid, Learning video representations using contrastive bidirectional transformer. arXiv:1906.05743 (13 June 2019).
- 68 D. J. Willshaw, O. P. Buneman, H. C. Longuet-Higgins, Non-holographic associative memory. *Nature* **222**, 960–962 (1969).
- 69 J. Weston, S. Chopra, A. Bordes, “Memory networks” in *International Conference on Learning Representations*, Y. Bengio, Y. LeCun, Eds. arXiv:1410.3916v11 (29 November 2015).
- 70 A. Graves et al., Hybrid computing using a neural network with dynamic external memory. *Nature* **538**, 471–476 (2016).
- 71 J. Russin, R. C. O’Reilly, Y. Bengio, “Deep learning needs a prefrontal cortex” in *ICLR Workshop on Bridging AI and Cognitive Science (ICLR, 2020)*.
- 72 J. Piaget, *The Origins of Intelligence in Children* (International Universities Press Inc., New York, NY, 1952).
- 73 K. M. Hermann et al., Grounded language learning in a simulated 3D world. arXiv:1706.06551 (20 June 2017).
- 74 R. Das, M. Zaheer, S. Reddy, A. McCallum, “Question answering on knowledge bases and text using universal schema and memory networks” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, R. Barzilay, M.-Y. Kan, Eds. (Association for Computational Linguistics, Stroudsburg, PA, 2017), pp. 358–365.
- 75 D. S. Chaplot, K. M. Sathyendra, R. K. Pasumarthi, D. Rajagopal, R. Salakhutdinov, “Gated-attention architectures for task-oriented language grounding” in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*, S. A. McIlraith, K. Q. Weinberger, Eds. (AAAI Press, Cambridge, MA, 2018), pp. 2819–2826.
- 76 J. Oh, S. Singh, H. Lee, P. Kohli, “Zero-shot task generalization with multi-task deep reinforcement learning” in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, D. Precup, Y. W. Teh, Eds. (JMLR.org, 2017), pp. 2661–2670.
- 77 F. Hill et al., “Environmental drivers of systematicity and generalization in a situated agent” in *International Conference on Learning Representations (ICLR, 2020)*.
- 78 J. Bryson, Embodiment vs. memetics. *Mind Soc.* **7**, 77–94 (2008).
- 79 G. Lakoff, M. Johnson, *Metaphors We Live By* (University of Chicago, Chicago, IL, 1980).
- 80 J. Feldman, S. Narayanan, Embodied meaning in a neural theory of language. *Brain Lang.* **89**, 385–392 (2004).