An Analytical Theory of Curriculum Learning in Teacher-Student Networks

Anonymous Author(s)
Affiliation
Address
email

Abstract

In animals and humans, curriculum learning—presenting data in a curated order—is critical to rapid learning and effective pedagogy. A long history of experiments has demonstrated the impact of curricula in a variety of animals but, despite its ubiquitous presence, a theoretical understanding of the phenomenon is still lacking. Surprisingly, in contrast to animal learning, curricula strategies are not widely used in machine learning and recent simulation studies reach the conclusion that curricula are moderately effective or even ineffective in most cases. This stark difference in the importance of curriculum raises a fundamental theoretical question: when and why does curriculum learning help? In this work, we analyse a prototypical neural network model of curriculum learning in the high-dimensional limit, employing statistical physics methods. We study a task in which a sparse set of informative features are embedded amidst a large set of noisy features. We analytically derive average learning trajectories for simple neural networks on this task, which establish a clear speed benefit for curriculum learning in the online setting. However, when training experiences can be stored and replayed (for instance, during sleep), the advantage of curriculum in standard neural networks disappears, in line with observations from the deep learning literature. Inspired by synaptic consolidation techniques developed to combat catastrophic forgetting, we propose curriculum-aware algorithms that consolidate synapses at curriculum change points and investigate whether this can boost the benefits of curricula. We derive generalisation performance as a function of consolidation strength (implemented as an L_2 regularisation/elastic coupling connecting learning phases), and show that curriculum-aware algorithms can yield a large improvement in test performance. Our reduced analytical descriptions help reconcile apparently conflicting empirical results, trace regimes where curriculum learning yields the largest gains, and provide experimentally-accessible predictions for the impact of task parameters on curriculum benefits. More broadly, our results suggest that fully exploiting a curriculum may require explicit consolidation at curriculum boundaries.

1 Introduction

Presenting learning materials in a meaningful order according to a curriculum greatly helps learning in animals and humans [1, 2, 3, 4], and is considered an essential aspect of good pedagogy [5]. For example, humans have been shown to learn visual discriminations faster when presented with examples that exaggerate the relevant difference between classes, a phenomenon known as "fading" [6, 7, 8]. Beyond humans, curricula in the form of "shaping" or "staircase" procedures are a near-universal feature of task designs in animal studies, without which training often fails entirely. For

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

instance, the International Brain Laboratory task, a standardised perceptual decision-making training paradigm in mice, involves six stages of increasing difficulty before reaching final performance [9].

Building from this intuition, a seminal series of papers proposed a similar curriculum learning approach for machine learning (ML) [10, 11, 12]. In striking contrast to the clear benefits of curriculum in biological systems, however, curriculum learning has generally yielded equivocal benefits in artificial systems. Experiments in a variety of domains [13, 14] have found usually modest speed and generalisation improvements from curricula. Recent extensive empirical analyses have found minimal benefits on standard datasets [15]. Indeed, a common intuition in deep learning practice holds that training distributions should ideally be as close as possible to testing distributions, a notion which runs counter to curriculum. Perhaps the only areas where curricula are actively used are in large language models [16] and certain reinforcement learning settings [17].

This gap between the effect of curriculum in biological and artificial learning systems poses a puzzle for theory. When and why is curriculum learning useful? What properties of a task determine the extent of possible benefits? What ordering of learning material is most beneficial? And can new learning algorithms better exploit curricula? Compared to the empirical investigations of curriculum learning, theoretical results on curriculum learning remain sparse. Most notably, [18, 19] show that curriculum can lead to faster learning in a simple setting, but the effects of curriculum on asymptotic generalisation and the dependence on task structure remain unclear. A hint that indeed curriculum learning might lead to statistically different minima comes from a connection between constraint-satisfaction problems and physics results on flow networks [20], but to our knowledge no direct result has been reported in the modern theoretical ML literature.

In this work we study the impact of curriculum using the analytically tractable teacher-student framework and the tools of statistical physics [21, 22, 23, 24]. High-dimensional teacher-student models are a popular approach for systematically studying learning behaviour in neural networks [25, 26, 22], and have recently been leveraged to analyse a variety of phenomena [27, 28, 29, 30, 31, 32]. Using a simple model to build structured data [12], we examine the impact of ordering examples by increasing difficulty (curriculum), decreasing difficulty (anti-curriculum), or standard shuffled training. We derive exact expressions for the online learning dynamics and the performance of batch learning. However, in the latter, curriculum confers no benefit under standard training. Motivated by theories of synaptic consolidation and elastic weight consolidation [33, 34], we introduce elastic penalties (Gaussian priors) that regularise training toward solutions obtained in prior curriculum phases. With these priors, curriculum yields benefits both in the online 3 and in the batch 4 settings.

2 Model definition and overview of approach

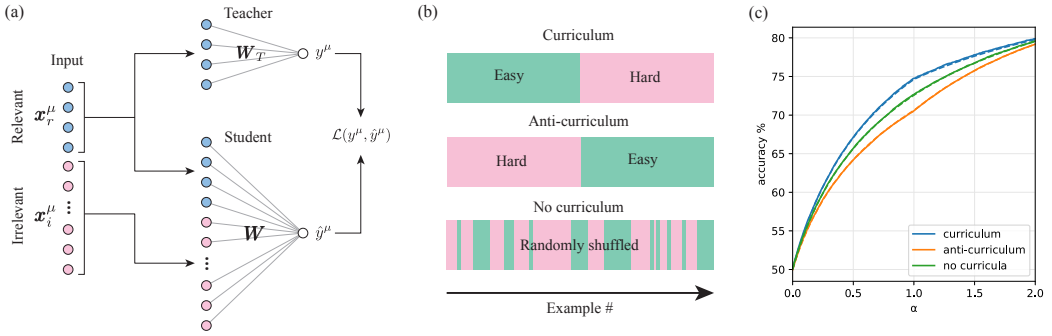

Figure 1: **Teacher-student setting for curriculum learning.** (a) Illustration of teacher-student setting in which a "student" network is trained from *i.i.d.* inputs with labels from a "teacher" network. Since the teacher network is sparse, its output depends only on a subset of *relevant* input features. (b) We consider curricula which order examples by difficulty, here taken to be the variance in the irrelevant feature dimensions. We refer to increasing, decreasing, and random difficulty order as curriculum, anti-curriculum, and no curriculum, respectively. (c) Example test error on hard examples for the student over training. The curriculum boundary lies at $\alpha = 1/2$. Solid lines show numerical simulations, while dashed lines show theoretical predictions derived in Section 3. For this particular parameter setting, curriculum speeds learning but only modestly improves final performance at $\alpha = 1$. Parameters: $\alpha_1 = 1$, $\alpha_2 = 1$, $\Delta_1 = 0$, $\Delta_2 = 1$, $\gamma = 10^{-5}$, $\eta = 3$.

In the following, we revisit a prototypical model of curriculum learning from [12] that finds correspondence to the fading literature [6] as highlighted in Sec. 4. Our setting is summarised in Fig. 1. The model entails a simple teacher-student setup, where teacher and student are each shallow 1-layer neural networks of size N (also known as perceptrons). The learning task for the student is a binary classification problem, with dataset $\mathcal{D} = \{(y^\mu, \boldsymbol{x}^\mu)\}_{\mu=1}^M$, where the ground-truth labels are produced by the teacher network $y^\mu = \text{sign}\, \boldsymbol{W}_T \cdot \boldsymbol{x}^\mu$. A key feature of this model is that the teacher network is sparse, with only a fraction $\rho < 1$ of $\sim \mathcal{N}(0,1)$ non-zero components. Therefore, in order to achieve a good test accuracy, the student has to guess which components should be set to zero and align the relevant weights in the correct direction. A large range of $0 < \rho < 1$ could give rise to the phenomenology we seek to analyse. In the remainder of the paper we will focus on the case $\rho = 0.5$.

We model the variable degree of difficulty in the samples by decomposing each input vector as $\boldsymbol{x}^\mu = [\boldsymbol{x}_r^\mu, \boldsymbol{x}_i^\mu] \in \mathbb{R}^N$, where $\boldsymbol{x}_r^\mu \in \mathbb{R}^{\rho N}$ denotes the relevant components of the input, and $\boldsymbol{x}_i^\mu \in \mathbb{R}^{(1-\rho)N}$ the irrelevant ones. Note that, crucially, the sparse teacher network is completely blind to the irrelevant part of the input: $y^\mu = \text{sign} \sum_{j=1}^{\rho N} W_{T,j} x_{r,j}^\mu$. While $x_{r,j}^\mu$ i.i.d. $\mathcal{N}(0,1)\,, \forall \mu$,[1] we consider the variance for the irrelevant components to be sample-dependent $x_{i,j}^\mu \sim \mathcal{N}(0, \Delta^\mu)$. Note that a smaller variance in the irrelevant part induces a higher SNR in the student learning problem.

The dataset is partitioned according to difficulty levels given by the variances of the irrelevant inputs. For simplicity we consider only two partitions in most of our analysis, but generalisations to more difficulty levels follow straightforwardly. We have a dataset with $M = (\alpha_1 + \alpha_2)N$ total samples, in which the irrelevant inputs of the first $\alpha_1 N$ samples have variance Δ_1, and the remaining $\alpha_2 N$ samples have variance $\Delta_2 > \Delta_1$. In the curriculum learning condition we present the easy examples first, while in the anti-curriculum condition we present the hard examples first. Standard learning presents examples shuffled in random order.

3 Online dynamical solution in the large input limit

We start by focusing on the same online learning setting explored in [12]. We consider a 1-layer student network with sigmoidal activation function, $\sigma(\cdot) = \text{erf}(\cdot/\sqrt{2})$, that learns to minimise a mean square error loss with L_2 regularisation of intensity γ, using gradient descent. This yields the updates

$$\boldsymbol{W}^{\mu+1} = \boldsymbol{W}^\mu - \frac{\eta}{\sqrt{N}} \sigma'\left(\frac{\boldsymbol{W}^\mu \cdot \boldsymbol{x}^\mu}{\sqrt{N}}\right)\left(\sigma\left(\frac{\boldsymbol{W}^\mu \cdot \boldsymbol{x}^\mu}{\sqrt{N}}\right) - y^\mu\right)\boldsymbol{x}^\mu - \gamma \boldsymbol{W}^\mu. \tag{1}$$

The dynamics of the model can be analysed in the high-dimensional limit $N \to \infty$. Generalising the results of [26, 35] on the online stochastic gradient descent dynamics in single-layer regression problems, we obtain a precise description of the performance at all times, as a function of several order parameters: the squared norm of the relevant and irrelevant part of the student weights $Q_r = \frac{1}{N}\boldsymbol{W}^r \cdot \boldsymbol{W}^r$ and $Q_i = \frac{1}{N}\boldsymbol{W}^i \cdot \boldsymbol{W}^i$, respectively; the overlap of the relevant weights of the student and teacher $R = \frac{1}{N}\boldsymbol{W}^r \cdot \boldsymbol{W}_T$; and the squared norm of the teacher vector $T = \frac{1}{N}\boldsymbol{W}_T \cdot \boldsymbol{W}_T$. In particular, given Q_r, Q_i, R and T, the test loss (i.e. average loss on a new example) on a dataset with variance Δ in the irrelevant inputs is given by

$$\mathcal{L}_{\text{MSE}} = \frac{1}{2} + \frac{1}{\pi}\sin^{-1}\frac{Q_r + \Delta Q_i}{1 + Q_r + \Delta Q_i} - \frac{2}{\pi}\sin^{-1}\frac{R/\sqrt{T}}{\sqrt{Q_r + \Delta Q_i + 1}},$$

the accuracy by

$$\mathcal{A} = \mathbb{E}\left[\frac{1}{2}(y\ \text{sign}\hat{y} + 1)\right] = \frac{1}{2} + \frac{1}{\pi}\sin^{-1}\left(\frac{R}{\sqrt{T(Q_r + \Delta Q_i)}}\right). \tag{2}$$

If the dataset contains a random mixture of different difficulty levels $\Delta_1, \Delta_2, \ldots$, the loss and accuracy can be obtained by taking a weighted average over the partitions.

To understand how test performance changes through learning, we study the evolution of the order parameters. Combining their definition with the definition of the dynamics (1) and the fact that the

[1] In [12] the distribution of relevant and irrelevant inputs is uniform between 0 and 1, but this difference does not qualitatively change the results.

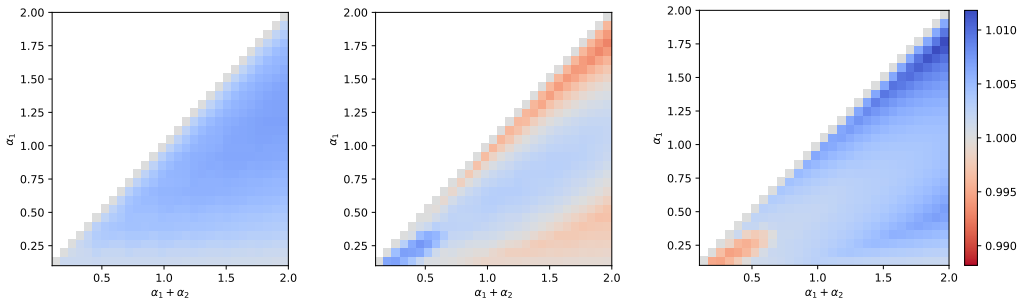

(a) Curriculum learning. (b) Anti-curriculum learning. (c) Curriculum vs anti-curriculum.

Figure 2: **Phase diagram of online learning performance gap with optimal parameters.** The colour scale shows the ratio of the accuracy on hard instances reached by curriculum over no-curriculum (left), anti-curriculum over no-curriculum (center), and curriculum over anti-curriculum (right), as a function of the total dataset size ($\alpha_1 + \alpha_2$) and easy dataset size (α_1). Curriculum broadly benefits performance and anti-curriculum is effective in certain regions, but the size of the improvement is modest. Parameters: $\rho = 0.50, \Delta_1 = 0, \Delta_2 = 1$.

random variables concentrate in the high-dimension as $N \to \infty$, we obtain an analytic form for the updates: $Q_r \leftarrow f_{Q_r}\big(Q_r, Q_i, R, T\big), Q_i \leftarrow f_{Q_i}\big(Q_r, Q_i, R, T\big), R \leftarrow f_R\big(Q_r, Q_i, R, T\big)$; where f_{Q_r}, f_{Q_i} and f_R are long but explicit expressions that are reported in the supplementary material (SM).

Dynamical advantages of curriculum. With these theoretical results in hand, we now characterise the performance of curricula in the online setting. The dynamical equations have two key advantages relative to simulating models in this setting. First, they are free of finite size effects and stochastic fluctuations. And second, their evaluation is very fast, enabling systematic exploration of the parameter space of the problem, along with fine-grained optimisation over hyper-parameters such as learning rate, weight decay and scaling in the initialisation.

Optimising final test accuracy separately for each curriculum strategy, we find that curriculum learning is the optimal strategy, followed by baseline (no-curriculum) and lastly anti-curriculum. In Fig. 1c we show typical learning trajectories for a dataset with equal numbers of easy and hard samples. The results of the simulations (solid lines) are well-described by our theoretical equations (dashed lines), and show that the curriculum strategy leads to better performance throughout training. Fig. 1c shows the evolution during training of the test accuracy computed on the whole dataset.

Next we systematically trace the effect of curriculum for a range of total dataset sizes ($\alpha_1 + \alpha_2$) and number of easy examples α_1 in the phase diagram in Fig. 2. This diagram shows on the left (centre) the accuracy on hard instances reached at the end of training by curriculum learning (anti-curriculum learning respectively) normalised by the accuracy reached by the standard strategy. The two heatmaps show that curriculum learning always outperforms standard learning and that, on the other hand, anti-curriculum learning outperforms standard learning only in part of the diagram. Comparing the two strategies, Fig. 2 (right), we can observe that there is a region for small α and α_1 where anti-curriculum learning is the best strategy, while in the majority of the situations curriculum learning is the best strategy. Interestingly, there is a sizeable region of the diagram in which *both* curriculum and anti-curriculum help, possibly explaining why both have been recommended in prior work [12, 14, 36, 37, 38].

Further, we find that our setting, in which a small task-relevant signal is embedded in large task-irrelevant variation, is critical to the benefit of curriculum. Fig. 5 shows performance as a function of sparsity ρ. Non-sparse tasks do not benefit. Hence curriculum aids tasks with many irrelevant factors of variation. Interestingly, the literature from human psychology shows precisely this: no curriculum benefits for low-D tasks or tasks with no variation in irrelevant dimensions [6].

Our results also highlight the intricate dependence of curriculum on parameters of the learning setup. If not all parameters are correctly optimised, we can observe more complex scenarios. For instance, anti-curriculum learning is always the best strategy starting from a large initialisation, as Fig. 3 shows for weights of order 1. In this case, curriculum learning shows an advantage only in the first phase

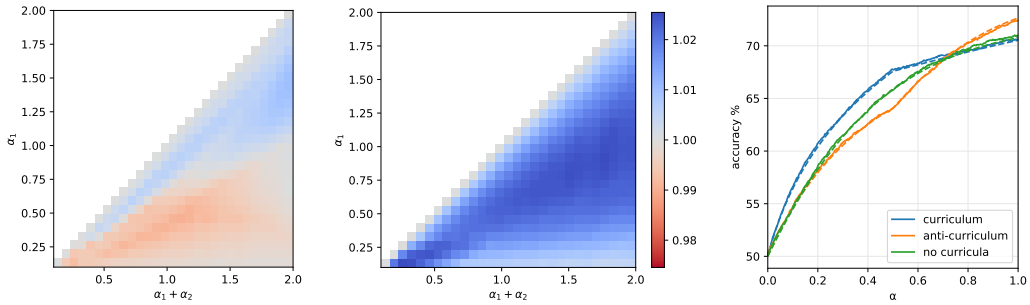

(a) Curriculum learning. (b) Anti-curriculum learning. (c) Accuracy hard samples.

Figure 3: **Performance gap starting from high initialisation.** The first two figures show the accuracy-gap on hard instances between curriculum learning and the baseline (a) and anti-curriculum learning and the baseline (b). Contrary to the phase diagram in Fig. 2, curriculum learning is not always the optimal and anti-curriculum is not always the worst strategy. The right panel shows the accuracy for $\alpha_1 = \alpha_2 = 0.5$.

when easy examples are shown, which is consistent with the results of [19]. However, in the next phase when hard examples are shown, the curriculum strategy does not extract enough information and it is outperformed by the other two strategies. The fact that curriculum or anti-curriculum can look better depending on other learning parameters like initialisation might help explain the confusion in the literature over the best protocol [12, 14, 36, 37, 38]. At least in this model, better performance from anti-curriculum is a signature of a sub-optimal choice of the parameters. To summarise our findings in this online learning setting, curriculum mainly offers a *dynamical advantage*: it speeds learning, with only minimal impact on asymptotic performance.

4 Batch learning solution

The previous section discussed the online case where each example is used once and then discarded. However, in common machine learning practice, neural networks typically revisit each sample repeatedly until convergence. Therefore an important question is: *can curricula lead to a generalisation improvement when trained on the same dataset until convergence?*

We investigate this question by considering a student that learns from slices of a dataset in distinct optimisation phases, where in each phase the student optimises a L_2-regularised logistic loss. Without further modification, curriculum can have no effect in this setting: due to the convex nature of the teacher-student setup [22], the network is bound to converge to a minimum uniquely determined by the final slice of data, with no memory of the progress made at intermediate steps. This simple observation may help explain empirical observations on real data, such as [15], which find no benefit of curriculum in standard settings (though we note that curriculum could still influence non-convex problems [12]). This observation raises the theoretical question of how curriculum learning could induce a non-vanishing effect in batch learning settings.

To instantiate a memory effect in our model, we propose biasing the optimisation landscape via a Gaussian prior, centred around the optimiser of the previous learning phase. The additional term in the loss acts as an elastic coupling between the successive phases, and the associated intensity γ_{12} is then an additional hyper-parameter of the model. This scheme is similar to regularisation methods proposed against catastrophic interference in continual learning, such as Synaptic Intelligence [39].

Tools from statistical physics can be used to analytically compute test performance under this scheme. In order to simplify the presentation, we first consider just two learning phases. It is natural to frame this setting as a 2-level problem, involving two systems with independent copies of the network weights $\boldsymbol{W}_1$ and $\boldsymbol{W}_2$. In a typical statistical physics approach, we associate a Boltzmann-Gibbs measure to the systems, with an energy function determined by the regularised logistic loss $\mathcal{L}_\gamma$. While the statistical properties of the first system can be determined self-consistently, the added elastic interaction creates a dependence of the second measure on the configurations of the first system. In

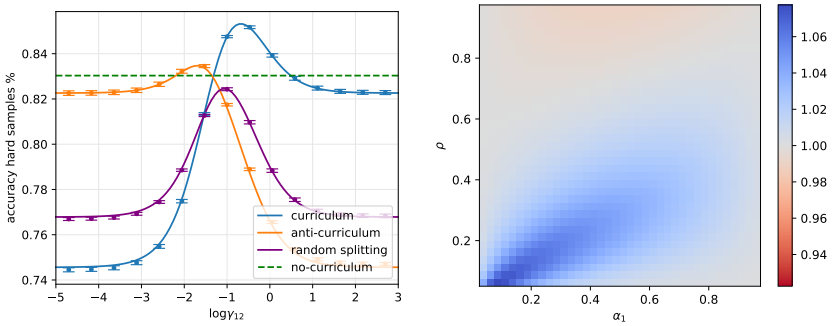

(a) Learning with curricula (b) Curriculum learning vs no-curriculum

Figure 4: **Effect of elastic coupling (Gaussian prior) between curriculum phases.** (Left) comparison between asymptotic performance of curricula (full lines) and single batch learning, at $\alpha_1 = 1$ $\alpha_2 = 1$, with a regularisation γ_1 that yields the best generalisation when learning the entire dataset (in principle not optimal for the other strategies). The points represent the results from 10 numerical simulations at size $N = 2000$. Parameters: $\rho = 0.50$, $\Delta_1 = 0$ and $\Delta_2 = 1$. (Right) ratio between the accuracy reached by curriculum learning over anti-curriculum as a function of the number of easy samples in a dataset of dimension $\alpha_1 + \alpha_2 = 1$, and of the sparsity level of the teacher ρ. Note that ρ can also be seen as the fraction of relevant components in the inputs. $\Delta_1 = 0$ and $\Delta_1 = 1$. $\gamma_1 = \gamma_2$ and γ_{12} where set the values that optimise test performance.

mathematical terms, the coupled system is represented by the following partition function:

$$\langle Z(\boldsymbol{W}_2, \boldsymbol{W}_1; \mathcal{D}_1, \mathcal{D}_2)\rangle_{\boldsymbol{W}_1} = \int d\boldsymbol{W}_1 \frac{e^{-\beta_1 \mathcal{L}_{\gamma_1}(\boldsymbol{W}_1, \mathcal{D}_1)}}{Z_1(\boldsymbol{W}_1)} \log \int d\boldsymbol{W}_2 \, e^{-\beta_2 \left(\mathcal{L}_{\gamma_2}(\boldsymbol{W}_2, \mathcal{D}_2) + \frac{\gamma_{12}}{2} \|\boldsymbol{W}_2 - \boldsymbol{W}_1\|_2^2\right)} \tag{3}$$

where $\mathcal{D}_1, \mathcal{D}_2$ denote the two dataset slices. This object represents the normalisation of the Boltzmann-Gibbs measure, and allows one to extract relevant information on the asymptotic behaviour of our model. The optimisations entailed in each learning phase can be described in the "low noise" limit of $\beta_1, \beta_2 \to \infty$, where the measures focus on the minimisers of the respective losses. In order to study a self-averaging quantity that does not depend on a specific realisation of the dataset, we aim to compute the associated average free-energy:

$$\Phi = \lim_{N\to\infty} \lim_{\beta_1, \beta_2 \to \infty} \frac{1}{\beta_2 N} \left\langle \log \langle Z(\boldsymbol{W}_2, \boldsymbol{W}_1; \mathcal{D}_1, \mathcal{D}_2)\rangle_{\boldsymbol{W}_1} \right\rangle_{\mathcal{D}_1, \mathcal{D}_2} . \tag{4}$$

This quantity can be seen as a special case of the so-called Franz-Parisi potential computation [40, 41], and the entailed double average can be evaluated through the replica method. Refer to SM for details.

Similar to the online case, in high-dimensions the free-entropy concentrates on a deterministic function that depends on several order parameters that capture the geometrical distribution of teacher and student configurations. In addition to those already introduced in Sec. 3, we also have δQ, which is linked to the variance of the student norm. Moreover, for each order parameter we also need to introduce a conjugate parameter, denoted in the following with the hat symbol. The final expression for the free-energy reads:

$$\Phi = \operatorname{extr}\Big[- \left(\hat{R}R + \frac{1}{2} \left(\left(\hat{Q}\delta Q - \delta\hat{Q}Q \right)_{r+i} \right) \right) + g_S(\gamma_1, \gamma_2, \gamma_{12}) + \alpha_1 \, g_E\left(\Delta_1\right) + \alpha_2 \, g_E\left(\Delta_2\right) \Big] \tag{5}$$

where g_S and g_E are two scalar functions, often called entropic and energetic channels, that encode the dependence of the optimisation problem on the Gaussian prior and the logistic loss respectively. The extremum condition for the free-energy yields a system of fixed-point equations that converge to an asymptotic prediction for the order parameters, comparable with the results of numerical simulations on large instances, Fig. 4. At convergence, the order parameters can be inserted again in Eq. 2 to obtain an estimate of the test accuracy. Note that this formalism is not limited to two phases, but can be extended to the case of a discrete number of sequential stages.

The importance of sparsity. Sparsity is a key ingredient in determining the impact of curriculum strategies. It naturally introduces a notion of relevant and irrelevant input components, and defines a secondary learning goal, i.e. identifying what part of the presented data should be disregarded by the model. Curriculum learning can be extremely helpful in this identification process, since the easy samples are more transparent to this structure. This is also observed in human experiments [6]. However, the relative difficulty of the problem of inferring the support of the teacher and the problem of aligning with its non-zero components depends on the degree of sparsity ρ, so the effectiveness of curriculum can vary with it.

In the right panel of Fig. 4, we explore the interplay between the sparsity of the teacher ρ and the fraction of easy samples in the dataset α_1, comparing curriculum with the no-curriculum baseline. The phase diagram highlights the variability in the impact of the curriculum ordering:

- Curriculum is most effective at low values of ρ and close to the diagonal, where the fraction of easy examples in the dataset is comparable to the fraction of relevant dimensions.
- When $\rho > 0.5$, the possible gain from ordering the samples according to difficulty is counterbalanced by the instrinsic cost of splitting the information content into two blocks, thus curriculum can become detrimental.
- When α_1 is too small compared to ρ (above diagonal), the first stage in the curriculum strategy can only help in the support identification problem, but will not allow a good estimation of the direction of the teacher. Because of the elastic prior, the second stage cannot improve too much over it and the effect of curriculum is small.
- When α is larger than the sparsity (below diagonal), the easy examples contain sufficient information for solving both the support and the teacher estimation problems, and this information is also exploited by the baseline. Thus the improvement of curriculum becomes negligible.

We refer to the SM for an in-depth comparison with anti-curriculum.

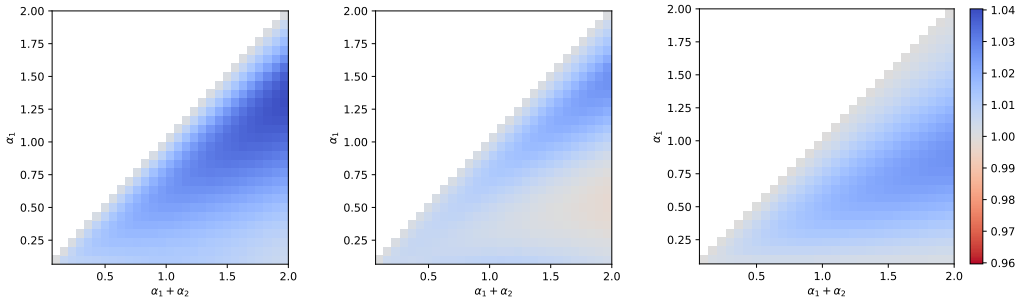

(a) Curriculum learning. (b) Anti-curriculum learning. (c) Curriculum vs anti-curriculum.

Figure 5: **Phase diagram for the performance gap in the batch setting.** The colour scale shows the ratio of the accuracy on hard instances for curriculum over no-curriculum (left), anti-curriculum over no-curriculum (center), and curriculum over anti-curriculum (right), as a function of the total dataset size ($\alpha_1 + \alpha_2$) and easy dataset size (α_1). In contrast to the online case, performance benefits are greater and curriculum is strictly better than anti-curriculum. Both $\gamma_1 = \gamma_2$ and γ_{12} are optimised point-wise, in order to yield the best test accuracy. Parameters: $\rho = 0.50, \Delta_1 = 0, \Delta_2 = 1$.

Asymptotic advantages of curriculum. Contrary to the online SGD case, if the fraction of relevant directions is small, batch learning with elastic coupling notably improves test accuracy of both curriculum and anti-curriculum above the baseline. This confirms the utility of curriculum strategies when the signal is partially "hidden in clutter" [42]. Fig. 5 shows similar phase diagrams to Fig. 2 but for the batch setting. At each point in the phase diagram the regularisation level $\gamma_1 = \gamma_2$ and the coupling γ_{12} are optimised to yield the best accuracy. In batch learning the performance order appears to be nearly always preserved: curriculum followed by anti-curriculum followed by baseline. We remark that this result is not trivial as splitting the learning process in two stages is not advantageous per se. Note that this In the SM we observe similar improvement applying the elastic coupling on the online setting and on real data.

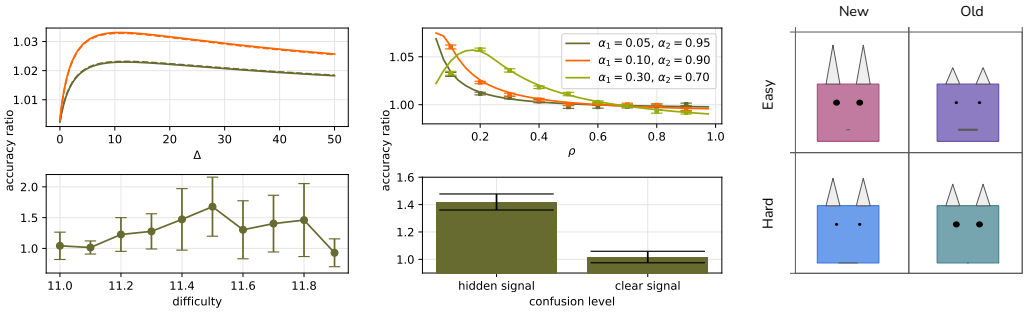

(a) Generalisation gap per difficulty. (b) Role of sparsity. (c) Fading experiment [6].

Figure 6: **Connection with psychology experiments.** (a) Top: Accuracy ratio of different strategies in the model, with curriculum/no-curriculum in green and curriculum/anti-curriculum in orange. The ratio shows non-monotonic behaviour. Bottom: The accuracy ratio obtained by [43]. Parameters $\rho = 0.5$, $\Delta_1 = 0.0$, $\Delta_2 = 1.0$, $\alpha_1 = 1$, $\alpha_2 = 1$ and optimal learning rate, variance at initialisation and weight decay. (b) Top: Dependence on the sparsity of the generalisation gain of curriculum over no-curriculum, measured as ratio between final accuracy, for fixed total dataset size ($\alpha_1 + \alpha_2 = 1$). Bottom: The ratio obtained from experiments 3 and 4 of [6]. (c) Example cartoon stimuli from the "fading" paradigm used in [6], where participants distinguish daemons of the old world from daemons of the new world. The distinguishing feature (horn length) is diluted among many irrelevant features (colour, eye size, mouth size). Highlighting the relevant feature to participants leads to better and faster learning.

Connection with experimental literature

Recent work has suggested that curriculum learning could provide an important window into the learning algorithms at work in biology [44]. Our analysis makes several predictions for curriculum effects. In this section we assess these predictions based on connections to extant experiments and propose future experimental tests.

First, we find that a curriculum strategy yields a speed up in learning in all the tested settings (see Fig. 1c). This acceleration is broadly consistent with the findings from cognitive science [1, 2, 6]. By contrast, our results show that the speed improvement does not necessarily translate into a sizeable generalisation error improvement, and the performance achieved at the end of training can even deteriorate when learning hyperparameters are not fully optimised (c.f. Fig. 3). Deterioration due to curricula has generally not been reported in the psychology literature, though it has been observed in ML [15]. This fact may suggest that animals naturally learn with near-optimal hyperparameters such that curricula generally confer benefits.

A more specific observation concerns the performance on different difficulties after learning. As reported in [43], human and rodent subjects trained in an auditory task using curricula showed the greatest improvement for intermediate level of difficulties as depicted in Fig. 6a bottom panel. The same conclusion can be drawn from the experiment of [7, 8], where, surprisingly, subjects trained with curricula to classify medical images showed poor performance in hard tasks compared to the control group. To address this phenomenon, we calculate accuracy as a function of difficulty in the model in Fig. 6a top panel. Consistent with these experiments, we find regimes where the gap between curriculum learning and the baseline is non-monotonic, with the largest performance gain for intermediate difficulties. Contrary to [7, 8], however, we do not observe negative effects of curriculum for high difficulties. Further experiments that more systematically manipulate training and transfer difficulties could provide a stronger test of these predictions.

A key ingredient in our model is the role of sparsity, such that a small signal is embedded amidst many irrelevant features. Experimentally, the importance of having many factors of variation to obtaining a curriculum effect has been documented in the "fading" experiments of [6]. Human subjects were trained on classification tasks involving stimuli with one task-relevant feature dimension and a variable number of task-irrelevant feature dimensions. Example cartoon "daemon" stimuli are depicted in Fig. 6c, where for instance horn height might be the distinguishing feature while colour,

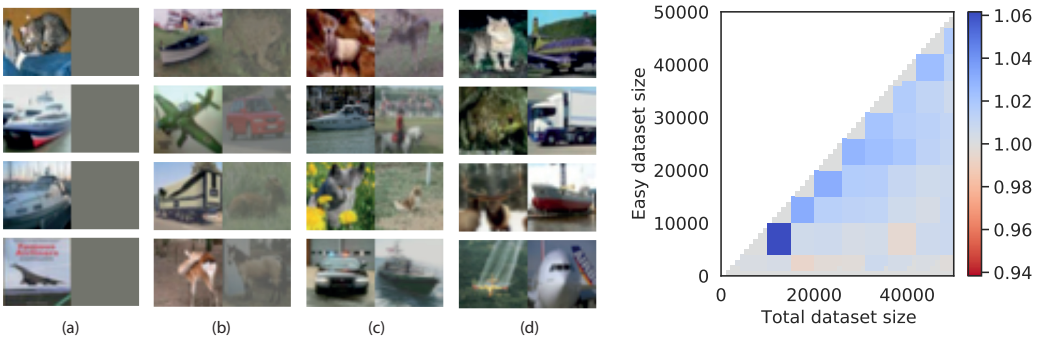

Figure 7: **Experimental setting on CIFAR10-derived data.** (Left) Input samples combine a task-relevant image with a distractor image, and become progressively harder from left to right. (Right) Ratio between final accuracy on hard instances for curriculum learning versus no curriculum. $\eta, \gamma, \gamma_{12}$, init, and stopping time are optimised.

eye size, and mouth size might constitute task-irrelevant features. Without any irrelevant factors of variation ($\rho = 1$), they report no curriculum benefit. By contrast when 75% of features are irrelevant ($\rho = .25$), they record a strong curriculum effect, as shown in Fig. 6b bottom. This qualitative trend is also observed in our model (Fig. 6b top). While these experiments tested only two sparsity levels, further experiments could sample this dimension more extensively and test for interactions with the fraction of easy and hard examples. We note that while the connectionist literature has addressed the effect of curriculum in several settings [45, 46, 10, 11], we found that easy-to-hard effects appear even in a simple setup without need for complex networks and/or dynamics.

Finally, our results may shed light on self-generated curricula during human development [47, 48]. Children undergo a spurt of vocabulary development that coincides with their ability to grasp and centre objects in the visual field [48]. Quantitative estimates of the amount of clutter (irrelevant objects) in self-generated views decrease due to this grasping ability, yielding a self-generated curriculum [42, 49]. Our model similarly predicts that reducing clutter should improve learning speed and performance. To verify this in a richer visual setting, we apply our curriculum-aware algorithm on real data. We construct a simple cluttered object recognition task from the CIFAR10 dataset [50] by patching two images together into a 32×64 input image (Fig. 7a). The network has to learn that the classification depends only on the left image, while the right image is a distractor that is irrelevant to the classification. Easy and hard instances are obtained by reducing the brightness of the distractor. We train a single-layer network with the cross-entropy loss and the curriculum protocol with Gaussian prior between the two stages. Weights are optimised using SGD and momentum, with an annealed learning rate. All training parameters are optimised, and full details are presented in the SM. Fig. 7 shows a robust curriculum advantage in this setting, suggesting a possible functional benefit for children's self-generated curriculum.

Conclusions

We analysed a model of curriculum learning introduced by [12] and amenable of analytical treatment. This simple setting sheds light on results observed in the cognitive science and machine learning literature, and the theoretical tractability allows for exploration of a wide range of parameters that would be costly to obtain through experiments. Future work will need to move beyond models with simple loss landscapes to address the impact of curricula in complex tasks like reinforcement learning. Nevertheless, the model recapitulates a variety of observations in the literature [43, 51, 52], revealing that easy-to-hard effects can appear when a sparse signal is embedded in many irrelevant dimensions of variation. We find that curriculum-aware algorithms can better exploit curricula, offering a potential route for improved practical algorithms. On the psychology side, some of our predictions can help in designing new experiments. The benefit of anti-curriculum learning for intermediate sparsity, is a counter-intuitive result testable in animal experiments.

References

[1] Douglas H Lawrence. The transfer of a discrimination along a continuum. *Journal of Comparative and Physiological Psychology*, 45(6):511, 1952.

[2] Robert A Baker and Stanley W Osgood. Discrimination transfer along a pitch continuum. *Journal of Experimental Psychology*, 48(4):241, 1954.

[3] Renee Elio and John Anderson. The effects of information order and learning mode on schema abstraction. *Memory & Cognition*, 12:20–30, January 1984.

[4] Robert C. Wilson, Amitai Shenhav, Mark Straccia, and Jonathan D. Cohen. The Eighty Five Percent Rule for optimal learning. *Nature Communications*, 10(1):4646, November 2019.

[5] Judith Avrahami, Yaakov Kareev, Yonatan Bogot, Ruth Caspi, Salomka Dunaevsky, and Sharon Lerner. Teaching by Examples: Implications for the Process of Category Acquisition. *The Quarterly Journal of Experimental Psychology Section A*, 50(3):586–606, August 1997. Publisher: SAGE Publications.

[6] Harold Pashler and Michael C. Mozer. When does fading enhance perceptual category learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(4):1162–1173, 2013.

[7] Adam N. Hornsby and Bradley C. Love. Improved classification of mammograms following idealized training. *Journal of Applied Research in Memory and Cognition*, 3(2):72–76, June 2014.

[8] Brett D. Roads, Buyun Xu, June K. Robinson, and James W. Tanaka. The easy-to-hard training advantage with real-world medical images. *Cognitive Research: Principles and Implications*, 3, October 2018.

[9] The International Brain Laboratory, Valeria Aguillon-Rodriguez, Dora Angelaki, Hannah Bayer, Niccolo Bonacchi, Matteo Carandini, Fanny Cazettes, Gaelle Chapuis, Anne K Churchland, Yang Dan, Eric Dewitt, Mayo Faulkner, Hamish Forrest, Laura Haetzel, Michael Häusser, Sonja B Hofer, Fei Hu, Anup Khanal, Christopher Krasniak, Ines Laranjeira, Zachary F Mainen, Guido Meijer, Nathaniel J Miska, Thomas D Mrsic-Flogel, Masayoshi Murakami, Jean-Paul Noel, Alejandro Pan-Vazquez, Cyrille Rossant, Joshua Sanders, Karolina Socha, Rebecca Terry, Anne E Urai, Hernando Vergara, Miles Wells, Christian J Wilson, Ilana B Witten, Lauren E Wool, and Anthony M Zador. Standardized and reproducible measurement of decision-making in mice. *eLife*, 10:e63711, May 2021. Publisher: eLife Sciences Publications, Ltd.

[10] Jeffrey L. Elman. Learning and development in neural networks: the importance of starting small. *Cognition*, 48(1):71–99, July 1993.

[11] Kai A. Krueger and Peter Dayan. Flexible shaping: how learning in small steps helps. *Cognition*, 110(3):380–394, March 2009.

[12] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009.

[13] Anastasia Pentina, Viktoriia Sharmanska, and Christoph H. Lampert. Curriculum learning of multiple tasks. pages 5492–5500. IEEE Computer Society, June 2015. ISSN: 1063-6919.

[14] Guy Hacohen and Daphna Weinshall. On the power of curriculum learning in training deep networks. In *ICML*, volume 97, pages 2535–2544. PMLR, 2019.

[15] Xiaoxia Wu, Ethan Dyer, and Behnam Neyshabur. When do curricula work? *ICLR*, 2020.

[16] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott

Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.

[17] Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. Prioritized Level Replay. *arXiv:2010.03934 [cs]*, January 2021. arXiv: 2010.03934.

[18] Daphna Weinshall, Gad Cohen, and Dan Amir. Curriculum learning by transfer learning: Theory and experiments with deep networks. In *International Conference on Machine Learning*, pages 5238–5246. PMLR, 2018.

[19] Daphna Weinshall and Dan Amir. Theory of curriculum learning, with convex loss functions. *Journal of Machine Learning Research*, 21(222):1–19, 2020.

[20] Miguel Ruiz-García, Andrea J Liu, and Eleni Katifori. Tuning and jamming reduced to their minima. *Physical Review E*, 100(5):052608, 2019.

[21] Marc Mézard, Giorgio Parisi, and Miguel Angel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987.

[22] Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001.

[23] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.

[24] Yasaman Bahri, Jonathan Kadmon, Jeffrey Pennington, Sam S Schoenholz, Jascha Sohl-Dickstein, and Surya Ganguli. Statistical mechanics of deep learning. *Annual Review of Condensed Matter Physics*, 2020.

[25] Leticia F Cugliandolo and Jorge Kurchan. Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model. *Physical Review Letters*, 71(1):173, 1993.

[26] Michael Biehl and Holm Schwarze. Learning by on-line gradient descent. *Journal of Physics A: Mathematical and general*, 28(3):643, 1995.

[27] M.S. Advani, A.M. Saxe, and H. Sompolinsky. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132:428 – 446, 2020.

[28] Sebastian Goldt, Madhu Advani, Andrew M Saxe, Florent Krzakala, and Lenka Zdeborová. Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[29] Stefano Sarao Mannelli, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborova. Passed & spurious: Descent algorithms and local minima in spiked matrix-tensor models. 97:4333–4342, 09–15 Jun 2019.

[30] Stefano Sarao Mannelli, Giulio Biroli, Chiara Cammarota, Florent Krzakala, and Lenka Zdeborová. Who is afraid of big bad minima? analysis of gradient-flow in spiked matrix-tensor models. In *Advances in Neural Information Processing Systems*, pages 8679–8689, 2019.

[31] Stefano Sarao Mannelli, Giulio Biroli, Chiara Cammarota, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Complex dynamics in simple neural networks: Understanding gradient flow in phase retrieval. 33:3265–3274, 2020.

[32] Hugo Cui, Luca Saglietti, and Lenka Zdeborová. Large deviations for the perceptron model and consequences for active learning. In *Mathematical and Scientific Machine Learning*, pages 390–430. PMLR, 2020.

[33] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017.

[34] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

[35] David Saad and Sara A Solla. Exact solution for on-line learning in multilayer neural networks. *Physical Review Letters*, 74(21):4337, 1995.

[36] Tom Kocmi and Ondřej Bojar. Curriculum Learning and Minibatch Bucketing in Neural Machine Translation. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 379–386, Varna, Bulgaria, September 2017. INCOMA Ltd.

[37] Nathan Schneider, Dirk Hovy, Anders Johannsen, and Marine Carpuat. Semeval-2016 task 10: Detecting minimal semantic units and their meanings (dimsum). In Steven Bethard, Daniel M. Cer, Marine Carpuat, David Jurgens, Preslav Nakov, and Torsten Zesch, editors, *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016*, pages 546–559. The Association for Computer Linguistics, 2016.

[38] Xuan Zhang, Pamela Shapiro, Gaurav Kumar, Paul McNamee, Marine Carpuat, and Kevin Duh. Curriculum Learning for Domain Adaptation in Neural Machine Translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1903–1915, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

[39] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3987–3995, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.

[40] Silvio Franz and Giorgio Parisi. Phase diagram of coupled glassy systems: A mean-field study. *Physical review letters*, 79(13):2486, 1997.

[41] Luca Saglietti and Lenka Zdeborová. Solvable model for inheriting the regularization through knowledge distillation. *CoRR*, abs/2012.00194, 2020.

[42] Elizabeth M. Clerkin, Elizabeth Hart, James M. Rehg, Chen Yu, and Linda B. Smith. Real-world visual statistics and infants' first-learned object names. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 372(1711), January 2017.

[43] Estella H Liu, Eduardo Mercado III, Barbara A Church, and Itzel Orduña. The easy-to-hard effect in human (homo sapiens) and rat (rattus norvegicus) auditory identification. *Journal of Comparative Psychology*, 122(2):132, 2008.

[44] Daniel R. Kepple, Rainer Engelken, and Rajan Kanaka. Curriculum learning as a tool to uncover learning principles in the brain. *ICLR*, 2022.

[45] Kim Plunkett, Virginia Marchman, and Steen Ladegaard Knudsen. From rote learning to system building: acquiring verb morphology in children and connectionist nets. In *Connectionist Models*, pages 201–219. Elsevier, 1991.

[46] Kim Plunkett and Virginia Marchman. U-shaped learning and frequency effects in a multi-layered perception: Implications for child language acquisition. *Cognition*, 38(1):43–102, 1991.

[47] Hadar Karmazyn Raz, Drew H. Abney, David Crandall, Chen Yu, and Linda B. Smith. How do infants start learning object names in a sea of clutter? *Annual Conference of the Cognitive Science Society*, 2019:521–526, July 2019.

[48] Linda B. Smith and Lauren K. Slone. A Developmental Approach to Machine Learning? *Frontiers in Psychology*, 8, 2017.

[49] Chen Yu and Linda B. Smith. Embodied attention and word learning by toddlers. *Cognition*, 125(2):244–262, November 2012.

[50] Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. 2009.

[51] Itzel Orduña, Estella H Liu, Barbara A Church, Ann C Eddins, and Eduardo Mercado III. Evoked-potential changes following discrimination learning involving complex sounds. *Clinical neurophysiology*, 123(4):711–719, 2012.

[52] Barbara A Church, Eduardo Mercado III, Matthew G Wisniewski, and Estella H Liu. Temporal dynamics in auditory perceptual learning: impact of sequencing and incidental learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(1):270, 2013.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [No] The model is theoretical in nature and we do not foresee any negative societal impact.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Refer to sections "model definition".
   (b) Did you include complete proofs of all theoretical results? [N/A] In the SM we provide details about the derivation.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The architecture used is very simple to simulate and we do not provide code.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] All the figures resulting from the experiments contain details to produce them. Full parameters for the experiments on real data are given in the SM.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We report estimated total amount of compute and type of compute in the SM for the experiment on real data (≈ 10000 GPU hours, ≈ 1110 kg CO2 eq). The data for each theory-based figure can be obtained in a few hours of laptop simulation.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We use and cite the CIFAR10 dataset.
   (b) Did you mention the license of the assets? [No] It is widely known that CIFAR10 is an MIT License dataset.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]