

ShapleyLaw: A Game-Theoretic Approach to Multilingual Scaling Laws

Xuyang Cao^{1,2*}, Qianying Liu^{1*}, Chuan Xiao², Yusuke Oda^{1,3},
Jiayi Wang^{1,4†}, Pontus Stenetorp^{1,4}, Daisuke Kawahara^{1,5},
Makoto Onizuka², Sadao Kurohashi¹, Shuyuan Zheng^{2‡}

¹NII LLMC, ²The University of Osaka,

³Nara Institute of Science and Technology, ⁴University College London, ⁵Waseda University

{xuyang.cao, chuanx, onizuka, zheng}@ist.osaka-u.ac.jp
{ying, odashi, pontus, kurohashi}@nii.ac.jp dkw@waseda.jp jiaywang@cs.ucl.ac.uk

Abstract

In multilingual pretraining, the test loss of a pretrained model is heavily influenced by the proportion of each language in the pretraining data, namely the *language mixture ratios*. Multilingual scaling laws can predict the test loss under different language mixture ratios and can therefore be used to estimate the optimal ratios. However, the current approaches to multilingual scaling laws do not explicitly quantify *cross-lingual transfer* between individual languages, resulting in suboptimal mixture-ratio optimization. In this paper, we consider multilingual pretraining as a cooperative game in which each language acts as a player that jointly contributes to pretraining, gaining the resulting reduction in test loss as the payoff. Consequently, from the perspective of cooperative game theory, we quantify the cross-lingual transfer from each language by its contribution in the game, and propose a game-theoretic multilingual scaling law called *ShapleyLaw*. Our experiments show that ShapleyLaw outperforms baseline methods in model performance prediction and language mixture optimization.¹

1 Introduction

As large language models (LLMs) serve increasingly diverse users, multilingual pretraining on mixed-language corpora has become standard practice for modern foundation models (Liu et al., 2020; Lin et al., 2020; Huang et al., 2023; Zhao et al., 2024; Xu et al., 2025; Wang et al., 2025a). The performance of a pretrained model on a given language j depends not only on the model size N and the pretraining data size D , but also critically on the proportion of each language i in the pretraining corpus, i.e., the mixture ratios $\mathbf{p}$. This is because

knowledge learned from language i can transfer to language j , a phenomenon known as *cross-lingual transfer* (Zhuang et al., 2021; Conneau and Lample, 2019; Schuster et al., 2019; Pfeiffer et al., 2020; Fernandes et al., 2023; Isonuma and Titov, 2024; He et al., 2025; Wang et al., 2025c). Therefore, identifying an optimal set of language mixture ratios that minimizes the test loss is a crucial problem for multilingual pretraining.

Recently, He et al. (2025) proposed using *multilingual scaling laws* to predict the test loss of pretrained models under different language mixture ratios, thereby optimizing language mixture. A central challenge in constructing such laws is how to measure cross-lingual transfer and relate mixture ratios to test loss. To simplify the complicated transfer effect among languages, the authors grouped languages by language family and assumed independence across families: *cross-lingual transfer occurs only among languages within the same family*. Then, they constructed multilingual scaling laws for each language family, referred to as *FamilyLaw*. However, FamilyLaw can merely be used to optimize mixture ratios across language families rather than across individual languages; the authors simply used the smoothed sampling strategy to allocate language mixture ratios within each family.


Moreover, the family-independence assumption underlying FamilyLaw does not hold. In our experiments, we find that cross-lingual transfer also exists between languages from different families, and that its strength varies substantially. For example, as Japanese and Chinese share substantial lexical overlap, although Chinese, Japanese, and Spanish belong to different language families, Japanese data improves performance on the Chinese task much more than Spanish data does (see Section 2.2). This result is consistent with the findings in Lin et al. (2019), which showed that the cross-lingual transfer effect from language i to language j largely depends on linguistic distances such as geographic distance

*Equal contribution.


†This work was done while the author was visiting NII LLMC.

‡Corresponding author.

¹Our code and data are available at <https://github.com/teijyogen/ShapleyLaw>.



(a) Cross-lingual transfer measured by the SV.



(b) Test loss prediction by ShapleyLaw.

Figure 1: **(a)** Cross-lingual transfer effects quantified by normalized SVs estimated on a small-scale configuration ($N = 50\text{M}$, $D = 50\text{B}$), shown as an asymmetric matrix. In contrast, the FamilyLaw (He et al., 2025) baseline yields zero off-diagonal entries. **(b)** Across diverse pre-training mixtures and model scales $N \in \{50\text{M}, 243\text{M}, 684\text{M}, 1460\text{M}\}$ with total data size $D = 100\text{B}$, the proposed ShapleyLaw $\mathcal{L}_j(N, D, \Theta_{ja})$ accurately fits and predicts Japanese performance.

and syntactic distance. By ignoring cross-family transfer, FamilyLaw may inform language mixtures that deviate substantially from the optimum.

In this work, we propose quantifying cross-lingual transfer from the perspective of cooperative game theory (CGT). We formulate multilingual pretraining as a cooperative game, in which each language is treated as a player and the payoff is defined by the reduction in test loss on a target evaluation language. We then use the Shapley value (SV) to measure the contribution of each training language to the target language, and interpret this contribution as cross-lingual transfer, as illustrated in Figure 1a.

Based on this transfer measure, we propose *ShapleyLaw*, a game-theoretic multilingual scaling law that predicts test loss from model size, data size, and language mixture ratios. To capture the effect of multilingual mixture ratios on a target language, we introduce an *aggregate cross-lingual transfer* quantity that summarizes the transfer effect of a mixture through SV-based transfer strengths. In this way, ShapleyLaw extends standard scaling law formulations with a language-level characterization of cross-lingual transfer. As shown in Figure 1b, it accurately predicts test loss on various mixtures and model scales, enabling mixture optimization in individual languages.

We summarize our contributions as follows.

- **SV-based measure of cross-lingual transfer.** We formulate multilingual pretraining as a cooperative game, which allows us to apply the SV to quantify the cross-lingual transfer effect across languages.

- **ShapleyLaw for multilingual scaling laws.** ShapleyLaw is the first framework to characterize the functional relationship between language-level mixture ratios $\mathbf{p}$ and test loss, thereby informing optimal mixture ratios across languages.
- **Empirical validation.** We evaluate ShapleyLaw on multilingual pretraining tasks covering 10 languages. The results show that ShapleyLaw consistently outperforms baseline methods in test loss prediction, mixture-ratio optimization, and downstream task performance prediction.

2 Preliminaries

2.1 Problem Statement

We consider pretraining a multilingual model $\theta_{\mathcal{K}}$ that supports a given set of languages $\mathcal{K} = \{1, \dots, K\}$. For each language $i \in \mathcal{K}$, we have a corpus of size $\bar{D}_i$. To construct a pretraining dataset $\mathcal{D}_{\mathcal{K}}$ of size D for pretraining $\theta_{\mathcal{K}}$, for each language i , we sample a training subset $\mathcal{D}_i$ from the corpus according to a mixture ratio p_i , where the subset size is $D_i = p_i \cdot D$. The union of all the languages' training subsets forms the pretraining dataset, i.e., $\mathcal{D}_{\mathcal{K}} = \bigcup_{i=1}^K \mathcal{D}_i$.

Performance prediction. Because pretraining a large language model is computationally expensive, we aim to predict the test loss of the pretrained model $\theta_{\mathcal{K}}$ on a target evaluation task $\mathcal{V}_{\mathcal{K}}$ by constructing a multilingual scaling law. Specifically, the target task $\mathcal{V}_{\mathcal{K}} = \{\mathcal{V}_1, \dots, \mathcal{V}_K\}$ consists of multiple language-specific tasks, each denoted by $\mathcal{V}_j$,

where $\mathcal{V}_j$ evaluates the model’s performance on language $j \in \mathcal{K}$. For each task $\mathcal{V}_j$, we estimate a language-specific scaling law $\mathcal{L}_j(N, D, \mathbf{p})$ to predict the loss of model θ_K on that task, where N is the model size, and $\mathbf{p} = [p_1, \dots, p_K]$ is the vector of all the language mixture ratios. Then, we can predict the model’s performance on the target task $\mathcal{V}_K$ by aggregating all the language-specific scaling laws into a multilingual scaling law:

$$\mathcal{L}(N, D, \mathbf{p}) = \sum_{j=1}^K w_j \cdot \mathcal{L}_j(N, D, \mathbf{p})$$

where w_j denotes a weight representing the relative importance of task $\mathcal{V}_j$. For example, we can determine $w_1, \dots, w_K$ by estimating the real-world distribution of language tokens across languages.

Mixture optimization. In addition to predicting test loss, which is the traditional function of scaling laws, multilingual scaling laws can further be used to inform optimal mixture ratios. Specifically, given a model size N , a total data size D , and importance weights $\mathbf{w} = [w_1, \dots, w_K]$, we can estimate the optimal mixture ratios $\mathbf{p}^*$ by solving the following optimization problem:

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \sum_{j=1}^K w_j \cdot \mathcal{L}_j(N, D, \mathbf{p}). \quad (1)$$

In summary, a multilingual scaling law should (1) accurately predict test losses for different parameter configurations $(N, D, \mathbf{p})$, and (2) enable the model pretrained with the optimized mixture ratios $\mathbf{p}^*$ to achieve a test loss as low as possible.

2.2 Limitations of the FamilyLaw

He et al. (2025) proposed the current SOTA multilingual scaling law, which we refer to as FamilyLaw. To simplify the modeling of cross-lingual transfer, FamilyLaw groups the languages into M language families and predicts the test loss $\mathcal{L}_m^{\text{fam}}$ for each family $m \in [M]$ as follows:

$$\mathcal{L}_m^{\text{fam}}(N, D, \mathbf{p}) = (E_m + \frac{A_m}{N^{\alpha_m}} + \frac{B_m}{D^{\beta_m}})(\mathbf{p}_m)^{-\gamma_m}$$

where $\mathbf{p}_m$ is the sum of the mixture ratios for languages in family m , and $A_m, B_m, \alpha_m, \beta_m$ are fitting parameters.

However, we find that FamilyLaw has the following three limitations. (1) For performance prediction, FamilyLaw does not construct language-specific scaling laws, making it difficult to predict




Figure 2: Consider pretraining two models for a Chinese task with training data mixtures in Chinese (zh), Japanese (ja), and Spanish (es), which belong to different language families. We first pretrain Model 1 by sampling Chinese, Spanish, and Japanese data with ratios [0.500, 0.495, 0.005], respectively. We then swap the mixture ratios of Japanese and Spanish to pretrain Model 2. Although Models 1 and 2 share the same model size $N=684\text{M}$ and data size $D=100\text{B}$, they exhibit significantly different losses on the Chinese task.

model performance on single-language tasks. (2) For mixture ratio optimization, FamilyLaw only optimizes mixture ratios across language families, but does not inform the optimal mixture ratios within each family (He et al. (2025) simply applied smoothed sampling to determine the intra-family ratios). (3) FamilyLaw relies on a core assumption that cross-lingual transfer occurs only within the same language family. However, as shown in Figure 2, this assumption does not hold in practice: Although both Japanese and Spanish belong to language families different from Chinese, Japanese data exhibits substantially stronger cross-lingual transfer for the Chinese task because Japanese and Chinese share substantial lexical overlap.

Our objectives. Considering the above limitations, we aim to develop a multilingual scaling law that (1) predicts language-specific test loss, (2) informs optimal mixture ratios across individual languages rather than language families, and (3) models cross-lingual transfer across languages.

3 Shapley Value-Based Scaling Law

3.1 Multilingual Pretraining Game

The core challenge in developing multilingual scaling laws lies in quantifying cross-lingual transfers across languages. We address this challenge using CGT, which studies how to fairly assess the contribution of each player when multiple players

jointly complete a task. In the context of multilingual pretraining, each language $i \in \mathcal{K}$ can likewise be viewed as a player. These players jointly pretrain a model, which is used to complete an evaluation task $\mathcal{V}_j$ in a target language j . This formulation allows us to use contribution metrics from CGT to quantify the contribution of each language i to the target language j , and thereby assess cross-lingual transfer. We refer to this cooperative game as a *multilingual pretraining game*, defined as follows.

Definition 3.1 (Multilingual Pretraining Game (MPG)). An MPG $\mathcal{G} = (\mathcal{K}, N, D, \mathcal{V})$ is specified by a language set $\mathcal{K}$, a model size N , a data size D , and a target evaluation task $\mathcal{V}$. For any language subset $S \subseteq \mathcal{K}$, the payoff $v_j^{\mathcal{G}}(S)$ of S on any language-specific task $\mathcal{V}_j \in \mathcal{V}$ is the decreased test loss by pretraining $\theta_{\mathcal{K}}$ using a pretraining dataset $\mathcal{D}_S^{\text{uni}}$, where $\mathcal{D}_S^{\text{uni}}$ consists of training data from languages S with uniform mixture ratios. A contribution metric ϕ to the game $\mathcal{G}$ returns a $K \times K$ contribution matrix $\phi(\mathcal{G}) = \{\phi_{i,j}(\mathcal{G})\}_{i,j \in \mathcal{K}}$ that distributes the payoff $v_j^{\mathcal{G}}(\mathcal{K})$ among all the languages $\mathcal{K}$ for all task $\mathcal{V}_j \in \mathcal{V}$, i.e., $\forall j, \sum_{i=1}^K \phi_{i,j}(\mathcal{G}) = v_j^{\mathcal{G}}(\mathcal{K})$.

Intuitively, in the MPG $\mathcal{G}$, all the languages $\mathcal{K}$ as players participate in pretraining with the same amount of data. For each task $\mathcal{V}_j$, they gain a payoff of $v_j^{\mathcal{G}}(\mathcal{K})$, which is the reduction in the test loss of the pretrained model. The contribution of language i to task $\mathcal{V}_j$, denoted as $\phi_{i,j}(\mathcal{G})$, measures the portion of this test loss reduction attributable to language i . Consequently, we can measure the cross-lingual transfer from language i to language j as the contribution $\phi_{i,j}$ in MPG.

3.2 Shapley Value for Measuring Cross-Lingual Transfer

The Shapley value (SV) is a widely-adopted contribution metric because it is the only metric that satisfies four important properties: *linearity*, *efficiency*, *symmetry*, and *null player* (see Appendix A.1 for details). In MPG, we can also employ the SV ϕ^{SV} to compute the contribution matrix as follows:

$$\phi_{i,j}^{SV}(\mathcal{G}) = \sum_{S \subseteq \mathcal{K} \setminus \{i\}} W(S) \cdot \Delta_{i,j}^{\mathcal{G}}(S)$$

where $W(S) = \frac{|S|!(K-|S|-1)!}{K!}$, and $\Delta_{i,j}^{\mathcal{G}}(S) = v_j^{\mathcal{G}}(S \cup \{i\}) - v_j^{\mathcal{G}}(S)$.

Intuitively, the SV measures the marginal contribution of language i to the reduction in test loss on task $\mathcal{V}_j$, averaged over all possible subsets of languages. Specifically, the SV (1) enumerates all

language subsets $S \subseteq \mathcal{K} \setminus \{i\}$, (2) calculates the marginal improvement $\Delta_{i,j}^{\mathcal{G}}(S)$ in the test loss by adding language i into S for pretraining, and (3) finally aggregates the improvements by language i across all subsets S as language i 's contribution to task $\mathcal{V}_j$, where $W(S)$ is the weight for aggregation. To ensure comparability across different tasks, we further apply exponential normalization to compute the normalized SV as follows:

$$\phi_{i,j}^{NSV}(\mathcal{G}) = \exp(\phi_{i,j}^{SV}(\mathcal{G}) - \max_{i'} \phi_{i',j}^{SV}(\mathcal{G}))$$

which ensures that the language i with the largest contribution to task $\mathcal{V}_j$ has a normalized SV of 1, while the values for other languages lie in the range $(0, 1]$. The rationale for choosing this normalization scheme is provided in Appendix C.5. Consequently, the normalized SV can serve as a measure of the cross-lingual transfer from any language i to any language j . For simplicity, we denote $\phi_{i,j}^{NSV}(\mathcal{G})$ as $\phi_{i,j}^{NSV}$ in the following, when no ambiguity arises.

3.3 Construction of ShapleyLaw

Finally, we construct a multilingual scaling law, named *ShapleyLaw*, based on the SV. Formally, given a set of parameter configurations $(N, D, \mathbf{p})$ and a set of normalized SVs $\{\phi_{i,j}^{NSV}\}$ that represent cross-lingual transfers among languages in $\mathcal{K}$, ShapleyLaw predicts test loss on each task $\mathcal{V}_j$ as:

$$\mathcal{L}_j(N, D, \mathbf{p} \mid \{\phi_{i,j}^{NSV}\}) = \mathcal{L}_j^{\text{chin}}(N, D) \cdot (\Theta_j)^{-\gamma_j}$$

where $E_j, A_j, B_j, \alpha_j, \beta_j, \gamma_j$ are fitting parameters; $\mathcal{L}_j^{\text{chin}}(N, D) = E_j + \frac{A_j}{N^{\alpha_j}} + \frac{B_j}{D^{\beta_j}}$ corresponds to the Chinchilla scaling law (Hoffmann et al., 2022) that does not consider multilingual mixture ratios; and $\Theta_j = \sum_{i \in \mathcal{K}} p_i \cdot \phi_{i,j}^{NSV}$ measures the *aggregate cross-lingual transfer* from languages $\mathcal{K}$ to language j weighted by their mixture ratios. For a fixed model size N and data budget D , different multilingual mixtures can yield different losses on task $\mathcal{V}_j$ due to differences in their cross-lingual transfer effects. Accordingly, $\mathcal{L}_j^{\text{chin}}(N, D)$ captures the base scaling behavior, while Θ_j characterizes the effective language-specific support induced by the data mixture; their multiplicative interaction models cross-lingual transfer as modulating how effectively a fixed scale budget benefits the target language. For example, in Figure 2, the aggregate transfer to Chinese is computed as $\Theta_{\text{zh}} = \sum_i p_i \phi_{i,\text{zh}}^{NSV}$. Swapping the mixture ratios of Japanese and Spanish changes Θ_{zh} according to their respective transfer contributions to Chinese, thereby yielding different Chinese losses despite identical N and D .




Figure 3: **Comparison of Shapley values across model sizes.** **Left:** normalized Shapley values estimated with $N = 50\text{M}$ and $D = 50\text{B}$. **Right:** normalized Shapley values estimated with $N = 243\text{M}$ and $D = 50\text{B}$. The two heatmaps exhibit highly consistent cross-lingual transfer patterns across model sizes.

SV transferability. A key assumption of Shapley-Law is that the SVs used to quantify cross-lingual transfer are transferable across model sizes. Under this assumption, the SVs estimated from a relatively small model can be reused as fixed transfer coefficients when fitting scaling laws for larger models. To empirically verify this property, we conduct a controlled comparison across two model sizes, $N = 50\text{M}$ and $N = 243\text{M}$, while keeping the pretraining data size fixed at $D = 50\text{B}$. We consider six language sets, $\mathcal{K} = \{\text{de}, \text{en}, \text{fr}, \text{ja}, \text{ko}, \text{zh}\}$. For each model size, we construct a multilingual pretraining game over the same language set and estimate the normalized SV matrix $\phi^{NSV} \in \mathbb{R}^{|\mathcal{K}| \times |\mathcal{K}|}$, where each entry $\phi_{i,j}^{NSV}$ measures the transfer contribution from training language i to evaluation language j . We then compare the two normalized SV matrices obtained at different model sizes. We use cosine similarity and Pearson correlation coefficient as the comparison metrics. Specifically, we vectorize the two SV matrices in the same language order and compute their similarity over all language pairs.

Figure 3 shows the normalized SV matrices estimated from the 50M and 243M models, which exhibit highly consistent cross-lingual transfer structures. In both cases, self-transfer dominates, while the off-diagonal entries reveal similar asymmetric patterns, such as strong transfer from English to several European languages and consistent interactions among Japanese, Korean, and Chinese. Quantitatively, the two vectorized SV matrices achieve a cosine similarity and Pearson correlation of 0.999 and 0.999, respectively; after excluding the diagonal entries, they remain high at 0.995 and 0.989. These results demonstrate that the cross-lingual transfer structure captured by SVs remains highly stable as

the model scales from 50M to 243M parameters. Overall, these results support the transferability of SVs across model sizes, substantially reducing the computational cost of estimating transfer effects while preserving the transfer structure required for accurate multilingual scaling-law fitting.

Because SV-based cross-lingual transfer exhibits transferability, we can construct an MPG $\mathcal{G} = (\mathcal{K}, N', D', \mathcal{V})$ using a small model size N' and data size D' to compute the normalized SVs. These normalized SVs can then be treated as fixed parameters when constructing ShapleyLaw for larger model sizes $N \gg N'$ and data sizes $D \gg D'$. As shown in Figure 1a, we compute the normalized SVs under the setting $N' = 50\text{M}$ and $D' = 50\text{B}$; ShapleyLaw accurately fits the test losses across settings with $N \in \{243\text{M}, 684\text{M}, 1460\text{M}\}$ and $D = 100\text{B}$.

In addition to predicting language-specific test loss, ShapleyLaw can also be used to estimate *optimal mixture ratios* by solving the optimization problem in Formula 1 (see Appendix A.2). Therefore, ShapleyLaw achieves our three objectives.

SV approximation. Since the SV $\phi_{i,j}^{SV}$ enumerates all non-empty language subsets $S \subseteq \mathcal{K}$ to evaluate their payoff $v_j^{\mathcal{G}}(S)$, we need to pretrain $2^K - 1$ models for exact SV computation, which is NP-hard. To accelerate SV computation, we employ the In-Run Data Shapley (IRDS) approximation technique (Wang et al., 2025b) (see Appendix A.3) in our experiments, which pretrains only one model to approximate the SVs. With IRDS, ShapleyLaw achieves $O(1)$ computational complexity and is therefore a tractable multilingual scaling law that can scale up to dozens or even hundreds of languages. Our experiments in Appendix C.3 show that IRDS accurately approximates the exact SV.




Figure 4: **Fitting ShapleyLaw across multilingual mixtures and target languages.** **Top row:** fits with respect to aggregate cross-lingual transfer Θ on curated multilingual mixtures for Chinese (ZH), Spanish (ES), and Japanese (JA) across four model scales $\{50M, 243M, 684M, 1460M\}$; circles denote fitting points and yellow diamonds denote held-out evaluation points. **Middle row:** fits under randomly sampled mixtures at model scale $N = 243M$; circles indicate fitting points, yellow diamonds indicate held-out points, and orange crosses denote removed outliers. **Bottom row:** fits with respect to aggregate cross-lingual transfer Θ on curated multilingual mixtures for low-resource languages Finnish (FI), Croatian (HR), and Standard Malay (MS).

4 Experiments

Setup. We describe the basic experimental setup here. More details are provided in Appendix B. We study $K = 10$ languages from eight language families, as summarized in Table 5. We sample 0.1% of the corpora as a fixed test set for evaluating all pretrained models. We consider four model scales with non-embedding parameter counts $N \in \{50M, 243M, 684M, 1460M\}$, following the open-source LLM-jp configurations. All models share the same architecture family and tokenizers. We consider the cross-entropy (CE) loss for evaluation. The SVs are computed in an MPG with $N = 50M$ and $D = 50B$ across all experiments.

Research questions. Our experiments address the following research questions. **Q1** (Performance Prediction): Can ShapleyLaw more accurately predict the model’s test loss? **Q2** (Mixture Optimization): Can ShapleyLaw achieve better language mixtures? **Q3** (Downstream Tasks): Can ShapleyLaw’s ability to predict model performance and optimal mixture ratios generalize to downstream tasks?




Figure 5: **Controlled comparison between FamilyLaw and ShapleyLaw on a Chinese evaluation task.** We pre-train multilingual models with $N = 50M$ and $D = 100B$ using Chinese, Japanese, and Spanish corpora. Starting from the mixture $[0.05, 0.55, 0.40]$, we vary the Spanish–Japanese proportions up to $[0.05, 0.945, 0.05]$ while keeping the Chinese ratio fixed. Performance is evaluated on the Chinese test set. **Left:** FamilyLaw, where x-axis stands for language ratio, yields a poor fit. **Right:** ShapleyLaw, where the x-axis represents different Chinese aggregate transfer Θ , fits accurately.

Q4 (SV Approximation): How is the fitting performance of ShapleyLaw based on approximate SVs? (See Appendix C.3.)

4.1 Experiments on Performance Prediction

Experiment design. We pretrain models with $N \in \{50\text{M}, 243\text{M}, 684\text{M}, 1460\text{M}\}$ and $D = 100\text{B}$. At each model scale, we consider both *controlled mixtures* and *empirical mixtures*. Controlled mixtures isolate the effect of language mixtures under fixed ratios, including monolingual training, uniform mixtures over selected languages (e.g., zh–ja–ko, fi–hr–ms and en–de–fr–es, abbreviated as *ZJK*, *FHM* and *EDFS*), and the uniform mixture over all languages (*Uni*). Empirical mixtures mirror multilingual data distributions adopted in prior LLM studies, including the language distributions of Common Crawl² and FineWeb-2 (Penedo et al., 2025), the English-centric multilingual model LLaMA-2 (Touvron et al., 2023), and the Japanese-focused pretraining corpus LLM-JP-V4³ (Aizawa et al., 2024) (abbreviated as *CC*, *FW2*, *L2*, and *JPv4*). To evaluate the fitting performance, we also use the normalized MAE to quantify the prediction error (PE).

In addition, we examine 50 randomly sampled language mixtures under $N = 243\text{M}$ to further evaluate generalization to unseen mixture ratios. To ensure reliable fitting, we exclude those random mixtures where a training language i has subset size $D_i \leq 1.5\text{B}$. This regime corresponds to extremely low data availability for language i , causing optimization to become unstable even for $N = 243\text{M}$ models. This threshold remains practically relevant: in FineWeb-scale corpora, a language budget of roughly 1.5B tokens already extends to many low-resource languages considered in current multilingual LLM studies, including Basque and Armenian. We then split the remaining runs into fitting and held-out evaluation subsets.

Results. Figure 4 shows the fitted results for Chinese, Spanish, and Japanese. On curated multilingual mixtures across four model scales, ShapleyLaw achieves strong held-out performance, with $\text{PE}_{\text{test}} = 0.021$ for Chinese, 0.010 for Spanish, and 0.006 for Japanese ($R^2 = 0.990, 0.997, 0.996$). On randomly sampled mixtures at $N = 243\text{M}$, the fitted curves remain smooth with respect to Θ , with training-set fits of $\text{PE}_{\text{test}} = 0.008, 0.005$, and 0.002 for Chinese, Spanish, and Japanese ($R^2 = 0.975, 0.942, 0.987$), respectively. These results indicate that the proposed method makes stable predictions across sampling regimes and typologically diverse languages. Additional results for French,

German and Korean are provided in Appendix C.1.

We next compare ShapleyLaw with the family-level baseline FamilyLaw in a controlled setting that isolates cross-family transfer. As shown in Figure 5, we run experiments on the *ZJS* mixture. We fix the Chinese mixture ratio at $p_{\text{zh}} = 0.05$ and vary only the Japanese–Spanish proportions, keeping the total token budget fixed. Under FamilyLaw, both Japanese and Spanish lie outside the Chinese language family, so changing their ratios should not affect Chinese performance; accordingly, the fitted relation is nearly constant and fails to explain the observed variation ($R^2 \leq 0$). In contrast, ShapleyLaw captures the cross-language transfer and fits the loss on Chinese much more accurately ($R^2 = 0.931$). This result shows that family-level aggregation is too coarse to model important cross-lingual interactions in multilingual pretraining, such as shared script and lexical overlap.

Moreover, Appendix C.7 further demonstrates that ShapleyLaw generalizes well to a held-out language when a related language from the same family is retained in the fitting set.

4.2 Experiments on Mixture Optimization

Experiment design. We compare the ShapleyLaw-based mixture optimization against eight representative baselines chosen to reflect the main strategies used in multilingual data allocation, including *Uniform Sampling*, *Smoothed Sampling*, *FamilyLaw Optimization*, *ODM* (Shukor et al., 2025), *DML* (Ye et al., 2025), *ATLAS* (Longpre et al., 2026), and *Empirical Mixtures* (see Appendix B for baseline introduction and Appendix A.4 for baseline selection clarification). We pretrain a model ($N = 243\text{M}$ and $D = 100\text{B}$) using the mixture ratios obtained by each method.

We consider two settings of w_j for weighting the importance of each language-specific task $\mathcal{V}_j$. **Unweighted Sum:** Each task $\mathcal{V}_j$ receives equal weight, so the objective is to minimize the average multilingual loss. **Normalized Sum:** Each task $\mathcal{V}_j$ is weighted by the inverse of its monolingual loss. This ensures that intrinsically harder languages do not dominate the aggregate objective in Formula 1 and provides a more balanced evaluation of multilingual trade-offs across languages.

Results. Table 1 shows the optimized mixtures and their resulting CE loss. Under both preference settings, the mixtures obtained from ShapleyLaw outperform all baselines in the final weighted loss.

²<https://commoncrawl.org>

³<https://gitlab.llm-jp.nii.ac.jp/datasets/llm-jp-corpus-v4>

Strategy	Language Mixture Ratio (%)										CE Loss ↓
	en	de	fr	es	zh	ja	ko	fi	hr	ms	$\mathcal{L}$
Uniform	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	24.440
$\alpha = 0.5$	13.2	14.5	12.6	13.6	19.1	11.4	4.9	4.7	3.7	2.4	24.683
Common Crawl	61.4	8.3	6.4	6.4	7.2	8.2	1.2	0.5	0.3	0.1	25.491
FineWeb-2	80.6	3.3	2.3	2.7	6.7	3.1	0.5	0.4	0.3	0.1	27.133
FamilyLaw	14.1	14.1	1.8	1.8	12.1	11.6	10.2	12.6	11.0	10.7	24.749
ODM	1.5	30.0	5.1	7.2	6.0	30.0	1.5	1.5	15.7	1.5	25.544
DML	1.5	36.7	4.7	17.5	5.9	5.0	1.5	5.3	8.9	13.0	25.330
ATLAS	1.5	1.5	5.1	1.7	3.4	28.0	23.3	12.4	12.5	10.5	25.353
ShapleyLaw	26.0	13.5	1.5	1.5	20.2	13.4	10.3	5.3	6.8	1.5	24.108

Strategy	Language Mixture Ratio (%)										CE Loss ↓
	en	de	fr	es	zh	ja	ko	fi	hr	ms	$\mathcal{L}$
Uniform	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	11.985
$\alpha = 0.5$	13.2	14.5	12.6	13.6	19.1	11.4	4.9	4.7	3.7	2.4	12.180
Common Crawl	61.4	8.3	6.4	6.4	7.2	8.2	1.2	0.5	0.3	0.1	12.715
FineWeb-2	80.6	3.3	2.3	2.7	6.7	3.1	0.5	0.4	0.3	0.1	13.613
FamilyLaw	9.9	9.9	1.5	1.5	11.6	8.6	9.3	17.4	14.9	15.4	12.153
ODM	1.5	30.0	3.7	9.1	5.7	30.0	1.5	1.5	15.5	1.5	12.592
DML	1.5	37.9	4.5	17.8	8.7	17.8	2.3	2.0	3.5	3.9	12.304
ATLAS	1.5	1.5	2.1	1.9	3.3	29.2	19.0	14.6	13.9	13.0	12.518
ShapleyLaw	18.1	15.4	1.5	1.5	19.7	9.3	9.6	10.3	12.0	2.6	11.813

Table 1: Comparison of mixture optimization strategies and their CE loss under the Unweighted Sum (**Left**) and Normalized Sum (**Right**) settings.




Figure 6: **Fitting Relationship between Downstream Performance and Test CE Loss.** Experiments are conducted on six downstream tasks for English, German, French, Spanish, Croatian, and Japanese. Across all tasks, we observe strong *negative correlations* ($|r| > 0.90$, with an average $R^2 = 0.933$), indicating that lower test loss consistently corresponds to higher downstream performance.

The gains are especially meaningful because they are obtained under actual pretraining rather than only from the fitted objective, indicating that the learned law transfers effectively from prediction to optimization. By contrast, the mixture obtained from FamilyLaw is consistently less effective in the per-language setting, suggesting that family-level aggregation is too coarse to capture the trade-offs needed for fine-grained multilingual allocation.

4.3 Experiments on Downstream Tasks

Experiment design. We focus on three multilingual benchmarks that cover complementary NLU abilities: **HellaSwag** (Zellers et al., 2019; Lai et al., 2023) for commonsense reasoning through contextually appropriate ending prediction, **XStoryCloze** (Lin et al., 2022; Mostafazadeh et al., 2017) for cross-lingual story ending prediction, and **XWinograd** for cross-lingual coreference resolution (Tikhonov and Ryabinin, 2021).

We evaluate pretrained models under the same setting as Section 4.1, using $N = 1.46\text{B}$ and $D = 100\text{B}$, and the same multilingual mixture families (controlled and empirical mixtures). We report 5-shot accuracy on the standard benchmark splits. Specifically, we evaluate HellaSwag in {en, fr, es, de, hr}, XWinograd in {en, fr, ja}, and XStoryCloze in {en, es}. For each task-language pair, we compare downstream accuracy against the corresponding pre-

training test cross-entropy on the same language, measured on the held-out test corpus. This yields one calibration line per task-language pair.

Results. We fit a simple calibration model $s = f(\ell) = a \cdot g(\ell) + b$ where ℓ is the test loss, s is downstream accuracy, and a, b are fitting parameters (Chen et al., 2025). We choose $g(\ell) = \exp(-\ell)$ as the transformation function $g(\cdot)$ (see more details in Appendix C.2). Figure 6 shows the resulting fits for the three multilingual benchmarks, with one line per task-language pair and points colored by language. Across all evaluated languages, lower pretraining loss consistently corresponds to higher downstream accuracy, and this relationship is generally well captured by a transformed linear fit. The fitted results show that this trend holds not only within a single benchmark, but across multilingual commonsense reasoning, story completion, and coreference resolution tasks. Table 2 reports the corresponding R^2 values for each task-language pair.

These results show that for the multilingual NLU benchmarks considered here, language-specific pretraining loss remains strongly informative for downstream accuracy across diverse mixtures. As ShapleyLaw predicts this loss accurately across language mixtures, it provides a useful intermediate signal for downstream-oriented mixture design.

Task	de	en	fr	es	ja	hr
HellaSwag	0.918	0.996	0.864	0.897	–	0.992
XWinograd	–	0.980	0.975	–	0.675	–
XStorycloze	–	0.947	–	0.905	–	–

Table 2: Corresponding R^2 values for the fitted relationship between downstream performance and test CE loss for each task–language pair. English achieves the highest R^2 , while Japanese is the lowest. The results consistently show a negative correlation between downstream performance and test CE loss.

Strategy	Language Mixture Ratio (%)					Overall Acc. $\uparrow$
	en	de	fr	es	hr	HellaSwag
Uniform	20.0	20.0	20.0	20.0	20.0	0.343
$\alpha = 0.5$	25.2	22.9	21.9	23.6	6.4	0.344
Common Crawl	73.9	10.0	7.7	7.8	0.6	0.354
FineWeb-2	90.3	3.7	2.5	3.0	0.5	0.355
FamilyLaw	31.7	31.7	4.1	4.1	28.4	0.346
ODM	0.1	88.2	0.1	11.5	0.1	0.321
DML	0.1	21.1	0.1	78.6	0.1	0.331
ATLAS	0.1	19.5	38.4	15.7	26.3	0.338
ShapleyLaw	95.7	0.1	0.1	0.1	4.0	0.356

Table 3: Comparison of mixture optimization strategies and average HellaSwag downstream-task accuracies across five languages under the unweighted-sum setting. ShapleyLaw achieves the best performance, with an average accuracy of 0.356.

Optimal Ratio Downstream Experiment for HellaSwag We further evaluated downstream mixture optimization for HellaSwag in {en, fr, es, de, hr}. To assess the effectiveness of ShapleyLaw on downstream tasks, we pretrained six models using different sampling strategies and evaluated them on HellaSwag. The overall performance was measured by the average accuracy across the five languages.

Results. As shown in Table 3, ShapleyLaw achieves the highest overall accuracy on the multilingual HellaSwag task, reaching 0.356 under the unweighted-sum setting. This result outperforms all baseline strategies. Notably, ShapleyLaw assigns most of the mixture to English while retaining a small proportion of Croatian, suggesting that the optimal downstream mixture ratio is highly English-dependent. These results demonstrate that ShapleyLaw can identify more effective language mixture ratios for downstream performance than baseline strategies, including FamilyLaw and ATLAS.

5 Related Work

Neural scaling laws describe predictable relationships between model performance and scale variables such as model parameters, training tokens, and compute (Kaplan et al., 2020; Hoffmann et al.,

2022), but standard formulations do not account for *data composition*. Recent work extends scaling-law analysis to data mixtures in domain-specific pretraining (Que et al., 2024; Shukor et al., 2025). These works highlight the importance of data composition, but they do not provide a multilingual scaling law that directly optimizes *language-level* mixture ratios. In multilingual settings, prior work either models mixture ratios at the family level (He et al., 2025) or studies bilingual transfer solely in multilingual machine translation (Fernandes et al., 2023; Isik et al., 2025). Concurrent to our work, ATLAS (Longpre et al., 2026) models cross-lingual transfer via empirically fitted coefficients over a target/transfer/remaining-language decomposition. To the best of our knowledge, ShapleyLaw is the first multilingual scaling law to ground cross-lingual transfer in cooperative game theory: rather than relying on ad hoc fitted coefficients, it derives each language’s contribution to target-language performance via SVs, thereby optimizing language-level mixture ratios through a theoretically grounded decomposition of cross-lingual transfer.

6 Conclusion

We proposed *ShapleyLaw*, a game-theoretic multilingual scaling law that explicitly incorporates aggregate cross-lingual transfer into multilingual pre-training. Across experiments on ten languages from eight language families, we show that the proposed law accurately predicts per-language pre-training loss across diverse model sizes, data budgets, and language mixtures, while capturing transfer effects that are missed by prior family-level approaches. We demonstrate that ShapleyLaw supports optimization of multilingual sampling ratios, producing better mixtures than heuristic and prior-law baselines under real pre-training. In addition, the predicted language-specific loss was shown to remain informative for downstream multilingual NLU performance, highlighting the practical utility of the framework beyond pre-training loss alone. We also examine the efficiency of approximate Shapley value estimation and show that it preserves strong empirical performance, highlighting the potential scalability of our method to larger multilingual settings. Taken together, these findings suggest that multilingual scaling behavior is better understood through fine-grained, transfer-aware language interactions, and that our method provides a useful foundation for efficient multilingual mixture design.

Limitations

Due to computational constraints, the largest model scale validated in the main text is $N = 1460\text{M}$ parameters, which already exceeds the largest scale ($N = 1.2\text{B}$) validated in FamilyLaw (He et al., 2025), the SOTA multilingual scaling law. Therefore, we believe that our empirical validation covers a sufficiently broad range of model scales to demonstrate the effectiveness of our method. To further address potential concerns regarding the extrapolation of ShapleyLaw to larger models, we additionally trained four models at $N = 3.2\text{B}$ parameters. ShapleyLaw continues to predict the test loss of these larger models with high accuracy, providing strong evidence that ShapleyLaw generalizes well beyond the scales used for fitting (see Appendix C.4.)

Ethical Statements

We anticipate no ethical concerns from this work.

Use of Generative AI

AI assistants, such as ChatGPT, were used in this work for grammar checking and language editing.

Acknowledgments

This work was financially supported by JST CREST [grant number JPMJCR22M2], JST BOOST [grant number JPMJBS2402], JSPS KAKENHI [grant numbers JP23K17456, JP23K28096, JP25H01117, JP25K21207, JP26K03246], and the United Kingdom Engineering and Physical Sciences Research Council [grant number UKRI5028], and was computationally supported by mdx and the "Development Acceleration Use" program of ABCI 3.0.

References

- Akiko Aizawa, Eiji Aramaki, Bowen Chen, Fei Cheng, Hiroyuki Deguchi, Rintaro Enomoto, Kazuki Fujii, Kensuke Fukumoto, Takuya Fukushima, Namgi Han, Yuto Harada, Chikara Hashimoto, Tatsuya Hiraoka, Shohei Hisada, Sosuke Hosokawa, Lu Jie, Keisuke Kamata, Teruhito Kanazawa, Hiroki Kanezashi, and 61 others. 2024. [LLM-jp: A cross-organizational project for the research and development of fully open japanese LLMs](#). *Preprint*, arXiv:2407.03963.
- Yangyi Chen, Binxuan Huang, Yifan Gao, Zhengyang Wang, Jingfeng Yang, and Heng Ji. 2025. [Scaling laws for predicting downstream performance in llms](#). *Trans. Mach. Learn. Res.*, 2025.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *NeurIPS 2019*, pages 7057–7067.
- Patrick Fernandes, Behrooz Ghorbani, Xavier Garcia, Markus Freitag, and Orhan Firat. 2023. Scaling laws for multilingual neural machine translation. In *ICML 2023*, pages 10053–10071.
- Yifei He, Alon Benhaim, Barun Patra, Praneetha Vaddamanu, Sanchit Ahuja, Parul Chopra, Vishrav Chaudhary, Han Zhao, and Xia Song. 2025. Scaling laws for multilingual language models. In *ACL 2025 (Findings)*, pages 4257–4273.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, and 3 others. 2022. Training compute-optimal large language models. In *NeurIPS 2022*, pages 30016–30030.
- Haoyang Huang, Tianyi Tang, Dongdong Zhang, Xin Zhao, Ting Song, Yan Xia, and Furu Wei. 2023. Not all languages are created equal in llms: Improving multilingual capability by cross-lingual-thought prompting. In *EMNLP 2023 (Findings)*, pages 12365–12394.
- Berivan Isik, Natalia Ponomareva, Hussein Hazimeh, Dimitris Paparas, Sergei Vassilvitskii, and Sanmi Koyejo. 2025. Scaling laws for downstream task performance in machine translation. In *ICLR 2025*, pages 88769–88790.
- Masaru Isonuma and Ivan Titov. 2024. Unlearning traces the influential training data of language models. In *ACL 2024*, pages 6312–6325.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *Preprint*, arXiv:2001.08361.
- Viet Lai, Chien Nguyen, Nghia Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan Rossi, and Thien Nguyen. 2023. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. In *EMNLP 2023: System Demonstrations*, pages 318–327.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona T. Diab, and 2 others. 2022. Few-shot learning with multilingual generative language models. In *EMNLP 2022*, pages 9019–9052.
- Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastopoulos, Patrick Littell, and Graham Neubig. 2019. Choosing transfer languages for cross-lingual learning. In *ACL 2019*, pages 3125–3135.

- Zehui Lin, Xiao Pan, Mingxuan Wang, Xipeng Qiu, Jiangtao Feng, Hao Zhou, and Lei Li. 2020. Pre-training multilingual neural machine translation by leveraging alignment information. In *EMNLP 2020*, pages 2649–2663.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Trans. Assoc. Comput. Linguistics*, 8:726–742.
- Shayne Longpre, Sneha Kudugunta, Niklas Muenighoff, I-Hung Hsu, Isaac Caswell, Alex Pentland, Serkan Arik, Chen-Yu Lee, and Sayna Ebrahimi. 2026. ATLAS: Adaptive transfer scaling laws for multilingual pretraining, finetuning, and decoding the curse of multilinguality. In *ICLR 2026*, pages 31839–31867.
- Nasrin Mostafazadeh, Michael Roth, Annie Louis, Nathanael Chambers, and James Allen. 2017. LSDSem 2017 shared task: The story cloze test. In *LSDSem 2017*, pages 46–51.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben alal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. The FineWeb datasets: Decanting the web for the finest text data at scale. In *NeurIPS 2024*, pages 30811–30849.
- Guilherme Penedo, Hynek Kydlíček, Vinko Sabolčec, Bettina Messmer, Negar Foroutan, Amir Hossein Kargaran, Colin Raffel, Martin Jaggi, Leandro Von Werra, and Thomas Wolf. 2025. [FineWeb2: One pipeline to scale them all – adapting pre-training data processing to every language](#). *Preprint*, arXiv:2506.20920.
- Jonas Pfeiffer, Ivan Vulic, Iryna Gurevych, and Sebastian Ruder. 2020. MAD-X: an adapter-based framework for multi-task cross-lingual transfer. In *EMNLP 2020*, pages 7654–7673.
- Haoran Que, Jiaheng Liu, Ge Zhang, Chenchen Zhang, Xingwei Qu, Yinghao Ma, Feiyu Duan, Zhiqi Bai, Jiakai Wang, Yuanxing Zhang, Xu Tan, Jie Fu, Jiamang Wang, Lin Qu, Wenbo Su, and Bo Zheng. 2024. D-CPT law: Domain-specific continual pre-training scaling law for large language models. In *NeurIPS 2024*, pages 90318–90354.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. Cross-lingual transfer learning for multilingual task oriented dialog. In *NAACL-HLT 2019*, pages 3795–3805.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. [Megatron-LM: Training multi-billion parameter language models using model parallelism](#). *Preprint*, arXiv:1909.08053.
- Mustafa Shukor, Louis Bethune, Dan Busbridge, David Grangier, Enrico Fini, Alaaeldin El-Nouby, and Pierre Ablin. 2025. Scaling laws for optimal data mixtures. In *NeurIPS 2025*, pages 129554–129579.
- Alexey Tikhonov and Max Ryabinin. 2021. It’s all in the heads: Using attention heads as a baseline for cross-lingual transfer in commonsense reasoning. In *ACL-IJCNLP 2021 (Findings)*, pages 3534–3546.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Fali Wang, Zhiwei Zhang, Xianren Zhang, Zongyu Wu, Tzuhao Mo, Qihao Lu, Wanjiang Wang, Rui Li, Junjie Xu, Xianfeng Tang, Qi He, Yao Ma, Ming Huang, and Suhang Wang. 2025a. [A comprehensive survey of small language models in the era of large language models: Techniques, enhancements, applications, collaboration with llms, and trustworthiness](#). *ACM Trans. Intell. Syst. Technol.*, 16(6):145:1–145:87.
- Jiachen T. Wang, Prateek Mittal, Dawn Song, and Ruoxi Jia. 2025b. Data Shapley in one training run. In *ICLR 2025*, pages 12358–12395.
- Jiayi Wang, Yao Lu, Maurice Weber, Max Ryabinin, David Ifeoluwa Adelani, Yihong Chen, Raphael Tang, and Pontus Stenetorp. 2025c. Multilingual language model pretraining using machine-translated data. In *EMNLP 2025*, pages 28087–28107.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V. Le, Tengyu Ma, and Adams Wei Yu. 2023. DoReMi: Optimizing data mixtures speeds up language model pretraining. In *NeurIPS 2023*, pages 69798–69818.
- Yuemei Xu, Ling Hu, Jiayi Zhao, Zihan Qiu, Kexin Xu, Yuqi Ye, and Hanwen Gu. 2025. [A survey on multilingual large language models: corpora, alignment, and bias](#). *Frontiers Comput. Sci.*, 19(11):1911362.
- Jiasheng Ye, Peiju Liu, Tianxiang Sun, Jun Zhan, Yunhua Zhou, and Xipeng Qiu. 2025. Data mixing laws: Optimizing data mixtures by predicting language modeling performance. In *ICLR 2025*, pages 82263–82287.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In *ACL 2019*, pages 4791–4800.
- Yiran Zhao, Wenxuan Zhang, Guizhen Chen, Kenji Kawaguchi, and Lidong Bing. 2024. How do large language models handle multilingualism? In *NeurIPS 2024*, pages 15296–15319.
- Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. 2021. [A comprehensive survey on transfer learning](#). *Proc. IEEE*, 109(1):43–76.

A More Technical Details

A.1 Theoretical Properties of the SV

The SV uniquely satisfies the following properties.

- *Efficiency*: The payoff (i.e., loss reduction) $v_j^G(\mathcal{K})$ for each task $\mathcal{V}_j$ is fully distributed among all the languages $\mathcal{K}$, i.e., $\forall j, \sum_{i=1}^K \phi_{i,j}^{SV}(\mathcal{G}) = v_j^G(\mathcal{K})$.
- *Linearity*: The SVs of each language for different evaluation tasks are linearly additive. Let $\phi_{i,j+j'}^{SV}$ denotes the Shapley value of language i for task $\mathcal{V}_{j+j'} = \mathcal{V}_j \cup \mathcal{V}_{j'}$. Then, for any language i , and for any two tasks $\mathcal{V}_j$ and $\mathcal{V}_{j'}$, we have $\phi_{i,j+j'}^{SV}(\mathcal{G}) = \phi_{i,j}^{SV}(\mathcal{G}) + \phi_{i,j'}^{SV}(\mathcal{G})$. That means that the contribution of language i to a combined task $\mathcal{V}_{j+j'}$ should be the sum of its contributions to the individual tasks $\mathcal{V}_j$ and $\mathcal{V}_{j'}$.
- *Symmetry*: If two languages i, i' are equivalent for improving model performance, then their SVs are identical. Formally, if $v_j^G(S \cup \{i\}) = v_j^G(S \cup \{i'\})$ for any language subset $S \subseteq \mathcal{K}$, then $\phi_{i,j}^{SV} = \phi_{i',j}^{SV}$.
- *Null player*: If a language i never improves the model performance when its data is included for pretraining, its contribution should be zero. That is, if $v_j^G(S \cup \{i\}) = v_j^G(S)$ for any language subset $S \subseteq \mathcal{K}$, then $\phi_{i,j}^{SV} = 0$.

A.2 ShapleyLaw Mixture Optimization

Optimal mixture ratios for language preferences.

In practice, multilingual pre-training is often guided by an application-specific evaluation suite spanning multiple languages, where different languages may carry different importance. To model this, suppose we have nonnegative preference weights $\{w_j\}_{j=1}^K$ (e.g., proportional to user demand or benchmark priority) and consider the weighted objective

$$\mathcal{J}(\mathbf{p}) = \sum_{j=1}^K w_j \mathcal{L}_j(N, D, \mathbf{p})$$

where $\mathcal{L}_j$ is the test loss for target language j . Introducing ShapleyLaw into Eq. A.2 yields

$$\mathcal{J}(\mathbf{p}) = \sum_{j=1}^K w_j \cdot C_j \cdot (\Theta_j)^{-\gamma_j} \quad (2)$$

where $C_j \triangleq E_j + \frac{A_j}{N^{\alpha_j}} + \frac{B_j}{D^{\beta_j}}$, which makes explicit how mixture influences the weighted loss through Θ_j .

Convexity and optimality conditions. For $\gamma_j > 0$, the function $x \mapsto x^{-\gamma_j}$ is convex over $x > 0$; since Θ_j is affine in $\mathbf{p}$, $\mathcal{J}(\mathbf{p})$ is a convex objective over the simplex whenever $\Theta_j > 0$ for all j . Therefore, the mixture selection problem

$$\min_{\mathbf{p} \in \Delta^{K-1}} \mathcal{J}(\mathbf{p})$$

admits a global optimum and can be solved efficiently. Using a Lagrangian with simplex constraints and Karush-Kuhn-Tucker (KKT) conditions, the gradient of Eq. 2 is

$$\frac{\partial \mathcal{J}}{\partial p_i} = - \sum_{j=1}^K w_j C_j \gamma_j \phi_{i,j}^{NSV} (\Theta_j)^{-(\gamma_j+1)}.$$

At an interior optimum (all $p_i > 0$), these gradients are equal up to a shared Lagrange multiplier, implying an intuitive *marginal utility equalization* principle: optimal mixtures allocate probability mass so that the expected reduction in weighted loss per unit increase of p_i is balanced across languages.

Closed-form solution in a diagonal (self-dominant) approximation. A useful and interpretable special case arises when each target language is primarily influenced by its own data, i.e., $\phi_{i,j}^{NSV}$ is approximately diagonal so that $\Theta_j \approx p_j r_j$ with $r_j \triangleq \phi_{j,j}^{NSV}$. Then Eq. 2 becomes separable:

$$\mathcal{J}(\mathbf{p}) \approx \sum_{j=1}^K w_j C_j (p_j r_j)^{-\gamma_j}.$$

The KKT conditions yield a closed-form optimal ratio:

$$p_i^* = \frac{\left(w_i C_i \gamma_i r_i^{-\gamma_i}\right)^{\frac{1}{\gamma_i+1}}}{\sum_{m=1}^K \left(w_m C_m \gamma_m r_m^{-\gamma_m}\right)^{\frac{1}{\gamma_m+1}}}. \quad (3)$$

First-order correction under cross-lingual transfer. For realistic multilingual settings, cross-lingual transfer introduces non-negligible off-diagonal terms in ϕ^{NSV} . Let the Shapley matrix be decomposed as

$$\phi^{NSV} = I + E$$

where E captures cross-lingual interactions with relatively small off-diagonal entries. Then

$$\Theta = \phi^{NSV} \mathbf{p} = \mathbf{p} + E\mathbf{p}.$$

Expanding $\Theta_i^{-(\gamma_i+1)}$ to first order around the diagonal solution $p^{(0)}$ from Eq. 3 yields

$$p_i^* \approx p_i^{(0)} \left(1 - \frac{1}{\gamma_i + 1} \sum_{j,j \neq i} \phi_{i,j}^{NSV} \right) \quad (4)$$

followed by renormalization to satisfy the simplex constraint. The correction term depends on the column sum $\sum_{j \neq i} \phi_{i,j}^{NSV}$, which measures how strongly language i contributes to other languages through cross-lingual transfer. Languages with larger column sums effectively provide training signal to many other languages, and therefore require less direct sampling mass in the optimal mixture. Conversely, languages with weaker transfer connections receive higher allocation to maintain balanced marginal utility across tasks.

In practice, the available corpus size for each language may be limited. Let $\bar{D}_i$ denote the maximum usable token budget for language i , and let the total training budget be D . Then the mixture ratio must satisfy the additional box constraint

$$p_i \leq \bar{p}_i \quad \text{where} \quad \bar{p}_i = \frac{\bar{D}_i}{D}. \quad (5)$$

The mixture optimization problem thus becomes

$$\min_{\mathbf{p} \in \Delta^{K-1}} \mathcal{J}(\mathbf{p}) \quad \text{s.t.} \quad 0 \leq p_i \leq \bar{p}_i.$$

Because $\mathcal{J}(\mathbf{p})$ is convex, the constrained problem remains convex and can be characterized via KKT conditions with complementary slackness. However, a simple and effective approximation can be derived by projecting the unconstrained optimum onto the feasible region.

Clipped closed-form approximation. Let $p_i^{(0)}$ denote the unconstrained solution from Eq. 3 (or its cross-lingual correction in Eq. 4). Define the *active set* of data-limited languages

$$\mathcal{A} = \{i \mid p_i^{(0)} > \bar{p}_i\}.$$

For languages in $\mathcal{A}$, the optimal solution saturates the constraint:

$$p_i^* = \bar{p}_i, \quad i \in \mathcal{A}.$$

The remaining probability mass is

$$M = 1 - \sum_{i \in \mathcal{A}} \bar{p}_i$$

which is redistributed among the unconstrained languages $\mathcal{F} = \{1, \dots, K\} \setminus \mathcal{A}$ proportionally to their original optimal ratios:

$$p_i^* \approx M \cdot \frac{p_i^{(0)}}{\sum_{m \in \mathcal{F}} p_m^{(0)}}, \quad i \in \mathcal{F}. \quad (6)$$

If redistribution causes additional violations ($p_i^* > \bar{p}_i$ for some $i \in \mathcal{F}$), the clipping-and-redistribution step can be repeated. Since each iteration adds at least one index to $\mathcal{A}$, the procedure converges in at most K steps.

This approximation corresponds to projecting the unconstrained optimum onto the truncated simplex defined by Eq. 5. Intuitively, languages with insufficient data are trained to exhaustion, while the remaining training budget is allocated among data-rich languages according to their relative marginal utility. Notably, the structure of Eq. 6 preserves the optimal trade-off encoded by Shapley-aware scaling laws while ensuring practical feasibility under corpus-size limitations.

A.3 In-Run Data Shapley for SV Approximation

Why In-Run Data Shapley approximation is needed. For a multilingual pretraining game (MPG), the Shapley value (SV) of language i on task j is defined as

$$\phi_{i,j}^{SV}(\mathcal{G}) = \sum_{S \subseteq \mathcal{K} \setminus \{i\}} W(S) \cdot \Delta_{i,j}^{\mathcal{G}}(S)$$

where $W(S) = \frac{|S|!(K-|S|-1)!}{K!}$, and $\Delta_{i,j}^{\mathcal{G}}(S) = v_j^{\mathcal{G}}(S \cup \{i\}) - v_j^{\mathcal{G}}(S)$. Exact computation requires evaluating the payoff of all language coalitions, which entails pre-training $2^K - 1$ models for K languages.

To avoid repeated pre-training, we replace coalition-level retraining with an In-Run Data Shapley (IRDS) approximation. Instead of asking how the final payoff changes when language i is added to an independently trained coalition, IRDS measures how language i contributes to the validation loss reduction along a specific multilingual pre-training trajectory. Thus, the resulting value directly attributes the cross-lingual transfer observed in the actual training run used for SV estimation.

Local payoff in IRDS. Consider a single multilingual pre-training run with checkpoints $\{w_0, w_1, \dots, w_T\}$. At training step t , let B_t be the sampled mini-batch and let $B_{t,i} \subseteq B_t$ denote the

subset of examples whose language is i . For a language subset $S \subseteq \mathcal{K}$, we define a counterfactual one-step update using only examples from languages in S :

$$\tilde{w}_{t+1}(S) = w_t - \eta_t \sum_{r \in S} \sum_{z \in B_{t,r}} \nabla \ell(w_t, z),$$

where η_t is the learning rate and $\ell(w, z)$ is the training loss.

For a target validation language j , let $\mathcal{V}_j$ be its validation set and let

$$\ell_j(w) = \frac{1}{|\mathcal{V}_j|} \sum_{z^{val} \in \mathcal{V}_j} \ell(w, z^{val})$$

denote the validation loss on task j . We define the local payoff of coalition S at step t as the one-step validation loss reduction:

$$u_j^{(t)}(S) = \ell_j(w_t) - \ell_j(\tilde{w}_{t+1}(S)). \quad (7)$$

This local payoff measures how much the current update, if performed using only languages in S , improves performance on target language j .

IRDS SV. The IRDS SV of language i on task j is obtained by computing the SV of i for each local payoff $u_j^{(t)}(\cdot)$ and accumulating it across training:

$$\hat{\phi}_{i,j}^{IRDS} = \sum_{t=0}^{T-1} \phi_i(u_j^{(t)}). \quad (8)$$

By the linearity of SV, this is equivalent to computing the SV of the cumulative payoff $\sum_{t=0}^{T-1} u_j^{(t)}(\cdot)$. Therefore, $\hat{\phi}_{i,j}^{IRDS}$ measures the cumulative contribution of language i to the loss reduction on target language j during the same pre-training run.

First-order approximation. Directly computing $\phi_i(u_j^{(t)})$ still requires enumerating language subsets within each step. We therefore approximate the local payoff in Eq. 7 by a first-order Taylor expansion. Define the validation gradient for target language j as

$$g_{t,j}^{val} = \nabla \ell_j(w_t),$$

and define the aggregated training gradient of language i at step t as

$$g_{t,i} = \sum_{z \in B_{t,i}} \nabla \ell(w_t, z).$$

Algorithm 1 In-Run Data Shapley Approximation

Require: Language set $\mathcal{K}$, uniform multilingual pre-training data, validation sets $\{\mathcal{V}_j\}_{j=1}^K$, learning rates $\{\eta_t\}_{t=0}^{T-1}$

Ensure: Approximate IRDS Shapley values $\{\hat{\phi}_{i,j}^{IRDS}\}_{i,j \in \mathcal{K}}$

```

1: for  $i = 1$  to  $K$  do
2:   for  $j = 1$  to  $K$  do
3:      $\hat{\phi}_{i,j}^{IRDS} \leftarrow 0$ 
4:   end for
5: end for
6: for  $t = 0$  to  $T - 1$  do
7:   Sample multilingual mini-batch  $B_t$ 
8:   for  $i = 1$  to  $K$  do
9:      $B_{t,i} \leftarrow \{z \in B_t : \text{lang}(z) = i\}$ 
10:     $g_{t,i} \leftarrow \sum_{z \in B_{t,i}} \nabla \ell(w_t, z)$ 
11:   end for
12:   for  $j = 1$  to  $K$  do
13:      $g_{t,j}^{val} \leftarrow \nabla \ell_j(w_t)$ 
14:     for  $i = 1$  to  $K$  do
15:        $\Delta_{i,j}^{(t)} \leftarrow \eta_t \langle g_{t,j}^{val}, g_{t,i} \rangle$   $\triangleright$  First-order IRDS
16:        $\hat{\phi}_{i,j}^{IRDS} \leftarrow \hat{\phi}_{i,j}^{IRDS} + \Delta_{i,j}^{(t)}$ 
17:     end for
18:   end for
19:   Update  $w_{t+1}$  using the full multilingual batch  $B_t$ 
20: end for
21: return  $\{\hat{\phi}_{i,j}^{IRDS}\}_{i,j \in \mathcal{K}}$ 

```

Then,

$$u_j^{(t)}(S) \approx \eta_t \sum_{r \in S} \langle g_{t,j}^{val}, g_{t,r} \rangle.$$

Since the approximated utility is additive over languages, its Shapley value has a closed form:

$$\phi_i(u_j^{(t)}) \approx \eta_t \langle g_{t,j}^{val}, g_{t,i} \rangle.$$

Thus, the first-order IRDS Shapley estimator is

$$\hat{\phi}_{i,j}^{IRDS-1} = \sum_{t=0}^{T-1} \eta_t \langle g_{t,j}^{val}, g_{t,i} \rangle. \quad (9)$$

Intuitively, language i receives a positive contribution at step t when its training gradient is aligned with the validation gradient of target language j , because this update decreases the validation loss on task j . Conversely, a negative value indicates that the update from language i increases the validation loss on task j .

The first term measures the direct alignment between language i and the target language j . IRDS also provides a second-order approximation that captures interactions between language i and the full multilingual batch under the local curvature of the validation loss. In practice, we use the first-order estimator by default because it is simpler, more stable, and incurs minimal additional cost. The second-order estimator can be used when within-batch language interactions are important.

Normalization. After obtaining the IRDS Shapley matrix $\{\hat{\phi}_{i,j}^{IRDS}\}_{i,j \in \mathcal{K}}$, we apply the same normalization as in the main text:

$$\hat{\phi}_{i,j}^{NSV} = \exp \left(\hat{\phi}_{i,j}^{IRDS} - \max_{i'} \hat{\phi}_{i',j}^{IRDS} \right).$$

The normalized values are then used as cross-lingual transfer strengths in ShapleyLaw:

$$\Theta_j = \sum_{i \in \mathcal{K}} p_i \cdot \hat{\phi}_{i,j}^{NSV}.$$

Computational complexity. Unlike exact SV computation, IRDS requires only one multilingual pre-training run for SV estimation. During this run, the method accumulates language-level gradient contributions at each update step. The first-order estimator requires gradient dot-products between the validation gradient of each target language and the aggregated training gradient of each source language. These dot-products can be computed efficiently during backpropagation and do not require pre-training additional coalition models. Therefore, the number of pretrained models is reduced from $O(2^K)$ for exact SV to $O(1)$ pre-training runs for IRDS.

A.4 Clarification on Baseline Selection

The goal of this paper is to develop a multilingual scaling law for performance prediction and language mixture optimization. Accordingly, appropriate baselines are methods that can optimize language mixtures for large-scale training configurations *from observations made under small-scale training configurations*. Uniform Sampling, Smoothed Sampling, FamilyLaw Optimization, and Empirical Mixtures all satisfy this criterion, and we therefore adopt them as our mixture optimization baselines.

In contrast, some recent data mixture optimization methods such as DoReMi (Xie et al., 2023) address a different problem: optimizing data mixtures under a fixed data budget. As a result, they cannot predict optimal mixtures for unseen data scales. Since the core contribution of ShapleyLaw is mixture optimization across both data and model scales, a direct comparison with these methods would be unfair and would not evaluate the capability that ShapleyLaw is designed to provide. We therefore do not include them as baselines.

B Experimental Setup Details

Corpora. We use the English *sample-350BT* subset⁴ from *FineWeb* (Penedo et al., 2024), and the German, French, Spanish, Chinese, Japanese, Korean, Finnish, Croatian, and Standard Malay subsets from *FineWeb-2 v2.0.0* (Penedo et al., 2025). The total tokens are summarized in the following Table 5.

Training Configuration. We pre-train multilingual LLMs with Megatron-LM (Shoeybi et al., 2019) on 8-node 8×H200 clusters. We adopt the open-source model architectures and training hyperparameters provided by the *llm-jp* project⁵.

Table 4 summarizes the configurations of five transformer language models with increasing non-embedding parameter scales (50M, 243M, 684M, 1460M, and 3171M). As the scale grows, the number of layers and hidden dimensions increase, while the number of attention heads and context length remain constant. Larger models allocate substantially more parameters to both embedding and non-embedding components, with non-embedding parameters growing more rapidly, indicating that most additional capacity comes from deeper and wider transformer blocks rather than the embedding layer.

We adopt a decoder-only Transformer with RMSNorm, RoPE, and SwiGLU. Tokenization uses the LLaMA-2-style `llm-jp-tokenizer-100k.ver3.0b1`⁶ tokenizer with a 100k multilingual vocabulary. Models are trained with learning rate 3×10^{-4} , sequence length 4,096, and global batch size 512 using Adam ($\beta_1=0.9$, $\beta_2=0.95$), cosine decay with warmup, and standard stabilization techniques (weight decay, gradient clipping, Z-loss). Training efficiency is improved through bf16 precision, FlashAttention, fused kernels, and Megatron-LM parallelism, thereby enabling stable large-scale multilingual pre-training.

Mixture Optimization Baselines. We consider the following baselines for mixture optimization.

- *Uniform Sampling:* All languages are assigned equal proportions. This balanced heuristic serves as the simplest non-parametric baseline.
- *Smoothed Sampling* (Conneau and Lample, 2019): This strategy downweights dominant lan-

⁴A subset randomly sampled from the whole dataset of around 350B gpt2 tokens (388GB).

⁵<https://huggingface.co/llm-jp>

⁶<https://github.com/llm-jp/llm-jp-tokenizer>

Params	Layers	Hidden size	Heads	Context length	Embedding parameters	Non-embedding parameters
150M (50M)	12	512	8	4,096	101,874,688	50,344,448
440M (243M)	16	1,024	8	4,096	203,749,376	243,303,424
980M (684M)	20	1,536	8	4,096	305,624,064	684,258,816
1.8B (1460M)	24	2,048	16	4,096	407,498,752	1,459,718,144
3.7B (3171M)	28	3,072	24	4,096	611,248,128	3,171,068,928

Table 4: Model configurations and parameter breakdown.

Family	Languages	Tokens (B)
Germanic	English / German	373 / 450
Romance	French / Spanish	340 / 397
Sino-Tibetan	Chinese	788
Japonic	Japanese	281
Korean	Korean	52
Uralic	Finnish	54
Slavic	Croatian	29
Austronesian	Standard Malay	12

Table 5: **Corpora used in our experiments.** The study covers ten languages from eight language families, the available token counts are listed for each language.

guages while still reflecting corpus size through a smoothing parameter α , i.e., $p_i = \frac{(\bar{D}_i/D)^\alpha}{\sum_{i'=1}^K (\bar{D}_{i'}/D)^\alpha}$. We adopt the common setting $\alpha = 0.5$ for all the related experiments.

- *FamilyLaw Optimization* (He et al., 2025): This strategy first optimizes family-level mixture ratios using FamilyLaw and then allocates language-level mixture ratios using the Smoothed Sampling (see Appendix F in (He et al., 2025) for details).
- *Empirical Mixtures*: As shown in Table 6, empirical mixture ratios are drawn based on statistics of existing open-source corpora, including *Common Crawl*⁷ and *FineWeb-2*. Note that the available Korean corpus is larger than the Finnish corpus in disk size (213.43,GB vs. 143.03,GB) of *FineWeb-2*, whereas the Finnish corpus contains slightly more tokens (54B vs. 52B), likely due to our tokenizer. We include this setting to evaluate whether the proposed method improves upon corpus-native sampling strategies.

Downstream Tasks. We evaluate all seven pre-training languages on publicly available benchmarks covering linguistic competence, natural language understanding, and commonsense reasoning. The tasks are: *HellaSwag* (commonsense completion) (Zellers et al., 2019; Lai et al., 2023),

⁷<https://commoncrawl.github.io/cc-crawl-statistics/plots/languages.html>

XStoryCloze (cross-lingual story ending prediction) (Lin et al., 2022; Mostafazadeh et al., 2017), and *XWinograd* (multilingual coreference resolution) (Tikhonov and Ryabinin, 2021). Following standard multilingual evaluation practice, we use a *five-shot* setting and report accuracy.

C Additional Experimental Results

C.1 Additional Experiments on Performance Prediction.

To further validate the robustness of the proposed Shapley-based scaling law, we report additional fitting results on four target languages: EN, FR, DE, and KO. For each target language, we fit the proposed relation using both curated multilingual mixtures and randomly sampled mixtures, and evaluate whether the aggregate cross-lingual transfer Θ consistently explains the variation in held-out test CE loss. Representative examples for EN, and FR are shown in Figure 7: the top row presents fits across predefined mixture settings and model scales, while the bottom row shows results under random mixture sampling. In both cases, empirical losses align closely with the predicted curves, and held-out points remain near the fitted trend, indicating that the scaling law generalizes beyond the mixtures used for fitting.

The same pattern is observed for DE, and KO (see Figure 8), where losses across different model sizes and sampling strategies can also be well described by the ShapleyLaw. Across these additional languages, the fitted curves remain smooth and monotonic with respect to Θ , and the resulting high R^2 values range from 0.985 to 0.998. These results provide further evidence that the proposed parameterization is not specific to a single language or mixture family, but instead captures a common regularity in multilingual pre-training. Overall, the additional results strengthen the conclusion that aggregate cross-lingual transfer serves as an effective mixture variable for modeling per-language scaling behavior across diverse target languages.

Dataset / Language	en	de	fr	es	zh	ja	ko	fi	hr	ms
Common Crawl	61.4%	8.3%	6.4%	6.4%	7.2%	8.2%	1.2%	0.5%	0.3%	0.1%
FineWeb-2	80.6%	3.3%	2.3%	2.7%	6.7%	3.1%	0.5%	0.4%	0.3%	0.1%

Table 6: **Normalized statistics across languages in open-source corpora.** The statistics are derived from the official Common Crawl statistics and, for FineWeb-2, estimated based on the reported disk sizes (Penedo et al., 2025). Specifically, the disk sizes of the FineWeb-2 subsets for {en, de, fr, es, zh, ja, ko, fi, hr, ms} are {33 TB, 1.51 TB, 1.11 TB, 1.32 TB, 2.42 TB, 1.50 TB, 213.43 GB, 143.03 GB, 35.91 GB, 31.94 GB}, respectively.




Figure 7: **Additional fitting results for Shapley-based scaling laws on <EN,FR>.** **Top row:** Fits on curated multilingual mixtures for EN, and FR across four model scales {50M, 243M, 684M, 1460M}; gray circles denote training points and yellow diamonds denote held-out points. **Bottom row:** Fits under randomly sampled mixtures under model scale $N=243M$. For both languages, the fitted curves remain smooth and monotonic with respect to Θ , and held-out points closely follow predicted trends, demonstrating strong generalization to unseen mixtures.

Task	ℓ	$-\ell$	$\log \ell$	$\exp(-\ell)$	$1/\ell$
HellaSwag	0.994	0.994	0.995	0.996	0.996
XWinograd	0.971	0.971	0.975	0.980	0.978
XStorycloze	0.947	0.947	0.943	0.936	0.939

Table 7: **Model Selection.** Model comparison for $s = a \cdot g(\ell) + b$ using different transformations of loss. Numbers denote R^2 . Across tasks except XStorycloze, $g(\ell) = \exp(-\ell)$ consistently achieves the highest explanatory power, suggesting that accuracy scales approximately linearly with token-level probability $p = \exp(-\ell)$.

C.2 Additional Experiments on Downstream Tasks

Results of Model Selection. To determine an appropriate form for $g(\ell)$, we evaluate five candidate transformations: $\{\ell, -\ell, \log \ell, \exp(-\ell), 1/\ell\}$. For

each candidate, we fit a linear model of the form $s = a \cdot g(\ell) + b$ that maps loss values to predicted scores. Table 7 reports the resulting model comparisons. Among these transformations, $\exp(-\ell)$ always exhibits the strongest explanatory power, achieving the highest R^2 and consistently providing the best fit to the observed relationship between loss and score. These results indicate that an exponential transformation of loss most effectively captures the empirical mapping between the two quantities, and we therefore adopt $g(\ell) = \exp(-\ell)$ in the subsequent experiments.

Fitting Results on Downstream Tasks. We next model the relationship between pre-training loss ℓ and downstream performance using the func-




Figure 8: **Additional fitting results for Shapley-based scaling laws on <DE,KO>.** **Top row:** Fits on curated multilingual mixtures for DE and KO across the same four model scales. **Bottom row:** Fits under randomly sampled mixtures with model scale $N=243M$. For both languages, the fitted curves remain smooth and monotonic with respect to aggregate cross-lingual transfer Θ , and held-out points closely follow predicted trends, demonstrating strong generalization to unseen mixtures.

Task	Pearson r	p -value	Best $g(\ell)$	R^2	a	b
HellaSwag	-0.997	8.15×10^{-7}	$\exp(-\ell)$	0.996	4.030	0.098
XWinograd	-0.985	4.92×10^{-5}	$\exp(-\ell)$	0.980	4.824	0.382
XStorycloze	-0.973	2.25×10^{-4}	$\exp(-\ell)$	0.980	3.034	0.433

Table 8: **Fitted Results (English Downstream Tasks).** Pearson correlation between pretraining loss ℓ and few-shot accuracy. All tasks exhibit strong negative correlation ($|r| > 0.90$), indicating that lower loss consistently corresponds to higher downstream performance. Correlation between pretraining loss ℓ and few-shot accuracy, and best-fit results for $\hat{s} = ag(\ell) + b$. Across all tasks, $\exp(-\ell)$ consistently achieves the highest R^2 .

tional form $s = a \cdot \exp(-\ell) + b$. Table 8 reports the fitting results across six English downstream tasks, comparing pre-training loss with few-shot (5-shot) accuracy. Across all tasks, we observe consistently strong *negative correlations*, with Pearson coefficients ranging from -0.909 to -0.997 (p -value < 0.01 in all cases), indicating that lower pre-training loss is strongly associated with higher downstream performance. The fitted models also achieve high explanatory power, with R^2 values between 0.843 and 0.996, suggesting that the exponential transformation $\exp(-\ell)$ captures the empirical relationship between loss and accuracy well. These results support the use of $s = a \cdot \exp(-\ell) + b$ as a simple yet effective mapping for predicting down-

stream task performance from pre-training loss.

Additional Fitting Results on Other Language Tasks. We conduct additional experiments on downstream tasks for Chinese and Japanese to explore the feasibility of ShapleyLaw. In addition, we fit results of downstream performance vs test loss of 50 models ($N = 243M$, $D = 100B$) with randomly sampled language mixtures. As shown in Figure 9, we report fitting results for Chinese, and Japanese.

C.3 Additional Experiments on SV Approximation

Experiment Design. Exact SV computation requires evaluating all non-empty language coalitions,




Figure 9: **Left.** Downstream performance versus test cross-entropy (CE) loss on Chinese tasks (*XCOPA_zh* and *XStoryCloze_zh*). **Middle.** Downstream performance versus test CE loss on the Japanese task (*XWinograd*). **Right.** Downstream performance for English versus test CE loss across 50 models ($N = 243\text{M}$, $D = 100\text{B}$) trained with randomly sampled data mixtures.




Figure 10: **Fitting results with exact and IRDS-approximated SVs.** To reduce the workload of SV computation, we use the IRDS approximation by pretraining a small-scale model ($N = 50\text{M}$, $D = 50\text{B}$) with a uniform mixture ratio over seven languages. The resulting IRDS SVs produce fitting curves that closely match those obtained with exact SVs across target languages.

Comparison	Cosine Similarity	Pearson
Exact SV vs. IRDS SV	0.9584	0.9390

Table 9: Similarity comparison between exact SV and IRDS SV.

Target	de	en	fr	es	ja	ko	zh
Exact	0.015	0.004	0.015	0.014	0.006	0.009	0.015
IRDS	0.015	0.003	0.036	0.017	0.008	0.023	0.013

Table 10: Comparison of corresponding PE_{fit} values for the fitted results with exact and IRDS SVs.

resulting in $2^K - 1$ pretrained models for K languages. This cost grows exponentially with K and becomes prohibitive in large multilingual settings. To assess whether ShapleyLaw remains effective without exact SVs, we evaluate IRDS as a scalable approximation that estimates language contributions from a single model pretrained with a uniform multilingual mixture. We use the same MPG as in the main experiments. Specifically, we pretrain a small-scale model with $N = 50\text{M}$ and $D = 50\text{B}$ using a uniform mixture over seven languages, estimate the IRDS matrix, and construct the aggregate transfer variable Θ_j . We then fit ShapleyLaw under the same settings as the exact-SV variant, isolating the effect of SV approximation on both transfer estimation and scaling-law fitting quality.

Similarity between Exact and IRDS SVs. Table 9 reports the similarity between exact SVs and IRDS-estimated SVs. The two SV matrices are highly aligned, achieving a cosine similarity of 0.9584 and a Pearson correlation of 0.9390. This indicates that IRDS preserves the global geometry of the cross-lingual transfer matrix: languages that receive large contributions under exact SV estimation also tend to receive large contributions under IRDS. Therefore, although IRDS does not explicitly enumerate all language coalitions, it provides a close approximation to the relative transfer strengths needed by ShapleyLaw.

Fitting Results. Figure 10 shows that the IRDS-based fitting curves closely match those obtained with exact SVs across target languages, indicating that ShapleyLaw is robust to approximation noise in the SV matrix. The close alignment of the curves suggests that IRDS preserves this transfer structure sufficiently for loss prediction. Table 10 further confirms this trend. IRDS achieves fitting errors comparable to exact SVs for most languages, matching or slightly improving the exact-SV results for German, Chinese, and English. The degradation is modest for Japanese and Spanish but more pronounced for French and Korean. Nevertheless, the




Figure 11: **Additional fitting results for ShapleyLaw at a larger model scale.** We extend the evaluation to $N = 3171M$ and $D = 100B$. Fits are shown for DE, EN, FR, ES, JA, and KO across model scales $\{50M, 243M, 684M, 1460M, 3171M\}$. Gray circles denote fitting points and yellow diamonds denote held-out points.

PEs remain low, and the fitted trends are broadly consistent with the exact-SV results.

Discussion. These results show that exact SV computation is not essential for ShapleyLaw’s practical effectiveness. In the seven-language setting, exact SV estimation requires $2^7 - 1 = 127$ pre-trained models, whereas IRDS requires only one. Despite this large cost reduction, IRDS achieves SV estimates with high similarity to exact SVs and yields comparable fitting quality for most languages. The remaining gaps mainly reflect approximation noise and do not materially affect the fitted scaling behavior, supporting ShapleyLaw’s scalability to larger multilingual settings.

C.4 Additional Experiments on Larger Model Scale

Experiment Design. To further evaluate the extrapolation ability of ShapleyLaw beyond the model scales used in the main experiments, we conduct additional experiments at a larger model size. Specifically, we pretrain models with $N = 3171M$ and $D = 100B$ under four empirical multilingual mixtures: *CC*, *FW2*, *L2*, and *JPv4*. These mixtures correspond to commonly used multilingual data distributions and provide a realistic evaluation setting for large-scale multilingual pretraining.

We treat all $N = 3171M$ results as held-out points and fit ShapleyLaw using the smaller-scale observations. This setting directly tests whether the learned scaling relationship can generalize to a larger model scale that is not used for fitting. We evaluate six target languages, including DE, EN, FR,

ES, JA, and KO. For each target language, the input variable remains the SV-based aggregate cross-lingual transfer Θ_j , while the scaling law jointly captures the effects of model size, data size, and language mixture.

Results. Figure 11 shows that ShapleyLaw accurately predicts the held-out $N = 3171M$ test losses across all six languages. The held-out points generally lie close to the fitted curves, indicating that the learned relationship between aggregate cross-lingual transfer and test loss remains stable when extrapolated to a larger model scale.

ShapleyLaw achieves low held-out prediction errors across languages: $PE_{test} = 0.008$ for German, 0.004 for English, 0.008 for French, 0.013 for Spanish, 0.008 for Japanese, and 0.018 for Korean. These values are comparable to the fitting errors obtained on smaller-scale models, suggesting that the proposed law does not merely interpolate among observed configurations but can also provide reliable predictions for larger-scale models.

The results are encouraging because the held-out points combine two sources of distribution shift: they use a larger model size and empirical multilingual mixtures. Despite this harder setting, ShapleyLaw preserves accurate predictions across typologically diverse target languages. This provides evidence that SV-based aggregate transfer is a stable predictor of multilingual test loss and that ShapleyLaw can support practical mixture evaluation at larger model scales without requiring exhaustive large-scale pretraining for every candidate mixture.

Method	de	en	fr	es	ja	ko	zh
Signed Normalization	0.039	0.004	0.037	0.023	0.029	0.042	0.037
Min–Max Normalization	0.039	0.006	0.039	0.024	0.026	0.035	0.035
Ours	0.016	0.004	0.016	0.011	0.005	0.009	0.013

Table 11: **Comparison of Shapley-value normalization schemes.** Prediction errors obtained with Signed Normalization, Min–Max Normalization, and the softmax-like normalization adopted in ShapleyLaw across different test languages. Lower values indicate better prediction performance. The proposed normalization achieves the lowest or tied-lowest prediction error for all evaluated languages.

	de	en	fr	es	ja	ko	zh
IRDS Error (ℓ_1 distance)	10.605	9.903	9.922	9.916	9.892	10.603	9.888
ΔPE_{fit}	0.000	-0.001	0.021	0.003	0.002	0.014	-0.002

Table 12: **IRDS approximation error and the corresponding change in fitting performance across test languages.** The first row reports the ℓ_1 distance between the IRDS-approximated and reference Shapley values, while the second row reports the change in fitting error after using the IRDS approximation. Positive values of ΔPE_{fit} indicate increased fitting error, whereas negative values indicate a slight improvement.

Excluded Language	Germanic/EN	Germanic/DE	Romance/FR	Romance/ES
PE_{fit}	0.006	0.004	0.006	0.009
PE_{test}	0.009	0.004	0.012	0.005

Table 13: **Leave-one-language-out evaluation within language families.** Experiments are conducted with $N = 50\text{M}$ and $D = 50\text{B}$. For each setting, all fitting mixtures containing the designated language are excluded, and the fitted ShapleyLaw is evaluated on the corresponding held-out within-family mixture.

C.5 Analysis of the Normalization Scheme in ShapleyLaw

ShapleyLaw applies a softmax-like transformation to raw Shapley values in order to obtain normalized transfer coefficients that are comparable across different test-language tasks. In particular, the normalization maps the transfer contribution of each language to $(0, 1]$ and anchors the self-transfer coefficient of each test language at 1. This provides a common scale for modeling and comparing cross-lingual transfer patterns across tasks.

Properties of the underlying cooperative game.

The cooperative game defined in Definition 3.1 assigns each language an independent and fixed amount of training data. Consequently, languages do not compete for a shared data budget when forming different coalitions. Under this formulation, adding a language that provides no additional benefit produces a zero marginal contribution, rather than a negative contribution induced by resource competition. Empirically, the raw Shapley values obtained in all of our experiments are positive, indicating that negative marginal contributions do not arise in the evaluated setting.

Effect of the normalization. The softmax-like transformation does not preserve the classical null-player property after normalization, since every normalized coefficient is strictly positive. This distinction concerns the normalized representation rather than the underlying cooperative game: the null-player property remains well defined at the level of raw marginal contributions and raw Shapley values. In the setting considered by ShapleyLaw, where the observed raw Shapley values are positive, the normalization primarily serves to place transfer effects from different test-language tasks on a common scale while preserving their relative magnitudes.

Comparison with alternative normalization schemes.

We additionally consider Signed Normalization and Min–Max Normalization, which preserve a zero-valued contribution under normalization. Table 11 reports the resulting prediction errors across seven test languages. The softmax-like normalization used in ShapleyLaw consistently yields lower prediction error than the alternatives, with the best or tied-best performance for every evaluated language. These results suggest that anchoring self-transfer at 1 while smoothly rescaling cross-lingual contributions provides a more effective representa-

tion for the predictive modeling objective considered in ShapleyLaw.

C.6 Analysis of Language-wise Variation in Fitting Performance

We further examine the language-wise variation in fitting performance when Shapley values are approximated using IRDS. Table 12 reports both the IRDS approximation error and the corresponding change in fitting error for each test language.

The results do not indicate a clear association between the magnitude of the IRDS approximation error and the resulting change in fitting performance. For example, German exhibits the largest approximation error (10.605) but essentially no increase in fitting error. Korean has a similarly large approximation error (10.603), yet its fitting error increases by 0.014. Conversely, French shows the largest increase in fitting error (0.021), although its approximation error (9.922) is close to those of several other languages. These observations suggest that the uneven language-wise changes in fitting performance cannot be directly explained by the IRDS approximation error alone.

A plausible source of this variation is the sensitivity of the fitting procedure to small perturbations in the fitted data points. The absolute fitting errors are already small across the evaluated languages; consequently, minor variations in the approximated Shapley values can induce relatively different changes in PE_{fit} across individual test languages, even when their overall approximation errors are similar. Importantly, these variations remain small in absolute magnitude and do not exhibit a systematic relationship with the IRDS approximation error.

C.7 Leave-One-Language-Out Generalization within Language Families

We further evaluate whether ShapleyLaw can generalize to a language that is absent from the fitting set while another language from the same family remains represented. We consider two language families: the Germanic family, consisting of English (EN) and German (DE), and the Romance family, consisting of French (FR) and Spanish (ES). All experiments use $N = 50\text{M}$ and $D = 50\text{B}$.

For each setting, we designate one language as the held-out language and remove from the fitting set all mixtures containing that language. The remaining language from the same family is retained during fitting. For example, in the *Germanic/EN* setting, all fitting points involving English are excluded, while

fitting points involving German remain available. The resulting model is then evaluated on the corresponding within-family DE–EN mixture. The other three settings are constructed analogously.

As shown in Table 13, ShapleyLaw maintains low prediction error across all four leave-one-language-out settings. When EN and DE are held out from the fitting set, respectively, the resulting PE_{test} values are 0.009 and 0.004. Similarly, holding out FR and ES yields PE_{test} values of 0.012 and 0.005, respectively. These values remain close to the corresponding fitting errors of 0.006, 0.004, 0.006, and 0.009.

These results provide evidence that ShapleyLaw can generalize to a language excluded from the fitting data when another language from the same family remains represented. This within-family generalization suggests that transfer patterns learned from related languages can provide useful predictive signals for languages with limited fitting observations. Broader evaluation across additional language families and more diverse linguistic relationships is required to characterize the generality of this behavior.