A mathematical theory of semantic development in deep neural networks

Andrew M. Saxe^{a,1}, James L. McClelland^b, and Surya Ganguli^{c,d}

^aDepartment of Experimental Psychology, University of Oxford, Oxford OX2 6GG, United Kingdom; ^bDepartment of Psychology, Stanford University, Stanford, CA 94305; ^cDepartment of Applied Physics, Stanford University, Stanford, CA 94305; and ^dGoogle Brain, Google, Mountain View, CA 94043

Edited by Terrence J. Sejnowski, Salk Institute for Biological Studies, La Jolla, CA, and approved April 9, 2019 (received for review December 6, 2018)

An extensive body of empirical research has revealed remarkable regularities in the acquisition, organization, deployment, and neural representation of human semantic knowledge, thereby raising a fundamental conceptual question: What are the theoretical principles governing the ability of neural networks to acquire, organize, and deploy abstract knowledge by integrating across many individual experiences? We address this question by mathematically analyzing the nonlinear dynamics of learning in deep linear networks. We find exact solutions to this learning dynamics that yield a conceptual explanation for the prevalence of many disparate phenomena in semantic cognition, including the hierarchical differentiation of concepts through rapid developmental transitions, the ubiquity of semantic illusions between such transitions, the emergence of item typicality and category coherence as factors controlling the speed of semantic processing, changing patterns of inductive projection over development, and the conservation of semantic similarity in neural representations across species. Thus, surprisingly, our simple neural model qualitatively recapitulates many diverse regularities underlying semantic development, while providing analytic insight into how the statistical structure of an environment can interact with nonlinear deep-learning dynamics to give rise to these regularities.

semantic cognition | deep learning | neural networks | generative models

Human cognition relies on a rich reservoir of semantic knowledge, enabling us to organize and reason about our complex sensory world (1–4). This semantic knowledge allows us to answer basic questions from memory (e.g., “Do birds have feathers?”) and relies fundamentally on neural mechanisms that can organize individual items, or entities (e.g., *canary* or *robin*), into higher-order conceptual categories (e.g., *birds*) that include items with similar features, or properties. This knowledge of individual entities and their conceptual groupings into categories or other ontologies is not present in infancy, but develops during childhood (1, 5), and in adults, it powerfully guides inductive generalizations.

The acquisition, organization, deployment, and neural representation of semantic knowledge has been intensively studied, yielding many well-documented empirical phenomena. For example, during acquisition, broader categorical distinctions are generally learned before finer-grained distinctions (1, 5), and periods of relative stasis can be followed by abrupt conceptual reorganization (6, 7). Intriguingly, during these periods of developmental stasis, children can believe illusory, incorrect facts about the world (2).

Also, many psychophysical studies of performance in semantic tasks have revealed empirical regularities governing the organization of semantic knowledge. In particular, category membership is a graded quantity, with some items being more or less typical members of a category (e.g., a sparrow is a more typical bird than a penguin). Item typicality is reproducible across individuals (8, 9) and correlates with performance on a diversity of semantic tasks (10–14). Moreover, certain categories themselves are thought to be highly coherent (e.g., the set of things that are *dogs*), in contrast to less coherent categories (e.g., the set of things that are *blue*). More coherent categories play a privileged role in the organization of our semantic knowledge; coherent categories are the ones that are most easily learned and represented (8, 15, 16). Also, the organization of semantic knowledge power-

fully guides its deployment in novel situations, where one must make inductive generalizations about novel items and properties (2, 3). Indeed, studies of children reveal that their inductive generalizations systematically change over development, often becoming more specific with age (2, 3, 17–19).

Finally, recent neuroscientific studies have probed the organization of semantic knowledge in the brain. The method of representational similarity analysis (20, 21) revealed that the similarity structure of cortical activity patterns often reflects the semantic similarity structure of stimuli (22, 23). And, strikingly, such neural similarity structure is preserved across humans and monkeys (24, 25).

This wealth of empirical phenomena raises a fundamental conceptual question about how neural circuits, upon experiencing many individual encounters with specific items, can, over developmental time scales, extract abstract semantic knowledge consisting of useful categories that can then guide our ability to reason about the world and inductively generalize. While several theories have been advanced to explain semantic development, there is currently no analytic, mathematical theory of neural circuits that can account for the diverse phenomena described above. Interesting nonneural accounts for the discovery of abstract semantic structure include the conceptual “theory-theory” (2, 17, 18) and computational Bayesian (26) approaches. However, neither currently proposes a neural implementation that can infer abstract concepts from a stream of examples. Also, they hold that specific domain theories or a set of candidate structural forms must be available a priori for learning to occur. In contrast, much prior work has shown, through simulations, that neural networks can gradually extract semantic

Significance

Over the course of development, humans learn myriad facts about items in the world, and naturally group these items into useful categories and structures. This semantic knowledge is essential for diverse behaviors and inferences in adulthood. How is this richly structured semantic knowledge acquired, organized, deployed, and represented by neuronal networks in the brain? We address this question by studying how the nonlinear learning dynamics of deep linear networks acquires information about complex environmental structures. Our results show that this deep learning dynamics can self-organize emergent hidden representations in a manner that recapitulates many empirical phenomena in human semantic development. Such deep networks thus provide a mathematically tractable window into the development of internal neural representations through experience.

Author contributions: A.M.S., J.L.M., and S.G. designed research; A.M.S. and S.G. performed research; and A.M.S., J.L.M., and S.G. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹To whom correspondence should be addressed. Email: andrew.saxe@psy.ox.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1820226116/-DCSupplemental.

Published online May 17, 2019.

structure by incrementally adjusting synaptic weights via error-corrective learning (4, 27–32). However, the theoretical principles governing how even simple artificial neural networks extract semantic knowledge from their ongoing stream of experience, embed this knowledge in their synaptic weights, and use these weights to perform inductive generalization, remain obscure.

In this work, our goal is to fill this gap by using a simple class of neural networks—namely, deep linear networks. Surprisingly, this model class can learn a wide range of distinct types of structure without requiring either initial domain theories or a prior set of candidate structural forms, and accounts for a diversity of phenomena involving semantic cognition described above. Indeed, we build upon a considerable neural network literature (27–32) addressing such phenomena through simulations of more complex nonlinear networks. We build particularly on the integrative, simulation-based treatment of semantic cognition in ref. 4, often using the same modeling strategy in a simpler linear setting, to obtain similar results but with additional analytical insight. In contrast to prior work, whether conceptual, Bayesian, or connectionist, our simple model permits exact analytical solutions describing the entire developmental trajectory of knowledge acquisition and organization and its subsequent impact on the deployment and neural representation of semantic structure. In the following, we describe semantic knowledge acquisition, organization, deployment, and neural representation in sequence, and we summarize our main findings in *Discussion*.

A Deep Linear Neural Network Model

Here, we consider a framework for analyzing how neural networks extract semantic knowledge by integrating across many individual experiences of items and their properties, across developmental time. In each experience, given an item as input, the network is trained to produce its associated properties or features as output. Consider, for example, the network's interaction with the semantic domain of living things, schematized in Fig. 1*A*. If the network encounters an item, such as a *canary*, perceptual neural circuits produce an activity vector $\mathbf{x} \in \mathbb{R}^{N_1}$ that identifies this item and serves as input to the semantic system. Simultaneously, the network observes some of the item's properties—for example, that a canary *can fly*. Neural circuits produce an activity feature vector $\mathbf{y} \in \mathbb{R}^{N_3}$ of that item's properties which serves as the desired output of the semantic network. Over time, the network experiences many such episodes with different items and their properties. The total collected experience furnished by the environment to the semantic system is thus a set of P examples $\{\mathbf{x}^i, \mathbf{y}^i\}$, $i = 1, \dots, P$, where the input vector $\mathbf{x}^i$ identifies item i , and the output vector $\mathbf{y}^i$ is a set of features to be associated with this item.

The network's task is to predict an item's properties $\mathbf{y}$ from its perceptual representation $\mathbf{x}$. These predictions are generated by propagating activity through a three-layer linear neural network (Fig. 1*B*). The input activity pattern $\mathbf{x}$ in the first layer propagates through a synaptic weight matrix $\mathbf{W}^1$ of size $N_2 \times N_1$, to

create an activity pattern $\mathbf{h} = \mathbf{W}^1 \mathbf{x}$ in the second layer of N_2 neurons. We call this layer the “hidden” layer because it corresponds neither to input nor output. The hidden-layer activity then propagates to the third layer through a second synaptic weight matrix $\mathbf{W}^2$ of size $N_3 \times N_2$, producing an activity vector $\hat{\mathbf{y}} = \mathbf{W}^2 \mathbf{h}$ which constitutes the output prediction of the network. The composite function from input to output is thus simply $\hat{\mathbf{y}} = \mathbf{W}^2 \mathbf{W}^1 \mathbf{x}$. For each input $\mathbf{x}$, the network compares its predicted output $\hat{\mathbf{y}}$ to the observed features $\mathbf{y}$, and it adjusts its weights so as to reduce the discrepancy between $\mathbf{y}$ and $\hat{\mathbf{y}}$.

To study the impact of depth, we will contrast the learning dynamics of this deep linear network to that of a shallow network that has just a single weight matrix, $\mathbf{W}^s$ of size $N_3 \times N_1$ linking input activities directly to the network's predictions $\hat{\mathbf{y}} = \mathbf{W}^s \mathbf{x}$. It may seem that there is no utility in considering deep linear networks, since the composition of linear functions remains linear. Indeed, the appeal of deep networks is thought to lie in the increasingly expressive functions they can represent by successively cascading many layers of nonlinear elements (33). In contrast, deep linear networks gain no expressive power from depth; a shallow network can compute any function that the deep network can, by simply taking $\mathbf{W}^s = \mathbf{W}^2 \mathbf{W}^1$. However, as we see below, the learning dynamics of the deep network is highly nonlinear, while the learning dynamics of the shallow network remains linear. Strikingly, many complex, nonlinear features of learning appear even in deep linear networks and do not require neuronal nonlinearities.

To illustrate the power of deep linear networks to capture learning dynamics, even in nonlinear networks, we compare the two learning dynamics in Fig. 2. Fig. 2*A* shows a low-dimensional visualization of the simulated learning dynamics of a multilayered nonlinear neural network trained to predict the properties of a set of items in a semantic domain of animals and plants (for details of the neural architecture and training data see ref. 4). The nonlinear network exhibits a striking, hierarchical, progressive differentiation of structure in its internal hidden representations, in which animals vs. plants are first distinguished, then birds vs. fish and trees vs. flowers, and, finally, individual items. This remarkable phenomenon raises important questions about the theoretical principles governing the hierarchical differentiation of structure in neural networks: How and why do the network's dynamics and the statistical structure of the input conspire to generate this phenomenon? In Fig. 2*B*, we mathematically derive this phenomenon by finding analytic solutions to the nonlinear dynamics of learning in a deep linear network, when that network is exposed to a hierarchically structured semantic domain, thereby shedding theoretical insight onto the origins of hierarchical differentiation in a deep network. We present the derivation below, but for now, we note that the resemblance in Fig. 2 suggests that deep linear networks can form an excellent, analytically tractable model for shedding conceptual insight into the learning dynamics, if not the expressive power, of their nonlinear counterparts.

Acquiring Knowledge

We now outline the derivation that leads to Fig. 2*B*. The incremental error-corrective process described above can be formalized as online stochastic gradient descent; each time an example i is presented, the weights $\mathbf{W}^2$ and $\mathbf{W}^1$ receive a small adjustment in the direction that most rapidly decreases the squared error $\|\mathbf{y}^i - \hat{\mathbf{y}}^i\|^2$, yielding the standard back propagation learning rule

$$\Delta \mathbf{W}^1 = \lambda \mathbf{W}^{2T} (\mathbf{y}^i - \hat{\mathbf{y}}^i) \mathbf{x}^{iT}, \quad \Delta \mathbf{W}^2 = \lambda (\mathbf{y}^i - \hat{\mathbf{y}}^i) \mathbf{h}^{iT}, \quad [1]$$

where λ is the learning rate. This incremental update depends only on experience with a single item, leaving open the fundamental question of how and when the accumulation of such incremental updates can extract over developmental time abstract structures, like hierarchical taxonomies, that are emergent properties of the entire domain of items and their features.

Fig. 1. (A) During development, the network experiences sequential episodes with items and their properties. (B) After each episode, the network adjusts its synaptic weights to reduce the discrepancy between actual observed properties $\mathbf{y}$ and predicted properties $\hat{\mathbf{y}}$.

Fig. 2. (A) A two-dimensional multidimensional scaling (MDS) visualization of the temporal evolution of internal representations, across developmental time, of a deep nonlinear neural network studied in ref. 4. Reprinted with permission from ref. 4. (B) An MDS visualization of analytically derived learning trajectories of the internal representations of a deep linear network exposed to a hierarchically structured domain.

We show that the extraction of such abstract domain structure is possible, provided learning is gradual, with a small learning rate λ . In this regime, many examples are seen before the weights appreciably change, so learning is driven by the statistical structure of the domain. We imagine training is divided into a sequence of learning epochs. In each epoch, the above rule is followed for all P examples in random order. Then, averaging [1] over all P examples and taking a continuous time limit gives the mean change in weights per learning epoch,

$$\tau \frac{d}{dt} \mathbf{W}^1 = \mathbf{W}^{2T} (\Sigma^{yx} - \mathbf{W}^1 \Sigma^x), \quad [2]$$

$$\tau \frac{d}{dt} \mathbf{W}^2 = (\Sigma^{yx} - \mathbf{W}^2 \mathbf{W}^1 \Sigma^x) \mathbf{W}^1 T, \quad [3]$$

where $\Sigma^x \equiv E[\mathbf{x}\mathbf{x}^T]$ is an $N_1 \times N_1$ input correlation matrix, $\Sigma^{yx} \equiv E[\mathbf{y}\mathbf{x}^T]$ is an $N_3 \times N_1$ input-output correlation matrix, and $\tau \equiv \frac{1}{P\lambda}$ (see SI Appendix). Here, t measures time in units of learning epochs. These equations reveal that learning dynamics, even in our simple linear network, can be highly complex: The second-order statistics of inputs and outputs drives synaptic weight changes through coupled nonlinear differential equations with up to cubic interactions in the weights.

Explicit Solutions from Tabula Rasa. These nonlinear dynamics are difficult to solve for arbitrary initial conditions. However, we are interested in a particular limit: learning from a state of essentially no knowledge, which we model as small random synaptic weights. To further ease the analysis, we shall assume that the influence of perceptual correlations is minimal ($\Sigma^x \approx \mathbf{I}$). Our goal, then, is to understand the dynamics of learning in [2] and [3] as a function of the input-output correlation matrix Σ^{yx} . The learning dynamics is tied to terms in the singular value decomposition (SVD) of Σ^{yx} (Fig. 3A),

$$\Sigma^{yx} = \mathbf{U}\mathbf{S}\mathbf{V}^T = \sum_{\alpha=1}^{\min(N_1, N_3)} s_{\alpha} \mathbf{u}^{\alpha} \mathbf{v}^{\alpha T}, \quad [4]$$

which decomposes any matrix into the product of three matrices. Each of these matrices has a distinct semantic interpretation.

For example, the α 'th column $\mathbf{v}^{\alpha}$ of the $N_1 \times N_1$ orthogonal matrix $\mathbf{V}$ can be thought of as an object-analyzer vector; it determines the position of item i along an important semantic dimension α in the training set through the component $\mathbf{v}_i^{\alpha}$. To illustrate this interpretation concretely, we consider a simple example dataset with four items (*canary*, *salmon*, *oak*, and *rose*) and five properties (*can fly*, *can swim*, *has bark*, and *has petals*). For this dataset, while the first row of $\mathbf{V}^T$ is a uniform mode, the second row, or the second object-analyzer vector $\mathbf{v}^2$, determines where items sit on an *animal-plant* dimension and hence

has positive values for the *canary* and *salmon* and negative values for the plants. The other dimensions identified by the SVD are a *bird-fish* dimension and a *flower-tree* dimension.

The corresponding α 'th column $\mathbf{u}^{\alpha}$ of the $N_3 \times N_3$ orthogonal matrix $\mathbf{U}$ can be thought of as a feature-synthesizer vector for semantic distinction α . Its components $\mathbf{u}_m^{\alpha}$ indicate the extent to which feature m is present or absent in distinction α . Hence, the feature synthesizer $\mathbf{u}^2$ associated with the *animal-plant* semantic dimension has positive values for *move* and negative values for *roots*, as animals typically can move and do not have roots, while plants behave oppositely. Finally, the $N_3 \times N_1$ singular value matrix $\mathbf{S}$ has nonzero elements $s_{\alpha}, \alpha = 1, \dots, N_1$ only on the diagonal, ordered so that $s_1 \geq s_2 \geq \dots \geq s_{N_1}$. s_{α} captures the overall strength of the α 'th dimension. The large singular value for the *animal-plant* dimension reflects the fact that this one dimension explains more of the training set than the finer-scale dimensions *bird-fish* and *flower-tree*.

Given the SVD of the training set's input-output correlation matrix in [4], we can now explicitly describe the network's learning dynamics. The network's overall input-output map at time t is a time-dependent version of this SVD (Fig. 3B); it shares the object-analyzer and feature-synthesizer matrices of the SVD of Σ^{yx} , but replaces the singular value matrix $\mathbf{S}$ with an effective singular value matrix $\mathbf{A}(t)$,

$$\mathbf{W}^2(t)\mathbf{W}^1(t) = \mathbf{U}\mathbf{A}(t)\mathbf{V}^T = \sum_{\alpha=1}^{N_2} a_{\alpha}(t) \mathbf{u}^{\alpha} \mathbf{v}^{\alpha T}, \quad [5]$$

where the trajectory of each effective singular value $a_{\alpha}(t)$ obeys

$$a_{\alpha}(t) = \frac{s_{\alpha} e^{2s_{\alpha} t / \tau}}{e^{2s_{\alpha} t / \tau} - 1 + s_{\alpha} / a_{\alpha}^0}. \quad [6]$$

Eq. 6 describes a sigmoidal trajectory that begins at some initial value a_{α}^0 at time $t=0$ and rises to s_{α} as $t \rightarrow \infty$, as plotted

Fig. 3. (A) SVD of input-output correlations. Associations between items and their properties are decomposed into modes. Each mode links a set of properties (a column of $\mathbf{U}$) with a set of items (a row of $\mathbf{V}^T$). The strength of the mode is encoded by the singular value (diagonal element of $\mathbf{S}$). (B) Network input-output map, analyzed via the SVD. The effective singular values (diagonal elements of $\mathbf{A}(t)$) evolve over time during learning. (C) Trajectories of the deep network's effective singular values $a_i(t)$. The black dashed line marks the point in time depicted in B. (D) Time-varying trajectories of a shallow network's effective singular values $b_i(t)$.

in Fig. 3C. This solution is applicable when the network begins from an undifferentiated state with little initial knowledge corresponding to small random initial weights (see *SI Appendix*), and it provides an accurate description of the learning dynamics in this regime, as confirmed by simulation in Fig. 3C. The dynamics in Eq. 6 for one hidden layer can be generalized to multiple hidden layers (*SI Appendix*). Moreover, while we have focused on the case of no perceptual correlations, in *SI Appendix*, we examine a special case of how perceptual correlations and semantic task structure interact to determine learning speed. We find that perceptual correlations can either speed up or slow down learning, depending on whether they are aligned or misaligned with the input dimensions governing task-relevant semantic distinctions.

This solution also gives insight into how the internal representations in the hidden layer of the deep network evolve. An exact solution for $\mathbf{W}^2$ and $\mathbf{W}^1$ is given by

$$\mathbf{W}^1(t) = \mathbf{Q}\sqrt{\mathbf{A}(t)}\mathbf{V}^T, \quad \mathbf{W}^2(t) = \mathbf{U}\sqrt{\mathbf{A}(t)}\mathbf{Q}^{-1}, \quad [7]$$

where $\mathbf{Q}$ is an arbitrary $N_2 \times N_2$ invertible matrix (*SI Appendix*). If initial weights are small, then the matrix $\mathbf{Q}$ will be close to orthogonal, i.e., $\mathbf{Q} \approx \mathbf{R}$, where $\mathbf{R}^T \mathbf{R} = \mathbf{I}$. Thus, the internal representations are specified up to an arbitrary rotation $\mathbf{R}$. Factoring out the rotation, the hidden representation of item i is

$$h_i^\alpha(t) = \sqrt{a^\alpha(t)} \mathbf{v}_i^\alpha. \quad [8]$$

Thus, internal representations develop by projecting inputs onto more and more input-output modes as they are learned.

The shallow network has an analogous solution, $\mathbf{W}^s(t) = \sum_{\alpha=1}^{\min(N_1, N_3)} b_\alpha(t) \mathbf{u}^\alpha \mathbf{v}^{\alpha T}$, but each singular value obeys

$$b_\alpha(t) = s_\alpha \left(1 - e^{-t/\tau}\right) + b_\alpha^0 e^{-t/\tau}. \quad [9]$$

In contrast to the deep network's sigmoidal trajectory, Eq. 9 describes a simple exponential approach to s_α , as plotted in Fig. 3D. Hence, depth fundamentally changes the learning dynamics, yielding several important consequences below.

Rapid Stage-Like Transitions Due to Depth. We first compare the time course of learning in deep vs. shallow networks as revealed in Eqs. 6 and 9. For the deep network, beginning from a small initial condition $a_\alpha^0 = \epsilon$, each mode's effective singular value $a_\alpha(t)$ rises to within ϵ of its final value s_α in time

$$t(s_\alpha, \epsilon) = \frac{\tau}{s_\alpha} \ln \frac{s_\alpha}{\epsilon}, \quad [10]$$

in the limit $\epsilon \rightarrow 0$ (*SI Appendix*). This is $O(1/s_\alpha)$ up to a logarithmic factor. Hence, modes with stronger explanatory power, as quantified by the singular value, are learned more quickly. Moreover, when starting from small initial weights, the sigmoidal transition from no knowledge of the mode to perfect knowledge can be arbitrarily sharp. Indeed, the ratio of time spent in the sigmoidal transition regime to time spent before making the transition can go to zero as the initial weights go to zero (*SI Appendix*). Thus, rapid stage-like transitions in learning can be prevalent, even in deep linear networks.

By contrast, the shallow-network learning timescale is

$$t(s_\alpha, \epsilon) = \tau \ln \frac{s_\alpha}{\epsilon}, \quad [11]$$

which is $O(1)$ up to a logarithmic factor. Here, the timescale of learning a mode depends only weakly on its associated singular value. Essentially all modes are learned simultaneously, with an exponential rather than sigmoidal learning curve.

Progressive Differentiation of Hierarchical Structure. We are now almost ready to explain how we analytically derive the result in

Fig. 2B. The only remaining ingredient is a mathematical description of the training data. Indeed, the numerical results in Fig. 2A arise from a toy training set, making it difficult to understand how structure in data drives learning. Here, we introduce a probabilistic generative model to reveal general principles of how statistical structure impacts learning. We begin with hierarchical structure, but subsequently show how diverse structures come to be learned by the network (compare Fig. 9).

We use a generative model corresponding to a branching diffusion process that mimics evolutionary dynamics to create an explicitly hierarchical dataset (*SI Appendix*). In the model, each feature diffuses down an evolutionary tree (Fig. 4A), with a small probability of mutating along each branch. The items lie at the leaves of the tree, and the generative process creates a hierarchical similarity matrix between items such that items with a more recent common ancestor on the tree are more similar to each other (Fig. 4B). We analytically compute the SVD of this dataset, and find that the object-analyzer vectors, which can be viewed as functions on the leaves of the tree in Fig. 4C, respect the hierarchical branches of the tree, with the larger singular values corresponding to broader distinctions. In Fig. 4A, we artificially label the leaves and branches of the evolutionary tree with organisms and categories that might reflect a natural realization of this evolutionary process.

Now, inserting the singular values in Fig. 4D (and *SI Appendix*) into the deep-learning dynamics in Eq. 6 to obtain the time-dependent singular values $a^\alpha(t)$, and then inserting these along with the object analyzers vectors $\mathbf{v}^\alpha$ in Fig. 4C into Eq. 8, we obtain a complete analytic derivation of the development of internal representations in the deep network. An MDS visualization of these evolving hidden representations then yields Fig. 2B, which qualitatively recapitulates the much more complex network and dataset that leads to Fig. 2A. In essence, this analysis completes a mathematical proof that the striking progressive differentiation of hierarchical structure observed in Fig. 2 is an inevitable consequence of deep-learning dynamics, even in linear networks, when exposed to hierarchically structured data. The essential intuition is that dimensions of feature variation across items corresponding to broader hierarchical distinctions have stronger statistical structure, as quantified by the singular values of the training data, and hence are learned faster, leading to waves of differentiation in a deep but not shallow network. This pattern of hierarchical differentiation has long been argued to apply to the conceptual development of infants and children (1, 5–7).

Illusory Correlations. Another intriguing aspect of semantic development is that children sometimes attest to false beliefs [e.g., worms have bones (2)] they could not have learned through experience. These errors challenge simple associationist accounts of semantic development that predict a steady, monotonic accumulation of information about individual properties (2, 16, 17, 34). Yet, the network's knowledge of individual properties exhibits

Fig. 4. Hierarchy and the SVD. (A) A domain of eight items with an underlying hierarchical structure. (B) The correlation matrix of the features of the items. (C) SVD of the correlations reveals semantic distinctions that mirror the hierarchical taxonomy. This is a general property of the SVD of hierarchical data. (D) The singular values of each semantic distinction reveal its strength in the dataset and control when it is learned.

complex, nonmonotonic trajectories over time (Fig. 5). The overall prediction for a property is a sum of contributions from each mode, where the contribution of mode α to an individual feature m for item i is $a_\alpha(t) \mathbf{u}_m^\alpha \mathbf{v}_i^\alpha$. In the example of Fig. 5A, the first two modes make a positive contribution, while the third makes a negative one, yielding the inverted U-shaped trajectory.

Indeed, any property–item combination for which $\mathbf{u}_m^\alpha \mathbf{v}_i^\alpha$ takes different signs across different α will exhibit a nonmonotonic learning curve, making such curves a frequent occurrence, even in a fixed, unchanging environment. In a deep network, two modes with singular values that differ by Δ will have an interval in which the first is learned, but the second is not. The length of this developmental interval is roughly $O(\Delta)$ (SI Appendix). Moreover, the rapidity of the deep network's stage-like transitions further accentuates the nonmonotonic learning of individual properties. This behavior, which may seem hard to reconcile with an incremental, error-corrective learning process, is a natural consequence of minimizing global rather than local prediction error: The fastest possible improvement in predicting all properties across all items sometimes results in a transient increase in the size of errors on specific items and properties. Every property in a shallow network, by contrast, monotonically approaches its correct value (Fig. 5B), and, therefore, shallow networks provably never exhibit illusory correlations (SI Appendix).

Organizing and Encoding Knowledge

We now turn from the dynamics of learning to its final outcome. When exposed to a variety of items and features interlinked by an underlying hierarchy, for instance, what categories naturally emerge? Which items are particularly representative of a categorical distinction? And how is the structure of the domain internally represented?

Category Membership, Typicality, and Prototypes. A long-observed empirical finding is that category membership is not simply a binary variable, but rather a graded quantity, with some objects being more or less typical members of a category (e.g., a sparrow is a more typical bird than a penguin) (8, 9). Indeed, graded judgments of category membership correlate with performance on a range of tasks: Subjects more quickly verify the category membership of more typical items (10, 11), more frequently recall typical examples of a category (12), and more readily extend new information about typical items to all members of a category (13, 14). Our theory begins to provide a natural mathematical definition of item typicality that explains how it emerges from the statistical structure of the environment and improves task performance. We note that all results in this section apply specifically to data generated by binary trees as in Fig. 4A, which exhibit a one-to-one correspondence between singular dimensions of the SVD and individual categorical distinctions (Fig. 4C). The theory characterizes the typicality of items with respect to individual categorical distinctions.

A natural notion of the typicality of an item i for a categorical distinction α is simply the quantity $\mathbf{v}_i^\alpha$ in the corresponding

object-analyzer vector. To see why this is natural, note that, after learning, the neural network's internal representation space has a semantic distinction axis α , and each object i is placed along this axis at a coordinate proportional to $\mathbf{v}_i^\alpha$, as seen in Eq. 8. For example, if α corresponds to the bird–fish axis, objects i with large positive $\mathbf{v}_i^\alpha$ are typical birds, while objects i with large negative $\mathbf{v}_i^\alpha$ are typical fish. Also, the contribution of the network's output to feature neuron m in response to object i , from the hidden representation axis α alone, is given by

$$\hat{\mathbf{y}}_m^\alpha \leftarrow \mathbf{u}_m^\alpha s_\alpha \mathbf{v}_i^\alpha. \quad [12]$$

Hence, under our definition of typicality, an item i that is more typical than another item j will have $|\mathbf{v}_i^\alpha| > |\mathbf{v}_j^\alpha|$, and thus will necessarily have a larger response magnitude under Eq. 12. Any performance measure which is monotonic in the response will therefore increase for more typical items under this definition. Thus, our definition of item typicality is both a mathematically well-defined function of the statistical structure of experience, through the SVD, and provably correlates with task performance in our network.

Several previous attempts at defining the typicality of an item involved computing a weighted sum of category-specific features present or absent in the item (8, 15, 35–37). For instance, a *sparrow* is a more typical bird than a *penguin* because it shares more relevant features (*can fly*) with other birds. However, the specific choice of which features are relevant—the weights in the weighted sum of features—has often been heuristically chosen and relied on prior knowledge of which items belong to each category (8, 36). Our definition of typicality can also be described in terms of a weighted sum of an object's features, but the weightings are uniquely fixed by the statistics of the entire environment through the SVD (SI Appendix):

$$\mathbf{v}_i^\alpha = \frac{1}{P s_\alpha} \sum_{m=1}^{N_3} \mathbf{u}_m^\alpha \mathbf{o}_m^i, \quad [13]$$

which holds for all i and α . Here, item i is defined by its feature vector $\mathbf{o}^i \in R^{N_3}$, where component $\mathbf{o}_m^i$ encodes the value of its m^{th} feature. Thus, the typicality $\mathbf{v}_i^\alpha$ of item i in distinction α can be computed by taking a weighted sum of the components of its feature vector $\mathbf{o}^i$, where the weightings are precisely the coefficients of the corresponding feature-synthesizer vector $\mathbf{u}^\alpha$ (scaled by the reciprocal of the singular value). The neural geometry of Eq. 13 is illustrated in Fig. 6 when α corresponds to the bird–fish categorical distinction.

In many theories of typicality, the particular weighting of object features corresponds to a prototypical object (3, 15), or the best example of a particular category. Such object prototypes are often obtained by a weighted average over the feature vectors for the objects in a category (averaging together the features of all birds, for instance, will give a set of features they share). However, such an average relies on prior knowledge of the extent to which an object belongs to a category. Our theory also yields a natural notion of object prototypes as a weighted average of object feature vectors, but, unlike many other frameworks, it yields a unique prescription for the object weightings in terms of environmental statistics through the SVD (SI Appendix):

$$\mathbf{u}_m^\alpha = \frac{1}{P s_\alpha} \sum_{i=1}^{N_1} \mathbf{v}_i^\alpha \mathbf{o}_m^i. \quad [14]$$

Thus, the feature synthesizer $\mathbf{u}^\alpha$ can be thought of as a category prototype for distinction α , as it is obtained from a weighted average of all object feature vectors $\mathbf{o}^i$, where the weighting of object i is simply the typicality $\mathbf{v}_i^\alpha$ of object i in distinction α . In essence, each element $\mathbf{u}_m^\alpha$ of the prototype vector signifies how important feature m is in distinction α (Fig. 6).

In summary, a beautifully simple duality between item typicality and category prototypes arises as an emergent property of

Fig. 5. Illusory correlations during learning. (A) Predicted value (blue) of feature “can fly” for item “salmon” during learning in a deep network (dataset as in Fig. 3). The contributions from each mode are shown in red. (B) The predicted value and modes for the same feature in a shallow network.

Fig. 9. Representation of explicit structural forms in a neural network. Each column shows a different structure. The first four columns correspond to pure structural forms, while the final column has cross-cutting structure. First row: The structure of the data generating PGM. Second row: The resulting item-covariance matrix arising from either data drawn from the PGM (first four columns) or designed by hand (final column). Third row: The input-analyzing singular vectors that will be learned by the linear neural network. Each vector is scaled by its singular value, showing its importance to representing the covariance matrix. Fourth row: MDS view of the development of internal representations.

progressively narrowing patterns of inductive projection across development.

Consider the act of learning that a familiar item has a novel feature (e.g., “a pine has property x ”). To accommodate this knowledge, new synaptic weights must be chosen between the familiar item *pine* and the novel property x (Fig. 10A), without disrupting prior knowledge of items and their properties already stored in the network. This may be accomplished by adjusting only the weights from the hidden layer to the novel feature so as to activate it appropriately. With these new weights established, inductive projections of the novel feature to other familiar items (e.g., “does a rose have property x ?”) naturally arise by querying the network with other inputs. If a novel property m is ascribed to a familiar item i , the inductive projection of this property to any other item j is given by (SI Appendix)

$$\hat{y}_m = \mathbf{h}_j^T \mathbf{h}_i / \|\mathbf{h}_i\|^2. \quad [16]$$

This equation implements a similarity-based inductive projection to other items, where the similarity metric is precisely the Euclidean similarity of hidden representations of pairs of items (Fig. 10B). In essence, being told “a pine has property x ,” the network will more readily project the novel property x to those familiar items whose hidden representations are close to that of the pine.

A parallel situation arises upon learning that a novel item possesses a familiar feature (e.g., “a blink can move”; Fig. 10C). Encoding this knowledge requires new synaptic weights between the item and the hidden layer. Appropriate weights may be found through standard gradient descent learning of the item-to-hidden weights for this novel item, while holding the hidden-to-output weights fixed to preserve prior knowledge about features. The network can then project other familiar properties to the novel item (e.g., “Does a blink have legs?”) by simply generating a feature output vector given the novel item as input. A novel item i with a familiar feature m will be assigned another familiar feature n through (SI Appendix)

$$\hat{y}_n = \mathbf{h}_n^T \mathbf{h}_m / \|\mathbf{h}_m\|^2, \quad [17]$$

where the α^{th} component of $\mathbf{h}_n$ is $\mathbf{h}_n^\alpha = \mathbf{u}_n^\alpha \sqrt{a_\alpha(t)}$. $\mathbf{h}_n \in \mathbb{R}^{N_2}$ can be thought of as the hidden representation of feature n

at developmental time t . In parallel to [16], this equation now implements similarity-based inductive projection of familiar features to a novel item. In essence, by being told “a blink can move,” the network will more readily project other familiar features to a blink, if those features have a similar internal representation as that of the feature move.

Thus, the hidden layer of the deep network furnishes a common, semantic representational space into which both features and items can be placed. When a novel feature m is assigned to a familiar item i , that novel feature is placed close to the familiar item in the hidden layer, and so the network will inductively project this novel feature to other items close to i in neural space. In parallel, when a novel item i is assigned a familiar feature m , that novel item is placed close to the familiar feature, and so the network will inductively project other features close to m in neural space onto the novel item.

This principle of similarity-based generalization encapsulated in Eqs. 16 and 17, when combined with the progressive differentiation of internal representations as the network learns from hierarchically structured data, as illustrated in Fig. 2B, then naturally explains the developmental shift in patterns of inductive projection from broad to specific, as shown in Fig. 10E. For example, consider specifically the inductive projection of a novel feature to familiar items (Fig. 10A and B). Earlier (later) in developmental time, neural representations of all items are more similar to (different from) each other, and so the network’s similarity-based inductive projection will extend the novel feature to many (fewer) items, thereby exhibiting progressively narrower patterns of projection that respect the hierarchical structure of the environment (Fig. 10E). Thus, remarkably, even a deep linear network can provably exhibit the same broad to specific changes in patterns of inductive projection, empirically observed in many works (2, 3, 17, 18).

Linking Behavior and Neural Representations

Compared with previous models which have primarily made behavioral predictions, our theory has a clear neural

Fig. 10. The neural geometry of inductive generalization. (A) A novel feature (property x) is observed for a familiar item (e.g., “a pine has property x ”). (B) Learning assigns the novel feature a neural representation in the hidden layer of the network that places it in semantic similarity space near the object which possesses the novel feature. The network then inductively projects that novel feature to other familiar items (e.g., “Does a rose have property x ?”) only if their hidden representation is close in neural space. (C) A novel item (a blink) possesses a familiar feature (e.g., “a blink can move”). (D) Learning assigns the novel item a neural representation in the hidden layer that places it in semantic similarity space near the feature possessed by the novel item. Other features are inductively projected to that item (e.g., “Does a blink have wings?”) only if their hidden representation is close in neural space. (E) Inductive projection of a novel property (“a pine has property x ”) over learning. As learning progresses, the neural representations of items become progressively differentiated, yielding progressively restricted projection of the novel feature to other items. Here, the pine can be thought of as the left-most item node in the tree.

Given that optimal learning is a prerequisite for neural similarity mirroring behavioral similarity, as in the previous section, there is a match between the two for pairs of networks trained from small random initial weights (Fig. 11 *A* and *B*, *Middle* and *Bottom*), but not for pairs of networks trained from large random

category's object analyzer, we can prove that typical items must enjoy enhanced performance in semantic tasks relative to atypical items (Eq. 12). Also, by defining the category prototype to be the associated feature synthesizer, we can prove that the most typical items for a category are those that have the most extremal projections onto the category prototype (Fig. 6 and Eq. 13). Finally, by defining the coherence of a category to be the associated singular value, we can prove that more coherent categories can be learned more easily and rapidly (Fig. 7) and explain how changes in the statistical structure of the environment determine what level of a category hierarchy is the most basic or important (Fig. 8). All our definitions of typicality, prototypes, and category coherence are broadly consistent with intuitions articulated in a wealth of psychology literature, but our definitions imbue these intuitions with enough mathematical precision to prove theorems connecting them to aspects of category learnability, learning speed, and semantic task performance in a neural network model.

More generally, beyond categorical structure, our analysis provides a principled framework for explaining how the statistical structure of diverse structural forms associated with different PGMs gradually becomes encoded in the weights of a neural network (Fig. 9). Remarkably, the network learns these structures without knowledge of the set of candidate structural forms, demonstrating that such forms need not be built in. Regarding neural representation, our theory reveals that, across different networks trained to solve a task, while there may be no correspondence at the level of single neurons, the similarity structure

of internal representations of any two networks will both be identical to each other and closely related to the similarity structure of behavior, provided that both networks solve the task optimally, with the smallest possible synaptic weights (Fig. 11).

While our simple neural network captures this diversity of semantic phenomena in a mathematically tractable manner, because of its linearity, the phenomena it can capture still barely scratch the surface of semantic cognition. Some fundamental semantic phenomena that require complex nonlinear processing include context-dependent computations, dementia in damaged networks, theory of mind, the deduction of causal structure, and the binding of items to roles in events and situations. While it is inevitably the case that biological neural circuits exhibit all of these phenomena, it is not clear how our current generation of artificial nonlinear neural networks can recapitulate all of them. However, we hope that a deeper mathematical understanding of even the simple network presented here can serve as a springboard for the theoretical analysis of more complex neural circuits, which, in turn, may eventually shed much-needed light on how the higher-level computations of the mind can emerge from the biological wetware of the brain.

ACKNOWLEDGMENTS. We thank Juan Gao and Madhu Advani for useful discussions. S.G. was supported by the Burroughs-Wellcome, Sloan, McKnight, James S. McDonnell, and Simons foundations. J.L.M. was supported by the Air Force Office of Scientific Research. A.M.S. was supported by Swartz, National Defense Science and Engineering Graduate, and Mind, Brain, & Computation fellowships.

- Keil F (1979) *Semantic and Conceptual Development: An Ontological Perspective* (Harvard Univ Press, Cambridge, MA).
- Carey S (1985) *Conceptual Change In Childhood* (MIT Press, Cambridge, MA).
- Murphy G (2002) *The Big Book of Concepts* (MIT, Cambridge, MA).
- Rogers TT, McClelland JL (2004) *Semantic Cognition: A Parallel Distributed Processing Approach* (MIT Press, Cambridge, MA).
- Mandler J, McDonough L (1993) Concept formation in infancy. *Cogn Dev* 8:291–318.
- Inhelder B, Piaget J (1958) *The Growth of Logical Thinking from Childhood to Adolescence* (Basic Books, New York).
- Siegler R (1976) Three aspects of cognitive development. *Cogn Psychol* 8:481–520.
- Rosch E, Mervis C (1975) Family resemblances: Studies in the internal structure of categories. *Cogn Psychol* 7:573–605.
- Barsalou L (1985) Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *J Exp Psychol Learn Mem Cogn* 11: 629–654.
- Rips L, Shoben E, Smith E (1973) Semantic distance and the verification of semantic relations. *J Verbal Learn Verbal Behav* 12:1–20.
- Murphy G, Brownell H (1985) Category differentiation in object recognition: Typicality constraints on the basic category advantage. *J Exp Psychol Learn Mem Cogn* 11:70–84.
- Mervis C, Catlin J, Rosch E (1976) Relationships among goodness-of-example, category norms, and word frequency. *Bull Psychon Soc* 7:283–284.
- Rips L (1975) Inductive judgments about natural categories. *J Verbal Learn Verbal Behav* 14:665–681.
- Osherson D, Smith E, Wilkie O, López A, Shafir E (1990) Category-based induction. *Psychol Rev* 97:185–200.
- Rosch E (1978) *Principles of Categorization in Cognition and Categorization*, eds Rosch E, Lloyd B (Lawrence Erlbaum, Hillsdale, NJ), pp 27–48.
- Murphy G, Medin D (1985) The role of theories in conceptual coherence. *Psychol Rev* 92:289–316.
- Keil F (1991) The emergence of theoretical beliefs as constraints on concepts. *The Epigenesis of Mind: Essays on Biology and Cognition*, eds Carey S, Gelman R (Psychology Press, New York).
- Carey S (2011) Précis of 'the origin of concepts'. *Behav Brain Sci* 34:113–124.
- Gelman S, Coley J (1990) The importance of knowing a dodo is a bird: Categories and inferences in 2-year-old children. *Dev Psychol* 26:796–804.
- Edelman S (1998) Representation is representation of similarities. *Behav Brain Sci* 21:449–467.
- Kriegeskorte N, Kievit R (2013) Representational geometry: Integrating cognition, computation, and the brain. *Trends Cogn Sci* 17:401–412.
- Carlson T, Simmons R, Kriegeskorte N, Slevc L (2014) The emergence of semantic meaning in the ventral temporal pathway. *J Cogn Neurosci* 26:120–131.
- Connolly A, et al. (2012) The representation of biological classes in the human brain. *J Neurosci* 32:2608–2618.
- Mur M, et al. (2013) Human object-similarity judgments reflect and transcend the primate-IT object representation. *Front Psychol* 4:128.
- Kriegeskorte N, et al. (2008) Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60:1126–1141.
- Kemp C, Tenenbaum J (2008) The discovery of structural form. *Proc Natl Acad Sci USA* 105:10687–10692.
- Hinton G (1986) Learning distributed representations of concepts. *Parallel Distributed Processing: Implications for Psychology and Neurobiology*, ed Morris R (Oxford Univ Press, Oxford).
- Rumelhart D, Todd P (1993) *Learning and Connectionist Representations in Attention and Performance XIV*, eds Meyer D, Kornblum S (MIT Press, Cambridge, MA).
- McClelland J (1995) A connectionist perspective on knowledge and development. *Developing Cognitive Competence: New Approaches to Process Modeling*, eds Simon T, Halford G (Erlbaum, Hillsdale, NJ).
- Plunkett K, Sinha C (1992) Connectionism and developmental theory. *Br J Dev Psychol* 10:209–254.
- Quinn P, Johnson M (1997) The emergence of perceptual category representations in young infants. *J Exp Child Psychol* 66:236–263.
- Colunga E, Smith L (2005) From the lexicon to expectations about kinds: A role for associative learning. *Psychol Rev* 112:347–382.
- Poole B, Lahiri S, Raghu M, Sohl-Dickstein J, Ganguli S (2016) Exponential expressivity in deep neural networks through transient chaos. *Advances in Neural Information Processing Systems*, eds Lee D, Sugiyama M, Luxburg U, Guyon I, Garnett R (MIT Press, Cambridge, MA), Vol 29, pp 3360–3368.
- Gelman R (1990) First principles organize attention to and learning about relevant data: Number and the animate-inanimate distinction as examples. *Cogn Sci* 14:79–106.
- Mervis C, Crisafi M (1982) Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Dev* 53:258–266.
- Davis T, Poldrack R (2014) Quantifying the internal structure of categories using a neural typicality measure. *Cereb Cortex* 24:1720–1737.
- Rosch E, Simpson C, Miller R (1976) Structural bases of typicality effects. *J Exp Psychol Hum Percept Perform* 2:491–502.
- Anglin J (1977) *Word, Object, and Conceptual Development* (Norton, New York), p 302.
- Rosch E, Mervis C, Gray W, Johnson D, Boyes-Braem P (1976) Basic objects in natural categories. *Cogn Psychol* 8:382–439.
- Mandler J, Bauer P, McDonough L (1991) Separating the sheep from the goats: Differentiating global categories. *Cogn Psychol* 23:263–298.
- Cortez J, Gluck M (1992) Explaining basic categories: Feature predictability and information. *Psychol Bull* 111:291–303.
- Khrennikov A, Kozyrev S (2005) Wavelets on ultrametric spaces. *Appl Comput Harmon Anal* 19:61–76.
- Murtagh F (2007) The Haar wavelet transform of a dendrogram. *J Classif* 24:3–32.
- Gray R (2005) Toeplitz and circulant matrices: A review. *Foundations Trends Communications Information Theory* 2:155–239.
- Laakso A, Cottrell G (2000) Content and cluster analysis: Assessing representational similarity in neural systems. *Philos Psychol* 13:47–76.
- Haxby J, et al. (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430.
- Shinkareva S, Malave V, Just M, Mitchell T (2012) Exploring commonalities across participants in the neural representation of objects. *Hum Brain Mapp* 33:1375–1383.
- Raizada R, Connolly A (2012) What makes different people's representations alike: Neural similarity space solves the problem of across-subject fMRI decoding. *J Cogn Neurosci* 24:868–877.