
Skill-it! A Data-Driven Skills Framework for Understanding and Training Language Models

Mayee F. Chen¹ Nicholas Roberts² Kush Bhatia¹ Jue Wang³ Ce Zhang⁴ Frederic Sala² Christopher Ré¹

Abstract

The quality of training data impacts the performance of pre-trained large language models (LMs). Given a fixed budget of tokens, it is unclear what data to best select for the model’s performance across tasks. To study this, we develop a new framework based on a simple hypothesis: similar to how humans acquire interdependent skills in a deliberate order, there exists a natural order in how the LM best learns a set of skills from its training data. If such order exists, it can be exploited for improved understanding of LMs and data-efficient training. Using this intuition, our framework formalizes the notion of a skill and of an ordered set of skills in terms of their associated data. We demonstrate that these ordered skill sets exist on synthetic and real data, and their existence enables skills to be learned with less data given that we train on their prerequisite skills. Using our proposed framework, we introduce an online data sampling algorithm, SKILL-IT, over mixtures of skills for both continuous pre-training and fine-tuning regimes, where we aim to efficiently learn multiple skills in the former and an individual skill in the latter. On the LEGO synthetic in the continual pre-training setting, SKILL-IT obtains 36.5 points higher accuracy than random sampling. On the Natural Instructions dataset in the fine-tuning setting, SKILL-IT reduces the validation loss on the target skill by 13.6% versus training on the skill itself. We apply our skills framework on the recent RedPajama dataset to continually pre-train a 3B-parameter LM, achieving higher accuracy on the LM Evaluation Harness with 1B tokens than uniform sampling over data sources with 3B tokens.

¹Department of Computer Science, Stanford University, USA
²Department of Computer Sciences, University of Wisconsin-Madison, USA
³Department of Computer Science, Zhejiang University, China
⁴Department of Computer Science, ETH-Zürich, Switzerland. Correspondence to: Mayee Chen <mfchen@stanford.edu>.

1. Introduction

Large language models (LMs) have astonishing capabilities, including producing creative content (Stevenson et al., 2022), writing or completing source code (Chen et al., 2021), or chatting with users (Brown et al., 2020). A key ingredient in enabling models to perform such tasks is the data on which they are trained (Google, 2023; Gururangan et al., 2020; Touvron et al., 2023). A natural goal is to improve the data in order to unlock particular capabilities; however, it is not clear how to select data for these capabilities given a fixed budget of training tokens. Existing methods rely on heuristics for filtering and mixing together different datasets (Lee et al., 2022; Touvron et al., 2023). We lack a formal framework for capturing how data influences the model’s capabilities, which is critical in understanding existing LMs and training improved LMs.

To develop such a framework, we take inspiration from how humans acquire knowledge. A classic idea in the education literature is the concept of *skills* that form a learning hierarchy (White, 1973). For example, one study found that students learned mathematical and scientific skills most quickly when these skills were presented in a particular order (Gagne, 1962). We seek to understand the extent that similar orderings characterize LM training. Such orderings may provide a better understanding of LMs as well as a mechanism for data-efficient training. For instance, to train an LM to perform Spanish question generation, we wish to know if training first on Spanish grammar and English question generation helps.

We aim to use the idea of skill orderings to build a framework that can relate data to LM training and behavior (Figure 1). This requires addressing two challenges revolving around the connection between skills and data. First, in order to show that there exist sets of skills that the LM learns most efficiently in some particular order, *an operational definition of LM skill and skill ordering must be developed and validated on data*. In initial experiments, we investigated if semantic groupings of data, such as metadata attributes or embedding clusters, were sufficient to represent a skill and characterize how models learn. For instance, we partitioned the Alpaca dataset (Taori et al., 2023) by instruction type—a technique used to capture dataset diversity (Wang

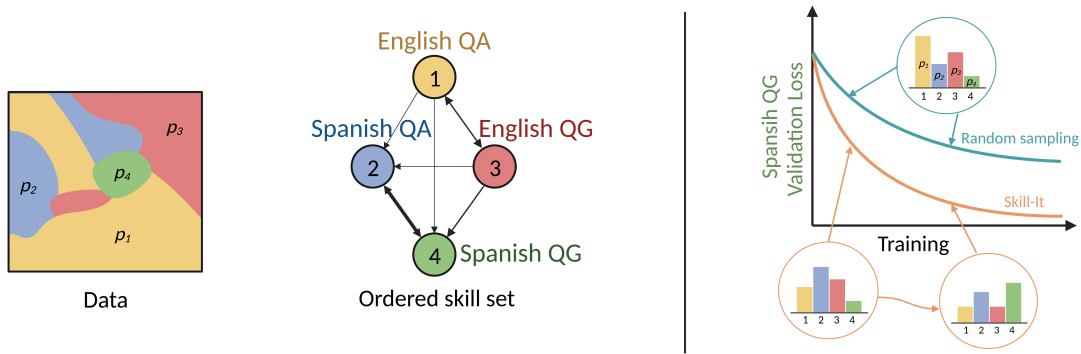


Figure 1. Inspired by how humans acquire knowledge, we hypothesize that LMs best learn skills in a particular order and that this can help improve our understanding and training of LMs. We show that these ordered skill sets exist in real data, which enables skills to be learned with less data given that we train on their prerequisite skills. We then propose SKILL-IT, an online data selection algorithm that learns skills quickly by exploiting their ordering.

et al., 2022a)—but we found that sampling based on instruction types and random sampling resulted in similar model performance, suggesting that not just any existing notion of data groups can characterize skills.

Second, once defined on data, *skills must be translated into means to improve data selection*. To develop criteria for a data selection algorithm that learns skills efficiently, we identify several challenges that naive selection approaches face. The standard approach of random uniform sampling over data fails to learn skills optimally for two reasons. First, skills can be distributed unevenly in the data, with more complex skills being rare—for instance, Spanish and question generation (QG) are 5% and 4% of the Natural Instructions dataset (Wang et al., 2022b), respectively, but Spanish QG is only 0.2%. Second, random sampling provides no mechanism for taking into account a particular training order and dependency structure on skills. More sophisticated techniques like curriculum learning account for ordering but not at the skill level. Our goal framework must account for these issues.

Skill-based framework We define a *skill* as a unit of behavior that a model can learn using an associated slice of data. An *ordered skill set* is a collection of skills with a directed *skills graph* that is neither complete nor empty (Figure 1), where an edge from a prerequisite skill to a skill exists if the amount of training it takes to learn the skill can be reduced if the prerequisite skill is also learned. We show that ordered skill sets exist in synthetic and real datasets using this operational definition. Interestingly, the existence of these ordered skill sets unveils that one can learn a skill quickly not by training solely on that skill, but on a mixture of that skill and prerequisite skills. For instance, in Figure 3 we observe that Spanish QG can be learned more efficiently when the model also learns English QG and Spanish—we can achieve 4% lower validation loss than training on only Spanish QG over a fixed budget of training steps.

Next, given an ordered skill set to train on, we use our framework to propose methods for how to select data so that the LM learn skills faster: skill-stratified sampling and an online generalization, SKILL-IT. We address the issue of unevenly distributed skills in datasets by proposing skill-stratified sampling, a simple approach that allows us to explicitly optimize for learning skills by uniformly sampling relevant skills (such as a target skill and its prerequisite skills in fine-tuning). Skill-stratified sampling uses the construction of the ordered skill set but is static, which does not incorporate the ordering as training proceeds and results in oversampling skills that may be already learned early on in training. We address this issue by proposing an online data selection algorithm, SKILL-IT, for selecting mixtures of training skills that allocates more weight towards learning skills that are not yet learned or towards influential skills. SKILL-IT is derived from an online optimization problem over the training skills for minimizing loss on a set of evaluation skills given a fixed budget of data and the skills graph. SKILL-IT is inspired by online mirror descent and can be adapted for continual pre-training, fine-tuning, or out-of-domain evaluation depending on the evaluation skill set.

We evaluate SKILL-IT on synthetic and real datasets at two model scales, 125M and 1.3B parameters. For the continual pre-training setting, we show on the LEGO synthetic (Zhang et al., 2022b) that we obtain a 35.8 point improvement in accuracy over randomly selecting training data and curriculum learning (Bengio et al., 2009). For the fine-tuning setting, we show that on the widely-used Natural Instructions dataset (Mishra et al., 2022; Wei et al., 2021), our algorithm over a mixture of skills is able to achieve up to 13.6% lower loss on that skill than solely training on that skill, given the same training budget. In the case that our training skills do not align perfectly with evaluation skills, our algorithm is able to achieve the lowest loss on 11 out of 12 evaluation skills corresponding to task categories in

the Natural Instructions test tasks dataset over random and skill-stratified sampling on the training data. We finally apply our framework to a case study on the recent RedPajama 1.2 trillion token dataset (Together, 2023). We use the data mixture produced by SKILL-IT to continually pre-train a 3B parameter model over 3B tokens. We find that SKILL-IT achieves higher accuracy with 1B tokens than uniform sampling over data sources with 3B tokens.

2. Related work

We present an abbreviated related work. A full treatment can be found in Appendix B. Existing work on data selection for LMs has generally ranged from more computationally expensive methods for dataset condensation on smaller datasets (Paul et al., 2021; Toneva et al., 2018; Phillips, 2016) to broader deduplication and filtering techniques for web-scale datasets (Abbas et al., 2023; Lee et al., 2022; Xie et al., 2023). Another way of improving model performance through choice of data is via curriculum learning (Bengio et al., 2009), which also draws inspiration from how humans learn and arranges data in order from easiest to hardest over samples or groups (Varshney et al., 2022). In contrast to existing works in both curriculum learning and data selection, our work focuses on selecting data for learning an ordered set of skills more efficiently. How LMs learn is also a topic of recent interest; there are hypotheses that models learn over quanta, discrete units of computation (Michaud et al., 2023), and that chain-of-thought reasoning emerges due to a structure of local clusters of variables that influence each other (Prystawski & Goodman, 2023). Lastly, the notion of skill has been studied in education, ranging from classical research on learning hierarchies (White & Gagné, 1974) to methods for decision-making over lesson sequences (Reddy et al., 2016).

3. Skills framework

First, we propose definitions of skills and ordered skill sets in order to formalize our intuition around how models learn skills, and we demonstrate that not just any existing notion of data groups can characterize an ordered skill set in the dataset. Then, we demonstrate the existence of ordered skill sets on synthetic and real data, which show how viewing data through a skills-based framework can help with training and understanding model performance. Finally, we explore unsupervised approaches for skill recovery from data, finding that embedding-based approaches do not adequately recover synthetic skills.

3.1. Definitions

We first present a definition of an individual skill. Let the input space of all possible text data be \mathcal{X} , where $x \in \mathcal{X}$ is an individual text sample that a next-token-prediction LM $f \in \mathcal{F} : \mathcal{X} \rightarrow \mathcal{X}$ is trained on. We quantify learning via

a metric $L : \mathcal{F} \times \mathcal{X} \rightarrow \mathbb{R}$, which maps from a model and evaluation data to a scalar quantity. In our setup, we use the cross-entropy validation loss applied over next-token predictions as our metric L .

Definition 3.1 (Skill). A *skill* s is a unit of behavior with associated data $\mathcal{X}_s \subseteq \mathcal{X}$ such that if f is trained on an dataset $\mathcal{D}_s \subset \mathcal{X}_s$, then f has improved metric L on samples belonging to $\mathcal{X}_s \setminus \mathcal{D}_s$ on average.

This definition of a skill is flexible—it simply means that given a training dataset associated with the skill, a model f has an improved metric when evaluated on validation data associated with this skill. Under this definition, a skill could be a granular task, such as writing elementary-level math word problems in Spanish, or can be defined over a data source, such as next-token prediction of legal data from tax court rulings. However, our next definition, the ordered skill set, has a more specific construction and provides a framework for how models learn across dependent skills.

Definition 3.2 (Ordered skill set, skills graph). An *ordered skill set* for f is a collection of skills $\mathcal{S} = \{s_1, \dots, s_k\}$ over which there is a directed *skills graph* $G = (\mathcal{S}, E)$ on the skill set that is neither complete or empty, where $(s_i, s_j) \in E$ if the amount of data needed to learn s_j when uniformly sampling from $\mathcal{D}_{s_i} \cup \mathcal{D}_{s_j}$ is no more than the amount of data needed when sampling only from \mathcal{D}_{s_j} . We equate learning a skill to f attaining a certain value of L .

This definition isolates complete and empty graphs as extrema that do not capture meaningful sets of skills. We discuss the three types of skill graphs—complete, empty, intermediate—and their implications for data selection. In particular, we present our motivating findings on constructing skills via various partitions of datasets, many of which did not help characterize how models learn.

- The complete graph demonstrates that all skills influence each other. A random partition is an example of a skill set that yields a complete graph. This graph suggests that the best approach for learning from data is random sampling on the dataset, and thus this is not a setting where we can gain much with skill-based sampling. For example, using instruction types as skills on the Alpaca dataset results in a nearly complete estimated skills graph (97.4% dense), and we find that stratified sampling on these skills only improves validation loss per skill by 0.007 points over random sampling (Figure 2 left).
- The empty graph demonstrates that each skill is independent. This can occur if skills are too granular; for instance, learning Spanish math problems is unlikely to help with English poem generation. This graph suggests that the best approach for learning each skill is to train on the skill itself. We see that empty graphs exist in real data; in Figure 2 (center), using data sources as skills on the

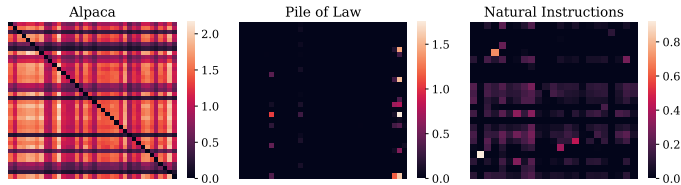


Figure 2. Heatmaps of adjacency matrices we compute for skill graphs for Alpaca, Pile of Law, and Natural Instructions (see Appendix D.2 for descriptions of how they were constructed and larger versions).

Pile of Law (Henderson* et al., 2022) results in a nearly empty skills graph (3.9% dense).

- Graphs that are neither empty nor complete thus suggest a nontrivial order of how skill influence each other. *This is the setting in which we expect that identifying skills and exploiting their ordering will help the most.* In Figure 2 right, we use task categories, which capture broader reasoning patterns, as skills on Natural Instructions and find that the estimated graph has intermediate density (42.7% dense). We show concrete examples of how skills can be learned more efficiently on this dataset in Section 3.2.

While intuitive notions of groupings result in ordered skill sets on some datasets (e.g., task categories on NI), this is not always the case (e.g., instruction types on Alpaca and sources on Pile of Law). Even though these attributes capture some notion of diversity in the dataset, our findings suggest that not just any semantic grouping induces an ordered skill set. We now empirically demonstrate that our definition of ordered skill sets align with how models learn and can be exploited for more data-efficient training.

3.2. Examples of skills and ordered skill sets

We provide examples of ordered skill sets on the LEGO synthetic dataset, an addition synthetic dataset, and the Natural Instructions dataset. On these datasets, we find that certain skills are better learned when trained along with their prerequisite skills rather than in isolation.

LEGO skills The LEGO synthetic, first introduced in Zhang et al. (2022b), can evaluate a model’s ability to follow a chain of reasoning. In this synthetic, the letters of the alphabet, \mathcal{A} , are variables each with some binary label in $\{0, 1\}$. An individual sample consists of k clauses for some fixed k across the dataset, each of the form $a = gx$ where $a, x \in \mathcal{A}$ and g is either a negation (“not”) or assertion (“val”), e.g. we assign a to the value of x , or we assign a to the opposite label. At the end of the sentence, we prompt the model for what the value of one of these variables is. Two samples $x \in \mathcal{X}$ are given below for $k = 5$:

Input: b = not y, r = val 1, m = val b, q = val m, y = not r. Output: b = 1.

Input: c = val x, p = val f, x = val k, f = not c, k =

val 0. Output: k = 0.

These samples each correspond to a chain of reasoning; for instance the first sample has the chain r, y, b, m, q , where knowing q ’s label requires the most reasoning steps. We define the i th skill s_i as the model’s ability to know the i th variable of the chain. From our example above, the first sample belongs to \mathcal{X}_{s_3} and the second sample belongs to \mathcal{X}_{s_1} . To demonstrate the existence of ordered skill sets, we continually pre-train the 125M parameter GPT-Neo model (Gao et al., 2020; Black et al., 2021) over various mixtures of LEGO skills with $k = 5$. In Figure 3 (left), we find that in 35.9% fewer training steps, training on a balanced mixture of $\mathcal{X}_{s_1}, \mathcal{X}_{s_2}$, and \mathcal{X}_{s_3} resulted in the same validation loss of 0.01 as training solely on \mathcal{X}_{s_3} . This suggests that s_1, s_2 helped unlock performance on s_3 and that there exist edges from s_1 or s_2 to s_3 in the skill graph. Additional observations are available in Appendix E.1, where we examine other edges as well as more complex reasoning chains.

Addition skills We consider a variant of a synthetic 5-digit addition dataset analyzed in Nanda et al. (2023). We show the existence of ordered skill sets for a simplified 3-digit addition dataset where we treat each digit prediction as a skill—the labels, in this case, are the integers $\{0, 1, \dots, 9\}$. Examples are of the following form:

Input: A = 1 0 6 + 0 7 1 , A 0 = ? Output: 7

Input: A = 6 0 6 + 8 7 9 , A 2 = ? Output: 4

where ‘A 0’ refers to the ones digit of the output (s_1) and ‘A 2’ refers to the hundreds digit (s_3). We find that in 32% fewer training steps, training on a balanced mixture of \mathcal{X}_{s_1} , and \mathcal{X}_{s_2} resulted in the same validation loss of 0.01 as training solely on \mathcal{X}_{s_3} . This is shown in Figure 3 (center). Once again, we find that performance on a skill can be improved via combined training with a prerequisite skill.

Natural Instructions (NI) skills We show that ordered skill sets exist in NI (Wang et al., 2022b). While we do not expect each text sample to pertain to one skill, we find that treating task categories as skills yields ordered skill sets:

- In Figure 3 (top right), we show that ordered skill sets exist over crosslingual task categories. Training on Spanish question generation (QG) along with equal parts of English QG, Spanish question answering (QA), and English

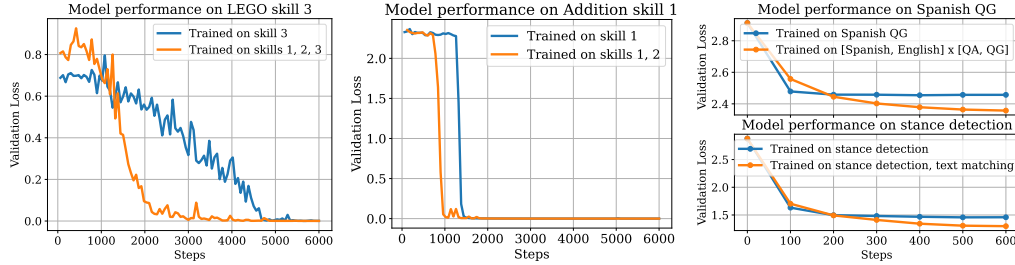


Figure 3. On the LEGO synthetic, 3-digit addition, and Natural Instructions, we identify examples of ordered skill sets in which training on a mixture of skills helps learn an individual skill faster than just training on that skill itself, given a fixed training budget.

QA results in 4.1% lower validation loss than training only on Spanish QG. Remarkably, the former only uses 25% of the latter’s Spanish QG data. This suggests that there are edges from Spanish QA, English QA, and English QG to Spanish QG.

- In Figure 3 (bottom right), we see that training on the task category Text Matching along with Stance Detection helps decrease the loss on Stance Detection by 11%. This suggests that these categories, which both involve understanding the relationship between two input texts, share an edge.

While using task category as an approximation of skill is imperfect, we see that there is signal within real data that suggests that ordered skill sets can improve data efficiency.

3.3. Skill recovery

A final component of characterizing skills is unsupervised recovery of ordered skill sets. We discuss embedding-based clustering approaches and propose a loss-based clustering mechanism that recovers LEGO skills. When clustering data using various trained and pre-trained embeddings, we find that they were unable to achieve above 39% accuracy on LEGO. Instead, we find that taking 10 random training runs and clustering data by their *loss* per timestep per run recovers the skills with 61% accuracy (Table 2). The intuition behind this method is that the validation losses on points from the same skill have similar trajectories. We discuss this approach more in Appendix E.2.

4. Skills-based data selection

We state the data selection problem for learning across skills in Section 4.1. We discuss how to learn the skills graph that will be exploited in our data selection methods in Section 4.2. We then introduce two sampling methods that utilize the graph, a simple skill-stratified sampling method and the online sampling method SKILL-IT, in Section 4.3.

4.1. Problem statement

We are given an ordered training skill set $\mathcal{S}_{\text{train}} = \{s_{\text{train},1}, \dots, s_{\text{train},k}\}$ on the training data, each with asso-

ciated support set $\mathcal{X}_{\text{train},1}, \dots, \mathcal{X}_{\text{train},k}$, and an ordered evaluation skill set $\mathcal{S}_{\text{eval}} = \{s_{\text{eval},1}, \dots, s_{\text{eval},m}\}$ of m evaluation skills. We aim to select n samples from $\mathcal{S}_{\text{train}}$ via a mixture of training skills, $p \in \Delta^{k-1}$, to achieve three goals depending on how $\mathcal{S}_{\text{eval}}$ is constructed:

- **Continual pre-training:** when $\mathcal{S}_{\text{eval}} = \mathcal{S}_{\text{train}}$, our goal is to select a mixture of training skills to learn all of them.
- **Fine-tuning:** when $\mathcal{S}_{\text{eval}} \subseteq \mathcal{S}_{\text{train}}$, our goal is to select a mixture of training skills to learn an individual target skill or subset of these skills.
- **Out-of-domain:** when $\mathcal{S}_{\text{eval}} \cap \mathcal{S}_{\text{train}} = \emptyset$, our goal is to select a mixture of training skills to learn a disjoint set of evaluation skills we cannot train on. This can arise when we have a separate downstream validation dataset or the skills identified in the training dataset are noisy.

Furthermore, we have a skills graph $G = (\mathcal{S}_{\text{train}} \cup \mathcal{S}_{\text{eval}}, E)$, where $E \subseteq \mathcal{S}_{\text{train}} \times \mathcal{S}_{\text{eval}}$ and $A \in \mathbb{R}^{k \times m}$ is a weighted adjacency submatrix, where A_{ij} describes the strength of the edge from $s_{\text{train},i}$ to $s_{\text{eval},j}$. We discuss how A can be estimated from the data.

4.2. Skills graph learning

The skills graph is important to utilize in sampling a mixture over the ordered skill set for training efficiently. We present two approaches for learning the skills graph—brute-force and linear approximation. Algorithms are provided in Appendix C.2. By definition 3.2, the brute-force way of identifying edges involves fixing a training budget of H steps and 1) training and evaluating the model on each s_i and 2) training the model on each pair of (s_i, s_j) and evaluating on s_i and s_j . If the loss on s_j when trained additionally on s_i is lower, there exists an edge from s_i to s_j . This approach has runtime $\mathcal{O}(Hk^2)$, which is feasible for small k . When k is large, we can approximate this approach in linear time by training on each s_i for $h < H$ steps and setting $A_{ij} > 0$ if the loss on s_j decreases for a runtime of $\mathcal{O}(hk)$. This linear approach is necessary in the out-of-domain setting when $\mathcal{S}_{\text{eval}}$ and $\mathcal{S}_{\text{train}}$ are disjoint, as we do not train on $\mathcal{S}_{\text{eval}}$. In addition, both graph learning approaches can be performed

Algorithm 1 SKILL-IT Online Data Selection Algorithm

- 1: **Input:** Ordered training skill set $\mathcal{S}_{\text{train}}$, ordered evaluation skill set $\mathcal{S}_{\text{eval}}$. Learning rate η , T rounds, n samples, H training steps per run for graph learning, model f_1 , window parameter w .
- 2: $A \leftarrow \text{LEARNGRAPH}(\mathcal{S}_{\text{train}}, \mathcal{S}_{\text{eval}}, H, f_1)$ (Alg. 2, 3).
- 3: Initialize $p_1^i = \exp(\eta \sum_{j=1}^m A_{ij})$ for all $i \in [k]$, the softmax over A .
- 4: **for** $t = 1, \dots, T - 1$ **do**
- 5: Observe losses $L_{\text{eval},j}(f_t)$ for all $s_{\text{eval},j} \in \mathcal{S}_{\text{eval}}$.
- 6: Train model f_t with n/T samples from mixture p_t over $\mathcal{S}_{\text{train}}$. Update model $f_{t+1} = \Phi(f_t, p_t)$.
- 7: Set $p_{t+1}^i = \exp(\eta \sum_{\tau=t-w+1}^t \sum_{j=1}^m A_{ij} L_{\text{eval},j}(f_\tau))$.
- 8: **end for**

on a smaller model, and the learned graph can be used for data selection for training a larger model (Appendix E.4).

4.3. Skills graph-aware sampling

We present two approaches for sampling over the mixture of training skills according to the skills graph: skill-stratified sampling, which samples uniformly over relevant training skills according to A , and SKILL-IT, which is an online generalization that incorporates knowledge of how skills are being learned throughout training.

4.3.1. SKILL-STRATIFIED SAMPLING

A straightforward sampling approach is to discard training skills that do not benefit the evaluation skills and sample uniformly over the set of relevant training skills, which we call *skill-stratified sampling*. For continual pre-training, the relevant skills are the entire training skill set; for each $s_{\text{train},i} \in \mathcal{S}_{\text{train}}$, $\Pr(s_{\text{train},i}) = \frac{1}{k}$. This enables each skill to have sufficient training data. For fine-tuning, the relevant skills are the target skills and prerequisite skills, which can be identified via positive entries of the i th column of A with $\mathcal{S}_{\text{prereq}} = \{s_{\text{train},i} : \exists s_{\text{eval},j} \text{ s.t. } A_{ij} > 0\}$. We then set $\Pr(s) = \frac{1}{|\mathcal{S}_{\text{prereq}} \cup \mathcal{S}_{\text{eval}}|}$ for $s \in \mathcal{S}_{\text{prereq}} \cup \mathcal{S}_{\text{eval}}$. For the out-of-domain setting, skill-stratified sampling is over the set of prerequisite skills. For each $s \in \mathcal{S}_{\text{prereq}}$, we set $\Pr(s) = \frac{1}{|\mathcal{S}_{\text{prereq}}|}$. Next, we propose our online algorithm that exploits the graph dynamically for more efficient training.

4.3.2. SKILL-IT ONLINE DATA SELECTION ALGORITHM

Despite accounting for prerequisite skills, one shortcoming of skill-stratified sampling is that even if a skill has already obtained sufficiently low validation loss early during training, we will continue to allocate the same weight to that skill throughout training. Therefore, we formulate our data selection problem as an online learning problem and propose SKILL-IT, which both prioritizes prerequisite skills

and skills that are not yet learned.

We are given a budget of T rounds and n total samples to train on. At round t , we select a mixture $p_t \in \Delta^{k-1}$ from the k -dimensional unit simplex, and for each training skill $s_{\text{train},i} \in \mathcal{S}_{\text{train}}$, we sample from $\mathcal{X}_{s_{\text{train},i}}$ with proportion p_t^i for a total of $\frac{n}{T}$ samples per round. Let f_t be the model at the start of round t . We can define f_t recursively as a function of the previous round’s model f_{t-1} and mixture p_{t-1} via a dynamics function $\Phi : \mathcal{F} \times \Delta^{k-1} \rightarrow \mathcal{F}$; that is, $f_t = \Phi(f_{t-1}, p_{t-1})$. Let $L_{\text{eval},j}(f_t)$ be the validation loss of f_t on $s_{\text{eval},j}$. Our goal is to select p_1, \dots, p_T to minimize loss per evaluation skill at the end of training:

$$\underset{p_1, \dots, p_T \in \Delta^{k-1}}{\text{minimize}} \quad \frac{1}{m} \sum_{j=1}^m L_{\text{eval},j}(f_T). \quad (1)$$

This optimization problem is challenging to solve without additional assumptions. In order to make the problem tractable, we impose an explicit dynamics rule for the each evaluation skill’s loss $L_{\text{eval},j}$ in terms of the current loss and data mixture. Assuming for simplicity that $\mathcal{S}_{\text{eval}} \subseteq \mathcal{S}_{\text{train}}$, a simple rule would be $L_{\text{eval},j}(f_t) = L_{\text{eval},j} \Phi(f_{t-1}, p_{t-1}) = L_{\text{eval},j}(f_{t-1})(1 - \alpha p_{t-1}^j)$ for $\alpha \in [0, 1]$. That is, we expect that allocating more data to skill j should result in the validation loss on skill j decreasing. However, such an expression assumes that only training on the j th skill will help learn the j th skill. Instead, Section 3.2 suggests that there are other skills that may help with the j th skill. We propose the following dynamics:

$$L_{\text{eval},j}(f_t) = L_{\text{eval},j}(f_{t-1})(1 - A_{:,j}^\top p_{t-1}), \quad (2)$$

where $A_{:,j}$ is the column with weights of all skills that influence $s_{\text{eval},j}$, and we absorb the scalar α into A . The optimization problem in (1) can thus be simplified as follows:

$$\underset{p_1, \dots, p_T \in \Delta^{k-1}}{\text{minimize}} \quad \frac{1}{m} \sum_{j=1}^m L_{\text{eval},j}(f_T) \quad (3)$$

$$\text{s.t. } f_t = \Phi(f_{t-1}, p_{t-1}) \quad \forall t = 1, \dots, T$$

$$L_{\text{eval},j}(f_t) = L_{\text{eval},j}(f_{t-1})(1 - A_{:,j}^\top p_{t-1}) \quad \forall j \in [m]$$

In Appendix C, we derive the following update rule via online mirror descent (Nemirovskij & Yudin, 1983) for learning rate $\eta > 0$:

$$p_{t+1}^i = p_t^i \exp \left(\eta \sum_{j=1}^m A_{ij} L_{\text{eval},j}(f_t) \right). \quad (4)$$

In addition, when equation 4 is expanded, we have that $p_{t+1}^i = p_1^i \exp \left(\eta \sum_{\tau=1}^t \sum_{j=1}^m A_{ij} L_{\text{eval},j}(f_\tau) \right)$. Since this summation over τ results in diminishing strength of updates,

we change it to a moving window of size w . Our full method is in Algorithm 1.

Intuitively, at each step we adjust the weight on skill i based on the losses of skills that i influences, with the assumption that more training data helps decrease loss. Note that when we use our algorithm with a complete graph or empty graph, we achieve expected behavior discussed in Section 3.1. For the complete graph, our algorithm reduces to stratified sampling. When we have a skill set with an empty graph, the update rule reduces to sampling proportional to each skill’s validation loss.

5. Experimental results

Given an ordered skill set, we aim to validate SKILL-IT’s ability to select data for efficiently learning skills in the continual pre-training, fine-tuning, and out-of-domain settings. We provide full tables of results in Appendix E.3.1 and results where we learn the skills graph on the 125M model and use it for the 1.3B parameter model in Appendix E.4. Skills graphs are in Appendix D.2, weight trajectories for SKILL-IT are in Appendix E.3.2, and ablations on the graph and online components of SKILL-IT are in Appendix E.5.

5.1. Continual pre-training

Baselines We compare SKILL-IT against three baselines that do not account for skills: random sampling, curriculum learning, and anticurriculum learning. Random sampling is a standard procedure for selecting samples given no additional information. Curriculum learning (Bengio et al., 2009) and anticurriculum learning (Wu et al., 2020) score the samples from easiest to hardest and vice versa, respectively, and sample over an expanding set of the lowest scored samples at every epoch; we use the pre-trained model’s loss to rank points. We evaluate skill-stratified sampling, which uses knowledge of the skills but is not online, and include an additional skills curriculum baseline in Appendix E.3.1

Setup We evaluate the ability of SKILL-IT to select data for efficiently learning over all skills. We measure average validation loss per skill after a fixed number of training steps. We construct the LEGO synthetic and addition synthetic with $k = 5$ and 3, respectively, and an imbalanced dataset over the skills. On the Natural Instructions dataset, we use 23 of the task categories as skills.

Analysis Our results are shown in Figure 4. Across our experiments we find that SKILL-IT outperforms baselines that do not use skills as well as skill-stratified sampling. On the LEGO dataset, all three baselines that do not utilize a notion of skills exhibit plateauing loss on four of the skills. Both skill-stratified sampling and SKILL-IT are able to significantly reduce loss on all skills, but the former is slower. Halfway through training, SKILL-IT exhibits an

accuracy improvement between 9.9 and 25.9 points over other approaches, reaching a final accuracy of 99.4 (Figure 18) SKILL-IT outperforms skill-stratified sampling by initially allocating more weight to prerequisite skills and eventually allocating more weight to skills that are learned more slowly (Figure 20). On the addition synthetic with $k = 3$, SKILL-IT converges to near-zero validation loss faster than the baselines on skills 1 and 2. While the random baseline may seem competitive at first glance, it fails to learn skill 1 (adding together the ones digits), which hurts its average loss per skill. On NI, the validation loss from SKILL-IT is 3.2% lower than from random sampling. Our results suggest that exploiting the construction and ordering of skills is critical to learning skills quickly.

5.2. Fine-tuning

Baselines We compare SKILL-IT against training on the target skill only and skill-stratified sampling over prerequisite skills and the target skill. The skill-stratified sampling approach uses the ordered skill set to identify prerequisite skills, but does not exploit them dynamically.

Setup We evaluate the ability of SKILL-IT to select data from an ordered training skill set for learning a target skill. Mirroring Figure 3, we evaluate on LEGO target skill 3 (third in reasoning chain), on the addition synthetic’s skill 1 (ones place digit addition), and on NI’s Spanish QG and Stance Detection.

Analysis Our results are shown in Figure 5. On LEGO, SKILL-IT results in the same validation loss of 0.01 as training only on the target skill in 38.1% fewer steps. We observe a similar trend on addition, with SKILL-IT converging to a validation loss of 0.01 in 59% fewer steps required to do so when training only on the target skill. Finally, on NI, SKILL-IT improves validation loss on Spanish question generation by 5.3% and Stance Detection by 13.6% over just training on the respective target skill only. A significant portion of the improvement over training only on the target skill comes from identification of prerequisite skills through the learned graph in the skill-stratified sampling method. SKILL-IT is further able to improve performance with finer-grained dynamic weighting on prerequisite skills.

5.3. Out-of-domain setting

Natural Instructions We compare SKILL-IT against random and skill-stratified sampling, both of which do not exploit the relationships between training skills and evaluation skills. We evaluate the ability of SKILL-IT to select data from a set of training skills for learning a disjoint set of evaluation skills that we cannot train on. We use all 59 task categories in the NI train tasks split as the training skills and the 12 task categories in the test tasks split as our evaluation skills. SKILL-IT achieves the lowest loss on 11 out of 12

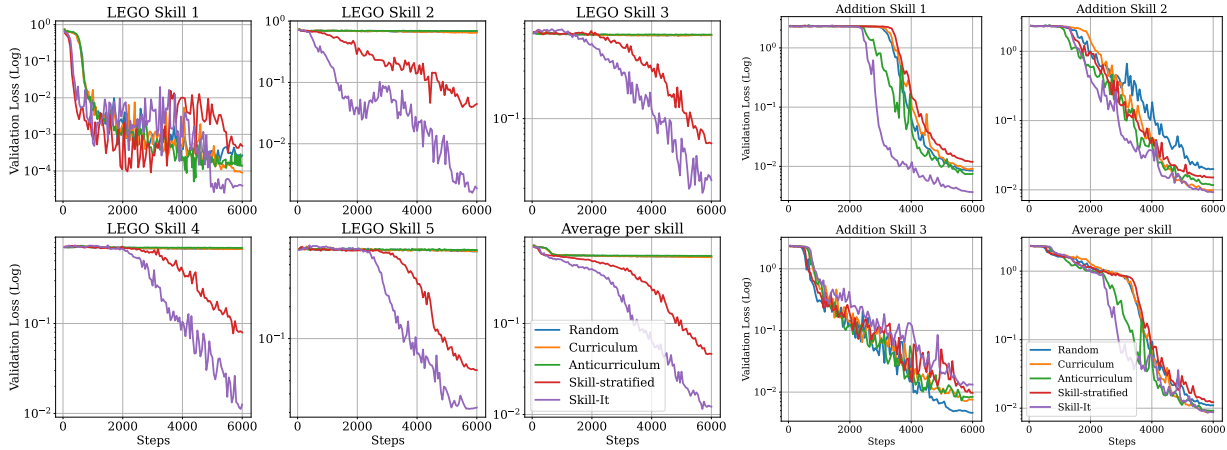


Figure 4. Performance of SKILL-IT on each skill in the continual pre-training setting (learning over all skills in the ordered training skill set) on the LEGO synthetic (left) and addition synthetic (right).

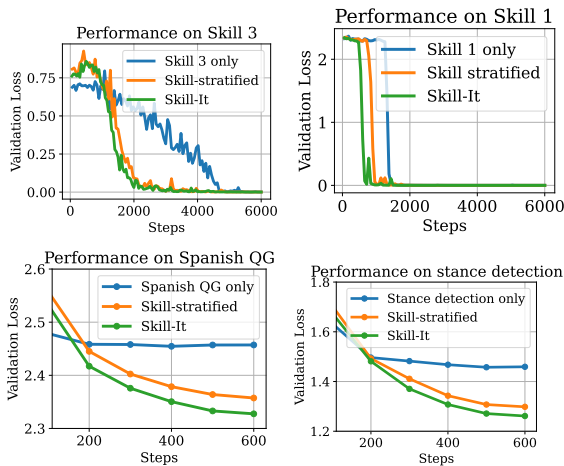


Figure 5. Performance of SKILL-IT in the fine-tuning setting (learning a target skill using the ordered training skill set) on LEGO, addition, and NL.

task categories over random and skill-stratified sampling (Figure 19, tables in Appendix).

RedPajama We use SKILL-IT to produce a data mixture on the RedPajama dataset. The training skills are the data sources comprising the dataset, and the evaluation skills are several tasks from the Language Model Evaluation Harness (Gao et al., 2021). SKILL-IT with $T = 1$ (i.e. a static, graph-based mixture) yields the mixture in Figure 6 (right). We continually pre-train a 3B parameter model trained on one trillion tokens for three billion additional tokens using this mixture, and see that it outperforms uniform sampling over the data sources (Figure 6 left). In particular, SKILL-IT achieves higher accuracy with 1B additional tokens than uniform with 3B additional tokens.

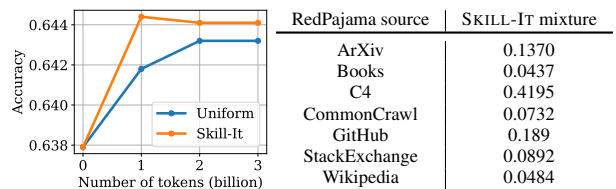


Figure 6. Left: Accuracy on LM Evaluation Harness for continual pre-training of a 3B parameter model using SKILL-IT on the RedPajama dataset. We achieve higher accuracy at 1B additional tokens than uniform at 3B tokens. Right: SKILL-IT mixture over RedPajama sources.

6. Conclusion

Given a fixed budget of data, knowing what data to train on to induce various capabilities in an LM is challenging. As LMs continue to improve, it will become increasingly important to extract as much signal as possible from the data and to direct that signal towards acquiring a broad variety of capabilities. In this paper, we introduce a skills-based framework for understanding how LMs learn and for selecting training data. We hope our study invites others to build on such a notion of skill and further explore how to align skills with data.

Acknowledgements

We thank Together Computer (<https://together.xyz/>) for providing portions of the compute used to train models in this paper. We thank Sabri Eyuboglu, Karan Goel, Arjun Desai, Neel Guha, Michael Zhang, Vishnu Sarrukai, Simran Arora, Ben Spector, Brandon Yang, Gautam Machiraju, and Sang Michael Xie for their helpful feedback and discussion.

We gratefully acknowledge the support of NIH under No. U54EB020405 (Mobilize), NSF under Nos. CCF1763315 (Beyond Sparsity), CCF1563078 (Volume to Velocity), and 1937301 (RTML); US DEVCOM ARL under No. W911NF-21-2-0251 (Interactive Human-AI Teaming); ONR under No. N000141712266 (Unifying Weak Supervision); ONR N00014-20-1-2480: Understanding and Applying Non-Euclidean Geometry in Machine Learning; N000142012275 (NEPTUNE); NXP, Xilinx, LETI-CEA, Intel, IBM, Microsoft, NEC, Toshiba, TSMC, ARM, Hitachi, BASF, Accenture, Ericsson, Qualcomm, Analog Devices, Google Cloud, Salesforce, Total, the HAI-GCP Cloud Credits for Research program, the Stanford Data Science Initiative (SDSI), and members of the Stanford DAWN project: Facebook, Google, and VMware. FS is supported by NSF CCF2106707 and the Wisconsin Alumni Research Foundation (WARF). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views, policies, or endorsements, either expressed or implied, of NIH, ONR, or the U.S. Government.

References

- Abbas, A., Tirumala, K., Simig, D., Ganguli, S., and Morcos, A. S. Semdedup: Data-efficient learning at web-scale through semantic deduplication, 2023.
- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., Das-Sarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., Joseph, N., Kadavath, S., Kernion, J., Conerly, T., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Hume, T., Johnston, S., Kravec, S., Lovitt, L., Nanda, N., Olsson, C., Amodei, D., Brown, T., Clark, J., McCandlish, S., Olah, C., Mann, B., and Kaplan, J. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.
- Biderman, S., Schoelkopf, H., Anthony, Q., Bradley, H., O’Brien, K., Hallahan, E., Khan, M. A., Purohit, S., Prashanth, U. S., Raff, E., Skowron, A., Sutawika, L., and van der Wal, O. Pythia: A suite for analyzing large language models across training and scaling, 2023.
- Black, S., Leo, G., Wang, P., Leahy, C., and Biderman, S. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow, March 2021. URL <https://doi.org/10.5281/zenodo.5297715>. If you use this software, please cite it using these metadata.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., Donahue, C., Doumbouya, M., Durmus, E., Ermon, S., Etchemendy, J., Ethayarajh, K., Fei-Fei, L., Finn, C., Gale, T., Gillespie, L., Goel, K., Goodman, N., Grossman, S., Guha, N., Hashimoto, T., Henderson, P., Hewitt, J., Ho, D. E., Hong, J., Hsu, K., Huang, J., Icard, T., Jain, S., Jurafsky, D., Kalluri, P., Karamcheti, S., Keeling, G., Khani, F., Khattab, O., Koh, P. W., Krass, M., Krishna, R., Kuditipudi, R., Kumar, A., Ladhak, F., Lee, M., Lee, T., Leskovec, J., Levent, I., Li, X. L., Li, X., Ma, T., Malik, A., Manning, C. D., Mirchandani, S., Mitchell, E., Munyikwa, Z., Nair, S., Narayan, A., Narayanan, D., Newman, B., Nie, A., Niebles, J. C., Nilforoshan, H., Nyarko, J., Ogut, G., Orr, L., Papadimitriou, I., Park, J. S., Piech, C., Portelance, E., Potts, C., Raghunathan, A., Reich, R., Ren, H., Rong, F., Roohani, Y., Ruiz, C., Ryan, J., Ré, C., Sadigh, D., Sagawa, S., Santhanam, K., Shih, A., Srinivasan, K., Tamkin, A., Taori, R., Thomas, A. W., Tramèr, F., Wang, R. E., Wang, W., Wu, B., Wu, J., Wu, Y., Xie, S. M., Yasunaga, M., You, J.,

- Zaharia, M., Zhang, M., Zhang, T., Zhang, X., Zhang, Y., Zheng, L., Zhou, K., and Liang, P. On the opportunities and risks of foundation models, 2021.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., de Oliveira Pinto, H. P., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F. P., Cummings, D., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W. H., Nichol, A., Paino, A., Tezak, N., Tang, J., Babuschkin, I., Balaji, S., Jain, S., Saunders, W., Hesse, C., Carr, A. N., Leike, J., Achiam, J., Misra, V., Morikawa, E., Radford, A., Knight, M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., and Zaremba, W. Evaluating large language models trained on code, 2021.
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J. E., Stoica, I., and Xing, E. P. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- Coleman, C., Yeh, C., Mussmann, S., Mirzasoleiman, B., Bailis, P., Liang, P., Leskovec, J., and Zaharia, M. Selection via proxy: Efficient data selection for deep learning, 2019.
- Gagne, R. M. The acquisition of knowledge. *Psychological review*, 69(4):355, 1962.
- Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- Gao, L., Tow, J., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., McDonell, K., Muennighoff, N., Phang, J., Reynolds, L., Tang, E., Thite, A., Wang, B., Wang, K., and Zou, A. A framework for few-shot language model evaluation, September 2021. URL <https://doi.org/10.5281/zenodo.5371628>.
- Geng, X., Gudibande, A., Liu, H., Wallace, E., Abbeel, P., Levine, S., and Song, D. Koala: A dialogue model for academic research. Blog post, April 2023. URL <https://bair.berkeley.edu/blog/2023/04/03/koala/>.
- Ghorbani, A. and Zou, J. Data shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning*, pp. 2242–2251. PMLR, 2019.
- Google. Palm2 technical report. Technical report, 2023. URL <https://ai.google/static/documents/palm2techreport.pdf>.
- Gupta, A. Advanced algorithms: Notes for cmu 15-850 (fall 2020), 2020.
- Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., and Smith, N. A. Don’t stop pretraining: Adapt language models to domains and tasks. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020. doi: 10.18653/v1/2020.acl-main.740. URL <http://dx.doi.org/10.18653/v1/2020.acl-main.740>.
- Guu, K., Webson, A., Pavlick, E., Dixon, L., Tenney, I., and Bolukbasi, T. Simfluence: Modeling the influence of individual training examples by simulating training runs, 2023.
- Henderson*, P., Krass*, M. S., Zheng, L., Guha, N., Manning, C. D., Jurafsky, D., and Ho, D. E. Pile of law: Learning responsible data filtering from the law and a 256gb open-source legal dataset, 2022. URL <https://arxiv.org/abs/2207.00220>.
- Hernandez, D., Brown, T., Conerly, T., DasSarma, N., Drain, D., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Henighan, T., Hume, T., Johnston, S., Mann, B., Olah, C., Olsson, C., Amodei, D., Joseph, N., Kaplan, J., and McCandlish, S. Scaling laws and interpretability of learning from repeated data, 2022.
- Ilyas, A., Park, S. M., Engstrom, L., Leclerc, G., and Madry, A. Datamodels: Predicting predictions from training data, 2022.
- Kirk, H. R., Jun, Y., Volpin, F., Iqbal, H., Benussi, E., Dreyer, F., Shtedritski, A., and Asano, Y. Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 2611–2624. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/1531beb762df4029513ebf9295e0d34f-Paper.pdf.
- Kitaev, N. and Klein, D. Constituency parsing with a self-attentive encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*

- (*Volume 1: Long Papers*), pp. 2676–2686, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1249. URL <https://www.aclweb.org/anthology/P18-1249>.
- Kitaev, N., Cao, S., and Klein, D. Multilingual constituency parsing with self-attention and pre-training. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3499–3505, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1340. URL <https://www.aclweb.org/anthology/P19-1340>.
- Koh, P. W. and Liang, P. Understanding black-box predictions via influence functions, 2017.
- Kung, P.-N., Yin, S.-S., Chen, Y.-C., Yang, T.-H., and Chen, Y.-N. Efficient multi-task auxiliary learning: Selecting auxiliary data by feature similarity. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 416–428, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.34. URL <https://aclanthology.org/2021.emnlp-main.34>.
- Langberg, M. and Schulman, L. J. *Universal approximators for integrals*, pp. 598–607. doi: 10.1137/1.9781611973075.50. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611973075.50>.
- Laurençon, H., Saulnier, L., Wang, T., Akiki, C., del Moral, A. V., Scao, T. L., Werra, L. V., Mou, C., Ponferrada, E. G., Nguyen, H., Frohberg, J., Šaško, M., Lhoest, Q., McMillan-Major, A., Dupont, G., Biderman, S., Rogers, A., allal, L. B., Toni, F. D., Pistilli, G., Nguyen, O., Nikpoor, S., Masoud, M., Colombo, P., de la Rosa, J., Villegas, P., Thrush, T., Longpre, S., Nagel, S., Weber, L., Muñoz, M., Zhu, J., Strien, D. V., Alyafeai, Z., Al-mubarak, K., Vu, M. C., Gonzalez-Dios, I., Soroa, A., Lo, K., Dey, M., Suarez, P. O., Gokaslan, A., Bose, S., Adelani, D., Phan, L., Tran, H., Yu, I., Pai, S., Chim, J., Lepercq, V., Ilic, S., Mitchell, M., Luccioni, S. A., and Jernite, Y. The bigscience roots corpus: A 1.6th composite multilingual dataset, 2023.
- Lee, K., Ippolito, D., Nystrom, A., Zhang, C., Eck, D., Callison-Burch, C., and Carlini, N. Deduplicating training data makes language models better. *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022. doi: 10.18653/v1/2022.acl-long.577. URL <http://dx.doi.org/10.18653/v1/2022.acl-long.577>.
- Lewis, D. D. A sequential algorithm for training text classifiers: Corrigendum and additional data. In *Acm Sigir Forum*, volume 29, pp. 13–19. ACM New York, NY, USA, 1995.
- Li, Y., Li, Y., and Risteski, A. How do transformers learn topic structure: Towards a mechanistic understanding, 2023.
- Liang, P. P., Wu, C., Morency, L.-P., and Salakhutdinov, R. Towards understanding and mitigating social biases in language models. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 6565–6576. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/liang21a.html>.
- Madaan, A., Zhou, S., Alon, U., Yang, Y., and Neubig, G. Language models of code are few-shot commonsense learners, 2022.
- Michaud, E. J., Liu, Z., Girit, U., and Tegmark, M. The quantization model of neural scaling, 2023.
- Mindermann, S., Razzak, M., Xu, W., Kirsch, A., Sharma, M., Morisot, A., Gomez, A. N., Farquhar, S., Brauner, J., and Gal, Y. Prioritized training on points that are learnable, worth learning, and not yet learned (workshop version), 2021.
- Mirzasoleiman, B., Bilmes, J., and Leskovec, J. Coresets for data-efficient training of machine learning models, 2019.
- Mishra, S., Khashabi, D., Baral, C., and Hajishirzi, H. Cross-task generalization via natural language crowdsourcing instructions. In *ACL*, 2022.
- Moore, R. C. and Lewis, W. Intelligent selection of language model training data. In *Proceedings of the ACL 2010 Conference Short Papers*, pp. 220–224, Uppsala, Sweden, July 2010. Association for Computational Linguistics. URL <https://aclanthology.org/P10-2041>.
- MosaicML. Introducing mpt-7b: A new standard for open-source, commercially usable llms, 2023. URL <https://www.mosaicml.com/blog/mpt-7b>.
- Nadeem, M., Bethke, A., and Reddy, S. Stereoset: Measuring stereotypical bias in pretrained language models. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021. doi: 10.18653/v1/2021.acl-long.416. URL <http://dx.doi.org/10.18653/v1/2021.acl-long.416>.
- Nanda, N., Chan, L., Lieberum, T., Smith, J., and Steinhardt, J. Progress measures for grokking via mechanistic

- interpretability. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=9XFSbDPmdW>.
- Nemirovskij, A. S. and Yudin, D. B. Problem complexity and method efficiency in optimization. 1983.
- Paul, M., Ganguli, S., and Dziugaite, G. K. Deep learning on a data diet: Finding important examples early in training, 2021.
- Phillips, J. M. Coresets and sketches, 2016.
- Prystawski, B. and Goodman, N. D. Why think step-by-step? reasoning emerges from the locality of experience, 2023.
- Reddy, S., Labutov, I., and Joachims, T. Latent skill embedding for personalized lesson sequence recommendation, 2016.
- Ruder, S. and Plank, B. Learning to select data for transfer learning with bayesian optimization. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017. doi: 10.18653/v1/d17-1038. URL <http://dx.doi.org/10.18653/v1/D17-1038>.
- Ruder, S., Ghaffari, P., and Breslin, J. G. Data selection strategies for multi-domain sentiment analysis, 2017.
- Saxena, S., Tuzel, O., and DeCoste, D. Data parameters: A new family of parameters for learning a differentiable curriculum. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/926ffc0ca56636b9e73c565cf994ea5a-Paper.pdf.
- Sorscher, B., Geirhos, R., Shekhar, S., Ganguli, S., and Morcos, A. S. Beyond neural scaling laws: beating power law scaling via data pruning, 2022.
- Soviany, P., Ionescu, R. T., Rota, P., and Sebe, N. Curriculum learning: A survey. *International Journal of Computer Vision*, 130(6):1526–1565, Apr 2022. ISSN 1573-1405. doi: 10.1007/s11263-022-01611-x. URL <http://dx.doi.org/10.1007/s11263-022-01611-x>.
- Stevenson, C., Smal, I., Baas, M., Grasman, R., and van der Maas, H. Putting gpt-3’s creativity to the (alternative uses) test, 2022.
- Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., and Hashimoto, T. B. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Together. Redpajama-data: An open source recipe to reproduce llama training dataset, 2023. URL <https://github.com/togethercomputer/RedPajama-Data>.
- Toneva, M., Sordoni, A., des Combes, R. T., Trischler, A., Bengio, Y., and Gordon, G. J. An empirical study of example forgetting during deep neural network learning, 2018.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. Llama: Open and efficient foundation language models, 2023.
- Varshney, N., Mishra, S., and Baral, C. Let the model decide its curriculum for multitask learning, 2022.
- Wang, Y., Kordi, Y., Mishra, S., Liu, A., Smith, N. A., Khashabi, D., and Hajishirzi, H. Self-instruct: Aligning language model with self generated instructions, 2022a.
- Wang, Y., Mishra, S., Alipoormolabashi, P., Kordi, Y., Mirzaei, A., Arunkumar, A., Ashok, A., Dhanasekaran, A. S., Naik, A., Stap, D., Pathak, E., Karamanolakis, G., Lai, H. G., Purohit, I., Mondal, I., Anderson, J., Kuznia, K., Doshi, K., Patel, M., Pal, K. K., Moradshahi, M., Parmar, M., Purohit, M., Varshney, N., Kaza, P. R., Verma, P., Puri, R. S., Karia, R., Sampat, S. K., Doshi, S., Mishra, S., Reddy, S., Patro, S., Dixit, T., Shen, X., Baral, C., Choi, Y., Smith, N. A., Hajishirzi, H., and Khashabi, D. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks, 2022b.
- Wei, J., Bosma, M., Zhao, V. Y., Guu, K., Yu, A. W., Lester, B., Du, N., Dai, A. M., and Le, Q. V. Finetuned language models are zero-shot learners, 2021.
- White, R. T. Research into learning hierarchies. *Review of Educational Research*, 43(3):361–375, 1973.
- White, R. T. and Gagné, R. M. Past and future research on learning hierarchies. *Educational Psychologist*, 11(1):19–28, 1974. doi: 10.1080/00461527409529119. URL <https://doi.org/10.1080/00461527409529119>.
- Wu, X., Dyer, E., and Neyshabur, B. When do curricula work?, 2020.

Xie, S. M., Santurkar, S., Ma, T., and Liang, P. Data selection for language models via importance resampling, 2023.

Zamir, A. R., Sax, A., Shen, W., Guibas, L., Malik, J., and Savarese, S. Taskonomy: Disentangling task transfer learning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun 2018. doi: 10.1109/cvpr.2018.00391. URL <http://dx.doi.org/10.1109/CVPR.2018.00391>.

Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X. V., Mihaylov, T., Ott, M., Shleifer, S., Shuster, K., Simig, D., Koura, P. S., Sridhar, A., Wang, T., and Zettlemoyer, L. Opt: Open pre-trained transformer language models, 2022a.

Zhang, Y., Backurs, A., Bubeck, S., Eldan, R., Gunasekar, S., and Wagner, T. Unveiling transformers with lego: a synthetic reasoning task, 2022b.

A. Broader Impacts and Limitations

Broader Impacts As more LMs are developed, a key criteria for their adoption and utility is if they exhibit a wide array of useful capabilities, such as generating harmless content, summarizing essays, and being conversational with the user. While improvements in other parts of the LM development pipeline such as training and architecture are important, many recent advances in building LMs with a wide array of useful capabilities have come from the data itself (Google, 2023; Taori et al., 2023; Chiang et al., 2023; Geng et al., 2023; MosaicML, 2023). Our work is fundamental in investigating how LMs learn and how to select data to learn skills more efficiently. However, we recognize that data selection methods can always be utilized to optimize for particular skills that may be considered malicious or negatively target or exclude specific groups (Bai et al., 2022). Furthermore, pre-trained LMs have been found to have various biases (Kirk et al., 2021; Nadeem et al., 2021; Liang et al., 2021; Bommasani et al., 2021).

Limitations The skills graph can either be provided (e.g., using a knowledge graph) or learned. Our work learns the skills graph using Algorithm 2 or Algorithm 3, which requires initial training runs on pairs of skills or each skill, respectively. This can be made more efficient by performing these training runs on a smaller model and for fewer number of steps, but tradeoffs here have yet to be thoroughly investigated. SKILL-IT also assumes that the ordered skill set is provided; as discussed in sections 3.1 and 3.3, it is challenging to recover ordered skill sets simply via metadata attributes or embedding clustering. Otherwise, the best way to sample over collections of skills that form a complete or empty graph is random or stratified sampling with no ordering to exploit. Our loss-based clustering approach presented in section 3.3 demonstrates that grouping by losses can provide an explanation for how skills are defined over data. An important direction for future work is to use such a clustering approach or other unsupervised algorithms in an end-to-end pipeline for skill discovery, skill graph learning, and data selection based on such skills.

Code release An anonymized repository can be found at <https://anonymous.4open.science/r/skillit-submission-7EDF/>.

B. Extended Related Work

We provide an extended related work below.

Data selection for LMs There have been several studies of large-scale data selection for LMs, Data deduplication (Lee et al., 2022; Abbas et al., 2023; Hernandez et al., 2022), in which identical or nearly identical samples are removed, is a method that enables LMs to be trained on smaller, cleaned datasets and has been increasingly used as a pre-processing step for training data (Touvron et al., 2023; Biderman et al., 2023; Zhang et al., 2022a). Other methods applied at scale involve ensuring high quality of data by explicitly filtering out samples or comparing the training dataset with a cleaned reference dataset (Brown et al., 2020; Touvron et al., 2023; Laurençon et al., 2023). Importance reweighting approaches have also been proposed for identifying training data from a large corpus that best approximates a smaller target distribution (Xie et al., 2023). These approaches can handle data selection for a particular target ability or topic, but do not provide much understanding of how the choice of data is related to the numerous skills that LMs learn.

Recent LMs have shifted focus from emphasizing the scale of the model to prioritizing the training data utilized. For example, models like Alpaca (Taori et al., 2023), Vicuna (Chiang et al., 2023), and Koala (Geng et al., 2023) are all based on the LLaMA model combined with instruction data generated by an existing LM. Palm 2’s technical report states that the data mixture was a critical component of the final model (Google, 2023), and Mosaic ML’s recent MPT model was trained on a hand-engineered mixture of the RedPajama dataset (MosaicML, 2023). However, these works lack rigorous explanation for why their datasets were constructed in this way.

Data selection methods Many data selection methods have been proposed for supervised, task-specific settings. In this setting, the most typical objective is dataset condensation, which aims to determine a small subset of data that captures the larger dataset’s properties with respect to the model. Some approaches include constructing coresets (Langberg & Schulman; Phillips, 2016), identifying samples that the model forgets during training (Toneva et al., 2018); identifying samples with the largest gradients (Paul et al., 2021) or gradients that approximate the overall gradient (Mirzasoleiman et al., 2019); clustering in embedding space and selecting points farthest from cluster centers (Sorscher et al., 2022); and selecting samples with the highest uncertainty or entropy (Lewis, 1995). These approaches have also been shown to transfer from smaller models to larger models (Coleman et al., 2019). Unlike these methods, we study how to select data for learning one or many skills at

the mixture level for LMs instead of the instance level.

Another area of interest is data selection for domain adaptation and multitask learning. For domain adaptation, there are a wide range of methods that select data to best match the target distribution. For example, the Moore-Lewis method matches data based on the difference in cross-entropy using a model trained on the target versus a model trained on the source data (Moore & Lewis, 2010). Several other approaches suggest training a model to distinguish between source and target and selecting points with high uncertainty (Ruder et al., 2017), or selecting points based on some divergence in an embedding space (Ruder & Plank, 2017). In comparison to these approaches, our work focuses on learning one or many skills and also finds that embedding-based heuristics do not fully identify skills.

Data attribution Another perspective on understanding training data is data attribution, which seeks to identify what data is responsible for particular model behaviors. Influence functions (Koh & Liang, 2017) and shapley values (Ghorbani & Zou, 2019) are two ways to quantify the role of individual samples. Datamodels (Ilyas et al., 2022) fit a model to predict behavior given a subset of training data, providing a framework for understanding individual samples as well as dataset counterfactuals. Simfluence (Guu et al., 2023) fits a Markov process to a set of training trajectories for finer-grained understanding of how data impacts training. We focus on understanding how groups of data associated with skills elicit broader model capabilities, and utilize this understanding to select data for more efficient training.

Curriculum learning Curriculum learning (Bengio et al., 2009) proposes to show the model data in order from easy samples to hard ones. Various criteria have been used to determine hardness, and anticurriculum as well as various pacing functions and mixing rates have been explored (Soviany et al., 2022). Curriculum learning can also be performed at the group level (Varshney et al., 2022). More sophisticated approaches include parametrizing each sample with a dynamic importance (Saxena et al., 2019), and also accounting for irrelevant and noisy data (Mindermann et al., 2021). Our approach similarly utilizes a curriculum, but it is defined over a skills graph and does not necessarily align with training on easiest to hardest skills.

How LMs learn Many different explanations for how LMs learn from data have been proposed. One hypothesis is that there exist discrete, universal building blocks of LM knowledge called quanta, and power law scaling emerges from a learning over a particular distribution of quanta in the right order (Michaud et al., 2023). Another is that chain of thought reasoning emerges due to local clusters of latent variables that influence each other, which can be validated by studying the LM’s ability to do conditional inference given intermediate variables (Prystawski & Goodman, 2023). Others have provided theoretical analysis of how transformers learn topics by studying co-occurrences of words in the training data (Li et al., 2023). Empirically, how models learn is still a mystery—for instance, models trained on code are found to perform fairly well at commonsense reasoning (Madaan et al., 2022). Our work initiates a study on how LMs learn various skills and how to exploit this for better data selection.

Task selection In multitask auxiliary learning, the goal is to train a model to perform well on one task by selecting the most beneficial auxiliary tasks to jointly train on. Using feature similarity has been proposed to select these tasks (Kung et al., 2021). In contrast, we find that feature similarity does not always capture dependent skills, and we focus on learning across one or many skills. Another line of work is Taskonomy (Zamir et al., 2018), in which a hypergraph over a set of tasks is learned and used to select tasks. The methods used to develop the taxonomy can be applied to further expand our graph learning (e.g., studying transitive and higher-order properties). However, their focus is on task selection in computer vision rather than data selection for LMs to learn skills.

C. Additional Algorithmic Details

C.1. Derivation of SKILL-IT Update Rule

First, we provide the derivation of our update rule from online mirror descent using the proximal point view (Gupta, 2020). We restate our optimization problem from (3):

$$\begin{aligned} & \underset{p_1, \dots, p_T \in \Delta^{k-1}}{\text{minimize}} \quad \frac{1}{m} \sum_{j=1}^m L_{\text{eval},j}(f_T) \\ & \text{s.t.} \quad L_{\text{eval},j}(f_t) = L_{\text{eval},j}(f_{t-1})(1 - \alpha A_{:,j}^\top p_{t-1}) \quad \forall j \in [m], t = 1, \dots, T \\ & \quad \quad f_t = \Phi(f_{t-1}, p_{t-1}) \quad \forall t = 1, \dots, T \end{aligned} \quad (5)$$

Let $\bar{L}_t(p) = \frac{1}{m} \sum_{i=j}^m L_{\text{eval},j}(f_{t+1}) = \frac{1}{m} \sum_{i=j}^m L_{\text{eval},j}(\Phi(f_t, p))$; that is, p is the mixture we must choose at time t and \bar{L}_t is the average loss per skill of the model after it is trained on p at round t . A greedy approximation of (5) is minimize $\bar{L}_t(p)$, given the model and mixtures at previous rounds. A linear approximation of $\bar{L}_t(p)$ is

$$\bar{L}_t(p) \approx \bar{L}_t(p_{t-1}) + \langle \nabla \bar{L}_{t-1}(p_{t-1}), p - p_{t-1} \rangle \quad (6)$$

Then, the problem of minimizing $\bar{L}_t(p)$ becomes

$$\underset{p \in \Delta^{k-1}}{\text{argmin}} \langle \eta \nabla \bar{L}_{t-1}(p_{t-1}), p \rangle \quad (7)$$

after we drop terms from (6) that do not depend on p . Note that the η is a constant and does not impact the solution. The optimal solution to this problem is selecting the p that has the most weight on the slice with the largest gradient. To improve stability and prevent overfitting, we introduce regularization via a Bregman divergence $D_h(p||p_{t-1}) = h(p) - h(p_{t-1}) - \langle \nabla h(p_{t-1}), p - p_{t-1} \rangle$. After dropping terms that do not contain p , our problem is now

$$\underset{p \in \Delta^{k-1}}{\text{argmin}} \langle \eta \nabla \bar{L}_{t-1}(p_{t-1}), p \rangle + h(p) - \langle \nabla h(p_{t-1}), p \rangle \quad (8)$$

Taking the gradient and setting it equal to 0 gives us

$$\eta \nabla \bar{L}_{t-1}(p_{t-1}) + \nabla h(p) - \nabla h(p_{t-1}) = 0 \quad (9)$$

Similar to in standard multiplicative weights, we set $h(p) = \sum_i p_i \ln p_i$ and $\nabla h(p) = [\ln p_i + 1]_i$. Then,

$$\begin{aligned} \ln p^i &= \ln p_{t-1}^i - \eta \nabla_i \bar{L}_{t-1}(p_{t-1}) \\ \Rightarrow p_{t+1}^i &= p_t^i \exp(-\eta \nabla_i \bar{L}_t(p_t)) \end{aligned} \quad (10)$$

where ∇_i is the i th element of the gradient. Now we wish to compute $\nabla_i \bar{L}_t(p_t) = \frac{1}{m} \sum_{j=1}^m \nabla_i [L_{\text{eval},j}(f_{t+1})] = \frac{1}{m} \sum_{j=1}^m \nabla_i [L_{\text{eval},j}(\Phi(f_t, p_t))]$. Recall the dynamics model for L_{eval} :

$$L_{\text{eval},j}(f_{t+1}) = L_{\text{eval},j}(f_t)(1 - A_{:,j}^\top p_t), \quad (11)$$

The gradient of this model with respect to each training skill s_i is

$$\begin{aligned} \nabla_i L_{\text{eval},j}(f_{t+1}) &= -A_{ij} L_{\text{eval},j}(f_t) \\ \Rightarrow \nabla_i \bar{L}_t(p_t) &= \frac{1}{m} \sum_{j=1}^m -A_{ij} L_{\text{eval},j}(f_t) \end{aligned} \quad (12)$$

Plugging this back into (10),

$$p_{t+1}^i = p_t^i \exp \left(\eta \sum_{j=1}^m A_{ij} L_{\text{eval},j}(f_t) \right), \quad (13)$$

where we can absorb the $\frac{1}{m}$ into η .

Algorithm 2 LEARNGRAPH (Brute-Force)

```

1: Input: Ordered skill set  $\mathcal{S} = \{s_1, \dots, s_k\}$ . Number of training steps  $H$ , base model  $f$ .
2: for  $j \in [k]$  do
3:   Train  $f$  on samples from  $\mathcal{X}_{s_j}$  for  $H$  steps and denote  $f_{H,j}$  to be the model after training.
4:   Observe change in loss,  $\delta_j^j = L_{\text{eval},j}(f) - L_{\text{eval},j}(f_{H,j})$ .
5: end for
6: for  $i, j \in [k]$  do
7:   Train  $f$  on samples from  $\mathcal{X}_{s_i} \cup \mathcal{X}_{s_j}$  for  $H$  steps and denote  $f_{H,i,j}$  to be the model after training.
8:   Observe change in loss,  $\delta_j^{i,j} = L_{\text{eval},j}(f) - L_{\text{eval},j}(f_{H,i,j})$ .
9:   if  $\delta_j^{i,j} > \delta_j^j$  then
10:    Draw edge  $s_i \rightarrow s_j$  and set  $A_{ij} > 0$ .
11:   end if
12: end for
13: Return Adjacency matrix  $A \in \mathbb{R}^{k \times k}$ 

```

Algorithm 3 LEARNGRAPH (Approximate)

```

1: Input: Ordered skill sets  $\mathcal{S}_{\text{train}}$  and  $\mathcal{S}_{\text{eval}}$ . Number of training steps  $H$ , base model  $f$ .
2: for  $i \in [k]$  do
3:   Train  $f$  on samples from  $\mathcal{X}_{s_{\text{train},i}}$  for  $H$  steps and denote  $f_{H,i}$  to be the model after training.
4:   for  $j \in [m]$  do
5:     Observe change in loss,  $\delta_j^i = L_{\text{eval},j}(f) - L_{\text{eval},j}(f_{H,i})$ .
6:     If  $\delta_j^i > 0$ , draw edge  $s_{\text{train},i} \rightarrow s_{\text{train},j}$  and set  $A_{ij} > 0$ .
7:   end for
8: end for
9: Return Bipartite adjacency submatrix  $A \in \mathbb{R}^{k \times m}$ 

```

C.2. Graph Learning Method

We provide algorithms for learning the graph over an ordered skill set. In Algorithm 2, we discuss the brute-force approach for learning the adjacency matrix. This approach only works when $\mathcal{S}_{\text{eval}} \subseteq \mathcal{S}_{\text{train}}$ (e.g. pre-training and fine-tuning cases), so we denote $\mathcal{S} = \mathcal{S}_{\text{train}}$ in the algorithm box. In Algorithm 3, we discuss the linear approach for learning the adjacency matrix. This approach works even in the out-of-domain case when $\mathcal{S}_{\text{eval}}$ and $\mathcal{S}_{\text{train}}$ are disjoint.

In both approaches, the exact value of A_{ij} can vary, but we can typically set it proportional to $\delta_j^{i,j} - \delta_j^j$ in the brute-force case or δ_j^i in the approximate case. The exact constructions and methods for learning each A in our experiments are in Appendix D.2.

D. Additional Experimental Details

D.1. Datasets

We present details about each dataset used, including information on the skills and the validation dataset. A summary is presented in Table 1.

- Alpaca dataset (Taori et al., 2023): the Alpaca dataset consists of 52K instruction examples that were generated from text-davinci-003. We applied the Berkeley Neural Parser (Kitaev et al., 2019; Kitaev & Klein, 2018) to each instruction, keeping 40777 samples it was able to parse successfully. If the sample began with a question, we annotated it with the skill “question”, and otherwise we annotated it with the verb identified from the parser. We grouped the data into a total, such as “list”, “edit”, “calculate”, “describe” and “identify”.
- Pile of Law (Henderson* et al., 2022): the Pile of Law dataset consists of various sources of legal and administrative data, ranging from tax rulings to the world’s constitutions. We evaluate on a subset of the Pile of Law validation dataset consisting of 13883 samples, where we selected $\max(645, \text{source size})$ samples per source. We truncated each sample to be 100K characters.

Table 1. We list each dataset used as well as its corresponding skill. We include the number of skills in the training dataset, as well as details on how the validation dataset is constructed.

Dataset	Skill	# skills	Validation data
Alpaca	Instruction type	38	50 samples per skill
Pile of Law	Legal data source	31	645 samples per skill
LEGO	Reasoning chain depth	5	100 samples per skill
Addition	Digit	3	100 samples per skill
NI (pre-training)	Task category	23	50 samples per task
NI (Spanish QG)	Task category \times language	4	100 samples per task
NI (stance detection)	Task category	2	50 samples per task
NI (out-of-domain)	Task category	59, 12	400 samples per task
RedPajama	Data source	7	LM eval harness

- LEGO (Zhang et al., 2022b): for the LEGO synthetic, we created a dataset of 192000 samples with $k = 5$. Our validation dataset consisted of 100 samples per skill.
- Addition: for the 3-digit addition synthetic, we created a dataset of 192000 samples and used a validation dataset of 100 samples per skill.
- Natural Instructions (Wang et al., 2022b; Mishra et al., 2022): the natural instructions dataset is a large collection of tasks and their definitions in natural language. For the pre-training setting, we used a set of 23 task categories that had the largest degree (in-degree + out-degree) in the learned skills graph, for a total of 1, 232, 437 samples and 425 tasks to select from. We evaluated on 50 samples per task.

For the fine-tuning setting with Spanish question generation, we select data over 4 skills (Spanish question generation, Spanish question answering, English question generation, English question answering) for a total of 513210 samples and 212 tasks to select from. We evaluated on 100 samples per task.

For the fine-tuning setting with stance detection, we select data over 2 skills (stance detection, text matching) for a total of 50990 samples and 19 tasks to select from. We evaluated on 50 samples per task.

For the out-of-domain setting, we select data over all 59 task categories for a total of 2, 417, 867 samples and 753 tasks to select from. The test split consisted of 12 task categories and 119 tasks, and we evaluated on $\min(400, \text{task size})$ samples per task.

- RedPajama (Together, 2023): the RedPajama dataset is a 1-trillion token dataset that aims to reproduce the LLaMA (Touvron et al., 2023) training dataset. We select over the 7 data sources and evaluate using the LM evaluation harness (Gao et al., 2021).

D.2. Graph Learning Details

We describe how the skills graph was learned on each dataset.

- Alpaca (Figure 7): we use Algorithm 3 and train for $K = 150$ steps per skill. Each edge $i \rightarrow j$ has a weight of δ_j^i , the difference in loss on skill j before and after training on i .
- Pile of Law (Figure 8): we use Algorithm 3 and train for $K = 150$ steps. Each edge $i \rightarrow j$ has a weight of δ_j^i , the difference in loss on skill j before and after training on i .
- LEGO (Figure 9): we use both Algorithm 2 and Algorithm 3 and train for $K = 6000$ steps each. Each edge $i \rightarrow j$ has a weight of 0.5 if the amount of data associated with skill j that is needed to reach 0.01 validation loss is less when training on (i, j) than on j (edges are set to 0 if 0.01 validation loss is not reached, even if loss is decreasing). Each edge $i \rightarrow j$ is also set to 0.5 if training on i decreases loss directly on j . We set each diagonal entry of A to be 1.
- Addition (Figure 10): we use Algorithm 2 and train for $K = 6000$ steps. Each edge $i \rightarrow j$ has a weight of 0.5 if the amount of data associated with skill j that is needed to reach 0.01 validation loss is less when training on (i, j) than on j (edges are set to 0 if 0.01 validation loss is not reached, even if loss is decreasing). We set each diagonal entry of A to be 1.
- Natural Instructions (Figure 11, 12, 13): we use Algorithm 3. For the pre-training setting, we train for $K = 600$ steps and assign each edge $i \rightarrow j$ a weight δ_j^i equal to the change in loss on j in the first 100 steps for all $i, j \in [k]$, including diagonal entries. For the fine-tuning setting, we train for $K = 600$ steps and assign each edge $i \rightarrow j$ a weight δ_j^i equal to

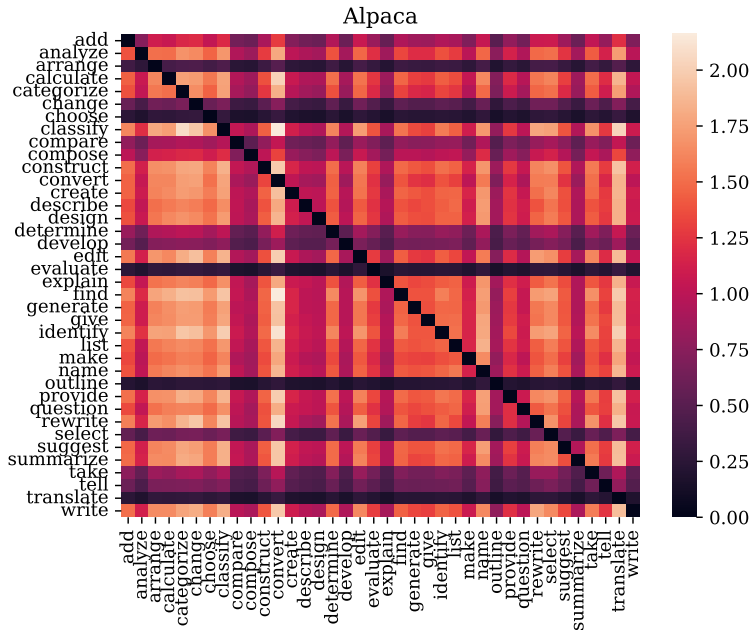


Figure 7. Alpaca heatmap where i, j th entry is $\max(0, \delta_j^i)$ (the change in loss on s_j after training on s_i for 150 steps). Diagonal entries are set to 0 for clearer visualization.

the change in loss before and after training. For the out-of-domain setting, we train for $K = 600$ steps and assign each edge $i \rightarrow j$ a weight δ_j^i equal to the change in loss before and after training in the first 100 steps.

- RedPajama (Figure 14): we use Algorithm 3 and train for 1-billion tokens per data source. We assign each edge $i \rightarrow j$ a weight δ_j^i equal to the change in perplexity on the validation data also before and after training.

D.3. Training Details

We describe the parameters used for SKILL-IT.

SKILL-IT pre-training

- LEGO: $\eta = 0.5, T = 6, w = 3$. We train for 6000 steps.
- Addition: $\eta = 0.1, T = 5, w = 3$. We train for 6000 steps.
- Natural Instructions (pre-training): $\eta = 0.2, T = 1$. We train for 5000 steps.

SKILL-IT fine-tuning

- LEGO: $\eta = 0.5, T = 10, w = 3$. We train for 6000 steps.
- Addition: $\eta = 0.1, T = 5, w = 3$. We train for 6000 steps.
- Natural Instructions (Spanish QG): $\eta = 0.8, T = 6, w = 3$. We train for 600 steps.
- Natural Instructions (stance detection): $\eta = 0.2, T = 6, w = 3$. We train for 600 steps.

SKILL-IT out-of-domain

- Natural Instructions: $\eta = 0.2, T = 10, w = 3$. We train for 5000 steps.
- RedPajama: $\eta = 100, T = 1$. We train for 3 billion tokens.

All results are computed over 5 random seeds.

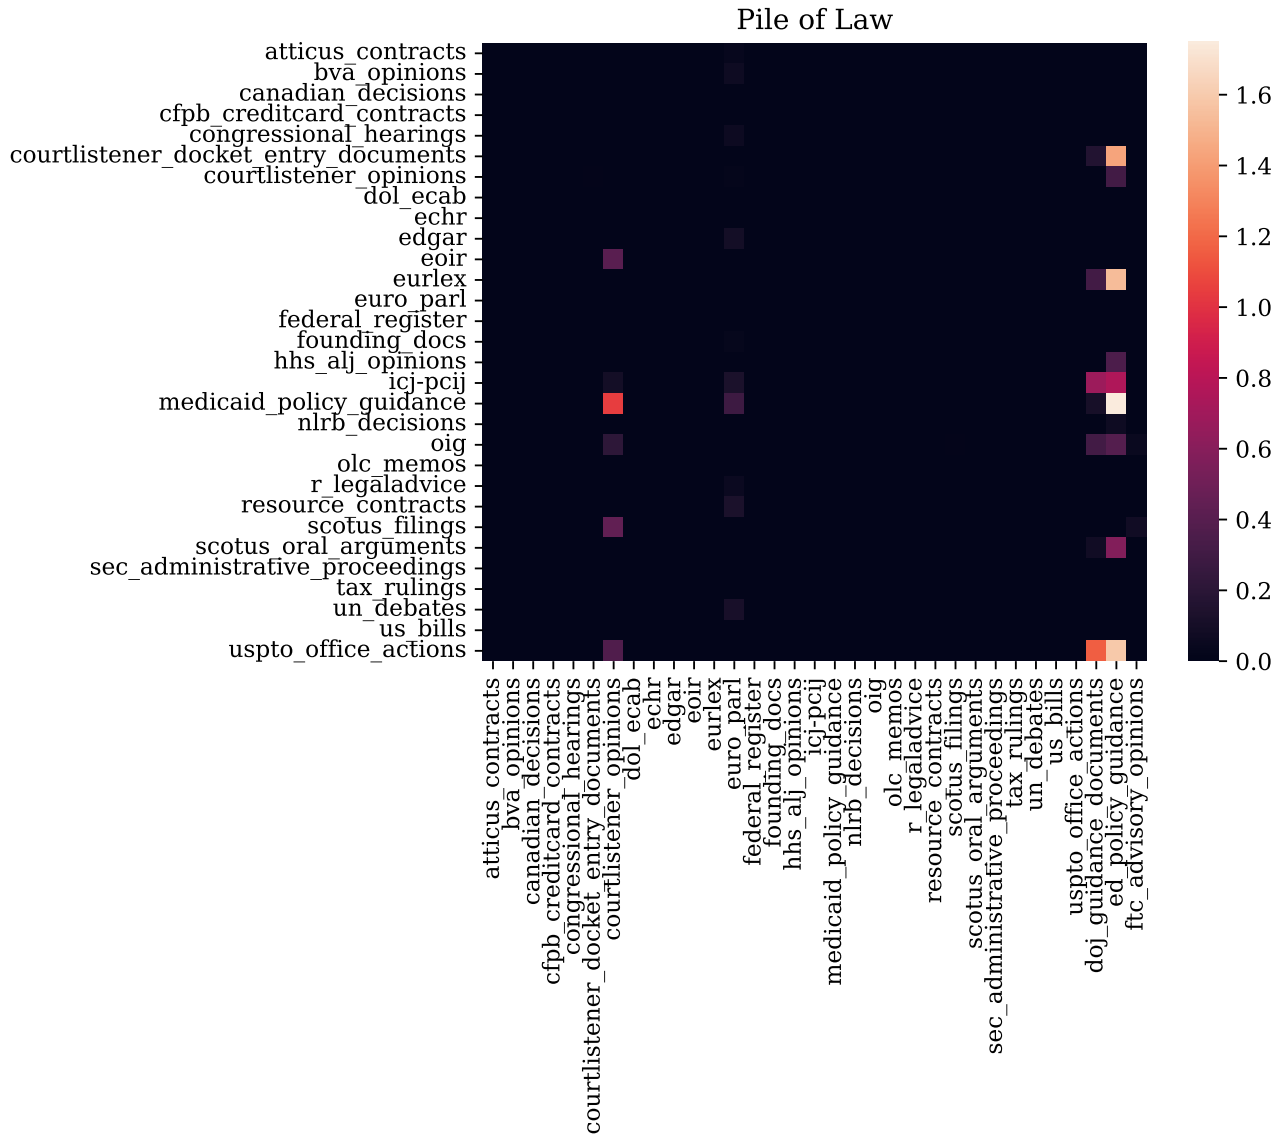


Figure 8. Pile of Law heatmap where i, j th entry is $\max(0, \delta_j^i)$ (the change in loss on s_j after training on s_i for 150 steps). Diagonal entries are set to 0 for clearer visualization.

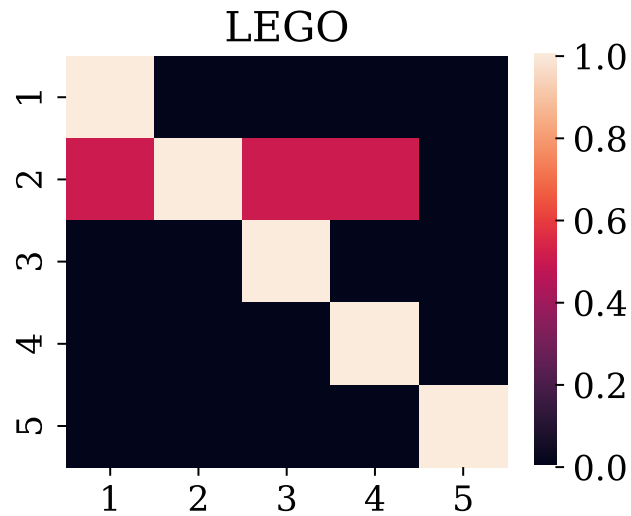


Figure 9. LEGO heatmap with $k = 5$ where i, j th entry is set to 0.5 if the number of steps needed to reach 0.01 loss on skill j when training on a balanced mixture of skills i and j is less than when training on skill j only.

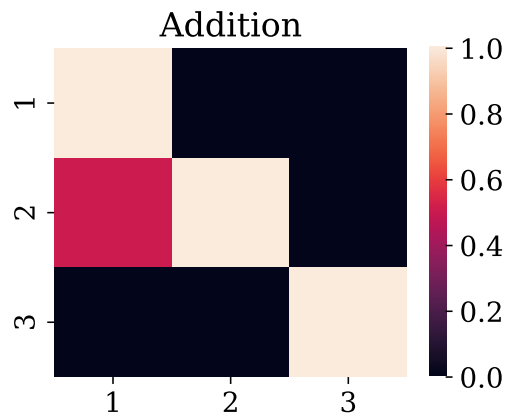


Figure 10. Addition heatmap with $k = 3$ where i, j th entry is set to 0.5 if the number of steps needed to reach 0.01 loss on skill j when training on a balanced mixture of skills i and j is less than when training on skill j only.

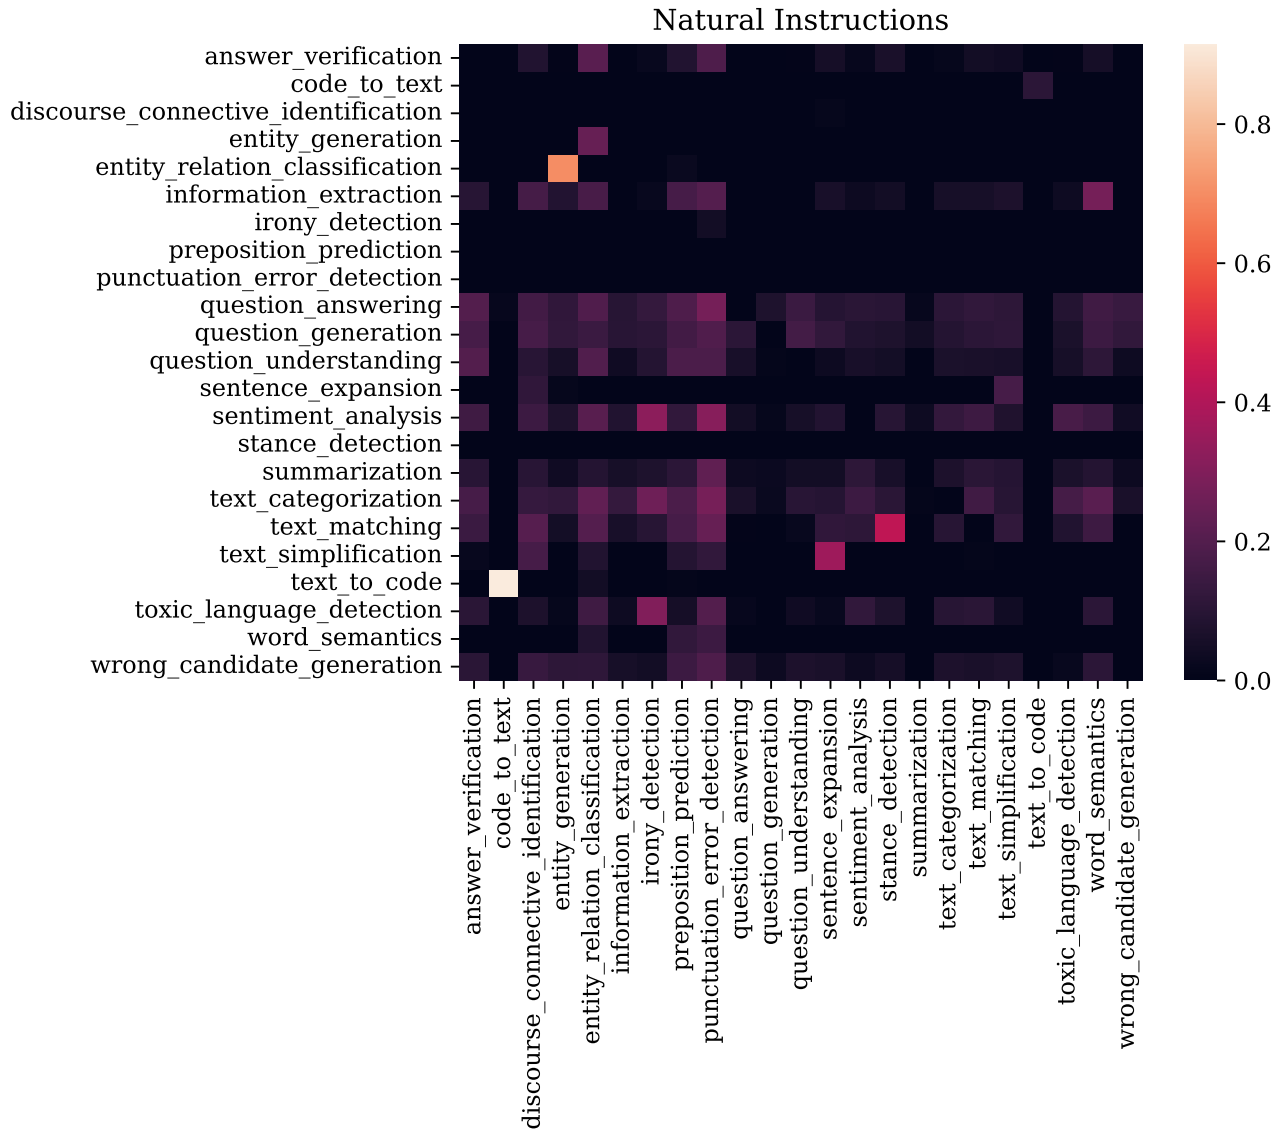


Figure 11. Natural Instructions heatmap where i, j th entry is $\max(0, \delta_j^i)$ (the change in loss on s_j after training on s_i for 100 steps). Diagonal entries are set to 0 for clearer visualization.

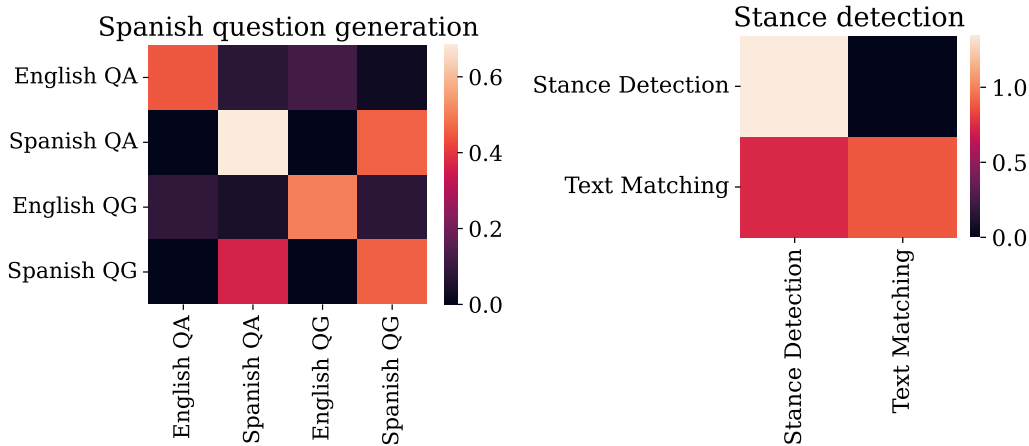


Figure 12. Spanish question generation and stance detection heatmaps where i, j th entry is $\max(0, \delta_j^i)$ (the change in loss on s_j after training on s_i for 100 steps).

Batch sizes of 32 and 64 were used for the LEGO and addition synthetic on the 125M and 1.3B parameter model, respectively. Batch sizes of 4 and 16 were used for the Natural Instructions experiments on the 125M and 1.3B parameter model.

For the out-of-domain Natural Instructions experiment and Alpaca graph learning experiments, a learning rate of $5e-6$ with linear scheduler and 50 warmup steps was used. For the Natural Instructions continual pre-training experiment on the 1.3B parameter model, a learning rate of $1e-6$ was used. All other experiments used a learning rate of $5e-5$. All experiments used AdamW with betas = 0.9, 0.999, eps = $1e-8$, and weight decay = 0.01. A context window of 512 was used.

Experiments with the Addition dataset were run using an Nvidia RTX A6000. Other experiments using the GPT-Neo 125M parameter model were run on an Nvidia Tesla P100. Experiments using the GPT-Neo 1.3B parameter model were run on an Nvidia Tesla A100.

E. Additional Experimental Results

E.1. Additional examples of LEGO ordered skill sets

For the LEGO synthetic, it may appear obvious that the skills graph is equivalent to the reasoning chain over the variables. However, in Figure 15 we see that this is not the case. Training on skills 2 and 4 together results in lower loss on skill 4 than when trained on skill 4 alone. However, training on skills 3 and 4 together results in roughly the same loss on skill 4 as when training on skill 4 alone, even though skill 3 and skill 4 share an edge in the LEGO synthetic’s underlying reasoning chain. This suggests that our intuition for how skills influence each other does not always match how the model learns skills.

Next, we consider a slightly more complex reasoning pattern on the LEGO synthetic. Instead of a chain, we construct a tree, where two variables in the LEGO synthetic are both defined in terms of the same parent variable. For example,

Input: $c = \text{val } 1, y = \text{not } w, v = \text{val } c, w = \text{not } c$. Output: $y = 1$.

In this example, $k = 4$ and both v and w are written in terms of c , and the reasoning graph has edges $1 \rightarrow 2, 1 \rightarrow 3, 2 \rightarrow 4$. In this case, we see that training on skill 2 or skill 3 both improve losses on skills 2 and 3 (Figure 16). However, unlike the previous figures, training on skills 2 and 4 or skills 3 and 4 do not significantly help reduce loss on skill 4 (Figure 17). Again, these measurements demonstrate that the reasoning graph does not necessarily equal the skills graph.

E.2. Unsupervised skill recovery

We explore several clustering techniques for recovering the skills in the LEGO synthetic on the validation dataset. Our results are shown in Table 2.

We first cluster based on the pre-trained model embeddings of the last token and the average token. We also report accuracies

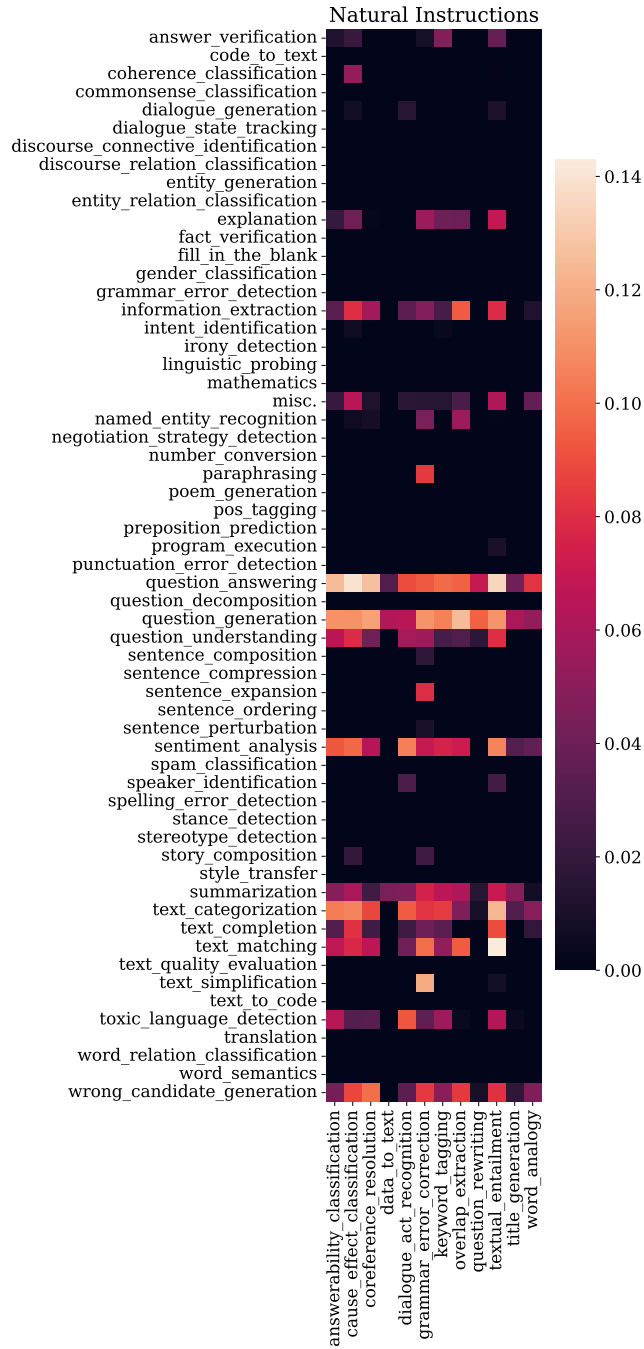


Figure 13. Natural Instructions heatmap for out-of-domain setting where rows are for the training skills and columns are for the evaluation skills. The i, j th entry is $\max(0, \delta_j^i)$ (the change in loss on s_j after training on s_i for 100 steps).

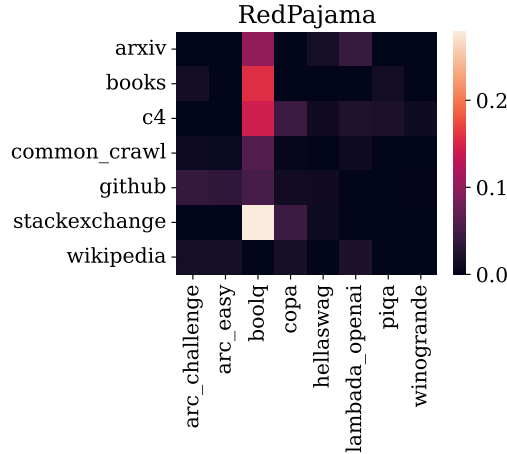


Figure 14. RedPajama heatmap for out-of-domain setting where rows are for the training skills and columns are for the evaluation skills. The i, j th entry is $\max(0, \delta_j^i)$ (the change in perplexity on s_j after training on s_i for 1B tokens).

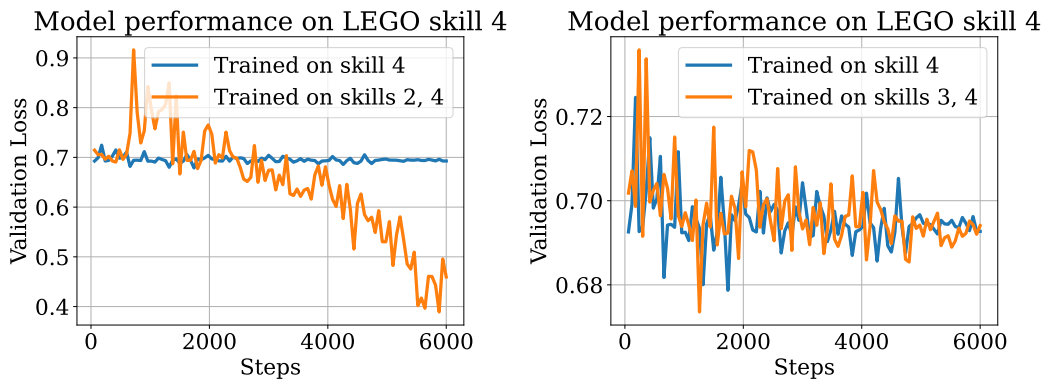


Figure 15. Performance on LEGO skill 4 when training on skill 4, skills 2 and 4, and skills 3 and 4. Even though skill 3 and skill 4 share an edge in the LEGO synthetic’s underlying reasoning chain (i.e. a model predicting correct for the fourth variable is one extra step beyond predicting correct for the third variable), we find that training on skills 2 and 4 helps improve performance on skill 4 more.

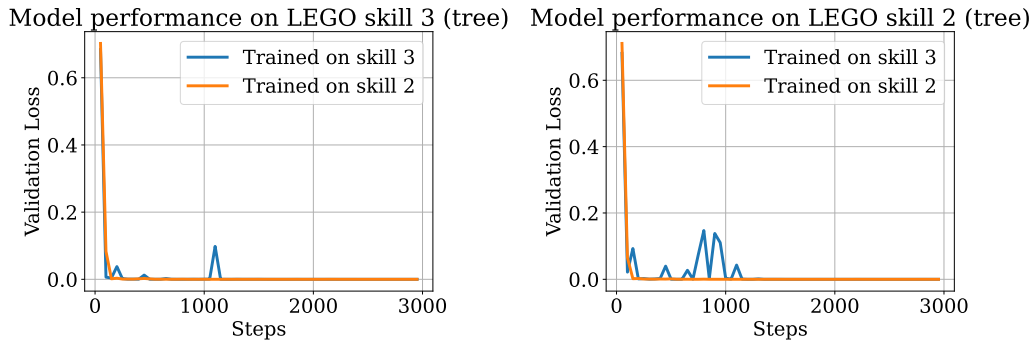


Figure 16. Performance on LEGO skill 2 and 3 when training on skills 2 and 3. The reasoning pattern is a tree rather than a chain over $k = 4$ variables. Skills 2 and 3 are at the same “depth” in the graph and both depend on skill 1, so there is positive influence between the skills despite there being no edge between 2 and 3 in the LEGO reasoning graph.



Figure 17. Performance on LEGO skill 4 when training on skills 2, 4 and skills 3, 4. We find that in both cases, the benefit from training on additional skills is minor. For instance, training on 2 and 4 reaches 0.01 loss in 2700 steps, while training on 4 only reaches it in 2100 steps.

Table 2. Clustering-based skill recovery methods on the LEGO dataset. The validation dataset we cluster consists of 500 points with $k = 5$, and results are reported over 10 runs of k-means.

Cluster method	Accuracy
Pretrained embedding of last token	24.8 ± 0.5
Pretrained embedding of average token	25.2 ± 1.1
Trained model embedding of last token	38.4 ± 0.8
Sentence-BERT embedding	23.9 ± 0.7
Losses over multiple runs	61.0 ± 1.6

of clustering based on the trained model embedding’s last token, where we train the model using random sampling for 6000 steps, and clustering based on Sentence-BERT embeddings. Among these four methods, using the trained model embeddings has the highest accuracy of 38.4 points.

Next, we cluster points based on losses. In particular, we do 10 runs, each for 6000 steps and with a randomly sampled mixture of skills. For each run, we evaluate the model on the validation dataset at 120 checkpoints. Then, each sample in the validation dataset has 1200 losses associated with it, comprising a feature vector for that sample. We perform k-means clustering on these features, which has an accuracy of 61.0 points, significantly higher than the second best accuracy of 38.4.

E.3. Full results for Section 5

E.3.1. PER-SKILL PERFORMANCE

We provide tables containing the per skill break-down of our results in Section 5.

Continual Pre-training In the continual pre-training setting, we report two additional baselines that combine curriculum learning with skills. Curriculum learning has been proposed for multitask learning (Varshney et al., 2022), in which groups of data are ranked by their average score and then trained in order of this ranking (with mixing of previously seen groups to avoid forgetting). We construct two baselines, Skill-curriculum and Skill-anticurriculum, using Algorithm 1 from (Varshney et al., 2022) with $\text{frac} = 0.4$.

We report loss per skill for the LEGO synthetic in Table 3, which corresponds to the results in Figure 4. We report accuracy per skill in Table 4 and Figure 18. For the random baseline in both, we use an imbalanced dataset with proportions 1:1:1:3:5.

We report the loss per skill for the Addition synthetic in Table 5, which also correspond to to the results in Figure 4. For the random baseline, we use an imbalanced dataset with randomly selected proportions: 13:14:18.

We report validation loss per task category for the Natural Instructions continual pre-training experiment in Table 6, where we find that SKILL-IT outperforms random sampling by 3.2% on average across skills.

Table 3. Results on validation loss per skill for LEGO pre-training experiment, averaged over 5 random seeds.

	Skill 1	Skill 2	Skill 3	Skill 4	Skill 5	Average
Random	0 ± 0.000	0.675 ± 0.041	0.688 ± 0.008	0.673 ± 0.049	0.667 ± 0.056	0.541 ± 0.031
Curriculum	0 ± 0.000	0.645 ± 0.052	0.686 ± 0.018	0.674 ± 0.042	0.671 ± 0.0459	0.535 ± 0.029
Anticurriculum	0 ± 0.000	0.690 ± 0.003	0.695 ± 0.004	0.693 ± 0.003	0.689 ± 0.004	0.554 ± 0.001
Skill-stratified	0 ± 0.000	0.045 ± 0.036	0.056 ± 0.029	0.079 ± 0.044	0.050 ± 0.025	0.046 ± 0.022
Skill-curriculum	0 ± 0.001	0.492 ± 0.282	0.476 ± 0.250	0.494 ± 0.196	0.484 ± 0.285	0.389 ± 0.154
Skill-anticurriculum	0 ± 0.000	0.373 ± 0.299	0.599 ± 0.099	0.660 ± 0.064	0.679 ± 0.031	0.462 ± 0.087
SKILL-IT	0 ± 0.000	0.002 ± 0.002	0.024 ± 0.031	0.013 ± 0.010	0.022 ± 0.021	0.012 ± 0.008

Table 4. Results on accuracy per skill (binary classification) for LEGO pre-training experiment, averaged over 5 random seeds.

	Skill 1	Skill 2	Skill 3	Skill 4	Skill 5	Average
Random	100.0 ± 0.0	54.2 ± 5.9	58.0 ± 3.1	48.0 ± 6.3	54.4 ± 7.3	62.9 ± 3.5
Curriculum	100.0 ± 0.0	60.0 ± 10.6	55.2 ± 5.8	51.2 ± 6.3	51.8 ± 6.1	63.6 ± 3.6
Anticurriculum	100.0 ± 0.0	53.4 ± 2.3	49.0 ± 4.8	48.2 ± 6.4	56.0 ± 5.7	61.3 ± 2.2
Skill-stratified	100.0 ± 0.0	98.2 ± 1.8	98.2 ± 1.3	97.8 ± 1.6	98.2 ± 1.3	98.5 ± 0.9
Skill-curriculum	100.0 ± 0.0	68.8 ± 19.3	67.4 ± 21.8	65.4 ± 20.8	66.2 ± 22.6	73.6 ± 14.8
Skill-anticurriculum	100.0 ± 0.0	78.4 ± 24.1	60.4 ± 11.5	50.0 ± 9.3	54.6 ± 5.0	68.7 ± 8.1
SKILL-IT	100.0 ± 0.0	99.2 ± 0.8	99.0 ± 1.0	99.4 ± 0.5	99.6 ± 0.5	99.4 ± 0.2

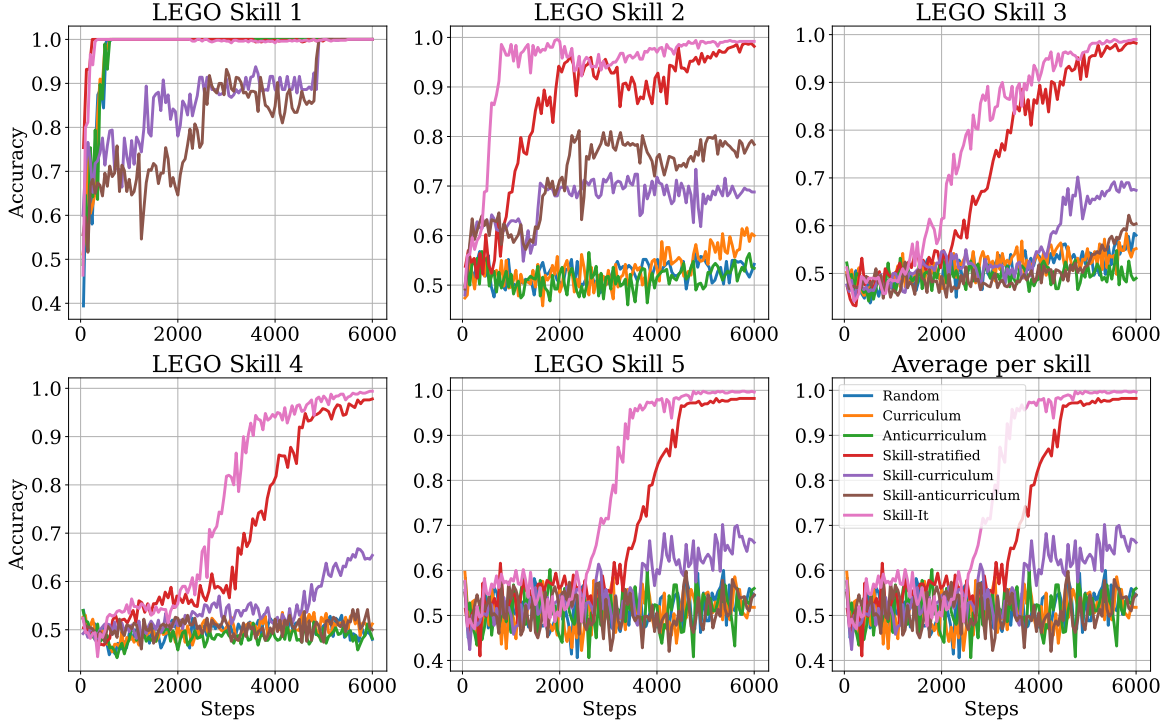


Figure 18. Accuracy of SKILL-IT on each skill (binary classification) on the LEGO synthetic in the continual pre-training setting. SKILL-IT attains higher accuracy more quickly than baselines that both do and do not utilize the notion of skills.

Table 5. Results on validation loss per skill for Addition pre-training experiment, averaged over 5 random seeds.

	Skill 1	Skill 2	Skill 3	Average
Random	0.008 ± 0.007	0.020 ± 0.019	0.005 ± 0.005	0.011 ± 0.014
Curriculum	0.009 ± 0.011	0.010 ± 0.008	0.008 ± 0.010	0.009 ± 0.010
Anticurriculum	0.007 ± 0.010	0.012 ± 0.013	0.008 ± 0.017	0.009 ± 0.014
Skill-stratified	0.012 ± 0.011	0.015 ± 0.015	0.010 ± 0.020	0.012 ± 0.016
SKILL-IT	0.004 ± 0.003	0.009 ± 0.007	0.013 ± 0.017	0.009 ± 0.011

Table 6. Validation loss per skill for data selection in continual pre-training setting on a subset of the Natural Instructions Dataset.

Skill	Random	Curriculum	Anticurriculum	Skill-stratified	Skill-curriculum	Skill-anticurriculum	SKILL-IT
Answer Verification	2.297 \pm 0.058	2.368 \pm 0.055	2.391 \pm 0.061	2.180 \pm 0.059	2.249 \pm 0.116	2.325 \pm 0.085	2.158 \pm 0.059
Code to Text	0.246 \pm 0.021	0.203 \pm 0.019	1.099 \pm 0.115	0.178 \pm 0.016	0.126 \pm 0.009	1.232 \pm 0.070	0.223 \pm 0.017
Discourse Connective Identification	2.927 \pm 0.069	3.084 \pm 0.067	2.932 \pm 0.058	2.805 \pm 0.071	2.891 \pm 0.001	2.925 \pm 0.011	2.784 \pm 0.068
Entity Generation	2.033 \pm 0.421	2.012 \pm 0.437	2.363 \pm 0.234	1.803 \pm 0.384	1.853 \pm 0.483	2.068 \pm 0.719	1.863 \pm 0.418
Entity Relation Classification	1.020 \pm 0.147	1.014 \pm 0.140	1.533 \pm 0.138	0.859 \pm 0.131	0.825 \pm 0.022	0.959 \pm 0.009	0.908 \pm 0.146
Information Extraction	2.154 \pm 0.040	2.247 \pm 0.037	2.352 \pm 0.042	2.140 \pm 0.037	2.286 \pm 0.022	2.338 \pm 0.025	2.073 \pm 0.042
Irony Detection	3.024 \pm 0.154	3.798 \pm 0.095	2.942 \pm 0.158	2.680 \pm 0.146	3.889 \pm 0.066	2.099 \pm 0.152	2.797 \pm 0.155
Preposition Prediction	0.979 \pm 0.124	0.887 \pm 0.147	1.488 \pm 0.213	0.845 \pm 0.152	0.941 \pm 0.019	1.044 \pm 0.029	0.876 \pm 0.173
Punctuation Error Detection	2.950 \pm 0.065	3.120 \pm 0.052	2.961 \pm 0.064	3.264 \pm 0.061	3.019 \pm 0.010	3.360 \pm 0.013	3.216 \pm 0.055
Question Answering	2.277 \pm 0.005	2.367 \pm 0.006	2.398 \pm 0.006	2.542 \pm 0.004	2.689 \pm 0.001	2.707 \pm 0.016	2.448 \pm 0.008
Question Generation	2.617 \pm 0.005	2.777 \pm 0.015	2.695 \pm 0.008	2.783 \pm 0.021	3.062 \pm 0.006	2.876 \pm 0.032	2.666 \pm 0.012
Question Understanding	1.965 \pm 0.051	2.199 \pm 0.059	2.060 \pm 0.033	1.958 \pm 0.051	2.385 \pm 0.022	2.100 \pm 0.054	1.895 \pm 0.043
Sentence Expansion	2.501 \pm 0.095	2.598 \pm 0.097	2.583 \pm 0.074	2.225 \pm 0.095	2.311 \pm 0.076	2.408 \pm 0.074	2.236 \pm 0.083
Sentiment Analysis	3.203 \pm 0.012	3.415 \pm 0.016	3.209 \pm 0.010	3.278 \pm 0.014	3.607 \pm 0.012	3.308 \pm 0.015	3.213 \pm 0.012
Stance Detection	1.810 \pm 0.100	1.775 \pm 0.120	2.231 \pm 0.128	1.385 \pm 0.070	1.361 \pm 0.114	1.823 \pm 0.189	1.556 \pm 0.125
Summarization	2.961 \pm 0.015	3.149 \pm 0.023	3.041 \pm 0.014	2.960 \pm 0.019	3.323 \pm 0.028	3.021 \pm 0.013	2.907 \pm 0.012
Text Categorization	2.488 \pm 0.023	2.692 \pm 0.029	2.553 \pm 0.006	2.570 \pm 0.015	3.001 \pm 0.007	2.635 \pm 0.014	2.448 \pm 0.017
Text Matching	2.177 \pm 0.059	2.232 \pm 0.055	2.316 \pm 0.048	2.152 \pm 0.061	2.324 \pm 0.004	2.304 \pm 0.035	2.093 \pm 0.054
Text Simplification	2.155 \pm 0.023	2.193 \pm 0.039	2.325 \pm 0.033	1.926 \pm 0.026	2.037 \pm 0.005	2.156 \pm 0.011	1.952 \pm 0.026
Text to Code	0.560 \pm 0.037	0.495 \pm 0.036	1.215 \pm 0.052	0.490 \pm 0.029	0.433 \pm 0.014	1.455 \pm 0.086	0.553 \pm 0.042
Toxic Language Detection	3.106 \pm 0.027	3.496 \pm 0.017	3.058 \pm 0.029	3.199 \pm 0.024	3.758 \pm 0.025	3.155 \pm 0.050	3.129 \pm 0.020
Word Semantics	2.092 \pm 0.027	2.334 \pm 0.034	2.156 \pm 0.064	1.916 \pm 0.043	1.784 \pm 0.048	2.424 \pm 0.038	1.952 \pm 0.019
Wrong Candidate Generation	2.438 \pm 0.021	2.606 \pm 0.039	2.519 \pm 0.027	2.506 \pm 0.026	2.849 \pm 0.029	2.574 \pm 0.018	2.432 \pm 0.025
Average	2.173 \pm 0.028	2.307 \pm 0.025	2.366 \pm 0.026	2.115 \pm 0.027	2.304 \pm 0.031	2.317 \pm 0.052	2.103 \pm 0.032

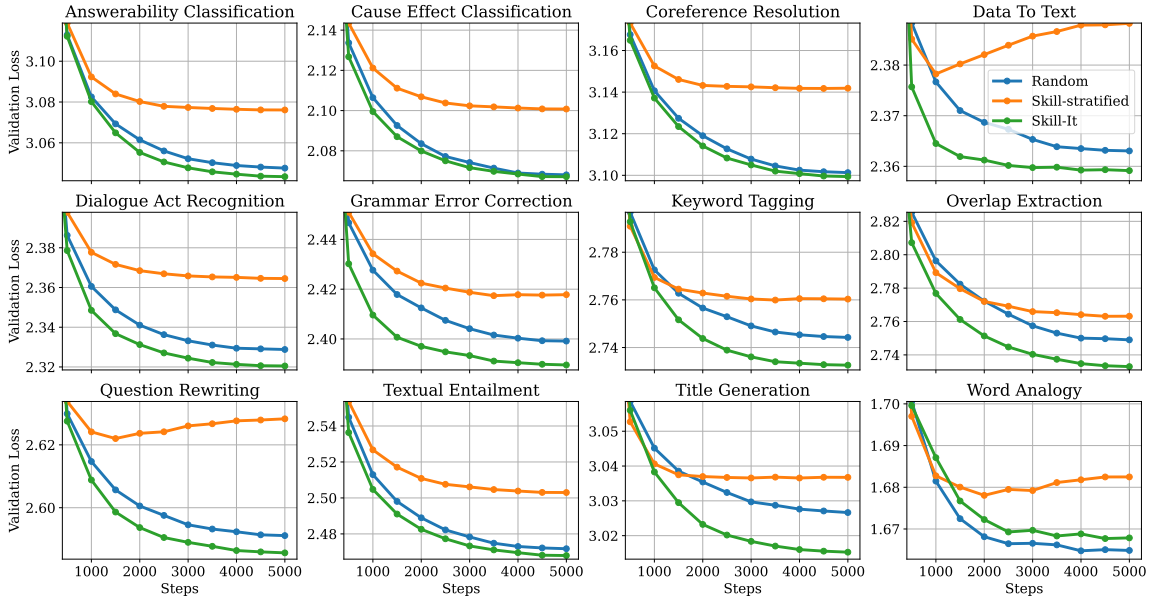


Figure 19. Performance of SKILL-IT in the out-of-domain setting for the NI test task split. SKILL-IT uses the graph between the train and evaluation skills to produce an online mixture on the training dataset.

Out-of-domain In Figure 19, we show the validation loss curves for each evaluation skill when using random sampling on the training data, skill-stratified sampling over prerequisite skills (e.g., the nonzero rows in Figure 13), and SKILL-IT. In Table 7, we provide a breakdown of validation loss per evaluation skill under these three sampling methods.

Table 7. Validation loss per skill for data selection in out-of-domain setting over Natural Instructions train task split and test task split.

Skill	Random	Skill-stratified	SKILL-IT
Answerability Classification	3.048 ± 0.003	3.076 ± 0.002	3.043 ± 0.003
Cause Effect Classification	2.068 ± 0.004	2.101 ± 0.005	2.067 ± 0.006
Coreference Resolution	3.101 ± 0.003	3.142 ± 0.004	3.099 ± 0.004
Data to Text	2.363 ± 0.004	2.388 ± 0.005	2.359 ± 0.005
Dialogue Act Recognition	2.329 ± 0.009	2.364 ± 0.010	2.320 ± 0.009
Grammar Error Correction	2.399 ± 0.008	2.418 ± 0.009	2.389 ± 0.007
Keyword Tagging	2.744 ± 0.005	2.760 ± 0.007	2.733 ± 0.006
Overlap Extraction	2.749 ± 0.011	2.763 ± 0.012	2.733 ± 0.010
Question Rewriting	2.591 ± 0.009	2.628 ± 0.011	2.586 ± 0.010
Textual Entailment	2.472 ± 0.002	2.503 ± 0.003	2.468 ± 0.002
Title Generation	3.027 ± 0.002	3.037 ± 0.002	3.015 ± 0.002
Word Analogy	1.665 ± 0.016	1.682 ± 0.015	1.668 ± 0.016
Average	2.546 ± 0.003	2.572 ± 0.003	2.540 ± 0.003

In Table 8 we provide a breakdown of the RedPajama experiment’s accuracy per evaluation skill, corresponding to the results in Figure 6.

E.3.2. WEIGHT TRAJECTORIES

We provide SKILL-IT’s weight trajectories for each result. The weight per skill across training steps for the LEGO pre-training experiment corresponding to Figure 4 (left) is shown in Figure 20. We see that SKILL-IT initially allocates more weight to skill 2 and less to 1, 3, 4, 5. Since skill 1 is learned quickly, the weight on skill 1 immediately drops to below 0.1 at 1000 steps. The weight on skills 3, 4, and 5 increase from around 0 to 3000 steps, during which their respective validation

Table 8. Performance of model trained on RedPajama with uniform sampling and SKILL-IT on LM evaluation harness. Unless otherwise noted, accuracy is reported for each task.

	1 Billion Tokens		2 Billion Tokens		3 Billion Tokens	
	Uniform	SKILL-IT	Uniform	SKILL-IT	Uniform	SKILL-IT
ARC Challenge (acc norm)	35.4	34.6	35.3	34.9	34.6	34.8
ARC Easy (acc norm)	62.2	61.2	62.4	61.7	62.5	62.0
BoolQ	68.9	68.2	67.7	68.6	67.2	68.7
COPA	81.0	82.0	80.0	81.0	81.0	81.0
HellaSwag (acc norm)	63.9	63.7	63.8	63.9	64.0	63.9
LAMBADA OpenAI	64.4	67.0	65.9	66.7	66.8	66.0
PIQA (acc norm)	74.8	75.0	75.5	75.2	75.0	75.7
Winogrande	62.8	63.9	63.9	63.2	63.4	63.1
Average accuracy	64.2	64.4	64.3	64.4	64.3	64.4

losses are higher than those of skills 1 and 2. Near the end of training, all losses are converging to 0, and so the weight per skill is roughly uniform.

The weight per skill across training steps for the addition pre-training experiment corresponding to Figure 4 (right) is shown in Figure 21. SKILL-IT allocates more weight to skill 2, which has an edge to skill 1 as shown in Figure 10. It also allocates very little weight to skill 3, which is learned faster than the other two skills. Eventually, it puts more weight on skill 1, the hardest skill, and then converges to uniform sampling as all validation losses approach 0.

The weight per skill across training steps for the LEGO fine-tuning experiment and the Spanish question generation and stance detection experiments corresponding to Figure 5 is shown in Figure 22. Since there is only one target skill in these experiments, the mixture of weights approaches uniform as the loss on the target skill approaches 0. It is interesting to explore how to reduce edge weights and regularization so that the mixture approaches the target skill instead, although preliminary experiments where we decayed the edge weight and the strength of the Bregman divergence term did not appear better. We hypothesize that since training on a uniform mixture (as in Figure 3) did strictly better than training on the target skill and their loss curves did not intersect during the training run, it is better to allocate non-negligible weight on all skills throughout the training run.

The weight per skill across training steps for the Natural Instructions out-of-domain experiment corresponding to Figure 19 is shown in Figure 23, where the legend is provided for the top 10 task categories with the largest weights. While the initial weights based on the skills graph roughly establishes the order of weight magnitude, the differences among the losses on the evaluation skills increases the range of weights as training continues. As validation losses saturate, the weights also converge to fixed values.

E.4. Experiments on 1.3B parameter model

We demonstrate that the skills graph learned on the 125M parameter model can be used for data selection when using the 1.3B parameter model. We present results in the continual pre-training setting on the LEGO synthetic and Natural Instructions.

All results are reported over 3 random seeds. For the LEGO experiment, we train for 1500 steps with $\eta = 0.5$, $T = 30$, $w = 3$. For the NI experiment, we train for 5000 steps with $\eta = 0.2$, and $T = 1$. The skill graphs were learned using the 125M parameter model as described in section D.2.

In Figure 24, we train the 1.3B model using SKILL-IT for the LEGO synthetic and find that it still outperforms random and skill-stratified sampling on average. In particular, while performance across sampling methods is similar for early skills, the discrepancy is larger for skill 5, for which SKILL-IT allocates more weight to dynamically. In Figure 25, we provide the weight trajectories of SKILL-IT. We observe that the weight trajectories are similar to that on the 125M parameter model, where initial weight is allocated towards skill 2. Later on, more weight is allocated towards skills 4 and 5, whose losses are higher, and eventually the weight mixture converges to uniform as all losses converge to near 0.

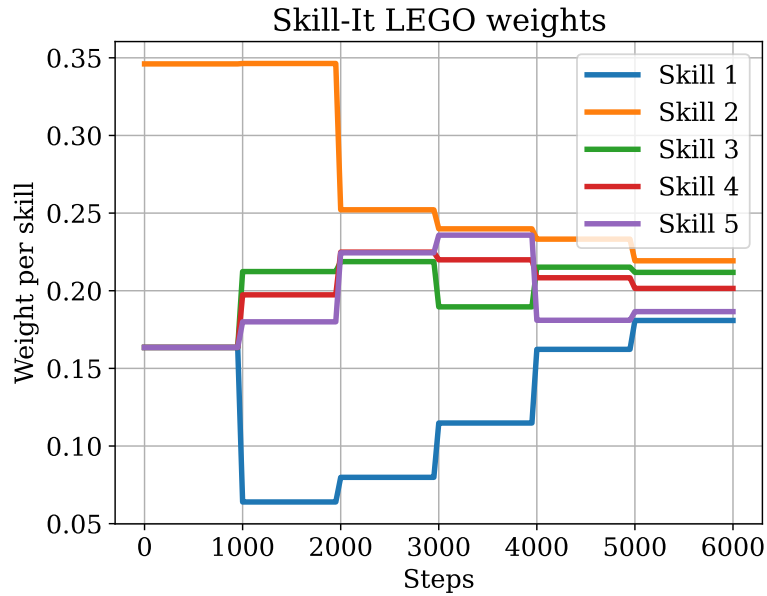


Figure 20. Weight per skill for LEGO pre-training experiment. SKILL-IT initially allocates more weight to skill 2, but eventually puts more weight on harder skills (3, 4, 5) before converging to uniform sampling when all losses converge roughly to 0.

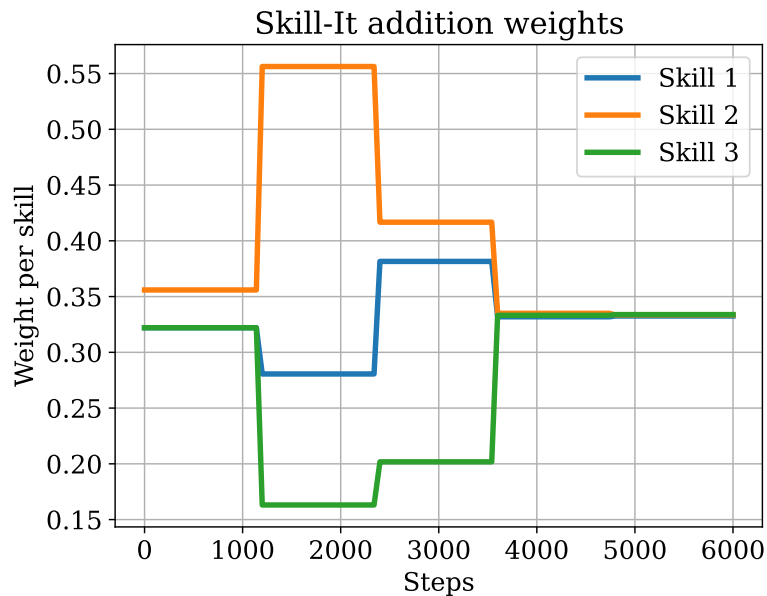


Figure 21. Weight per skill for addition pre-training experiment. SKILL-IT initially allocates more weight to skill 2, which has an edge to skill 1, while allocating little weight to skill 3 which is learned quickly. Eventually, SKILL-IT puts more weight on the harder skill 1 before converging to uniform sampling when all losses roughly approach 0.

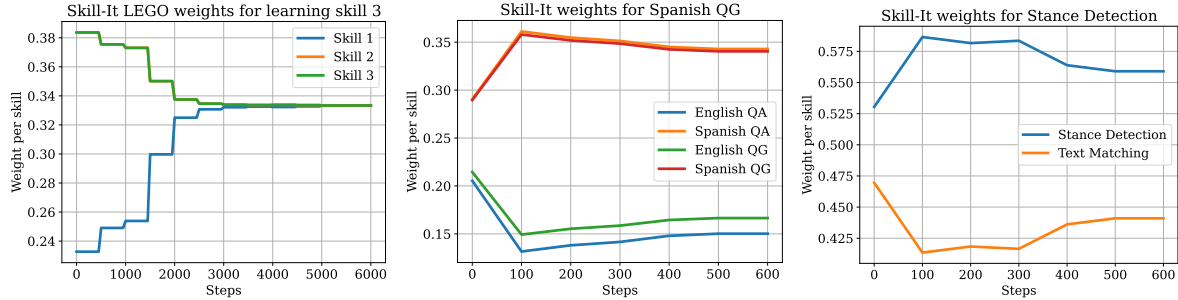


Figure 22. Weight per skill for fine-tuning experiments. Left: LEGO; Center: Spanish question generation; Right: stance detection.

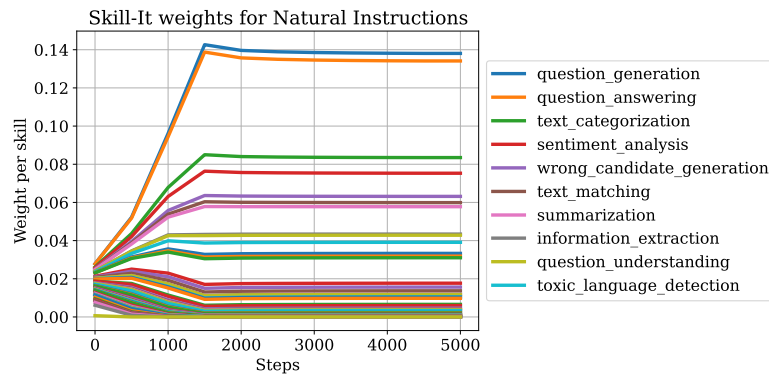


Figure 23. Weight per skill for Natural Instructions out-of-domain experiment. The legend shows the top 10 skills with the largest weight. While the relative order of weight magnitude does not change significantly across training, the incorporation of loss dramatically increases the range of the weights, showing the importance of an online algorithm.

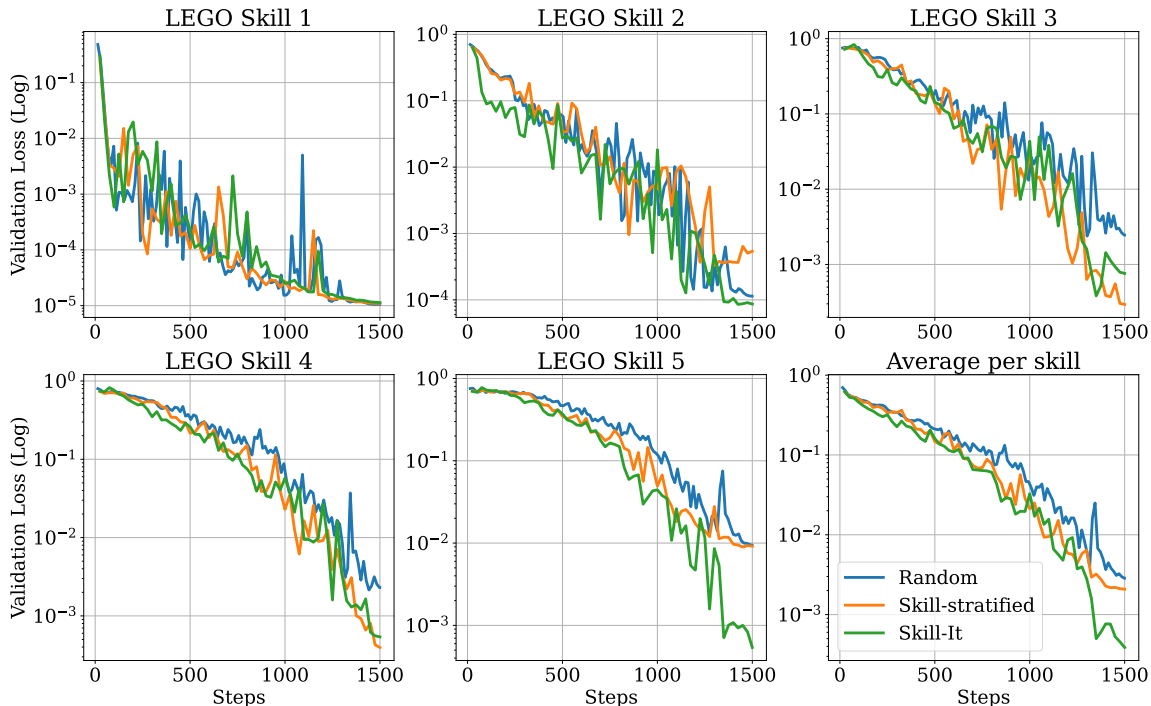


Figure 24. Performance of SKILL-IT for LEGO pre-training setting when skills graph is learned on a 125M parameter model and used for data selection with a 1.3B model. SKILL-IT on average still outperforms random and skill-stratified sampling, suggesting that findings on ordered skill sets can transfer from small models to large models.

In Table 9, we report performance of SKILL-IT with the 1.3B model on the Natural Instructions pre-training experiment and find that the trends from the smaller model hold—SKILL-IT outperforms random and skill-stratified sampling on average.

E.5. Ablations

We report ablations on the skills graph and the online component of SKILL-IT. Instead of using A in Algorithm 1, we study the performance when the identity matrix is used instead; intuitively, this corresponds to a misspecified skills graph where no skill influences another skill. We refer this approach as “No graph”. Note that the opposite case of a complete graph recovers skill-stratified sampling, and that this approach does not apply to the out-of-domain setting, since the submatrix $A \in \mathbb{R}^{k \times m}$ is not square.

Second, instead of sampling over multiple rounds and weighting according to the loss of each skill, we study the effect of setting $T = 1$, which only uses a softmax on A to yield static weights on the skills. We refer to this approach as “Static”. We omit results on Natural Instructions continual pre-training, since SKILL-IT uses $T = 1$ and using no graph with $T = 1$ recovers skill-stratified sampling. Intuitively, we expect the static version of SKILL-IT to perform somewhat well unless there is significant discrepancy among the losses (e.g. in synthetics where the loss on one skill can be close to 0 while the other is not, versus in Natural Instructions where all losses decrease consistently). For both ablations, we sweep over values of $\eta = [0.1, 0.2, 0.5, 0.8]$.

Figure 26 shows the comparison between SKILL-IT and no graph on the continual pre-training LEGO experiment, and Figure 27 shows the comparison between SKILL-IT and a static approach. We see that both the graph and the online dynamics of SKILL-IT are important for its performance. In particular, using no graph results in allocating significant weight to harder skills early on, even though many of them have easier prerequisite skills (such as skill 3 having edges to skills 1 and 2). Using a static graph results in consistent allocation of significant weight to prerequisite skills even after their validation losses converge to near 0, and thus the harder skills that have higher loss are not learned quickly afterwards.

We perform the same ablation on the Addition dataset—the results for this are shown in Figures 28 and Figure 29. We find that these simple baselines, including using a static graph and no graph perform similarly to SKILL-IT on average across all

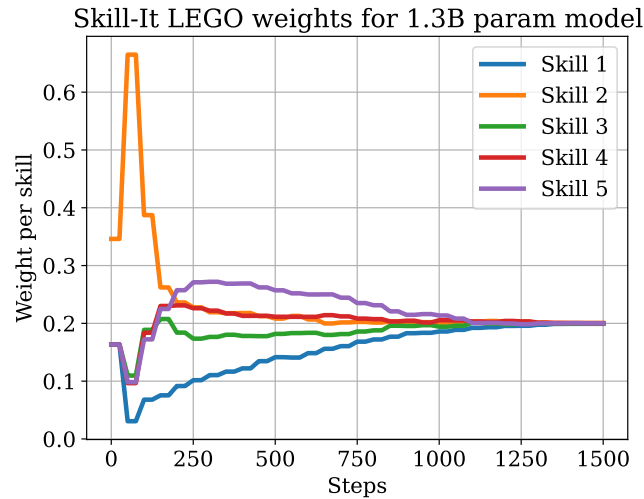


Figure 25. Weight per skill for LEGO pre-training experiment on 1.3B parameter model. The trajectories are similar to those of the 125M parameter model in Figure 20. SKILL-IT initially allocates more weight to skill 2, but eventually puts more weight on skills 4 and 5 before converging to uniform sampling when all losses converge to near 0.

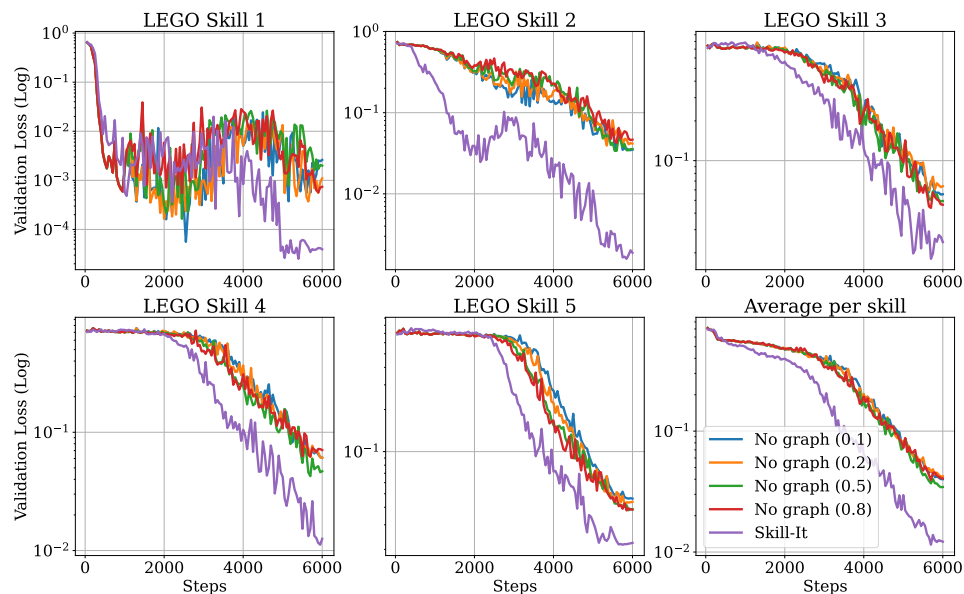


Figure 26. Comparison of SKILL-IT versus using the identity adjacency matrix (no skills graph) with $\eta = 0.1, 0.2, 0.5, 0.8$ on the LEGO continual pre-training experiment. The latter does not capture the relationship between skills, and we find that SKILL-IT attains lower loss on all skills.

Table 9. Results when skills graph for Natural Instructions learned on 125M parameter model is used for data selection with a 1.3B model. We see that SKILL-IT on average still outperforms random and skill-stratified sampling, even though the edges used by SKILL-IT are not derived from the larger model.

Skill	Random	Skill-stratified	SKILL-IT
Answer Verification	2.005 \pm 0.059	1.903 \pm 0.069	1.890 \pm 0.072
Code to Text	0.302 \pm 0.032	0.204 \pm 0.022	0.269 \pm 0.032
Discourse Connective Identification	2.529 \pm 0.046	2.372 \pm 0.054	2.393 \pm 0.056
Entity Generation	2.108 \pm 0.328	1.788 \pm 0.429	1.885 \pm 0.461
Entity Relation Classification	1.130 \pm 0.048	0.836 \pm 0.006	0.841 \pm 0.010
Information Extraction	2.032 \pm 0.013	1.992 \pm 0.006	1.933 \pm 0.013
Irony Detection	2.802 \pm 0.125	2.528 \pm 0.146	2.585 \pm 0.149
Preposition Prediction	1.095 \pm 0.040	0.686 \pm 0.041	0.774 \pm 0.029
Punctuation Error Detection	2.633 \pm 0.027	3.188 \pm 0.055	2.726 \pm 0.025
Question Answering	1.947 \pm 0.003	2.119 \pm 0.003	2.073 \pm 0.001
Question Generation	2.214 \pm 0.007	2.345 \pm 0.008	2.263 \pm 0.010
Question Understanding	1.928 \pm 0.020	1.837 \pm 0.031	1.700 \pm 0.042
Sentence Expansion	2.054 \pm 0.018	1.828 \pm 0.060	1.853 \pm 0.058
Sentiment Analysis	2.771 \pm 0.009	2.818 \pm 0.006	2.774 \pm 0.007
Stance Detection	1.814 \pm 0.151	1.500 \pm 0.117	1.628 \pm 0.149
Summarization	2.531 \pm 0.009	2.472 \pm 0.012	2.440 \pm 0.013
Text Categorization	2.289 \pm 0.016	2.341 \pm 0.021	2.231 \pm 0.022
Text Matching	1.967 \pm 0.008	1.913 \pm 0.005	1.872 \pm 0.005
Text Simplification	1.861 \pm 0.003	1.692 \pm 0.023	1.698 \pm 0.022
Text to Code	0.614 \pm 0.030	0.518 \pm 0.030	0.585 \pm 0.022
Toxic Language Detection	2.853 \pm 0.020	2.911 \pm 0.019	2.862 \pm 0.018
Word Semantics	1.999 \pm 0.023	1.870 \pm 0.039	1.902 \pm 0.024
Wrong Candidate Generation	2.187 \pm 0.028	2.192 \pm 0.023	2.140 \pm 0.020
Average	1.985 \pm 0.022	1.907 \pm 0.027	1.883 \pm 0.032

skills—while SKILL-IT performs the best on skill 2 compared to vanilla multiplicative weights, and SKILL-IT performs the best on skill 1 compared to a static graph. This suggests that Addition is somewhat easier than the other datasets that we consider, as SKILL-IT still outperforms other baselines, as shown in Figure 4.

Figure 30 compares SKILL-IT, no graph, and static data selection for the LEGO fine-tuning experiment. No graph can be interpreted as allocating equal weight to all training skills not equal to the target skill, and varying this weight versus the weight on the target skill. While SKILL-IT and setting $T = 1$ behave similarly, we see that SKILL-IT is slightly better than using no graph. For instance, SKILL-IT obtains a validation loss of 0.05 in 2000 steps, compared to 2050-2200 steps when using no graph.

Figure 31 and 32 compare SKILL-IT, no graph, and static data selection for the Natural Instructions fine-tuning experiments. For both Spanish QG and stance detection, SKILL-IT attains lower loss than using no graph or using $T = 1$ round.

Figure 33 compares SKILL-IT and static data selection for the Natural Instructions out-of-domain experiment. SKILL-IT attains the lowest validation loss on 7 out of 12 evaluation skills. It has an average loss of 2.540 compared to a range of 2.541-2.551 for static data selection.

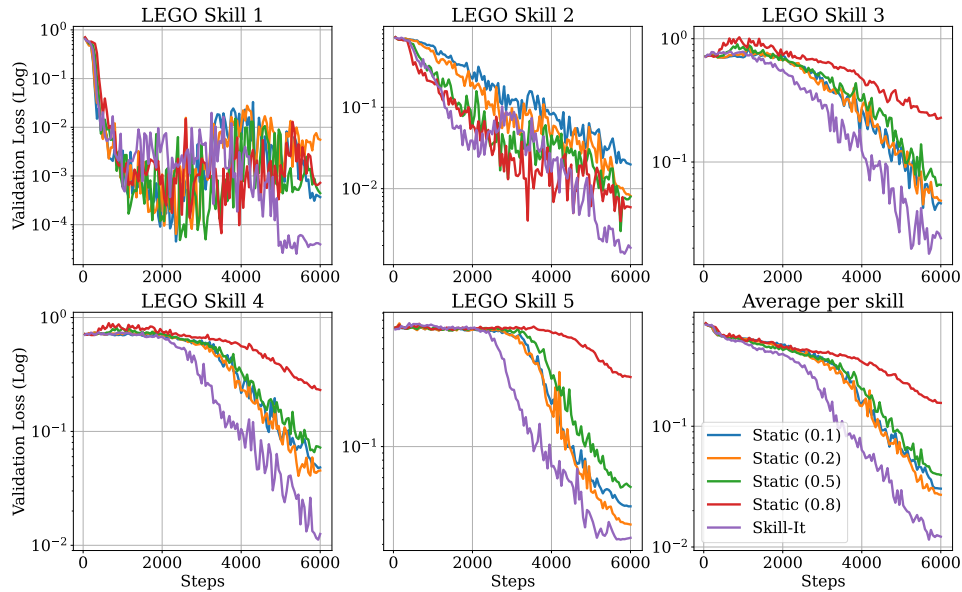


Figure 27. Comparison of SKILL-IT versus using static data selection ($T = 1$) with $\eta = 0.1, 0.2, 0.5, 0.8$ on the LEGO continual pre-training experiment. While SKILL-IT eventually allocates more weights to skills 3, 4, 5, which have higher loss, the static approach is not able to do this. We find that SKILL-IT attains lower loss on all skills.

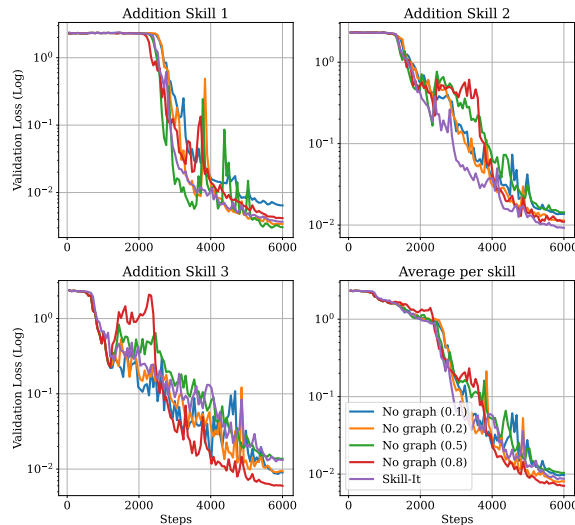


Figure 28. Comparison of SKILL-IT versus using the identity adjacency matrix (no skills graph) with $\eta = 0.1, 0.2, 0.5, 0.8$ on the Addition continual pre-training experiment. The latter does not capture the relationship between skills, and we find that SKILL-IT attains lower loss on skill 2, but attains similar performance to methods that do not use the skills graph.

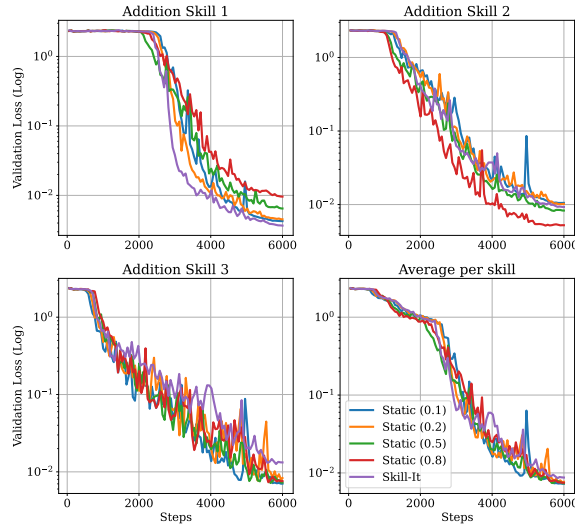


Figure 29. Comparison of SKILL-IT versus using static data selection ($T = 1$) with $\eta = 0.1, 0.2, 0.5, 0.8$ on the Addition continual pre-training experiment. We find that SKILL-IT attains lower loss on skill 1, but attains similar performance to the static methods.

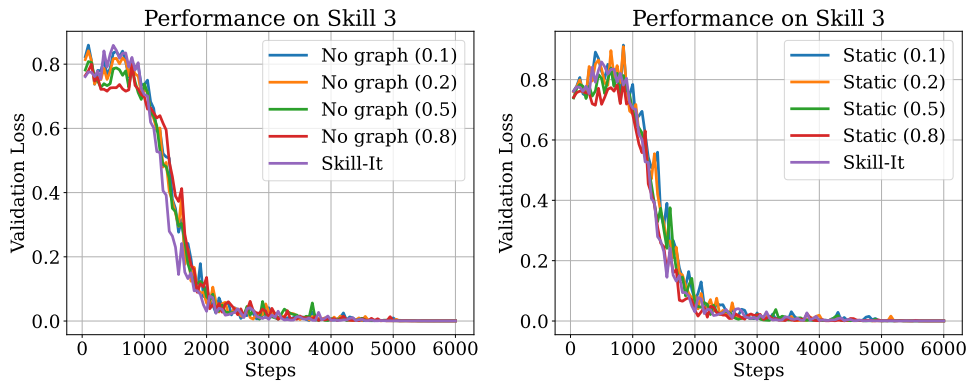


Figure 30. Comparison of SKILL-IT versus using no graph (left) and static data selection (right) with $\eta = 0.1, 0.2, 0.5, 0.8$ on the LEGO fine-tuning experiment. All approaches have roughly the same loss trajectories, but SKILL-IT is slightly lower than using no graph.

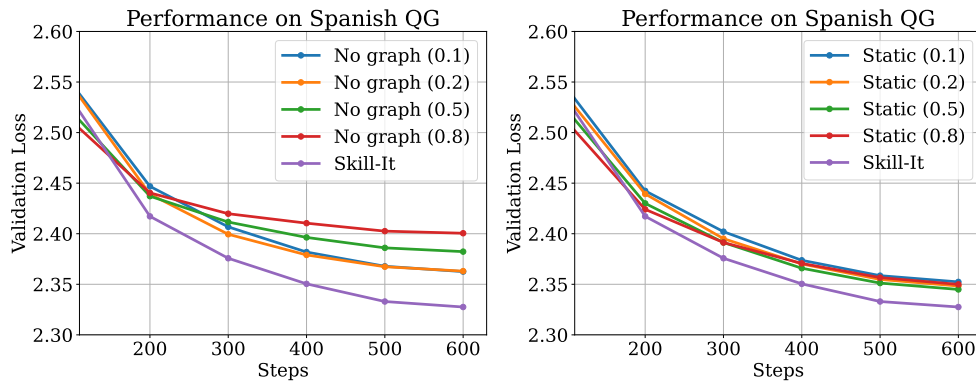


Figure 31. Comparison of SKILL-IT versus using no graph (left) and static data selection (right) with $\eta = 0.1, 0.2, 0.5, 0.8$ on the Natural Instructions Spanish QG fine-tuning experiment. SKILL-IT attains lower validation loss than both no graph and static data selection.

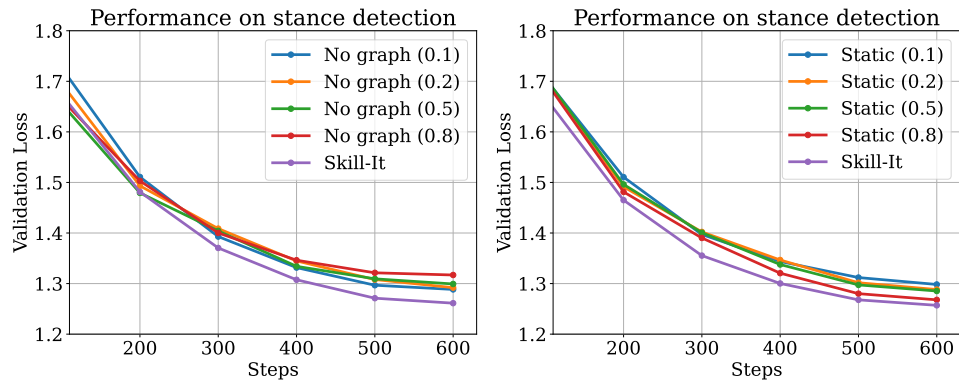


Figure 32. Comparison of SKILL-IT versus using no graph (left) and static data selection (right) with $\eta = 0.1, 0.2, 0.5, 0.8$ on the Natural Instructions stance detection fine-tuning experiment. SKILL-IT attains lower validation loss than both no graph and static data selection.

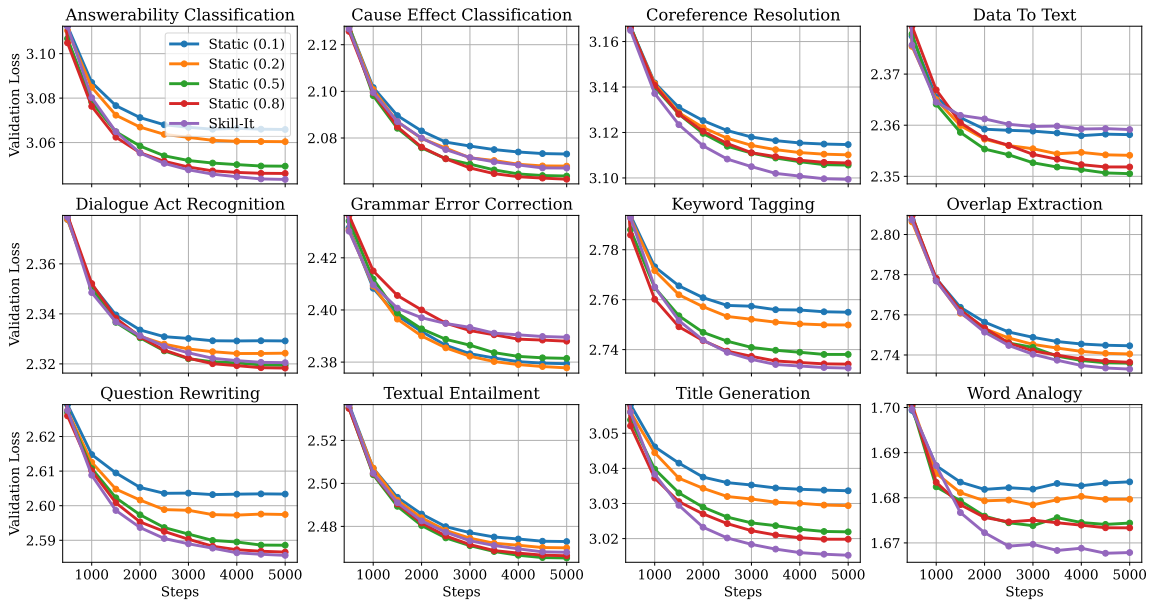


Figure 33. Comparison of SKILL-IT versus using static data selection with $\eta = 0.1, 0.2, 0.5, 0.8$ on the Natural Instructions out-of-domain experiment. SKILL-IT attains the lowest validation loss on 7 out of 12 evaluation skills, and an average loss of 2.540 compared to a range of 2.541-2.551 for static data selection.