
Languages You Know Influence Those You Learn: Impact of Language Characteristics on Multi-Lingual Text-to-Text Transfer

Benjamin Muller* Deepanshu Gupta† Siddharth Patwardhan† Jean-Philippe Fauconnier†

David Vandyke† Sachin Agarwal†

*INRIA, Paris, France

†Apple, Cupertino, USA

benjamin.muller@inria.fr

{dkg, patwardhan.s, jfauconnier, dvandyke, sachin_agarwal}@apple.com

Abstract

In this work, we analyze a pre-trained mT5 to discover the attributes of cross-lingual connections learned by this model. Through a statistical interpretation framework over 90 language pairs across three tasks, we show that transfer performance can be modeled by a few linguistic and data-derived features. These observations enable us to interpret cross-lingual understanding of the mT5 model. Through these observations, one can favorably choose the best source language for a task, and can anticipate its training data demands. A key finding of this work is that similarity of syntax, morphology and phonology are good predictors of cross-lingual transfer, significantly more than just the lexical similarity of languages. For a given language, we are able to predict zero-shot performance, that increases on a logarithmic scale with the number of few-shot target language data points.

1 Introduction

Multi-lingual language models (LM), such as mBERT [DCLT19], XLM-R [CKG⁺20], mT5 [XCR⁺20], mBART [LGG⁺20], have been remarkably successful in enabling natural language tasks in low-resource languages through cross-lingual transfer from high-resource languages. LM based pre-training and fine-tuning, combined with transfer learning resulted in state-of-art performance across various tasks [PSG19, LRF19].

In a typical cross-lingual transfer scenario, a single multi-lingual language model is pre-trained with large quantities of (unannotated) text from multiple languages. It is then fine-tuned for a given natural language understanding task using human-labeled examples of that task in a *source* language. Cross-lingual transfer occurs when this fine-tuned model can effectively perform this task on another language – the *target* language – without human-labeled data (called *zero-shot transfer*), or with only a few human-labeled examples in the target language (called *few-shot transfer*).

Recently, a line of work by [HNA⁺17, KMH⁺20] has analyzed the scaling effects of parameters, corpus size and number of training steps on pre-training loss in language models [DCLT19]. [Hut21] extended this analysis to the out-of-distribution transfer setting and showed that the effective amount of data transferred from the training distribution to the target distribution follows a power law of the number of parameters and the amount of training data. Similarly, [XAX⁺20] showed that the performance of a wide range of language tasks could be predicted with relatively good accuracy. Their approach consists of parameterizing the experimental setting with both data-driven features, and linguistic features fed to a gradient-boosting model [Fri01] to predict downstream performance.

Probing studies from [PSG19] and [XCR⁺20] suggest that large multi-lingual language models exhibit zero-shot transfer ability and can deliver state-of-art performance for low-resource languages. [KWMR20] have suggested that “structural similarity” between the source and target languages is one of the most important factors regardless of the lexical overlap or word frequency similarity. Along similar lines, [LRVG20] introduce a meta-regression framework and use it to predict cross-lingual task performance. [LCL⁺19] combined multiple features into a gradient-boosting model to predict zero-shot cross-lingual transfer performance. Finally, [dVWN22] combined multiple typological features in a single regression model to predict the cross-lingual transfer performance of XLM-R [CKG⁺20] in POS tagging. Our work extends their findings by presenting an interpretable statistical framework to explain zero-shot and few-shot cross-lingual transfer. We do it across three tasks and 90 language pairs. While language similarity is a critical factor in effective cross-lingual transfer, we show that corpora size or language model performance in pre-trained models plays an equally important role.

In our work, we try to better understand how multi-lingual pre-trained language models, such as mT5 [XCR⁺20], transfer *any* linguistic and semantic knowledge across languages. There are no explicit cross-lingual signals provided to the model during pre-training. Rather, unannotated texts from each language are presented to the model separately and independently of one another, and the model appears to implicitly learn cross-lingual connections. The fact that this model exhibits cross-lingual transfer may suggest that it is somehow aligning its learned “semantic spaces” of different languages [LRF19, MESS21]. But, *are the cross-lingual connections between every language pair equally strong? What properties of the source and the target language impact cross-lingual transfer performance? Can we quantify the impact of those properties on the cross-lingual transfer?* These are some of the key questions regarding effectiveness of cross-lingual transfer that naturally follow, and are the central theme of this work.

We posit that transfer between some languages is more dominant than others, based on the premise that not all language pairs are born equal [WD20]. As highlighted by [Rud20], designing an NLP system by mirroring what has been done on some high-resource languages (e.g., English) can lead to poor assumptions (e.g., ignoring the rich morphological connections between certain languages). This approach is sub-optimal as it ignores specific properties of, say, Swahili or Arabic that could potentially see larger transfer benefits from “non-traditional” source languages.

Our contributions are three-fold: First, we establish an interpretable statistical framework to enable introspection into cross-lingual transfer in mT5. Next, using the above framework, we assess the impact of various factors on cross-lingual transfer. Finally, we derive linear connections between language similarity features, language model performance and number of target training samples for transfer learning. A key finding of this work is that syntactic similarity, morphological similarity and phonological similarity are good predictors of cross-lingual transfer, significantly more so than just lexical similarity of language pairs. For a given $\{source, target\}$ language pair, we have the ability to predict zero-shot performance on the target language (for a given task), that is shown to increase on a logarithmic scale with the number of few-shot target language data points.

2 Analysis Framework

We start by first establishing an empirical framework to enable our analysis of cross-lingual transfer. One of the things we want to understand is how the “strength” of cross-lingual transfer for a given language pair can be linked to a characteristic (or combination of characteristics) of that language pair. In other words, can we ascertain if certain language pair characteristics will lead to a better cross-lingual transfer in mT5. There is no universally agreed upon methodology for such a study, and a theoretical model of transfer learning across languages is not obvious.

For our analysis framework we draw inspiration from transfer learning literature [MMR09, PY09, DL13] as our starting point. We analyze a pre-trained mT5 for pairs of languages (*source* language (S) and *target* language (T)) through observations of its performance (S_T) in the target language on NLP tasks (e.g., NER, question answering, etc.), after fine-tuning it for the tasks using source language training data (D_S), optionally fine-tuning with target language training data (D_T), and evaluating the task on target language test data. As such, cross lingual transfer can be captured through the function f :

$$S_T = f(n_T, n_S, S_S, D_S, D_T, A) \tag{1}$$

where, (S_S) is the performance on the source distribution, (n_S) is the number of samples used for fine-tuning on the source distribution, (n_T) is the number of samples used for fine-tuning on the target distribution, and (\mathcal{A}) is the learning algorithm with specific hyperparameter choices. Note that when $n_T = 0$, it is zero-shot transfer, while $0 < n_T \ll n_S$ it is few-shot transfer.

Since the data distribution for NLP tasks along with algorithmic complexity is hard to measure and observe, we make some reasonable simplifying assumptions to study cross-lingual transfer. We assume n_S to be constant and much larger than n_T (i.e. $n_T \ll n_S$), and we assume that we observe the performance of the model on the source distribution S_S so we discard n_S from equation 1. We utilize a measurable language similarity metric ($LS_{(T,S)}$) and a language model performance metric (LM) instead of D_S, D_T and \mathcal{A} respectively in equation 1. With this, we can update our cross-lingual transfer equation 1 to:

$$S_T = f(n_T, S_S, LS_{(T,S)}, LM) \tag{2}$$

where function f captures the relationship between various language features and target language performance. If we are able to observe and measure the inputs and outputs of equation 2, then this can provide us with insights into the factors that enable cross-lingual transfer in mT5. We search for an optimal combination of linguistic and/or data-driven features (the inputs to f) that can accurately estimate target language performance (the output from f) for *any* given language pair. The implication is that features that are good predictors of target language performance are important for cross-lingual transfer. In our experimentation framework, we consider several reasonable possibilities for $LS_{(T,S)}$ and modeling LM .

2.1 Language Similarity

Similarity of languages ($LS_{(T,S)}$) can be assessed in many different ways. A *historical* linguistic approach [Hoc09] defines language relatedness through common parent or ancestor languages. In a *typological* approach [DeL83], language similarity is viewed through similarities in phonological, morphological and syntactic properties of languages. Additionally, similarity between languages can be measured through *statistical* means, using large corpora via tokens and character sequence overlaps. Here, we are able to use some aspects of all of these similarity measures for model introspection. We model language similarity through their lexical, morphological, phonological, and syntactic properties, which enables us to assess the impact of these on cross-lingual transfer.

Lexical: We define lexical language similarity by first computing the distribution of character n-grams for each source and target languages¹. To capture the similarities of the dataset that are involved in each experiment, we compute those distributions using the training dataset of each task. We then compute a normalized Jensen-Shanon divergence (JSD) of the source distribution against the target distribution².

$$LEX_{(T,S)} = 1 - \frac{JSD(X_S, X_T)}{\max_{T,S} JSD(X_S, X_T)}$$

with X_L defining the character 3-gram frequency distribution of the language L . We also define the vocabulary size ratio (V_r) between the source and target languages by dividing the vocabulary length of the target language by the vocabulary length of the source language. Finally, we measure the SENTLEN ratio by dividing the average sentence length of the target language with the average sentence length of the source language.

Morphological: Following [XAX⁺20], we use the Type-Token-Ratio (TTR) as a measure of how morphologically-rich a language is. Based on this metric, we derive a Type-Token-Ratio similarity with:

$$MORPH_{(T,S)} = \frac{1}{K} \frac{TTR_T}{TTR_S}$$

where K , a normalization constant, is defined as $K = \max_{(T,S)} \frac{TTR_T}{TTR_S}$.

¹We note that n-grams level lexical similarity is only a proxy of lexical similarity defined as the word-level vocabulary overlap. To avoid relying on imperfect tokenization, we define it using character level 3-grams.

²Note that Jensen-Shanon divergence is a symmetric and smoothed version of the Kullback–Leibler divergence.

Phonological and Syntactic: We extract syntactic and phonological features from the World Atlas of Language Structures (WALS) database³. For each type of property, we compute the intersection over union of the list of properties of the source language with the list of properties of the target language. We refer to those metrics as PHONO for phonological similarity and as SYNT for syntactic similarity. For instance, for a given source language with the properties GENITIVE-NOUN-ORDER and SUBJECT-OBJECT-VERB ORDER (e.g. Japanese) and a target language which has also the property GENITIVE-NOUN-ORDER but a SUBJECT-VERB-OBJECT ORDER structure (e.g. French), SYNT would equal $\frac{1}{3}$ (1 shared property out of a union of 3 properties).

Embedding Driven: Multilingual language models capture rich information about language similarities. [LRF19] showed that language families can be retrieved from embedding representations computed with mBERT. [MESS21] showed that the centered kernel alignment (CKA) [KNLH19] of the hidden representations across languages correlates strongly with downstream zero-shot cross-lingual performance. Finally, [RBE20] showed that a cosine-based embedding driven metric based on mBERT can be statistically explained by genetic and typological signals.

Using those insights, we define an embedding-based language similarity. (a) We compute a language centroid vector by average-pooling the hidden states of the mT5 encoder across a large sample of sentences in a given language. (b) Then, based on these language centroid vectors, we compute a language similarity metric as:

$$\text{EMB}_{(\mathcal{T}, S)} = \cos(\bar{e}_s, \bar{e}_t)$$

with $\bar{e}_l = \frac{1}{k} \sum_i e_{i,l}$ and $(e_{i,l})_i$, k sentence embedding vectors in the language l .

2.2 Language Modeling

The LM term in equation (2) refers to performance of mT5 as a language model before any task-specific fine-tuning is done. We define two metrics related to the pre-training mechanism and analyze their relationship with downstream task-specific performance. First, we adapt the language model score defined by [SLNK20] to the denoising objective used to pre-train T5. We compute the output log-likelihood of the model on span-masked sentences. Formally, for a collection of sentences $X = x_1, \dots, x_{|X|}$ with $x_s = (x_1, \dots, x_{|x|})$ in the language L , we define the language model performance as:

$$LM_{\mathcal{L}(L)} = \frac{1}{|X|} \sum_x \sum_{i \in s} \log(p(x_i | x \setminus x_i))$$

For our second approach, we define an Exact-Match accuracy metric as follows:

$$LM_{EM(L)} = \frac{1}{|X|} \sum_{x,s} \prod_{i \in s} \mathbb{1}(\hat{x}_i = x_i)$$

with \hat{x}_i is the greedy-decoded model prediction of a sequence x after masking spans indexed by s . $LM_{EM(L)}$ simply captures how accurate the span predictions of the model are when feeding it masked sentences. We follow strictly the pre-training span masking procedure defined in [XCR⁺20]. We estimate both these statistics on the training dataset of each task.

3 Experiments

Based our framework made of equation 2 and the features presented in the previous section, we now present our experiments. We start by presenting a bi-variate analysis between each feature and the cross-lingual transfer performance. Then, we present a meta-regression that combines the multivariate effect of the features on the cross-lingual transfer performance.

As mentioned before, we focus on the mT5 framework, a multi-lingual adaptation of T5 [RSR⁺19]. T5, *Text-To-Text Transfer Transformer*, formulates any NLP tasks as sequence generation. If the task is a classification or regression, we generate the label token by token as if it were natural language.

³The WALS is a database of linguistic properties collected for a large number of languages <https://wals.info/>. We extract them using the lang2vec python package from [LML⁺17]

Task	$LM_{\mathcal{L}(S)}$	$LM_{\mathcal{L}(T)}$	$LM_{EM(S)}$	$LM_{EM(T)}$	SYNTAX	PHONO	MORPH	LEX	EMB	V_r	SENTLEN
QA	6.6	7.3*	-15.4*	-18.1*	38.4*	29.6*	-40.7*	19.3*	8.9*	-40.7*	3.4
NER	22.1*	1.3	1.3	6.6	2.4	14.7*	-2.4	23.6*	16.2*	-2.4	-27.1*
XNLI	6.5	23.9*	-26.5*	4.1	-0.1	0.1	10.8*	-6.3	1.3	10.8*	-12.4*

Table 1: Pearson correlation between the features introduced and the cross-lingual transfer performance in the zero-shot setting - measured as $S_T - S_S$ - for XNLI, QA and NER across source languages (S) and target languages (T) (* indicates statistical significance).

In a nutshell, the T5 framework abstracts away the output feature engineering from meaningless indexes to meaningful language tokens. Its architecture is a Transformer [VSP⁺17] encoder-decoder, pre-trained with a span-masking objective closely inspired by the BERT model [DCLT19]. We run our cross-lingual analysis on the base version of mT5.

Our analysis is conducted on Arabic, Bengali, English, Finnish, Indonesian, Russian, Swahili, Spanish, German, Hindi. Not all languages have training data for all the three tasks we work with but each task gets at least 7 languages. We report the detailed list of the languages used for each task in the Appendix in Table 4. Each language is used both as a source language (S) and as a target language (T) leading to up to 90 language pairs.⁴

Additionally, we focus on three tasks: Natural Language Inference (NLI), Name-Entity Recognition (NER), and Question Answering (QA). For NLI, we use the XNLI dataset [CRL⁺18], for NER the PANX dataset [GL17] and for QA the TyDiQA (for Typologically Diverse Question-Answering) dataset [CCC⁺20].

We report the standard evaluation score for each task: for XNLI we report the accuracy, for NER the F1 score and for QA the exact-match of the predicted answers with the gold answers. To allow comparison across languages, for each task, we control for the number of training samples in the source languages (100k samples for XNLI, 10k for NER, and 2.2k for QA for each source language). For the target language, we run experiments with $n \in \{0, 10, 30, 50, 100, 250\}$ for all three tasks and include $n \in \{500, 750, 1000\}$ for NER and XNLI.

3.1 Bivariate Correlation Analysis

We start this study with bi-variate correlation analysis over predictors introduced in section 2.1. We report in Table 1 the Pearson correlation of each predictors with cross-lingual transfer performance for each task. We indicate with * statistically significant results.⁵ Among others, we find that LEX, PHONO and EMB have significant correlation for QA and NER. For XNLI, we find that $LM_{\mathcal{L}(T)}$ and MORPH are correlated significantly with cross-lingual transfer. This suggests that those predictors are good candidates to be used in a predictive regression model.

Overall, the correlations have similar trends across tasks (correlation signs are the same across tasks in most cases) but with significant differences in strengths. We observe a few key differences for some features. For instance, the syntactic similarity correlates strongly with cross-lingual transfer

⁴Having access to many more languages for XNLI than for the other tasks, we extend our zero-shot cross-lingual transfer experiment to 7 extra languages, namely (Bulgarian, Greek, French, Turkish, Urdu, Vietnamese, and Chinese).

⁵We run a two-tailed statistical significance test on the Pearson correlation and consider correlation to be significant when the p-value $\leq 5\%$

Transformation of n	QA	NER	XNLI
n	67.3	58.1	14.2
$\log(n)$	83.9	69.0	14.8
$\log^2(n)$	76.9	67.8	18.0

Table 2: Pearson Correlation between S_T and various transformation of the number of samples in the target language n used for fine-tuning mT5 for XNLI, QA and NER. (with $n \in \{0, 10, 30, 50, 100, 250, 500, 750, 1000\}$).

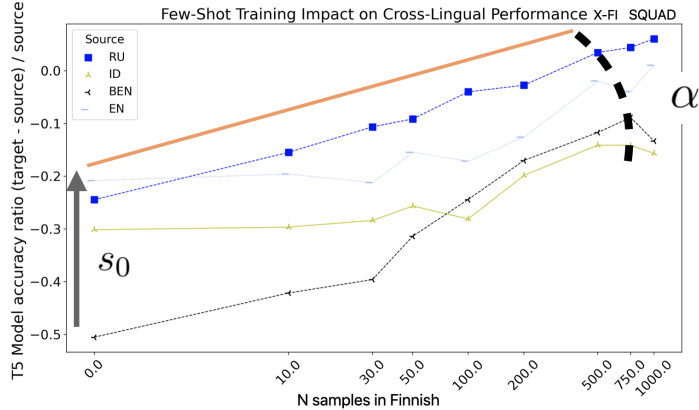


Figure 1: Zero-shot vs. N-shot Cross-Lingual Transfer for QA when transferring to Finnish from Russian, Indonesian, Bengali and English.

for Question Answering but does not with NER and XNLI. These results suggest that a multivariate linear model should be task-specific⁶.

Still, we note that despite the observed correlations, each of these metrics provides an incomplete view of the relationship between language distance and cross-lingual transfer. Indeed, the correlation is never close to 1 which means that none of these features alone is informative enough to predict cross-lingual transfer performance.

For instance, character level 3-grams is only limited to languages that share the same script. Languages with different script have very high lexical divergence. We note that the divergence is never undefined (infinite) even between languages written in different scripts. This is due to the numbers and residual Latin tokens used in the Arabic, Russian and Korean datasets.

For instance, the divergence between Swahili and Arabic is very close to the one from Swahili to Hindi but Arabic leads to a much better transfer. Additionally, in some cases our language similarity metrics fail to explain cross-lingual transfer. For instance, for QA, the embedding-based similarity between Indonesian and Swahili is significantly higher than the one between Finnish and Swahili. Still, the transfer to Swahili is better when the source is Finnish.

We run the same analysis with language model performance measured with log-likelihood and exact-match accuracy (cf. Section 2.2). Language model performance exhibits strong correlation, but surprisingly, the correlation is counter-intuitive. For instance, for QA, the higher the language model accuracy on the target language the lower the cross-lingual transfer.

To gain more insights into f , from Equation (2) in Section 2, we analyze the bivariate relationship between the target performance $S_{\mathcal{T}}$ and the number of samples n in the target language. As illustrated in Figure 1 with the n-shot cross-lingual performance from various languages to Finnish for Question Answering, there is a strong linear relationship between the log of the number of samples in the target language and downstream accuracy. We report in Table 2, the Pearson correlation between transformation of the number of samples with downstream performance in the target language for the three tasks. We find that downstream n -shot cross-lingual transfer performance correlates strongly with $\log(n)$ and the most for two of the three tasks. However, for XNLI, we note that $\log^2(n)$ is a stronger predictor and that the absolute correlation is weaker than for the two other tasks. We explain this by the fact that the absolute performance of XNLI in the few-shot setting increases only moderately compared to the zero-shot setting.

In summary, we introduced several predictors that, overall, correlate linearly and strongly with cross-lingual performance. In the following sections, we show that these predictors can be combined linearly to predict cross-lingual transfer reasonably well.

⁶We note that our results extend previous findings from [PSG19].

3.2 Meta-Regression of Cross-Lingual Performance

The bivariate analysis in the prior section principally motivates us to hypothesize that cross-lingual transfer has two components to itself. A source-target similarity component that dominates the zero-shot learning of a target language. While the logarithmic relation in number of target samples governs n-shot learning.

In consequence, we simplify f as follows:

$$S_{\mathcal{T}} = S_{\mathcal{T}}(0) + \alpha \log(n + 1)$$

with:

$$S_{\mathcal{T}}(0) = f_0(S_S, T, S, LM, \mathcal{T})$$

$$\alpha = f_1(S_S, T, S, LM, \mathcal{T})$$

Note that $S_{\mathcal{T}}(0)$ corresponds to the performance of the model in the zero-shot setting. While α is a function of source-target language, algorithm and task difficulty that governs the slope of the performance curve. As such, we will break our analysis into two separate components and study them as such.

Feature Selection: While we study an array of diverse feature transformation and combination, however, we do filter out most of them. To keep only the most simple and relevant features, we use a Lasso Regression [Tib96]. We run feature selection with recursive feature elimination using the absolute value of each coefficient. We start with all the features, we fit the regression and we iteratively remove the feature with the smallest coefficient.

Evaluation: Following [XAX⁺20] we fit our meta-regression model with l -folds cross-validation and define each fold so that they include only the observations for a single target language. We fit the model on the concatenation of $l - 1$ language-folds and we evaluate on the l^{th} fold. We report the average l -cross-validation score, which corresponds to the average computed on all the target languages. We evaluate the performance of the regression using the Root Mean Squared Error (RMSE), a standard metric to evaluate a regression model.

Additionally, for a given language and a number of annotated samples, we predict which source language leads to the best performance according to our regression. We do by solving argmax over all possible language pair for the Equation 2. We then measure the accuracy of the prediction by comparing it to the actual best source language. We denote this accuracy metric A_{src} .

3.2.1 Zero-Shot Transfer Meta-Regression

We find that for the three tasks, the zero-shot cross-lingual transfer can be effectively modeled with a linear combination of afore mentioned features. We summarize in Table 3a the performance of the regression. Note, a simple model with most relevant features is able to explain a lot of the variance in the QA performance for zero-shot transfer with RMSE as low as 5.48. Additionally, our model can be used to predict the best source language with a relatively good accuracy. For instance, on QA, the model can predict the correct source language in 64% of the cases.

We present in Figure 2 the final linear relationships between the cross-lingual performance and the selected predictors after recursive feature elimination for the three tasks. For interpretation purposes, we report the coefficients of the regression fitted on all the language-folds. We find that few features such as syntactic, morphological and lexical similarity are strong predictors of transfer learning performance. Surprisingly, despite the strong bivariate relationship, the language centroid similarity does not offer helpful signals in the presence of other features. Additionally, we find that language model performance as measured with the exact match accuracy is a useful predictor of task-specific performance.

For all the three tasks, the language model performance on the target language has a significant and positive coefficient. This suggests that a better language modeling of the target language leads to better task-specific prediction.

From the equations in Figure 2, we discover the following relationships between predictors and performance on the target language.

Task	RMSE	Top-1 Source Prediction Accuracy	Task	RMSE
QA	5.48	64.3%	SQUAD	5.69
NER	9.80	50.0%	NER	7.60
XNLI	4.26	16.4%	XNLI	8.13

(a) Zero-shot meta-regression fit measured with RMSE of f_1 for each task. and A_{src} top-1 source language prediction accuracy. (b) n -shot meta-regression RMSE of f_1 for each task.

Table 3: Meta-regression results in zero-shot and n -shot settings. The Root Mean-Square-Error (RMSE) corresponds to an average error of our regression in absolute points.

$$(QA) S_{\mathcal{T}}(0) = -65.38 + 0.62 S_S + 56.49 S_{yn} + 156.4 P_{hono} - 29.73 M_{orph} + 129.3 LM_{EM(\mathcal{T})}$$

$$(NER) S_{\mathcal{T}}(0) = -42.8 + 1.07 S_S + 14.63 S_{yn} + 68.22 LM_{EM(\mathcal{T})} + 4.09 Lex$$

$$(XNLI) S_{\mathcal{T}}(0) = 27.62 + 0.64 S_S + 10.12 P_{hono} - 1.2 M_{orph} + 46.87 LM_{EM(\mathcal{T})} + 1.2 Lex$$

Figure 2: Regression models after feature selection for the three tasks to predict zero-shot cross-lingual transfer. All the coefficients are non-null (with a p-value ≤ 0.05). All the variables are scaled between 0 and 1, allowing comparison across coefficients.

- For every 1% improvement in language modeling accuracy on the target language, we can expect about 1.3% improvement in target language performance for QA, 0.7% for NER and 0.5% for XNLI.
- Every 10 points increase in syntactic similarity leads to a corresponding increase of around 5 points in downstream accuracy for QA, 1.4 for NER.
- For QA and XNLI, we find that downstream cross-lingual performance is strongly related to phonological similarity: 10 points increase in language similarity can result in up to 15.6 points increase in cross-lingual performance in QA and 10 points in XNLI.
- Surprisingly, morphology has large negative effect on cross-lingual transfer for QA and more moderate one for XNLI.

3.2.2 N-Shot Cross-Lingual Meta-Regression

For n -shot cross-lingual transfer performance prediction, we run the same feature selection for f_1 . We found that the two main factors that are reliable predictors for each QA, NER and XNLI task performance: (i) number of samples available for training target language and the (ii) the zero-shot performance. We notice that the influence of the language similarity metrics and language model performance usually dies out once more target language samples are available for training. These results are consistent over all language pairs and tasks. We report the performance of the regression for the n -shot setting in Figure 3b.

We enumerate here several observations from our experiments. For QA, we notice that an increase in the number of data samples in the target language by a factor of 10 has a corresponding increase in target language performance by an average of over 5%. The slope of this increase, however, also depends on the zero-shot performance and language similarity. Observe that higher language similarity equates to an increase in zero-shot performance, which in turn reduces the influence of target language annotated data. For XNLI, we learn that the slope due to number of samples is pretty low and consistent over languages. The slope is very highly dependent on the few shot accuracy For NER, we find that a 10x increase in target training data has on average a 7% increase in the task performance. The broad takeaway from these observations is that target language performance is heavily predictable by training samples for all tasks.

4 Conclusion and Future Work

In this work, we analyze a pre-trained mT5 to study the nature of cross-lingual transfer. Through model interpretation experiments over multiple language pairs and tasks, we show that transfer can be modeled statistically by a few linguistic and data-derived features. We show that syntax, morphology, and phonology of languages are good predictors of cross-lingual transfer (significantly better than lexical similarity of languages). We also demonstrate that language model performance is an informative predictor of cross-lingual performance, providing an off-the-shelf metric to inform cross-lingual transfer.

Since, reproducing these relations for all language pairs and all tasks is a huge drain of resources, a natural future work for this line of research is to generalize these findings and relations over all tasks instead of the 3 that we did. Similarly, we focused on mT5 only due to its simplicity of architecture, an interesting follow-up would be to derive more general relations for other architectures like mBART, mBERT etc. Additionally, we would like to better understand influence of *multiple* source languages on each target language – i.e., *is there a combination of several source languages that could enable better cross-lingual transfer?* Finally, we would like to leverage this knowledge to influence the mT5 pre-training methodology to better facilitate few and zero-shot transfer across tasks.

References

- [CCC⁺20] Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470, 2020.
- [CKG⁺20] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online, July 2020. Association for Computational Linguistics.
- [CRL⁺18] Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium, October–November 2018. Association for Computational Linguistics.
- [DCLT19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [DeL83] Scott DeLancey. Language universals and linguistic typology: Syntax and morphology, 1983.
- [DL13] Li Deng and Xiao Li. Machine learning paradigms for speech recognition: An overview. *IEEE Transactions on Audio, Speech, and Language Processing*, 21(5):1060–1089, 2013.
- [dVWN22] Wietse de Vries, Martijn Wieling, and Malvina Nissim. Make the best of cross-lingual transfer: Evidence from POS tagging with over 100 languages. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7676–7685, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [Fri01] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.

- [GL17] Abbas Ghaddar and Phillippe Langlais. WiNER: A Wikipedia annotated corpus for named entity recognition. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 413–422, Taipei, Taiwan, November 2017. Asian Federation of Natural Language Processing.
- [HNA⁺17] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Patwary, Mostofa Ali, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017.
- [Hoc09] Hans Henrich Hock. *Principles of historical linguistics*. Walter de Gruyter, 2009.
- [HRS⁺20] Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. *CoRR*, abs/2003.11080, 2020.
- [Hut21] Marcus Hutter. Learning curve theory, 2021.
- [KB14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [KLSB20] Phillip Keung, Yichao Lu, Julian Salazar, and Vikas Bhardwaj. Don’t use English dev: On the zero-shot cross-lingual evaluation of contextual embeddings. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 549–554, Online, November 2020. Association for Computational Linguistics.
- [KMH⁺20] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [KNLH19] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. *arXiv preprint arXiv:1905.00414*, 2019.
- [KWMR20] Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. Cross-lingual ability of multilingual bert: An empirical study. In *International Conference on Learning Representations*, 2020.
- [LCL⁺19] Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastasopoulos, Patrick Littell, and Graham Neubig. Choosing transfer languages for cross-lingual learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3125–3135, Florence, Italy, July 2019. Association for Computational Linguistics.
- [LGG⁺20] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742, 2020.
- [LML⁺17] Patrick Littell, David R. Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori Levin. URIEL and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 8–14, Valencia, Spain, April 2017. Association for Computational Linguistics.
- [LRF19] Jindřich Libovický, Rudolf Rosa, and Alexander Fraser. How language-neutral is multilingual bert? *arXiv preprint arXiv:1911.03310*, 2019.
- [LRVG20] Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online, November 2020. Association for Computational Linguistics.

- [MESS21] Benjamin Muller, Yanai Elazar, Benoît Sagot, and Djamé Seddah. First align, then predict: Understanding the cross-lingual ability of multilingual BERT. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2214–2231, Online, April 2021. Association for Computational Linguistics.
- [MMR09] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms, 2009.
- [PSG19] Telmo Pires, Eva Schlinger, and Dan Garrette. How multilingual is multilingual BERT? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy, July 2019. Association for Computational Linguistics.
- [PY09] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- [RBE20] Taraka Rama, Lisa Beinborn, and Steffen Eger. Probing multilingual BERT for genetic and typological signals. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1214–1228, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics.
- [RSR⁺19] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.
- [Rud20] Sebastian Ruder. Why You Should Do NLP Beyond English. <http://ruder.io/nlp-beyond-english>, 2020.
- [SLNK20] Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. Masked language model scoring. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online, July 2020. Association for Computational Linguistics.
- [Tib96] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [VSP⁺17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [WD20] Shijie Wu and Mark Dredze. Are all languages created equal in multilingual BERT? In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 120–130, Online, July 2020. Association for Computational Linguistics.
- [WDS⁺20] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, October 2020. Association for Computational Linguistics.
- [XAX⁺20] Mengzhou Xia, Antonios Anastasopoulos, Ruochen Xu, Yiming Yang, and Graham Neubig. Predicting performance for natural language processing tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8625–8646, Online, July 2020. Association for Computational Linguistics.
- [XCR⁺20] Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*, 2020.

A Reproducibility

A.1 Fine-tuning and Evaluation Data

For NER, we use the data from [GL17], for XNLI we use the data from [CRL⁺18] and for Question Answering (QA), we use the TyDiQA dataset from [CCC⁺20]. All those datasets are downloaded and pre-processed using the scripts from the XTREME benchmark [HRS⁺20] available at <https://github.com/google-research/xtreme>. We work with the languages listed in Table 4.

Languages (iso)	Available Data Train/Test		
	XNLI	TyDiQA	NER
Arabic (AR)	1/1	1/1	1/1
Bengali (BN)	∅	1/1	∅
English (EN)	1/1	1/1	1/1
Finnish (FI)	∅	1/1	1/1
Indonesian (ID)	∅	1/1	1/1
Russian (RU)	1/1	1/1	1/1
Swahili (SW)	1/1	1/1	∅/1
Spanish (ES)	1/1	∅	1/1
German (DE)	1/1	∅	1/1
Hindi (HI)	1/1	∅	1/1

Table 4: Annotated Data available per studied tasks for languages included in mT5 pretraining. The selection was done starting on the TyDiQA languages and extended to have a better coverage across all tasks. 1/1 means that data is available for both training and evaluation. ∅/1 means it is only available for evaluation.

Each language is used both as a source language (fine-tuning) and as a target language (evaluation).

A.2 Implementation

Our experiments are based on the pre-trained Multilingual T5 models released by [XCR⁺20] and available in the Transformers library [WDS⁺20] based on which the fine-tuning experiments have been performed.

A.3 Sequence Prediction

With (m)T5, every NLP tasks is framed as a text generation task. For QA and XNLI, we follow closely what has been introduced by [RSR⁺19]. For NER, we use the gold tokenization provided in the dataset from [GL17] and we perform sequence labeling by generating the sequence of labels along with the original input tokens as follows:

Input		<i>It was named for Williams College in Williamstown.</i>
Output		It : O was : O named : O for : O Williams : B-ORG College : I-ORG in : O Williamstown : B-LOC . : O

A.4 Hyperparameters

Params.	QA	Bounds
batch size	128	[1, 8192]
Optimizer	Adam	-
learning rate	3e-4	[1e-6, 1e-3]
gradient clipping value	1.0	-
Max Sequence Length (token)	512	[1, 1024]

Table 5: Fine-tuning best hyper-parameters. Reported best epoch indicates the best selected epoch when English is the source language in the zero-shot cross-lingual setting.

We fine-tune mT5 using the same set of hyperparameters for each task. In contrast with the original implementation from [RSR⁺19], we use the Adam Optimizer [KB14]. We report in Table 5 the set of hyperparameters used.

As highlighted by [KLSB20], cross-lingual performance can be unstable from one run to another. To tackle this instability, each experiment is ran on at least 4 different random seeds. Our analysis is computed on the average of the successful runs.

B Absolute Performance and Prediction of mT5 in the Cross-Lingual Setting

We report in Table 3-7 the zero-shot and few-shot (using $n = 100$ as number of samples in the target language), for QA, NER and XNLI.

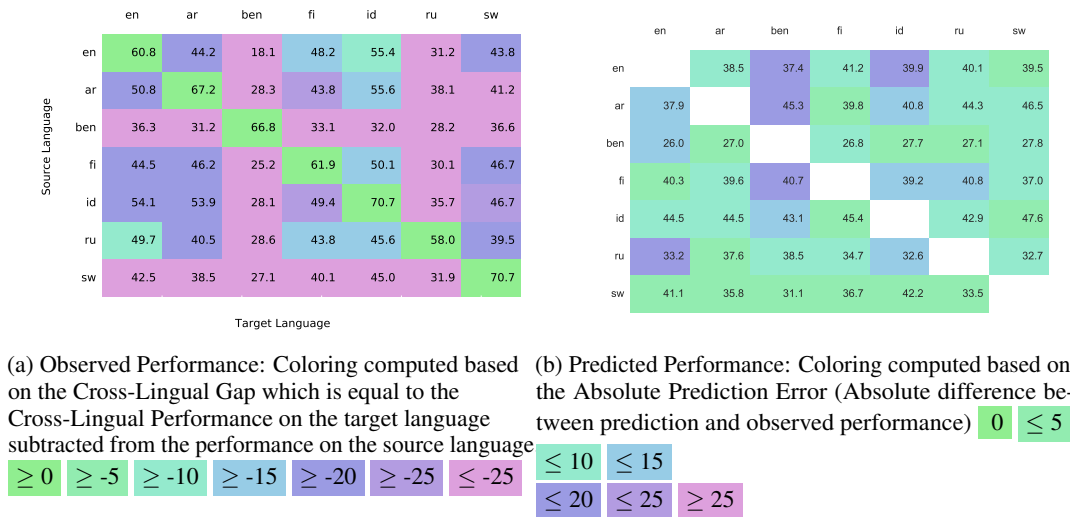


Figure 3: **Question Answering (QA) Zero-Shot** (0 annotated samples from the target language ($n = 0$)) Exact-Match Accuracy of mT5 base in the Cross-Lingual Setting.

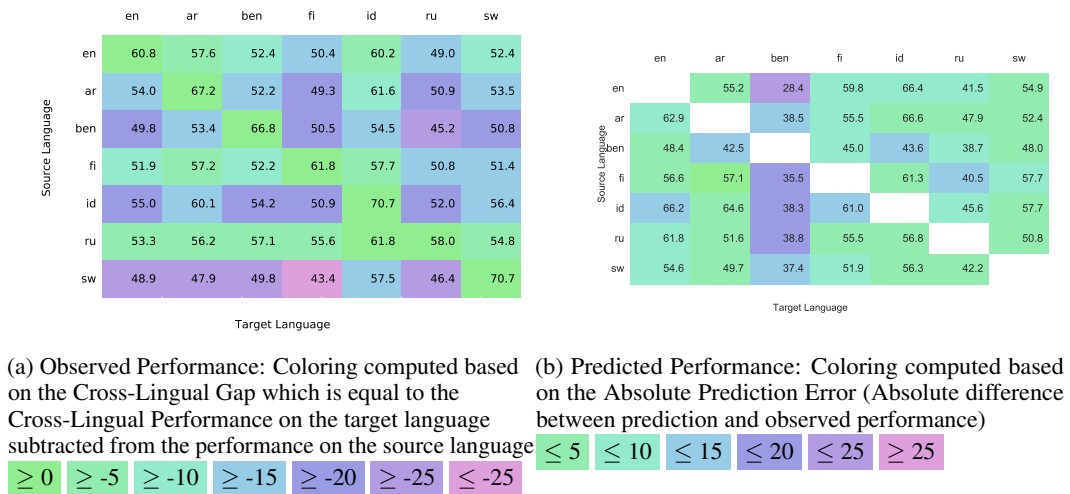




Figure 4: **Question Answering (QA) Few-shot** (100 annotated samples from the target language ($n = 100$)) Exact-Match Accuracy of mT5 base in the Cross-Lingual Setting.



(a) Observed Performance: Coloring computed based on the Cross-Lingual Gap which is equal to the Cross-Lingual Performance on the target language subtracted from the performance on the source language. (b) Predicted Performance: Coloring computed based on the Absolute Prediction Error (Absolute difference between prediction and observed performance)

≥ 0 ≥ -5 ≥ -10 ≥ -15 ≥ -20 ≥ -25 ≤ 5 ≤ 10 ≤ 15 ≤ 20 ≤ 25 ≥ 25

Figure 5: **NER Zero-shot** (0 annotated samples from the target language ($n = 0$)) F1 score of mT5 base in the Cross-Lingual Setting.



(a) Observed Performance: Coloring computed based on the Cross-Lingual Gap which is equal to the Cross-Lingual Performance on the target language subtracted from the performance on the source language. (b) Predicted Performance: Coloring computed based on the Absolute Prediction Error (Absolute difference between prediction and observed performance)

≥ 0 ≥ -5 ≥ -10 ≥ -15 ≥ -20 ≥ -25 ≤ 5 ≤ 10 ≤ 15 ≤ 20 ≤ 25 ≥ 25

Figure 6: **NER Few-shot** (100 annotated samples from the target language ($n = 100$)) F1 score of mT5 base in the Cross-Lingual Setting.




Figure 7: XNLI Zero-shot (0 annotated samples from the target language ($n = 0$)) F1 score of mT5 base in the Cross-Lingual Setting.

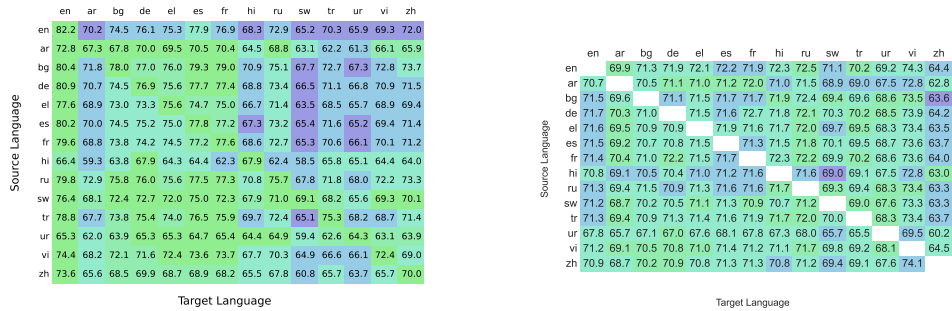


Figure 8: XNLI NEW Zero-shot (0 annotated samples from the target language ($n = 0$)) F1 score of mT5 base in the Cross-Lingual Setting.