

Evolving Culture vs Local Minima

Yoshua Bengio

Abstract We propose a theory that relates difficulty of learning in deep architectures to culture and language. It is articulated around the following hypotheses: (1) learning in an individual human brain is hampered by the presence of effective local minima; (2) this optimization difficulty is particularly important when it comes to learning higher-level abstractions, i.e., concepts that cover a vast and highly-nonlinear span of sensory configurations; (3) such high-level abstractions are best represented in brains by the composition of many levels of representation, i.e., by deep architectures; (4) a human brain can learn such high-level abstractions if guided by the signals produced by other humans, which act as hints or indirect supervision for these high-level abstractions; and (5), language and the combination of old ideas into new ideas provide an efficient evolutionary recombination operator, and this allows rapid search in the space of communicable ideas that help humans build up better high-level internal representations of their world. These hypotheses put together imply that human culture and the evolution of ideas have been crucial to counter an optimization difficulty: this optimization difficulty would otherwise make it very difficult for human brains to capture high-level knowledge of the world. The theory is grounded in experimental observations of the difficulties of training deep artificial neural networks. Plausible consequences of this theory for the efficiency of cultural evolutions are sketched.

1 Introduction

Interesting connections can sometimes be made at the interface between artificial intelligence research and the sciences that aim to understand human brains, cognition, language, or society. The aim of this paper is to propose and elaborate a theory at this interface, inspired by observations rooted in machine learning research, on so-called

Yoshua Bengio
Department of computer science and operations research, U. Montreal

Deep Learning¹. Deep Learning techniques aim at training models with many levels of representation, a hierarchy of features and concepts, such as can be implemented with artificial neural networks with many layers. The family of functions or distributions represented, e.g., by a neural network is called an architecture, and a *deep architecture* when it has more than 2 or 3 trained levels of representation. The visual cortex is believed to have between 5 and 10 such levels. Theoretical arguments have also been made to suggest that deep architectures are necessary to efficiently represent the kind of high-level concepts required for artificial intelligence (Bengio and LeCun, 2007). This paper starts from experimental observations of the difficulties in training deep architectures (Erhan *et al.*, 2010), and builds a theory of the role of cultural evolution to reduce the difficulty of learning high-level abstractions. The jist of this theory is that training deep architectures such as those found in the brain is difficult because of an optimization difficulty (local minima), but that the cultural evolution of ideas can serve as a way for a whole population of humans, over many generations, to efficiently discover better solutions to this optimization problem.

2 Neural Networks and Local Minima

2.1 Neural Networks

Artificial neural networks are computational architectures and learning algorithms that are inspired from the computations believed to take place in the biological neural networks of the brain (Arbib, 1995). The dominant and most successful approaches to training artificial neural networks are all based on the idea that *learning can proceed by gradually optimizing a criterion* (Rumelhart *et al.*, 1986). A neural network typically has free parameters, such as the synaptic strengths associated with connections between neurons. Learning algorithms formalize the computational mechanism for changing these parameters so as to take into account the evidence provided by observed (training) examples. Different learning algorithms for neural networks differ in the specifics of the criterion and how they optimize it, often approximately because no analytic and exact solution is possible. On-line learning, which is most plausible for biological organisms, involves changes in the parameters either after each example has been seen or after a small batch of examples has been seen (maybe corresponding to a day's worth of experience).

¹ See Bengio (2009) for a review of Deep Learning research, which had a breakthrough in 2006 (Hinton *et al.*, 2006; Bengio *et al.*, 2007; Ranzato *et al.*, 2007)

2.2 Training Criterion

In the case of biological organisms, one could imagine that the ultimate criterion involves the sum of expected future rewards (survival, reproduction, and other innately defined reward signals such as hunger, thirst, and the need to sleep). However, intermediate criteria typically involve modeling the observations from the senses, i.e., improving the prediction that could be made of any part of the observed sensory input given any other part, and improving the prediction of future observations given the past observations. Mathematically, this can often be captured by the statistical criterion of maximizing likelihood, i.e., of maximizing the probability that the model implicitly or explicitly assigns to new observations.

2.3 Learning

The learning problem can be formalized as follows. Let θ be a vector of parameters that are free to change while learning (such as the synaptic strengths of neurons in the brain). Let z represent an example, i.e., a measurement of the variables in the environment which are relevant to the learning agent. The agent is given a past history z_1, z_2, \dots, z_t , which in realistic cases also depends on the actions of the agent. Let $L(\theta, z)$ be a measurement of a loss to be minimized, whose future expected value is the criterion to be minimized. In the simple case where we ignore the effect of current actions on future rewards but only consider the value of a particular solution to the learning problem over the long term, the objective of the learner is to minimize the criterion

$$C(\theta) = \int P(z)L(\theta, z)dz = E[L(\theta, Z)] \quad (1)$$

which is the expected future loss, with $P(z)$ the unknown probability distribution from which the world generates examples for the learner. In the more realistic setting of reinforcement learning (Sutton and Barto, 1998), the objective of the learner is often formalized as the minimization of the expected value of the weighted sum of future rewards, with weights that decay as we go further in the future (food now is valued more than food tomorrow, in general). Although the training criterion cannot be computed exactly (because $P(\cdot)$ is unknown to the learner), the criterion $C(\cdot)$ can be approximately minimized by stochastic gradient descent (as well as other gradient-based optimization techniques): the learner just needs to estimate the gradient $\frac{\partial L(\theta, z)}{\partial \theta}$ of the loss L with respect to the parameters, i.e., estimate the effect of a change of the parameters on the loss. Let g be such an estimator (e.g., if it is unbiased then $E[g] = E[\frac{\partial L(\theta, z)}{\partial \theta}]$). For example, g could be based on a single example or a day's worth of examples. Then stochastic gradient descent proceeds by small steps of the form

$$\theta \leftarrow \theta - \epsilon g \quad (2)$$

where ϵ is a small constant called learning rate or gain. Applying these ideas to biological learners gives the following hypothesis.

Hypothesis H1. When the brain of a single biological agent learns, it performs an approximate optimization with respect to some endogenous objective.

Note that the criterion we have in mind here is not specialized to a single task, as is often the case in applications of machine learning. Instead, a biological learning agent must make good predictions in all the contexts that it encounters, and especially those that are more relevant to its survival. It needs to “solve” many tasks, corresponding to different types of contexts and decisions, something called *multi-task learning*, *transfer learning* or *self-taught learning* (Caruana, 1993; Raina *et al.*, 2007). All of these tasks share the same underlying “world” that surrounds the agent, and brains probably take advantage of these commonalities, allowing them to quickly generalize to new tasks (something that seems difficult with more standard single-task learning algorithms).

Note also that biological agents probably need to address multiple objectives together. However, in practice, since the same brain must take the decisions that can affect all of these criteria, these cannot be decoupled but they can be lumped into a single criterion with appropriate weightings (which may be innate and chosen by evolution). For example, it is very likely that biological learners must cater both to a “predictive” type of criterion (similar to the data-likelihood used in statistical models or in unsupervised learning algorithms) and a “reward” type of criterion (similar to the rewards used in reinforcement learning algorithms). The former explains curiosity and our ability to make sense and learn from things from which we derive no immediate or foreseeable benefit or loss. The latter is clearly crucial for survival, as biological brains need to focus its modeling efforts on what matters most to survival.

2.4 Local Minima

Stochastic gradient descent is one of many optimization techniques that perform a *local descent*: starting from a particular configuration of the parameters (e.g. a configuration of the brain’s synapses), one makes small gradual adjustments which in average tend to improve the expected loss, our training criterion. The theory proposed here relies on the following hypothesis:

Hypothesis H2. When the brain of a single biological agent learns, it relies on approximate local descent in order to gradually improve itself.

Here note that we refer to a *single* learning agent because we exclude the effect of interactions between learning agents, like those that occur because of communication between humans in a human society. Later we will advocate that in fact

when one takes into account the learning going on throughout a society, the optimization is not just a local descent but involves a global parallel search similar to that performed by evolution and sexual reproduction.

Let us now clarify what we mean by local descent. For example, if the learning algorithm is a form of stochastic gradient descent (eq. 2), where g approximates the gradient (it may even have a bias), and if ϵ is chosen small enough (compared to the largest second derivatives of C), then C will gradually decrease with high probability, and if ϵ is gradually decreased at an appropriate rate (such as $1/t$), then the learner will converge towards a *local minimum* of C . The proofs are usually for the unbiased case (Bottou, 2004), but a small bias is not necessarily very hurtful, as shown for Contrastive Divergence (Carreira-Perpiñan and Hinton, 2005; Yuille, 2005), especially if the magnitude of the bias also decreases as the gradient decreases (stochastic approximation convergence theorem (Yuille, 2005)).




Fig. 1 Illustration of learning that proceeds by local descent, and can get stuck near a local minimum (going from left figure to right figure). The horizontal axis represents the space of synaptic configurations (parameters of the learner), while the vertical axis represents the training criterion (expected future loss). The ball represents the learner's current state, which tends to go downwards (improving the expected loss). Note that the space of synaptic configurations is huge (number of synaptic connections on the order of 100 trillion in humans) but represented here schematically with a single dimension, the horizontal axis.

2.5 Effective Local Minima

As illustrated in Figure 1, a *local minimum* is a configuration of the parameters such that no small change can yield an improvement of the training criterion. A consequence of **Hypothesis H2**, if it is true, is therefore that biological brains would be likely to stop improving after some point, after they have sufficiently approached a local minimum. In practice, if the learner relies on a stochastic gradient estimator (which is the only plausible hypothesis we can see, because no biological learner has access to the full knowledge of the world required to directly estimate C), it will continue to change due to the stochastic nature of the gradient estimator (the training signal), hovering stochastically around a minimum. It is also quite possible

that biological learners do not have enough of a lifetime to really get very close to an actual local minimum, but what is plausible is that they get to a point where progress is very slow (so slow as to be indistinguishable from random hovering near a minimum). In practice, when one trains an artificial neural network with a learning algorithm based on stochastic gradient descent, one often observes that training *saturates*, i.e., no more observable progress is seen in spite of the additional examples being shown continuously. The learner appears stuck near a local minimum. Because it is difficult to verify that a learner is really near a local minimum, we call these *effective local minima*. It may happen that the training criterion is a complicated function of the parameters, such that stochastic gradient descent is sometimes practically stuck in a place in which it is not possible to improve in *most directions*, but from where other more powerful descent methods could escape (Martens, 2010).

3 High-Level Abstractions and Deep Architectures

Deep architectures (Bengio, 2009) are parametrized families of functions which can be used to model data using multiple levels of representation. In deep neural networks, each level is associated with a group of neurons (which in the brain could correspond to an area, such as areas V1, V2 or IT of the visual cortex). During sensory perception in animal brains, information travels quickly from lower (sensory) levels to higher (more abstract) levels, but there are also many feedback connections (going from higher to lower levels) as well as lateral connections (between neurons at the same level). Each neuron or group of neurons can be thought of as capturing a concept or feature or aspect, and being activated when that concept or feature or aspect is present in the sensory input, or when the model is generating an internal configuration (a “thought” or “mental image”) that includes that concept or feature or aspect. Note that very few of these features actually come to our consciousness, because most of the inner workings of our brains are not directly accessible (or rarely so) to our consciousness. Note also that a particular linguistic concept may be represented by many neurons or groups of neurons, activating in a particular pattern, and over different levels (in fact so many neurons are activated that we can see whole regions being activated with brain imaging, even when a single linguistic concept is presented as stimulus). This is called *distributed representation* because what would in most symbolic systems be represented by a single “on/off” bit (e.g., the symbol for ‘table’ is activated) is associated in the brain with a large number of neurons and groups of neurons being activated together in a particular pattern. In this way, concepts that are close semantically, i.e., share some attributes (e.g. represented by a group of neurons), can have an overlap in their brain representation, i.e., their corresponding patterns of activation have “on” bits in many of the same places.

3.1 Efficiency of Representation

Deeper architectures can be much more efficient in terms of representation of functions (or distributions) than shallow ones, as shown with theoretical results where for specific families of functions a too shallow architecture can require exponentially more resources than necessary (Yao, 1985; Håstad, 1986; Håstad and Goldmann, 1991; Bengio and LeCun, 2007; Bengio, 2009; Bengio *et al.*, 2010; Bengio and Delalleau, 2011). The basic intuition why this can be true is that in a deep architecture there is *re-use* of parameters and *sharing* of sub-functions to build functions. We do not write computer programs with a single main program: instead we write many subroutines (functions) that can call other subroutines, and this nested re-use provides not only flexibility but also great expressive power. However, this greater expressive power may come at the price of a making the learning task a more difficult optimization problem. Because the lower-level features can be used in many ways to define higher-level features, the interactions between parameters at all levels makes the optimization landscape much more complicated. At the other extreme, many shallow methods are associated with a convex optimization problem, i.e., with a single minimum of the training criterion.

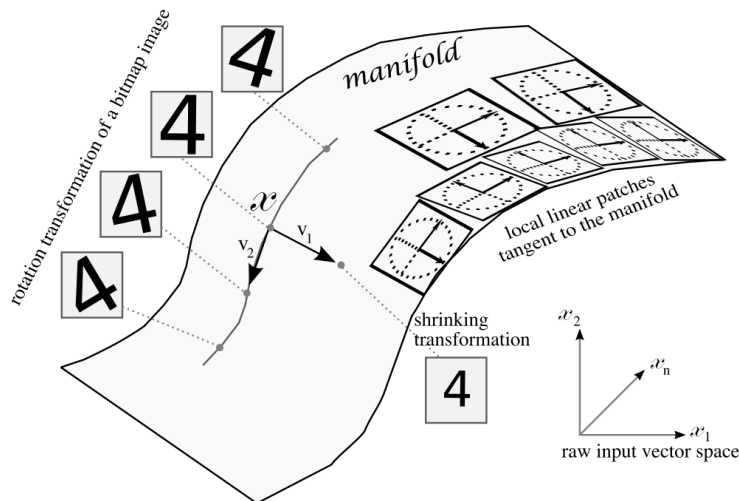


Fig. 2 Example of a simple manifold in the space of images, associated with a rather low-level concrete concept, corresponding to rotations and shrinking of a **specific** instance of the image of a drawn digit 4. Each point on the manifold corresponds to an image which is obtained by rotating or translating or scaling another image on the manifold. The set of points in the manifold defines a concrete concept associated with the drawing of a 4 of a particular shape irrespective of its position, angle and scale. Even learning such simple manifolds is difficult, but learning the much more convoluted and higher-dimensional manifolds of more abstract concepts is much harder.

3.2 High-Level Abstractions

We call *high-level abstraction* the kind of concept or feature that could be computed efficiently only through a deep structure in the brain (i.e., by the sequential application of several different transformations, each associated with an area of the brain or large group of neurons). An edge detector in an image seen by the eye can be computed by a single layer of neurons from raw pixels, using Gabor-like filters. This is a very low-level abstraction. Combining several such detectors to detect corners, straight line segments, curved line segments, and other very local but simple shapes can be done by one or two more layers of neurons, and these can be combined in such a way as to be locally insensitive to small changes in position or angle. Consider a hierarchy of gradually more complex features, constructing detectors for very abstract concepts which are activated whenever any stimulus within a very large set of possible input stimuli are presented. For a higher-level abstraction, this set of stimuli represents a highly-convoluted set of points, a highly curved manifold. We can picture such a manifold if we restrict ourselves to a very concrete concept, like the image of a specific object (the digit 4, as in Figure 2) on a uniform background. The only factors that can vary are due to the changes in geometry (location and orientation of the object with respect to the eye) and lighting, and we can use mathematics to help us make sense of such manifolds. Now think about all the images which can elicit a thought of a more abstract concept, such as “human”, or even more abstract, all the images which can elicit a thought of the concept “relaxing”. These images can be very different from each other, and in many complicated ways, for which scientists do not know how to construct the associated manifolds. Some concepts are clearly higher-level than others, and often we find that higher-level concepts can be defined in terms of lower-level ones, hence forming a hierarchy which is reminiscent of the kind of hierarchy that we find current deep learning algorithms to discover (Lee *et al.*, 2009). This discussion brings us to the formulation of a hypothesis about high-level abstractions and their representation in brains.

Hypothesis H3. Higher-level abstractions in brains are represented by deeper computations (going through more areas or more computational steps).

4 The Difficulty of Training Deep Architectures

There are a number of results in the machine learning literature that suggest that training a deeper architecture is often more difficult than training a shallow one, in the following sense. When trying to train all the layers together with respect to a joint criterion such as the likelihood of the inputs or the conditional likelihood of target classes given inputs, results can be worse than when training a shallow model,

unless some tricks are used that help provide a kind of training signal at each level in the hierarchy (each layer of the deep neural net).

4.1 *Unsupervised Layer-Wise Pre-training*

The first results of that nature appear in Bengio *et al.* (2007); Ranzato *et al.* (2007), where the same architecture gives very different results depending on the initialization of the network weights, either purely randomly, or based on *unsupervised layer-wise pre-training*. The idea of the layer-wise pre-training scheme (Hinton and Salakhutdinov, 2006; Hinton *et al.*, 2006; Bengio *et al.*, 2007; Ranzato *et al.*, 2007) is to train each layer with an unsupervised training criterion, so that it learns a new representation, taking as input the representation of the previous layer. Each layer is thus trained in sequence one after the other. Although this is probably not biologically plausible as such, what would be plausible is a mechanism for providing an unsupervised signal at each layer (group of neurons) that makes it learn to better capture the statistical dependencies in its inputs. That layer-local signal could still be combined with a global training criterion but might help to train deep networks if there is an optimization difficulty in coordinating the training of all layers simultaneously. Another indication that a layer-local signal can help to train deep networks came from the work of Weston *et al.* (2008), where the unsupervised layer-local signal was combined with a supervised global signal that was propagated through the whole network. This observation of the advantage brought by layer-local signals was also made in the context of purely unsupervised learning of a deep stochastic network, the Deep Boltzmann Machine (Salakhutdinov and Hinton, 2009a). By pre-training each layer as a Restricted Boltzmann Machine² before optimizing a Deep Boltzmann Machine (DBM) that comprises all the levels, the authors are able to train the DBM, whereas directly training it from random initialization was problematic. We summarize several of the above results in the deep learning literature with the following **Observation O1**: training deep architectures is easier if hints are provided about the function that intermediate levels should compute (Hinton *et al.*, 2006; Weston *et al.*, 2008; Salakhutdinov and Hinton, 2009a; Bengio, 2009). This is connected to an even more obvious **Observation O2**, from the work on artificial neural networks: it is much easier to teach a network with supervised learning (where we provide it examples of when a concept is present and when it is not present in a variety of examples) than to expect unsupervised learning to discover the concept (which may also happen but usually leads to poorer renditions of the concept).

² RBM, which ignores the interaction with the other levels, except for receiving input from the level below.




Fig. 3 Effect of depth on generalization error, with layer-wise unsupervised pre-training (bottom) and without (top). The training problem becomes more difficult for deeper nets, and using a layer-local cue to initialize each level helps to push the difficulty a bit farther.

4.2 More Difficult for Deeper Architectures and More Abstract Concepts

Another clue to this training difficulty came in later studies showing that directly training all the layers together would not only make it difficult to exploit all the extra modeling power of a deeper architecture but would actually get worse results *as the number of layers is increased* (Larochelle *et al.*, 2009; Erhan *et al.*, 2010), as illustrated in Figure 3. We call this **Observation O3**.

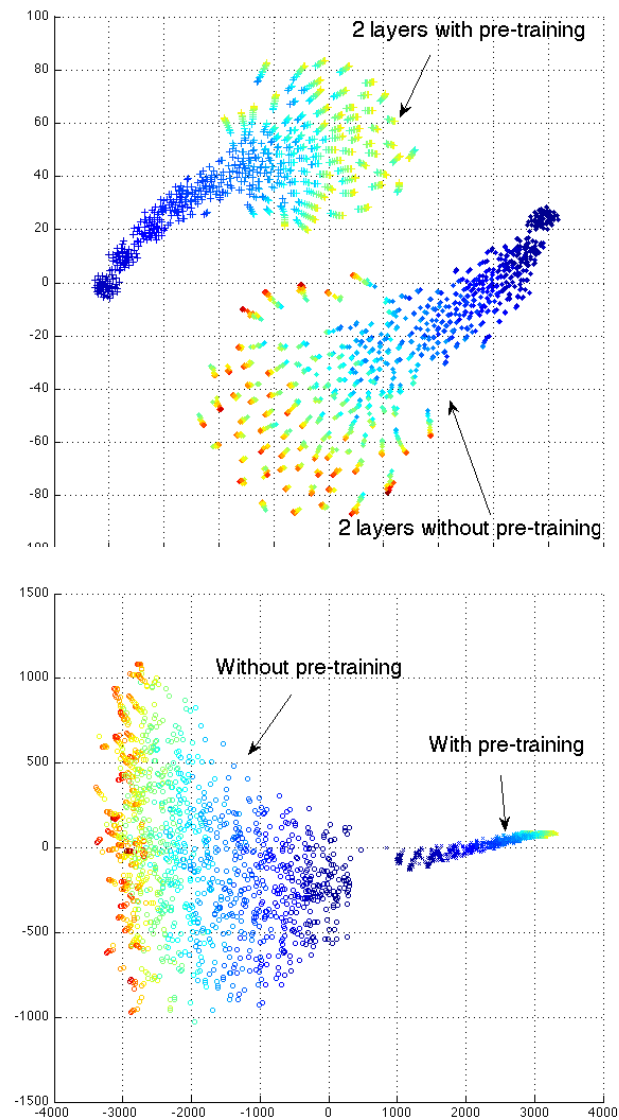


Fig. 4 Two-dimensional non-linear projection of the space of functions visited by artificial neural networks during training. Each cross or diamond or circle represents a neural network at some stage during its training, with color indicating its age (number of examples seen), starting from blue and moving towards red. Networks computing a similar function (with similar response to similar stimuli) are nearby on the graph. Top figure uses *t*-SNE for dimensionality reduction (insists on preserving local geometry) while the bottom figure uses Isomap (insists on preserving global geometry and volumes). The vertical crosses (top figure) and circles (bottom figure) are networks trained from random initialization, while the diamonds (top figure) and rotated crosses (bottom figure) are networks with unsupervised pre-training initialization.

In Erhan *et al.* (2010) we went further in an attempt to understand this training difficulty and studied the trajectory of deep neural networks during training, in function space. Such trajectories are illustrated in Figure 4. Each point in the trajectory corresponds to a particular neural network configuration and is visualized as a two-dimensional point as follows. First, we approximate the function computed by a neural network non-parametrically, i.e., by the outputs of the function over a large test set (of 10000 examples). We consider that two neural networks behave similarly if they provide similar answers on these test examples. We cannot use the network parameters to compare neural networks because the same function can be represented in many different ways (e.g., because permutations of the hidden neuron indices would yield the same network function). We therefore associate each network to a very long vector³, a point in a very high-dimensional space, and we compute these points for all the networks in the experiment. We then learn a mapping from these points to 2-dimensional approximations, so as to preserve local (and sometimes global) structure as much as possible, using non-linear dimensionality reduction methods such as t-SNE (van der Maaten and Hinton, 2008) or Isomap (Tenenbaum *et al.*, 2000). Figure 4 allows us to draw a number of interesting conclusions:

1. **Observation O4.** No two trajectories end up in the same local minimum. This suggests that the number of functional local minima (i.e. corresponding to different functions, each of which possibly corresponding to many instantiations in parameter space) must be huge.
2. **Observation O5.** A training trick (unsupervised pre-training) which changes the initial conditions of the descent procedure allows to reach much better local minima, and these better local minima do not appear to be reachable by chance alone (note how the regions in function space associated with the two “flowers” have no overlap at all, in fact being at nearly 90 degrees from each other in the high-dimensional function space).

Starting from **Hypothesis H2**, **Observation O4** and **Observation O5** bring us to the formulation of new hypothesis:

Hypothesis H4. Learning of a single human learner is limited by effective local minima.

We used the phrase “single human learner” because later in this paper we will hypothesize that a collection of human learners and the associated evolution of their culture can help to get out of what would otherwise be effective local minima.

Combining the above observations with the worse results sometimes observed when training deeper architectures (**Observation O3**, discussed above), we come to the following hypothesis.

³ containing in its elements the concatenation of the network outputs on the test examples.

Hypothesis H5. The detrimental effect of local minima tends to be more pronounced when training deeper architectures (by an optimization method based on iteratively descending the training criterion).

Finally, the presumed ability of deeper architectures to represent higher-level abstractions more easily than shallow ones (see Bengio (2009) and discussion at the beginning of this section) leads us to a human analogue of **Hypothesis H5**, which refines **Hypothesis H4**:

Hypothesis H6. A single human learner is unlikely to discover high-level abstractions by chance because these are represented by a deep sub-network in the brain.

5 Brain to Brain Transfer of Information to Escape Local Minima

If the above hypotheses are true, one should wonder how humans still manage to learn high-level abstractions. We have seen that much better solutions can be found by a learner if it is initialized in an area from which gradient descent leads to a good solution, and genetic material might provide enough of a good starting point and architectural constraints to allow learning of some abstractions. For example, this could be a plausible explanation for some visual abstractions (including simple face detection, which newborns do) and visual invariances, which could have had the chance to be discovered by evolution (since many of our evolutionary ancestors share a similar visual system). Recent work on learning algorithms for computer vision also suggest that architectural constraints can greatly help performance of a deep neural network (Jarrett *et al.*, 2011), to the point where even random parameters in the lower layers (along with appropriate connectivity) suffice to obtain reasonably good performance.

5.1 Labeled Examples as Hints

However, many of the abstractions that we master today have only recently (with respect to evolutionary scales) appeared in human cultures, so they could not have been evolved: each of them must have been discovered by at least one human at some point in the past and then been propagated or improved as they were passed from generation to generation. We will return later to the greater question of the evolution of ideas and abstractions in cultures, but let us first focus on the mechanics of communicating good synaptic configurations from one brain to another. Because we have a huge number of synapses and their values only make sense in the context

of the values of many others, it is difficult to imagine how the recipe for defining individual abstractions could be communicated from one individual to another in a direct way (i.e. by exchanging synaptic values). Furthermore, we need to ask how the hypothesized mechanism could help to escape effective local minima faced by a single learner.

The main insight to answering this question may come from **Observation O1** and **Observation O2**. Training a single layer neural network (supervised or unsupervised) is much easier than training a deeper one, so if one can provide a hint as to the function that deeper layers (corresponding to higher-level abstractions) should capture, then training would be much easier. In the extreme, specifying how particular neurons should respond in specific instances is akin to supervised learning.

The idea that we propose based on these premises relies on learning agents exchanging bits of information in the presence of a shared percept. Communicating about the presence of a concept in a sensory percept is something that humans do, and benefit from since their youngest age. The situation is illustrated in Figure 5.

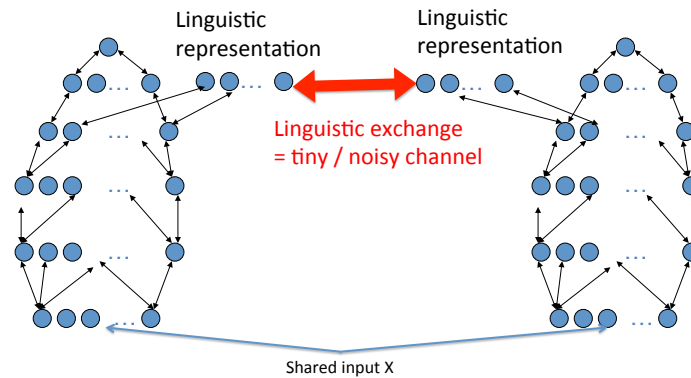


Fig. 5 Illustration of the communication between brains, typically through some language, in a way that can give hints to higher levels of one brain of how the concepts are represented in higher levels of another brain.

5.2 Language for Supervised Training

A very simple schema that would allow to communicate a concept⁴ from one brain to another is one in which there are many encounters between pairs of learners.

⁴ i.e., the characterization of the concept as a function that associates its presence with with all the sensory configurations that are compatible with it,

In each of them, two learners are faced with a similar percept (e.g., they both see the same scene) and they exchange bits of information about it. These bits can for example be indicators of the presence of high-level concepts in the scene. These indicators may reflect the neural activation associated with these high-level concepts. In humans, these bits of information could be encoded through a linguistic convention that helps the receiver of the message interpret them in terms of concepts that it already knows about. One of the most primitive cases of such a communication scenario could occur with animal and human non-verbal communication. For example, an adult animal sees a prey that could be dangerous and emits a danger signal (that could be innate) that a young animal could use as a supervised training signal to associate the prey to danger. Imitation is a very common form of learning and teaching, prevalent among primates, and by which the learner associates contexts with corresponding appropriate behavior. A richer form of communication, which would already be useful, would require simply *naming* objects in a scene. Humans have an innate understanding of the *pointing gesture* that can help identify which object in the scene is being named. In this way, the learner could develop a repertoire of object categories which could become handy (as intermediate concepts) to form theories about the world that would help the learner to survive better. Richer linguistic constructs involve the *combination of concepts* and allow the agents to describe relations between objects, actions and events, sequences of events (stories), causal links, etc., which are even more useful to help a learner form a powerful model of the environment.

This brings us to another hypothesis, supported by **Observation O2** and **Observation O1** and following from **Hypothesis H6**:

Hypothesis H7. A human brain can learn high-level abstractions if guided by the signals produced by other humans, which act as hints or indirect supervision for these high-level abstractions.

This hypothesis is related to much previous work in cognitive science, such as for example *cognitive imitation* (Subiaul *et al.*, 2004), which has been observed in monkeys, and where the learner imitates not just a vocalization or a behavior but something more abstract that corresponds to a cognitive rule.

5.3 *Learning by Predicting the Linguistic Output of Other Agents*

In the schema of Figure 5, it is not necessary for the emitter (who produces the utterance) to directly provide supervision to the high-level layers of the receiver (who receives the communication and can benefit from it). A similar effect can be achieved *indirectly* by simply making sure that the receiver's brain include in its training criterion the objective of *predicting what it observes*, which includes the linguistic output of the emitter in the context of the shared input percept. Since we have already assumed that the training criterion for human brains involves a term for

prediction or maximum likelihood, this could happen naturally, or be enhanced by innate reinforcement (e.g. children pay particular attention to the utterances of their parents). Hence the top-level hidden units of the receiver would receive a training signal that would encourage them to become good features as input to a sub-network that would map these concept-level features to units that represent the probability distribution of utterances that are received (see Figure ??). The same process could work for verbal or non-verbal communication, but using different groups of neurons to model the associated observations. In terms of existing learning algorithms one could for example imagine the case of a Deep Boltzmann Machine (Salakhutdinov and Hinton, 2009b): the linguistic units get 'clamped' by the external linguistic signal received by the learner, at the same time as the lower-level sensory input units get 'clamped' by the external sensory signal, and that conditions the likelihood gradient received by the hidden units, encouraging them to model the interactions between all of these, and in particular the joint distribution of linguistic units and sensory units.

One could imagine many more sophisticated communication schemes that go beyond the above scenario. For example, there could be a two-way exchange of information. It could be that both agents can potentially learn something from the other in the presence of the shared percept. It could be that there is no pre-assigned role of teacher (as emitter) and student (as receiver), but that depending on the confidence demonstrated by each agent for each particular percept, one pays more or less attention to the communicated output of the other. It could be that some aspects of the shared percept is well mastered by one agent but not the other, and vice-versa. Humans have the capability to know that some aspect of a situation is surprising (they would not have predicted it with high probability) and then they should rationally welcome "explanations" provided by others. A way to make the diffusion of useful knowledge more efficient is for the communicating agents to keep track of an estimated degree of "authority" or "credibility" of other agents. One would imagine that parents and older individuals in a human group would by default get more credit, and one of the products of human social systems is that different individuals acquire more or less authority and credibility. For example, scientists strive to maximize their credibility through very rigorous communication practices and a scientific method that insists on verifying hypotheses through experiments designed to test them.

5.4 Language to Evoke Training Examples at Will

Even more interesting scenarios that derive from linguistic abilities involve our ability to *evoke an input scene*. We do not need to be in front of danger to teach about it. We can describe a dangerous situation and mention what is dangerous about it. In this way, the diffusion of knowledge about the world from human brains to other human brains could be made even more efficient.

The fact that verbal and non-verbal communication between animals and humans happens through a noisy and very bandwidth-limited channel is important to keep in mind. Because very few bits of information can be communicated probably means that only the most useful elements should be communicated. If the objective is only to maximize collective learning, it seems that there is no point in communicating something that the receiver already knows. However, there may be other reasons why we communicate, such as for smoothing social interactions, acquiring status or trust, coordinating collective efforts, etc.

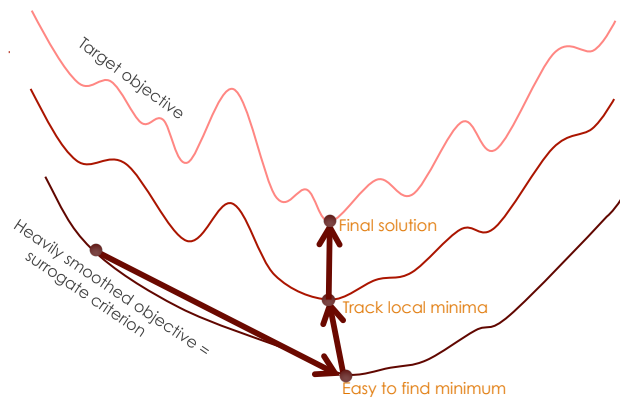


Fig. 6 A general strategy to reduce the impact of local minima is followed in continuation methods and simulated annealing. The idea is to consider a sequence of optimization problems that start with an easier one for which it is easy to find a global optimum (not corresponding to solving the actual problem of interest, though), and ending up in the problem of interest, each time starting at the solution previously found with an easier problem and tracking local minima along the way. It was hypothesized (Bengio, 2009) that following a curriculum could help learners thus find better solutions to the learning problem of interest.

5.5 Connection with Curriculum Learning

The idea that learning can be improved by guiding it, by properly choosing the sequence of examples seen by the learner, was already explored in the past. It was first proposed as a practical way to train animals through *shaping* (Skinner, 1958; Peterson, 2004), as a way to ease simulated learning of more complex tasks (Elman, 1993; Krueger and Dayan, 2009; Bengio *et al.*, 2009) by building on top of easier tasks. An interesting hypothesis introduced in Bengio *et al.* (2009) is that a proper choice of

training examples can be used to approximate a complex training criterion⁵ fraught with local minima with a smoother one (where, e.g., only prototypical examples need to be shown to illustrate the “big picture”). By gradually introducing more subtle examples and building on top of the already understood concepts (as is typically done in pedagogy), Bengio *et al.* (2009) propose that the learner goes through a sequence of gradually more difficult learning tasks, in a way that corresponds in the optimization literature to a *continuation method* or an *annealing method*, allowing to approximately discover global minima (or much better local minima), as illustrated in Figure 6. Interestingly, it was recently observed experimentally that humans use a form of curriculum learning strategy (starting from easier examples and building up) when they are asked to teach a concept to a robot (Khan *et al.*, 2011). (Khan *et al.*, 2011) also propose a statistical reason why a curriculum learning strategy can be more successful, having to do with the uncertainty that the learner has about the relevant factors explaining the variations seen in the data. If these theories are correct, an individual learner can be helped (to escape local minima or converge faster to better solutions) not only by showing examples of abstractions not yet mastered by the learner, but also by showing these examples in an appropriate order. This order corresponds to a curriculum that helps the learner build higher-level abstractions on top of lower-level ones, thus again defeating some of the difficulty believed to exist in training a learner to capture higher-level abstractions.

6 Memes, Cross-Over, and Cultural Evolution

In the previous section we have proposed a general mechanism by which knowledge can be *transmitted* between brains, without having to actually *copy synaptic strengths*. We hypothesized that such mechanisms could help an *individual learner* escape an effective local minimum and thus construct a better model of reality, when the learner is guided by the hints provided by other agents about relevant abstractions. But the knowledge had to come from another agent. Where did this knowledge arise in the first place? This is what we discuss here.

6.1 Memes and Evolution from Noisy Copies

Let us first step back and ask how “better”⁶ brains could arise. The most plausible explanation is that better brains arise due to some form of search or optimization (as stated in **Hypothesis H1**), in the huge space of brain configurations (architecture, function, synaptic strengths). Genetic evolution is a form of parallel search (with

⁵ the training criterion is here seen as a function of the learned parameters, as a sum of a loss over a distribution of possible contexts or examples.

⁶ “better” in the sense of the survival value they provide, and how well they allow their owner to understand the world around them.

each individual's genome representing a candidate solution) that occurs on a rather slow time-scale. Cultural evolution in humans is also a form of search, in the space of ideas or *memes* (Dawkins, 1976). A meme is a unit of selection for cultural evolution. It is something that can be copied from one mind to another. Like for genes, the copy can be imperfect. Memes are analogous to genes in the context of cultural evolution. Genes and memes have co-evolved, although it appears that cultural evolution occurs on a much faster scale than genetic evolution. Culture allows brains to modify their basic program and also go beyond their ability to learn from observing the nature that surrounds them. Culture allows brains to take advantage of knowledge acquired by other brains elsewhere and in previous generations.

To put it all together, the knowledge acquired by an individual brain combines three sources of adaptation: genetic evolution (over hundreds of thousands of years or more), cultural evolution (over dozens, hundreds or thousands of years), and individual learning and discovery (over minutes, hours and days). In all three cases, a form of adaptation is at play, which we hypothesize to be associated with a form of approximate optimization, in the same sense as stated in **Hypothesis H1**. One can also consider the union of all three adaptation processes as a global form of evolution and adaptation (see the work of Hinton and Nowlan (1989) on how learning can guide evolution in the style of Baldwinian evolution). Whereas genetic evolution is a form of parallel search (many individuals carry different combinations and variants of genes which are evaluated in parallel) and we have hypothesized that individual learning is a local search performing an approximate descent (**Hypothesis H2**), what about cultural evolution? Cultural evolution is based on individual learning (a learner trying to predict the behavior and speech output of another individual, as stated in **Hypothesis H7**), which is local, but considering the graph of interactions between humans in an evolving population, one must conclude that cultural evolution, like genetic evolution, is a form of parallel search, as illustrated in Figure 7.




Fig. 7 Illustration of parallel search in the space of synaptic configurations by a population of learners. Some learners start from configuration which happens to lead to a better solution when descending the training criterion.

The most basic working principle of evolution is the *noisy copy* and it is also at work in cultural evolution: a meme can be copied from one brain to another, and the meme can sometime be slightly modified in the process. Selective pressure then does the work of exponentially increasing the presence of successful memes in the population, by increasing the chances that a successful meme be copied in comparison with a competing less successful meme. This may happen simply because a useful meme allows its bearer to survive longer, communicate with more individuals, or because better ideas are promoted. With genetic evolution, it is necessary to copy a whole genome when the individual bearing it is successful. Instead, cultural evolution in humans has mechanisms to evaluate an individual meme and selectively promote it. Good ideas are more likely to be the subject of discussion in public communication, e.g., in the public media, or even better in scientific publications. Science involves powerful mechanisms to separate the worth of a scientist from the worth of his or her ideas (e.g. through independent replication of experimental results or theoretical proofs, or through blind reviewing). That may explain why the pace of evolution of ideas has rapidly increased since the mechanisms for scientific discovery and scientific dissemination of memes have been put in place. The fact that a good idea can stand on its own and be selected for its own value means that the selective pressure is more efficient because it is less hampered by the noisy evaluation that results when fitness is assigned to a whole individual, that integrates many memes and genes.

6.2 *Fast-Forward with Divide-and-Conquer From Recombination*

But if evolution only relied on the noisy copy principle, then it could only speed-up search at best linearly with respect to the number of individuals in a population. Instead of trying N random configurations with N individuals and picking the best by selective pressure, a population with $M > N$ individuals would discover a good selection M/N times faster in average. This is useful but we hypothesize that it would not be enough to make a real dent in the optimization difficulty due to a huge number of poor local minima in the space of synaptic configurations. In fact, evolution has discovered an evolutionary mechanism which can yield much larger speed-ups, and is based on *sexual reproduction* in the case of genetic evolution. With sexual reproduction, we have an interaction between two parent individuals (and their associated candidate configurations), and we mix some of the genes of one with some of the genes of the other in order to create new combinations that are not near-neighbors of either parent. This is very different from a simple parallel search because it allows to explore new configurations beyond local variations (i.e. it is not based on local search), and most importantly, it allows to *combine good, previously found, sub-solutions*. Maybe your father had exceptionally good genes for eyes and your mother exceptionally good genes for ears, and with about 25% probability you could get both, and this may confer you with an advantage that no one had had before. This kind of transformation of the population of candidate configurations is called a *cross-over operator* in the genetic algorithms literature (Holland, 1975). Cross-over is a *recombination mechanism*: it allows to create new candidate solutions by combining *parts* of previous candidate solutions. It has the potential for a much greater speed-up than simple parallelized search based on individual local search (noisy copy), because it can potentially exploit a form of *divide-and-conquer* (which, if well done, could yield exponential speed-up). For the divide-and-conquer aspect of the recombination strategy to work, it is best if *sub-solutions that can contribute as good parts to good solutions receive a high fitness score*. The success rate of recombination is also important i.e., what fraction of the recombination offsprings are viable? The *encoding* of information into genes has a great influence on this success rate as well as on the fitness assigned to good sub-solutions. We hypothesize that memes are particularly good units of selection in these two respects: they are by definition the units of cultural information that can be meaningfully recombined to form new knowledge. All these ideas are summarized in the following hypothesis.

Hypothesis H8. Language and the combination of old memes into new memes provide an efficient evolutionary recombination operator, and this allows rapid search in the space of communicable memes that help humans build up better high-level internal representations of their world.

6.3 *Where do New Ideas Come from?*

Where do completely new ideas (and memes) emerge? According to the views stated here, they emerge from two intertwined effects. On the one hand, our brain can easily combine different memes into new memes which it inherited from other humans, typically through linguistic communication and imitation. On the other hand, such recombination as well as other creations of new memes can arise from the optimization process taking place in a single brain, which tries to reconcile all the sources of evidence that it received into some kind of unifying theory. This search is local in parameter space (synaptic weights) but can involve a stochastic search in the space of neuronal firing patterns. For example, in a Boltzmann machine, neurons fire randomly but with a probability that depends on the activations of other connected neurons, and so as to explore and reach more plausible “interpretations” of the current and past observations (or “planning” for future actions in search for a sequence of decisions that would give rise to most beneficial outcomes), given the current synaptic strengths. In this stochastic exploration, new configurations of neuronal activation can randomly arise and if these do a better job of explaining the data (the observations made), then synaptic strengths will change slightly to make these configurations more likely in the future. This is already how some artificial neural networks learn and “discover” concepts that explain their input. In this way, we can see “concepts” of edges, parts of face, and faces emerge from a deep Boltzmann machine that “sees” images of faces (Lee *et al.*, 2009).

7 Future Work and Conclusion

A lot more needs to be done to connect the above hypotheses with the wealth of data and ideas arising in the social sciences. They can certainly be refined and expanded into more precise statements. Of central importance to future work following up on this paper is how one could go and *test these hypotheses*. Although many of these hypotheses agree with common sense, it might still be worthwhile verifying them empirically, to the extent this is possible. It is also quite plausible that many supporting experimental results from neuroscience, cognitive science, anthropology or primatology already exist that support these hypotheses, and future work should cleanly make the appropriate links.

To test **Hypothesis H1** would seem to require estimating a criterion (not an obvious task) and verifying that learning improves it in average. A proxy for this criterion (or its relative change, which is all we care about, here) might be measurable in the brain itself, for example by measuring the variation in the presence of reward-related molecules or the activity of neurons associated with reward. The effect of learning could be tested by allowing more or less training trials with respect to a rewarding task.

If **Hypothesis H1** is considered true, testing the additional assumptions of **Hypothesis H2** is less obvious because it is difficult to measure the change in synaptic

strengths in many places. However, a form of *stability* of synaptic strengths is a sufficient condition to guarantee that the optimization has to proceed by small changes.

There is already evidence for **Hypothesis H3** in the visual and auditory cortex, in the sense that neurons that belong to areas further away from the sensory neurons seem to perform a higher-level function. Another type of evidence comes from the time required to solve different cognitive tasks, since the hypothesis would predict that tasks requiring computation for the detection of more abstract concepts would require more “iterations” in the recurrent neural network of the brain.

Hypothesis H4 and **Hypothesis H6** are ethically difficult to test but are almost corollaries of the previous hypotheses. An indirect source of evidence may come from raising a primate without any contact with other primates nor any form of guidance from humans, and measure the effect on operational intelligence at different ages. One problem with such an experiment would be that other factors might also explain a poor performance (such as the effect of psychological deprivation from social support, which could lead to depression and other strong causes of poor decisions), so the experiment would require a human that provides warmth and caring, but no guidance whatsoever, even indirectly through imitation. Choosing a more solitary species such as the orangutan would make more sense here (to reduce the impacts due to lack of social support). The question is whether the tested primate could learn to survive as well in the wild as other primates of the same species.

Hypothesis H7 could already be supported by empirical evidence of the effect of education on intelligence, and possibly by observations of feral (wild) children. The important point here is that the intelligence tests chosen should not be about reproducing the academic knowledge acquired during education, but about decisions where having integrated knowledge of some learned high-level abstractions could be useful to properly interpret a situation and take correspondingly appropriate decisions. Using computational simulations with artificial neural networks and machine learning one should also test the validity of mechanisms for “escaping” local minima thanks to “hints” from another agent.

Hypothesis H8 could probably be best tested by computational models where we simulate learning of a population of agents that can share their discoveries (what they learn from data) by communicating the high-level abstractions corresponding to what they observe (as in the scenario of Section 5 and Figure 5). The question is whether one could set up a linguistic communication mechanism that would help this population of learners converge faster to good solutions, compared to a group of isolated learning individuals (where we just evaluate a group’s intelligence by the fitness, i.e. generalization performance, of the best-performing individual after training). If such algorithms would work, then they could also be useful to advance research in machine learning and artificial intelligence, and take advantage of the kind of massive and loose parallelism that is more and more available (to compensate for a decline in the rate of progress of the computing power accessible by a single computer core). This type of work is related to other research on algorithms inspired by the evolution of ideas and culture (see the *Wikipedia* entry on *Memetic Algorithms*).

If many of these hypotheses (and in particular the last one) is true, then we should also draw conclusions regarding the *efficiency of cultural evolution* and how different *social structures* may influence that efficiency, i.e., yielding greater group intelligence in the long run. Two main factors would seem to influence this efficiency: (1), the efficiency of exploration of new memes in the society, and (2), the rate of spread of good memes. Efficiency of exploration in meme-space would be boosted by a greater investment in scientific research, especially in high-risk high potential impact areas. It would also be boosted by encouraging diversity in all its forms because it would mean that individual humans explore a larger region of meme-space. For example, diversity would be boosted by a non-homogeneous education system, a general bias favoring openness to new ideas and multiple schools of thought (even if they disagree), and more generally to marginal beliefs and individual differences. The second factor, the rate of spread of good memes, would be boosted by communication tools such as the Internet, and in particular by open and free access to education, information in general, and scientific results in particular. The investment in education would probably be one of the strongest contributors of this factor, but other interesting contributors would be social structures making it easy for every individual to disseminate useful memes, .e.g., to publish on the web, and the operation of non-centralised systems of rating what is published (whether this is scientific output or individual blogs and posts on the Internet), helping the most interesting new ideas to bubble up and spread faster, and contributing both to diversity of new memes and more efficient dissemination of useful memes.

Acknowledgements

The author would like to thank Caglar Gulcehre, Myriam Côté, Aaron Courville and Olivier Delalleau for useful feedback, as well as NSERC and the Canada Research Chairs for funding.

References

- Arbib, M. A. (1995). *The Handbook of Brain Theory and Neural Networks*. MIT Press.
- Bengio, Y. (2009). Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, **2**(1), 1–127. Also published as a book. Now Publishers, 2009.
- Bengio, Y. and Delalleau, O. (2011). On the expressive power of deep architectures. In *Proceedings of the 22nd International Conference on Algorithmic Learning Theory*. Eds. Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann.
- Bengio, Y. and LeCun, Y. (2007). Scaling learning algorithms towards AI. In L. Bottou, O. Chapelle, D. DeCoste, and J. Weston, editors, *Large Scale Kernel*

- Machines*. MIT Press.
- Bengio, Y., Lamblin, P., Popovici, D., and Larochelle, H. (2007). Greedy layer-wise training of deep networks. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19 (NIPS'06)*, pages 153–160. MIT Press.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. (2009). Curriculum learning. In L. Bottou and M. Littman, editors, *Proceedings of the Twenty-sixth International Conference on Machine Learning (ICML'09)*. ACM.
- Bengio, Y., Delalleau, O., and Simard, C. (2010). Decision trees do not generalize to new variations. *Computational Intelligence*, **26**(4), 449–467.
- Bottou, L. (2004). Stochastic learning. In O. Bousquet and U. von Luxburg, editors, *Advanced Lectures on Machine Learning*, number LNAI 3176 in Lecture Notes in Artificial Intelligence, pages 146–168. Springer Verlag, Berlin.
- Carreira-Perpiñan, M. A. and Hinton, G. E. (2005). On contrastive divergence learning. In R. G. Cowell and Z. Ghahramani, editors, *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AISTATS'05)*, pages 33–40. Society for Artificial Intelligence and Statistics.
- Caruana, R. (1993). Multitask connectionist learning. In *Proceedings of the 1993 Connectionist Models Summer School*, pages 372–379.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.
- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, **48**, 781–799.
- Erhan, D., Bengio, Y., Courville, A., Manzagol, P.-A., Vincent, P., and Bengio, S. (2010). Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, **11**, 625–660.
- Håstad, J. (1986). Almost optimal lower bounds for small depth circuits. In *Proceedings of the 18th annual ACM Symposium on Theory of Computing*, pages 6–20, Berkeley, California. ACM Press.
- Håstad, J. and Goldmann, M. (1991). On the power of small-depth threshold circuits. *Computational Complexity*, **1**, 113–129.
- Hinton, G. E. and Nowlan, S. J. (1989). How learning can guide evolution. *Complex Systems*, **1**, 495–502.
- Hinton, G. E. and Salakhutdinov, R. (2006). Reducing the dimensionality of data with neural networks. *Science*, **313**(5786), 504–507.
- Hinton, G. E., Osindero, S., and Teh, Y. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, **18**, 1527–1554.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press.
- Jarrett, K., Kavukcuoglu, K., Ranzato, M., and LeCun, Y. (2011). What is the best multi-stage architecture for object recognition? In *Proc. International Conference on Computer Vision (ICCV'09)*, pages 2146–2153. IEEE.
- Khan, F., Zhu, X., and Mutlu, B. (2011). How do humans teach: On curriculum learning and teaching dimension. In *Advances in Neural Information Processing Systems (NIPS) 25*.

- Krueger, K. A. and Dayan, P. (2009). Flexible shaping: how learning in small steps helps. *Cognition*, **110**, 380–394.
- Larochelle, H., Bengio, Y., Louradour, J., and Lamblin, P. (2009). Exploring strategies for training deep neural networks. *Journal of Machine Learning Research*, **10**, 1–40.
- Lee, H., Grosse, R., Ranganath, R., and Ng, A. Y. (2009). Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In L. Bottou and M. Littman, editors, *Proceedings of the Twenty-sixth International Conference on Machine Learning (ICML'09)*. ACM, Montreal (Qc), Canada.
- Martens, J. (2010). Deep learning via hessian-free optimization. In L. Bottou and M. Littman, editors, *Proceedings of the Twenty-seventh International Conference on Machine Learning (ICML-10)*, pages 735–742. ACM.
- Peterson, G. B. (2004). A day of great illumination: B. F. Skinner’s discovery of shaping. *Journal of the Experimental Analysis of Behavior*, **82**(3), 317–328.
- Raina, R., Battle, A., Lee, H., Packer, B., and Ng, A. Y. (2007). Self-taught learning: transfer learning from unlabeled data. In Z. Ghahramani, editor, *Proceedings of the Twenty-fourth International Conference on Machine Learning (ICML'07)*, pages 759–766. ACM.
- Ranzato, M., Poultney, C., Chopra, S., and LeCun, Y. (2007). Efficient learning of sparse representations with an energy-based model. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19 (NIPS'06)*, pages 1137–1144. MIT Press.
- Rumelhart, D. E., McClelland, J. L., and the PDP Research Group (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. MIT Press, Cambridge.
- Salakhutdinov, R. and Hinton, G. (2009a). Deep Boltzmann machines. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics (AISTATS 2009)*, volume 8.
- Salakhutdinov, R. and Hinton, G. E. (2009b). Deep Boltzmann machines. In *Proceedings of The Twelfth International Conference on Artificial Intelligence and Statistics (AISTATS'09)*, volume 5, pages 448–455.
- Skinner, B. F. (1958). Reinforcement today. *American Psychologist*, **13**, 94–99.
- Subiaul, F., Cantlon, J., Holloway, R., and Terrace, H. (2004). Cognitive imitation in rhesus macaques. *Science*, **305**(5682), 407–410.
- Sutton, R. and Barto, A. (1998). *Reinforcement Learning: An Introduction*. MIT Press.
- Tenenbaum, J., de Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, **290**(5500), 2319–2323.
- van der Maaten, L. and Hinton, G. E. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research*, **9**, 2579–2605.
- Weston, J., Ratle, F., and Collobert, R. (2008). Deep learning via semi-supervised embedding. In W. W. Cohen, A. McCallum, and S. T. Roweis, editors, *Proceedings of the Twenty-fifth International Conference on Machine Learning (ICML'08)*, pages 1168–1175, New York, NY, USA. ACM.

- Yao, A. (1985). Separating the polynomial-time hierarchy by oracles. In *Proceedings of the 26th Annual IEEE Symposium on Foundations of Computer Science*, pages 1–10.
- Yuille, A. L. (2005). The convergence of contrastive divergences. In L. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17 (NIPS'04)*, pages 1593–1600. MIT Press.