

Linguistic Knowledge Can Enhance Encoder-Decoder Models (If You Let It)

Alessio Miaschi, Felice Dell’Orletta, Giulia Venturi

ItaliaNLP Lab, Institute for Computational Linguistics "A. Zampolli" (CNR-ILC), Pisa
{name.surname}@ilc.cnr.it

Abstract

In this paper, we explore the impact of augmenting pre-trained Encoder-Decoder models, specifically T5, with linguistic knowledge for the prediction of a target task. In particular, we investigate whether fine-tuning a T5 model on an intermediate task that predicts structural linguistic properties of sentences modifies its performance in the target task of predicting sentence-level complexity. Our study encompasses diverse experiments conducted on Italian and English datasets, employing both monolingual and multilingual T5 models at various sizes. Results obtained for both languages and in cross-lingual configurations show that linguistically motivated intermediate fine-tuning has generally a positive impact on target task performance, especially when applied to smaller models and in scenarios with limited data availability.

Keywords: encoder-decoder, intermediate fine-tuning, linguistic features, sentence complexity

1. Introduction

Understanding “how linguistic concepts that were common as features in NLP systems are captured in neural networks” (Belinkov and Glass, 2019) has been the focus of many studies in the recent NLP research. It has been extensively shown that pre-trained Neural Language Models (NLMs) are able to capture syntax- and semantic-sensitive phenomena (Hewitt and Manning, 2019; Pimentel et al., 2020; Li et al., 2022) and that there is a correlation between the degree of linguistic knowledge and its ability to solve correctly a downstream task (Miaschi et al., 2020; Sarti et al., 2021), although it is still highly debated (Ravichander et al., 2021). However, it has also been demonstrated that introducing additional linguistic information (Wang et al., 2019b; Zhou et al., 2020; Glavaš and Vulić, 2021) during the pre-training phase can enhance models’ performances. In addition, several works showed that transfer learning methods, such as fine-tuning on intermediate supporting tasks, are highly beneficial to improve pre-trained models’ performance in the resolution of multiple final target tasks (Phang et al., 2018; Wang et al., 2019a).

Starting from these premises, in this work we conducted multiple experiments devoted to evaluating the potential benefits of enriching a pre-trained NLM with multiple linguistic information that may enhance its performance in the resolution of a target task. Specifically, we defined a methodology based on an intermediate fine-tuning phase where the model is instructed to solve a set of raw, morpho-syntactic and syntactic tasks both in a multi- and single-task scenario. To this end, we tested our method on an Encoder-Decoder model, i.e. T5 (Raffel et al., 2020), that allows focusing the instructing process on specific linguistic tasks. Our target

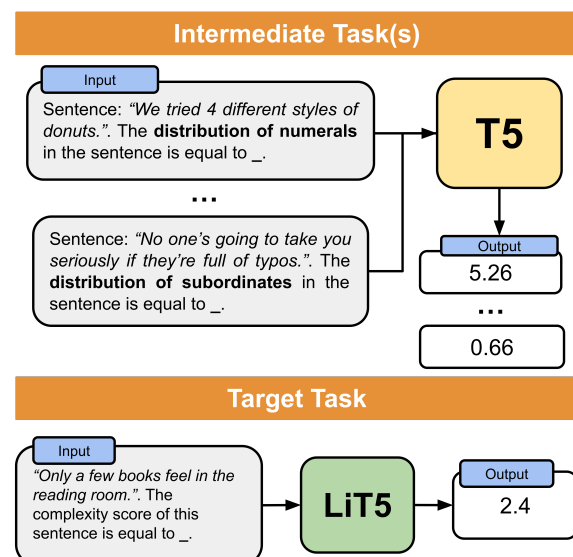


Figure 1: Illustrated example of the proposed methodology. T5 is previously fine-tuned on a subset of linguistic intermediate tasks in a multitask fashion. Then, the newly obtained model, LiT5, is tested on the target task.

task is one that strongly relies on the knowledge of the linguistic properties that characterize a sentence, namely, predicting the linguistic complexity of a sentence (see Figure 1 for an overview of the methodology).

The methodology was tested on two languages, Italian and English, employing both mono- and multi-lingual T5 models. Furthermore, we devised a cross-lingual evaluation scenario to assess the effectiveness of a model that underwent linguistic fine-tuning on data from a language different from that of the target task. The experiments on the target task were conducted by varying the training

data size to explore how linguistically informing the models impacts their performance across scenarios with limited data. Our purpose is not to propose a new framework for achieving state-of-the-art performance, but rather to inspect how this process of enhancing a model with linguistic knowledge scales across different languages and amounts of data.

Our main contributions are:

- We propose an intermediate fine-tuning approach to study the impact of enhancing pre-trained Encoder-Decoder models with knowledge of multi-level linguistic phenomena in solving a target task that has not been tested so far.
- We compare the effectiveness of linguistically informing both mono- and multi-lingual models, highlighting language-specific peculiarities.
- We test our approach on T5 models of increasing size and compare their performance varying the dimensions of the target training, thus uncovering the potential for resource-efficient, linguistically-informed small models, particularly in data-limited scenarios.
- We demonstrate the applicability of our method of linguistic fine-tuning across languages, offering valuable adaptability insights for cross-lingual scenarios¹.

2. Related Work

A large body of recent work focused on investigating whether performing multiple steps of fine-tuning on one or more intermediate tasks can enhance the performance of a pre-trained NLM before fine-tuning it on a target task. The core concept behind this approach is to specialize the pre-trained model by exposing it to tasks other than language modeling, which can enhance its capabilities through a transfer-learning approach. While conducting a comprehensive survey of the literature on this topic is beyond the scope of our work, we aim to concentrate on three key aspects that are frequently discussed about the effectiveness of the intermediate fine-tuning process. These aspects are interconnected, although we present them sequentially. The first one concerns the size of the intermediate fine-tuning and target datasets, and is generally related to a second aspect concerning the fine-tuning approach adopted (Vu et al., 2020; Chang and Lu, 2021). Among the others, e.g. Weller et al. (2022) showed that a multi-task learning approach (i.e. fine-tuning on a supporting task and the target task

¹Datasets and models are available at the following repository: https://github.com/alemiiaschi/linguistically_informed_t5.

simultaneously) to the intermediate fine-tuning on the GLUE benchmark (Wang et al., 2018) is more effective than conducting the two phases consecutively (a methodology often called STILTs) when the target dataset is smaller than the supporting dataset and vice-versa.

A third aspect highly discussed relates to the selection of intermediate tasks (Zhang and Zhang, 2021; Padmakumar et al., 2022; Goot, 2023). Generally, the effectiveness of the intermediate fine-tuning process is highly dependent on the choice of these tasks. For instance, Pruksachatkun et al. (2020) performed a study on the RoBERTa model (Liu et al., 2019) by fine-tuning it on 110 different combinations of intermediate tasks and evaluating the trained models via *probing tasks*. The experiments showed that intermediate tasks requiring high-level inference and reasoning abilities tend to work best. Differently from previous studies, our approach to intermediate fine-tuning is more ‘phenomenon-oriented’ than ‘task-oriented’. Specifically, we have devised a set of prediction tasks that reflect the progressive acquisition of multiple linguistic phenomena, which could prove beneficial in resolving the target task.

On a different side, different works focused on enhancing NLMs with linguistic knowledge during their pre-training phase (Wang et al., 2019b; Glavaš and Vulić, 2021). Among these, Zhou et al. (2020) enhanced a BERT model by training it in a multi-task learning technique with five syntax and semantics tasks, showing that the model can outperform a base one on several benchmarks.

In our work, we aim to bridge these perspectives by exploring an approach that combines the enhancement of a model with linguistic knowledge through intermediate-task fine-tuning, thereby exploiting the power of both transfer learning methods and linguistic knowledge integration.

3. Experimental Setting

We devised a two-step STILTs approach. Firstly, we fine-tuned the T5 models on several intermediate support tasks. They consist of a set of linguistic phenomena modeling aspects of a sentence related in the literature to complexity. They serve as input features for the models which are trained for a total of 25 epochs to predict their values in a multi-task setting. We saved model checkpoints at regular intervals, specifically every 5 epochs, thus resulting in 5 distinct snapshots of the models, each representing a different phase in their development². Secondly, we fine-tuned the “*Linguistically informed*” T5 models, LiT5s, on the prediction

²The LiT5 models are available at the following link: <https://huggingface.co/collections/alemiiaschi/lit5-65d5d480dd8841806fc0a2a0>.

of sentence complexity (one for each checkpoint), which we chose as the target task.

The two steps were tested using Italian and English T5 and the multilingual version. Additionally, we introduced a cross-lingual setting designed to evaluate the efficacy of multilingual models linguistically fine-tuned in a language other than that of the task. For both the intermediate fine-tuning and the target task phase, we verbalized each input sequence by adding a suffix that clarifies the task to be solved. For example, to instruct the models with linguistic knowledge we postpend each sentence with a set of verbalizations, as in the following example (the suffix is in bold): ‘*In 1982, he started a factory in Greece. **The distribution of numerals in the sentence is equal to _***’. While, for the target task, we postpend the string ‘*The complexity score of the sentence is equal to _*’³.

We also proposed a further evaluation scenario to investigate which linguistic feature is most informative for improving the models in solving the target task. To this end, we performed the same intermediate fine-tuning process, but in a single-task scenario, training the small monolingual models with each feature at a time.

3.1. Data

Target Task We considered the task of predicting the level of linguistic complexity of a sentence and we relied on the corpus introduced in Brunato et al. (2018). It contains Italian sentences taken from the newspaper section of the Italian Universal Dependency Treebank (IUDT) (Simi et al., 2014) and English sentences extracted from the Wall Street Journal section of the Penn Treebank (McDonald et al., 2013) manually rated by 20 crowdsourced workers for the corresponding level of perceived complexity on a 1-to-7-point scale. For our experiments, we decided to consider the average judgment of complexity given by all annotators to each sentence. For instance, the average complexity score associated with the following sentence ‘*Only a few books feel in the reading room.*’ is 2.4.

All the sentences contained in the two treebanks were grouped into 6 bins based on a different sentence length, i.e. 10, 15, 20, 25, 30, 35 tokens. Our intention in controlling sentence length was to create sets of sentences that feature comparable values for linguistic characteristics known to be associated with sentence length, such as parse tree depth or dependency links. The original corpus was divided in two: 50% was used as training and 50% as testing. We made sure to maintain a balanced distribution of sentences of 6 different lengths in order to expose the model to a diverse range of

³The complete list of the verbalizations can be found in Appendix A.

sentence lengths during training and testing. For the two languages, we further divided the training set into 5 bins containing an increasing number of sentences and balanced for the six sentence lengths. Specifically, we created 5 training sets each containing 120, 240, 360, 480, and 600 sentences for Italian and 72, 144, 216, 288, and 360 sentences for English. The test sets contain 600 and 360 sentences for Italian and English, respectively. We conducted experiments using training sets of increasing dimensions to explore the impact of enhancing the models with linguistic knowledge also in scenarios with limited data.

Intermediate Tasks They consist of predicting in a multi- and single-task setting the set of linguistic features selected with the approach described in 3.3. The values of the features were extracted from version 2.5 of the Italian and English Universal Dependency (UD) treebanks (Zeman et al., 2019). In total, we collected 16,000 and 4,000 sentences for the training and tests of the Italian treebank, and 17,600 and 4,400 for the English one.

3.2. Models

We relied on different versions of the T5 model (Rafael et al., 2020). For the experiments devised on English, we utilized three models of increasing size, all trained on the English language: *t5-small* (60M), *t5-base* (220M), and *t5-large* (770M). For experiments related to the Italian language, we employed IT5 (Sarti and Nissim, 2022), a T5 model trained on the Italian sentences extracted from a cleaned version of the mC4 corpus (Xue et al., 2021). Just like with the English models, we tested IT5 in three different sizes: *it5-small* (60M), *it5-base* (220M), and *it5-large* (738M).

In addition to these language-specific models, we also conducted cross-lingual experiments using the multilingual model. In this regard, we tested both *mt5-small* (300M) and *mt5-base* (580M). We used Huggingface’s transformers library (Wolf et al., 2020) for accessing all the models⁴.

Evaluation We used Spearman correlation score as evaluation metric. In particular, we computed the Spearman correlation between the gold value of each feature in the Italian or English treebank and the predicted value of the models for the intermediate tasks. For what concerns the target task, we computed the correlation between average judgments of complexity and the complexity scores obtained with the fine-tuned LiT5 models.

⁴For details regarding models, compute parameters and training details see Appendix B.

3.3. Linguistic Features

The set of linguistic features we used as intermediate tasks consists of a subset of those extracted with the ProfilingUD (Brunato et al., 2020), a tool that allows the extraction of more than 130 properties representative of the linguistic structure underlying a sentence and derived from raw, morpho-syntactic and syntactic levels of annotation based on the UD formalism (de Marneffe et al., 2021). The key advantage of relying on the UD formalism lies in the possibility of encoding different sentence properties consistently across various languages. This choice facilitates the adaptation of our methodology to diverse languages with relative ease. Additionally, these features have been shown to play a highly predictive role when leveraged by traditional learning models on various classification problems and can be also effectively used to profile the knowledge encoded in the internal representations of a pre-trained NLM (Miaschi et al., 2020).

We select the subset of linguistic features based on the degree of correlation between sentence-level complexity judgments and the values of linguistic features extracted by Profiling-UD from the target task datasets. Since, as pointed out by Brunato et al. (2018), linguistic complexity is strongly correlated with sentence length, we first calculated the correlation for each of the 6 bins, each containing sentences of different lengths, then averaged the correlation scores. This step was undertaken to minimize the influence of features directly related to sentence length, allowing us to concentrate on linguistic characteristics that contribute more indirectly to influencing the human perception of linguistic complexity. In the end, we obtained a list of the most correlated features for each language considered regardless of sentence length.

Selected features The list of the 10 most correlated ones, along with their Spearman correlation scores, are reported in Table 1. A noteworthy observation is that the scores for both languages appear relatively low. This can be attributed to our deliberate exclusion of features associated with sentence length, which typically exhibits the highest correlations with complexity. Interestingly, for each language, we find different linguistic phenomena, covering diverse aspects of sentence complexity. They include a **raw text feature** as the average length of words (*char_per_tok*); distribution of specific **parts-of-speech** (*upos_dist_**), i.e. adjectives (*ADJ*), numbers (*NUM*), punctuation marks (*PUNCT*), auxiliary verbs (*AUX*), determiners (*DET*), pronouns (*PRON*), symbols (*SYM*), ratio of content parts-of-speech over the total number of words (*lexical_density*), of **syntactic dependency relations** (*dep_dist_**), i.e. function words associated with a verb (*aux*), markers introducing a sub-

Features	Corr	Features	Corr
Italian		English	
char_per_tok	0.28	upos_dist_NUM	0.35
upos_dist_ADJ	0.21	dep_dist_nummod	0.31
upos_dist_NUM	0.19	upos_dist_SYM	0.27
lexical_density	0.17	upos_dist_AUX	0.25
dep_dist_aux	0.17	dep_dist_compound	0.25
dep_dist_mark	0.16	upos_dist_PRON	0.24
aux_mood_dist_Ind	0.14	upos_dist_DET	0.23
obj_post	0.14	subord_prop_dist	0.17
upos_dist_PUNCT	0.13	aux_form_dist_Fin	0.16
subord_prop_dist	0.12	aux_mood_dist_Ind	0.14

Table 1: Linguistic features selected as intermediate tasks along with their average correlation score (Spearman ρ coefficient) with the complexity judgments. All scores are statistically significant (p - value < 0.05).

ordinative clause (*mark*), multiword expressions (*compound*), numeric modifiers of a noun (*nummod*), of auxiliary verbs by **inflectional morphology traits** (*aux_**), i.e. finite forms (*form_dist_Fin*), indicative moods (*mood_dist_Ind*), and of **syntactic tree structures**, i.e. direct objects in post-verbal position (*obj_post*) and subordinate clauses (*subord_prop_dist*)⁵.

4. Results

In the following sections, we delve into the outcomes of our experiments. First, we present the results of the T5 models on the intermediate tasks. Subsequently, we analyze the results of the target task, comparing the performance of both the base and LiT5 models.

4.1. Enhancing T5 with Linguistic Features

Monolingual Models Figure 2 reports the results for the intermediate task (in terms of Spearman correlation coefficients) obtained with the monolingual *t5* models in a multi-task learning scenario. As a first observation, we can see that for both Italian and English, **all models tend to become progressively “linguistically-informed” over the 25 epochs of fine-tuning.**

Interestingly, it seems that **the size of the model impacts the results of the fine-tuning process.** On the one hand, large models consistently outperform smaller ones; on the other hand, it’s noteworthy that the improvement over epochs is particularly pronounced for smaller models. This observation possibly suggests that weights of smaller models even trained on the same amount of data and for the same number of epochs are more susceptible to modification. In addition, we can assume that smaller pre-trained models implicitly encode less

⁵A detailed description of the selected linguistic features is available in Appendix C.

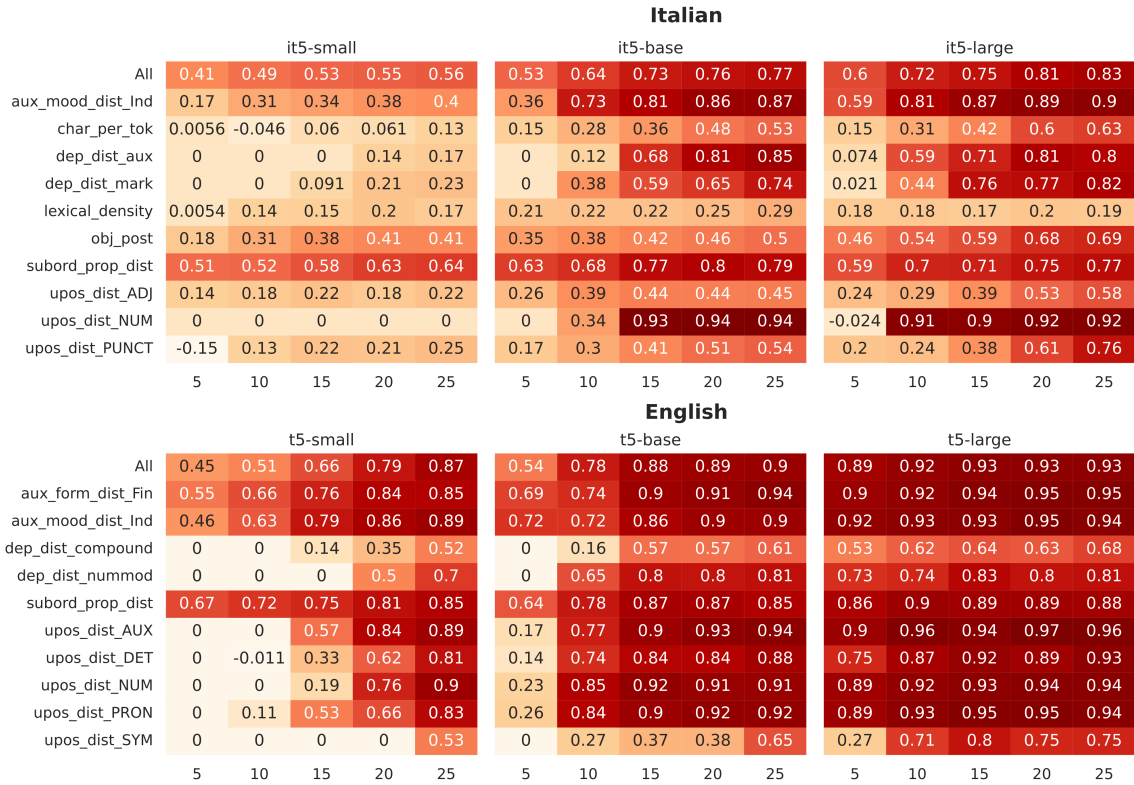


Figure 2: Spearman correlation coefficients for the intermediate tasks for the Italian (*top*) and English (*bottom*) datasets obtained with the monolingual T5 models. Each column in the heatmaps contains the results obtained by the models fine-tuned for multiple epochs (e.g. 5 = 5 epochs of fine-tuning).

linguistic knowledge than larger ones, thus making the addition of new explicit linguistic information more effective.

When we focus on the ranking of each learned linguistic feature, we can generally observe that there are multiple distinctions across models of different sizes. Specifically, it results that at the end of the fine-tuning process the *it5-small* model excels in mastering the distribution of subordinate clauses (*subord_prop_dist*) in comparison with other features learned by the same model, while the distribution of numerals (*upos_dist_NUM*) is the best-mastered property by *it5-base* and *large*. Conversely, the prediction of values of the lexical density of a sentence and of token length (*char_per_tok*) remains consistently a challenge: they are among the worst-mastered sentence characteristics for all the models.

Interestingly, model size also seems to have an impact on the learning speed of specific features. For example, the difference in the accuracy obtained between after 5 and 10 fine-tuning epochs in predicting the lexical density is much greater for *it5-small* than for larger models, for which we obtained quite similar scores across epochs. A quite peculiar case is represented by numbers (*upos_dist_NUM*): *it5-small* fails to master them even at the end of the entire fine-tuning process.

Similar observations hold for *t5*: the model size has an impact on the performance, even if **English models consistently outperform the Italian ones**. However, regardless of their dimensions, all three English models exhibit the lowest proficiency considering the same three features: the distribution of numerical modifiers (*dep_dist_nummod*), compounds (*dep_dist_compound*), and symbols (*upos_dist_SYM*). When we focus on the learning speed across epochs, we can observe that information about the distribution of syntactic dependencies (*dep_dist_**) is acquired by *t5-small* when it is fine-tuned for 10 epochs. Conversely, *t5-base* and *large* master these features earlier and consistently achieve superior results, in line with the general trend already noted. The prediction of symbols (*upos_dist_SYM*) represents a quite peculiar case: they start to be predicted by *t5-small* only after 20 epochs, and it is among the worst-predicted features by the bigger models.

Multilingual Models As we can observe from Figure 3, **multilingual models consistently outperform monolingual ones** for both languages. In addition, some differences in terms of learning speed can be highlighted. In particular, for all features, the improvement of *mt5-base* fine-tuned on the Italian treebank for 10 epochs is much higher

		Italian									
		mt5-small					mt5-base				
		5	10	15	20	25	5	10	15	20	25
All		0.51	0.53	0.65	0.73	0.74	0.56	0.83	0.88	0.89	0.9
aux_mood_dist_Ind		0.29	0.31	0.56	0.7	0.75	0.35	0.8	0.86	0.9	0.92
char_per_tok		0.057	-0.13	0.066	0.18	0.18	0.063	0.64	0.76	0.8	0.81
dep_dist_aux		0	0.037	0.49	0.74	0.77	0.13	0.92	0.92	0.94	0.94
dep_dist_mark		0	0	0.52	0.75	0.79	0.1	0.9	0.92	0.92	0.93
lexical_density		0	-0.02	0.079	0.11	0.083	0.1	0.43	0.58	0.64	0.68
obj_post		0.37	0.42	0.44	0.44	0.44	0.37	0.63	0.75	0.76	0.79
subord_prop_dist		0.61	0.77	0.81	0.83	0.83	0.73	0.85	0.87	0.86	0.86
upos_dist_ADJ		0.11	0.088	0.22	0.44	0.47	0.34	0.65	0.79	0.82	0.82
upos_dist_NUM		0	0.098	0.91	0.95	0.95	0.52	0.96	0.96	0.96	0.96
upos_dist_PUNCT		0.34	0.51	0.72	0.82	0.81	0.38	0.77	0.87	0.89	0.9

		English									
		mt5-small					mt5-base				
		5	10	15	20	25	5	10	15	20	25
All		0.66	0.84	0.88	0.89	0.89	0.88	0.93	0.93	0.94	0.94
aux_form_dist_Fin		0.77	0.87	0.89	0.9	0.89	0.93	0.96	0.96	0.96	0.96
aux_mood_dist_Ind		0.61	0.88	0.91	0.91	0.91	0.92	0.94	0.93	0.93	0.93
dep_dist_compound		0	0.32	0.48	0.53	0.53	0.51	0.6	0.63	0.65	0.65
dep_dist_nummod		0	0.8	0.8	0.79	0.8	0.7	0.84	0.87	0.88	0.89
subord_prop_dist		0.77	0.85	0.88	0.88	0.88	0.86	0.9	0.9	0.91	0.91
upos_dist_AUX		0.48	0.84	0.88	0.9	0.91	0.9	0.94	0.94	0.96	0.96
upos_dist_DET		0.31	0.67	0.74	0.79	0.8	0.75	0.89	0.93	0.93	0.93
upos_dist_NUM		0.2	0.9	0.91	0.93	0.92	0.93	0.94	0.93	0.94	0.94
upos_dist_PRON		0.43	0.77	0.86	0.89	0.9	0.84	0.95	0.95	0.97	0.97
upos_dist_SYM		0	0	0.45	0.45	0.59	0.37	0.72	0.76	0.84	0.8

Figure 3: Spearman correlation coefficients for the intermediate tasks for the Italian (*top*) and English (*bottom*) datasets obtained with the multilingual T5 models.

than for *it5-base* trained for the same numbers of epochs: this suggests that it becomes more rapidly informed on the same linguistic characteristics than the monolingual model. In general, the small model undergoes the main differences compared to its monolingual counterpart.

Similar differences in the learning speed can be highlighted for the English language, especially in the case of the small model. For example, symbols (*upos_dist_SYM*) are learned only after 20 epochs by *t5-small* but after 10 by *mt5-small*, or numerical modifiers (*upos_dist_NUM*) after 15 epochs by the monolingual model but only after 5 by the multilingual counterpart.

4.2. Predicting Complexity with Linguistically-Informed Models

Monolingual Models In Figure 4 we report the performance trend of the *Pre-Trained* and the LiT5 models when fine-tuned on the different portions of the target dataset and for increase numbers of epochs (*Intermediate_**).

The first notable observation is the varying trends observed across model dimensions and languages. **For the two languages, smaller models consistently outperform their pre-trained counterparts**, exhibiting a more pronounced impact from intermediate fine-tuning. This effect becomes even **more evident when working with limited training data**, highlighting the advantages of linguistic knowledge enhancement through a step of intermediate fine-tuning. Particularly evident is, in fact, the

performance gap when fine-tuning the models with a small portion of the training dataset. For instance, when the models are fine-tuned using only 1/5 of the entire training dataset, the disparities in performance of *it5-small* and *t5-small* are notably pronounced (0.29 and 0.30 between the *Pre-trained* and LiT5 correlation scores respectively). In contrast, when **considering bigger models (base and large)**, we notice **different behavior across languages**. For the Italian language, LiT5s outperform the pre-trained ones when fine-tuned on the first three portions of the dataset. However, their performance tends to converge with that of the pre-trained models as the fine-tuning process utilizes the majority of the available training data. In contrast, the trends for the English portion of the dataset are more nuanced, with variations between base and large models and distinctions emerging across different intermediate training epochs. These findings underscore the intricate interplay between model size, language, and the efficacy of intermediate fine-tuning in optimizing performance for specific tasks and datasets.

Focusing instead on the distinctions among the various LiT5 models, we uncover varying trends, particularly when considering different languages. For Italian, there appears to be no discernible difference among models or a consistent trend across increasing training sizes and model dimensions. In fact, scores among the LiT5 models remain generally quite similar. More pronounced variations are exhibited by the *it5-large* model: when trained for only 5 epochs (*Intermediate_5*) it yields generally lower scores compared to its counterpart trained for 25 epochs (*Intermediate_25*).

In contrast, surprisingly, **English smaller models benefit most from extended intermediate fine-tuning processes (> 5 epochs)**, whereas **the base and large models achieve their best performance when linguistically enhanced for only 5 epochs**. This trend becomes especially evident when evaluating performance on the entire target dataset, where the LiT5 models outperform the pre-trained ones. A possible explanation of this trend is that a smaller model may require a longer intermediate fine-tuning process to better adapt to the task, by gradually acquiring the necessary linguistic knowledge. On the other hand, larger models have already learned a wider range of linguistic features during their pre-training, reducing the need for extensive fine-tuning. This transfer learning advantage allows them to quickly adapt to specific tasks and perform well with shorter fine-tuning processes.

Multilingual Models Figure 5 shows the performance trend of the *Pre-Trained* and the LiT5 mul-

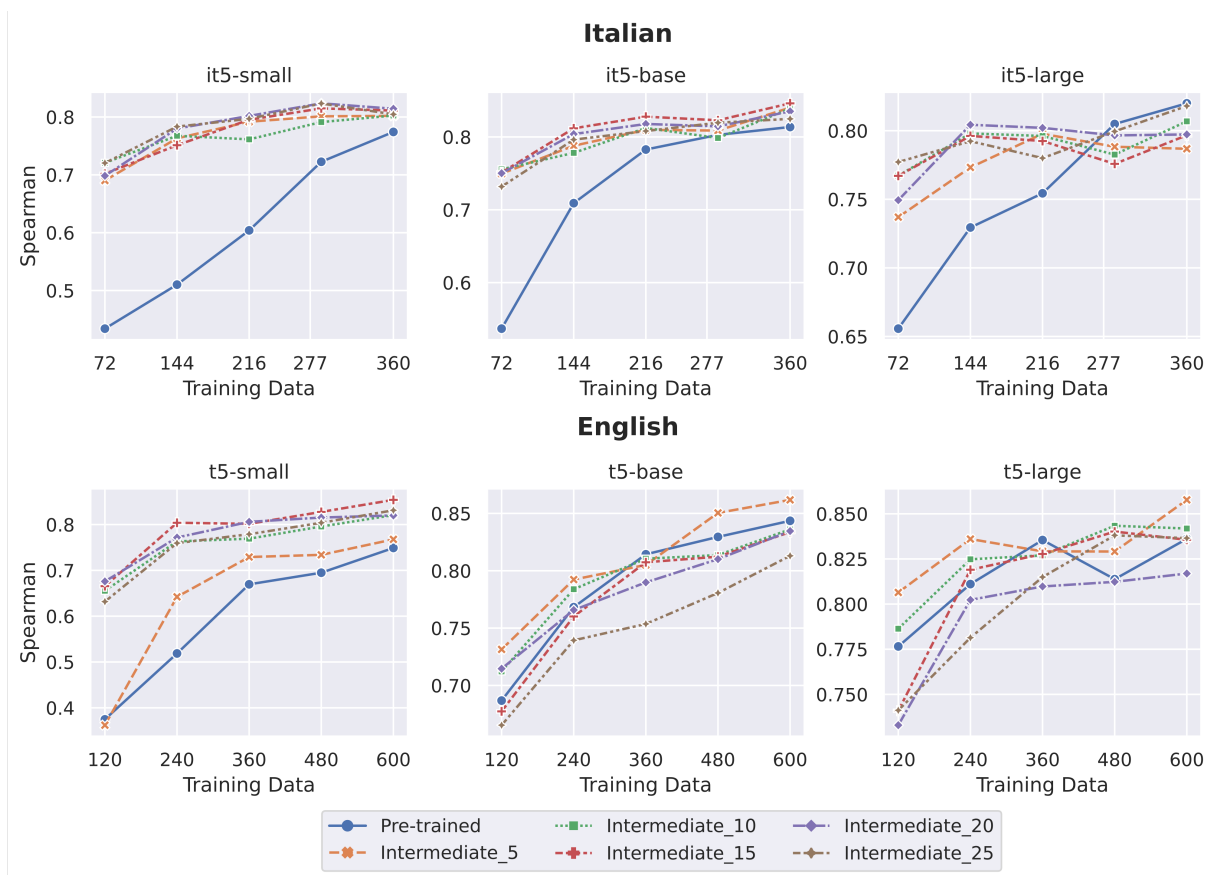


Figure 4: Spearman correlation coefficients for the target tasks for the Italian (*top*) and English (*bottom*) datasets obtained with the monolingual models using pre-trained and LiT5 models.

lingual models⁶. As it can be noted, there is a similar trend to what we observed with monolingual models. **In most configurations, multilingual LiT5 models intermediately fine-tuned either on Italian (**_Italian*) and English (**_English*) improve the performance on the target task, especially for smaller models and with limited data available.** If we focus instead more on the various LiT5 models, we notice that, in general, when training data is limited (1/5 or 2/5 of the entire dataset) the most significant impact for small models is seen in those that have completed the intermediate fine-tuning cycle (25 epochs), while training with larger portions of the dataset tends to reduce the difference between the models. In contrast, for *mt5-base* this tendency is more attenuated: the improvement is evident even when models are trained with only 1/5 of the dataset.

Additionally, we can see that, regardless of whether the LiT5 multilingual models were obtained through fine-tuning on a language different from the target task, their impact is still evident, suggest-

⁶We reported only the results obtained by the LiT5s fine-tuned for the minimum (5) and maximum (25) number of epochs. Complete results are reported in Appendix D.

ing the **effectiveness of the intermediate tuning even in cross-lingual configurations**. However, it is interesting to notice a substantial difference between the target tasks in the two languages. In fact, if we look at the results obtained on the Italian target task, we can see that having performed intermediate tasks in the same language as the target tends to be more effective than doing it in English. On the other hand, a different trend emerges with the English dataset: performing intermediate tasks on Italian data tends to increase the performance more than using a linguistically-informed model trained on English data. In fact, the performance of the LiT5 small model tested on Italian data and intermediately trained on English data increases by 0.31 points in terms of Spearman scores, while it increases by 0.52 when tested on English data and further trained on Italian data. Similar differences are observed for the base models, with a difference of 0.13 and 0.25 in the two scenarios⁷. This result seems to suggest that the enrichment of a multilingual model with language-specific linguistic knowledge is more effective for a language that is poorly represented in the original training of a

⁷Further details about performance differences are available in Appendix D.




Figure 5: Spearman correlation coefficients for the target tasks for the Italian (*top*) and English (*bottom*) datasets obtained with the *mt5*-*. Intermediate scores are reported for the models fine-tuned for Italian and English for the minimum (5) and maximum (25) number of epochs.

Features	Corr	Features	Corr
Italian		English	
upos_dist_NUM	0.65	aux_mood_dist_Ind	0.92
subord_prop_dist	0.60	aux_form_dist_Fin	0.86
aux_mood_dist_Ind	0.52	subord_prop_dist	0.81
obj_post	0.48	upos_dist_PRON	0.80
upos_dist_ADJ	0.22	upos_dist_NUM	0.65
char_per_tok	0.18	dep_dist_nummod	0.34
lexical_density	0.13	upos_dist_DET	0.24
upos_dist_PUNCT	0.09	upos_dist_AUX	-0.03
dep_dist_mark	-0.03	dep_dist_compound	-0.03
dep_dist_aux	-0.03	upos_dist_SYM	##

Table 2: Spearman correlation coefficients for the intermediate tasks for the Italian and English datasets obtained with the small models fine-tuned with one linguistic feature at a time.

multilingual model, such as Italian⁸. In contrast, a more represented language such as English, which might be widely informed of language-specific information, benefits more from an intermediate training on a different language.

4.3. Which linguistic features matter most?

To have a deeper understanding of the contribution of each linguistic feature to the target task, we propose a final evaluation by testing the Italian and English monolingual small models (*it5-small* and *t5-small*) when previously fine-tuned on each of the 10

⁸The Italian language covers around 2.43% of the mC4 corpus (Xue et al., 2021).

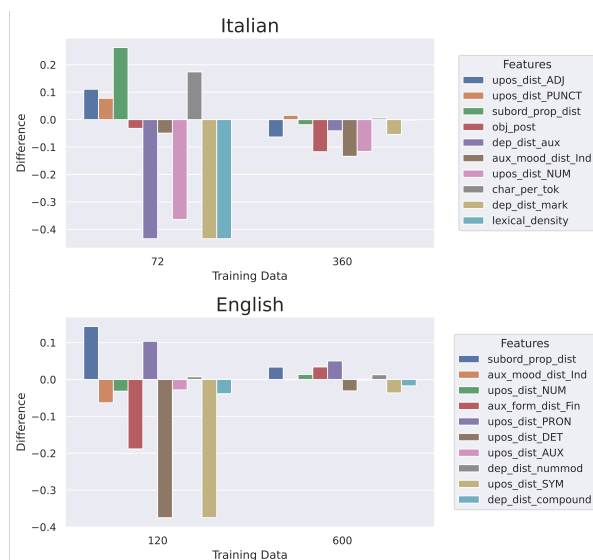


Figure 6: Performance differences between the *Pre-Trained* and the LiT5 models (trained on one linguistic feature at a time) in both scenarios of limited (1/5) and full training data (5/5) availability.

selected linguistic features at a time (i.e. one LiT5 model per property). Note that for this scenario, we have selected only one LiT5 model per feature⁹ instead of testing different checkpoints, as in Sec. 4.2. First of all, we report in Table 2 the results for the single-feature intermediate tasks. Overall, we can notice that the linguistic features are mastered by the two models with lower precision with respect to the multi-task learning setting of fine-tuning. This observation underscores the potential advantage of jointly predicting these sentence characteristics, potentially amplifying their interoperability and mutual reinforcement.

In Figure 6 we report the performance differences between the *Pre-Trained* and the LiT5 models in both scenarios of limited (1/5) and full training data (5/5) availability. The majority of negative differences indicate that, differently from the multi-task configuration, **only a few features effectively help the models in the resolution of the target task**. However, the impact varies depending on the language and the amount of training data. In the scenario **with limited resources, a higher number of features seem to be more informative for the Italian than for the English model and appear to contribute more significantly to enhancing the model performance** in the resolution of the target task. The two most effective ones for *it5* are a syntactic feature, i.e. the use of subordinative clauses (*subord_prop_dist*), and a raw text one, i.e. the length of tokens (*char_per_tok*). Interestingly, being informed about the presence of subordina-

⁹We kept the models that maximize the performance in the prediction of the linguistic features.

tive clauses enhances also the ability of the English model to predict the complexity of a sentence. Conversely, when more training data is available, the addition of linguistic information does not improve the Italian LiT5 models significantly, while it shows a slightly more beneficial effect for the English ones.

5. Conclusion

In this work, we study the effectiveness of enhancing pre-trained T5 models Encoder-Decoder models (i.e. T5), with linguistic knowledge for the prediction of sentence complexity. We conducted a series of experiments on Italian and English datasets using both monolingual and multilingual T5 models of various sizes. Our findings reveal several key insights.

First of all, we showed that informing models linguistically over several epochs allows them to progressively improve their degree of language proficiency, albeit with some differences between languages and model size. Interestingly, model size has an impact on the speed of learning specific linguistic features.

For what concerns the resolution of the target task, we found that our method of linguistic enhancement is particularly effective especially when applied to smaller models and in scenarios with limited availability of target training data. Interestingly, we observe that a linguistically-informed small model, refined through intermediate fine-tuning, can frequently surpass the performance of a larger model that has not undergone this intermediate refinement process. This highlights the potential for more resource-efficient models when linguistic knowledge is incorporated. This holds both for mono and multilingual models and when the linguistic fine-tuning process is conducted in a language other than that of the target task. The findings indicate that refining smaller models through efficient intermediate tuning phases could present a promising approach to building more sustainable models. Specifically, it may confer benefits compared to larger models trained for extended periods on more extensive datasets with a higher parameter count.

Finally, preliminary experiments conducted by investigating the impact of training linguistically-informed models with a single linguistic feature, highlighted that only a few of them seem to enhance the model performances. This result, on the one hand, underscores the evident benefits of intermediate fine-tuning within a multi-task framework. On the other hand, it poses questions for future research, emphasizing the need for a more comprehensive exploration of the relative impact of specific linguistic properties.

Looking ahead, there are several avenues for

future research, as detailed in the Limitations section. Additionally, it could be interesting to see if such an approach proves useful with generative Large Language Models (LLMs) in zero- and few-shot configurations. In fact, introducing linguistic knowledge via an instruction-tuning phase might enhance model performance, especially in tasks where linguistic competence plays a crucial role.

Limitations

In this section, we discuss the limitations of our work. 1) **Languages and Tasks.** The objective of our work was to investigate if enhancing a pre-trained model with linguistic knowledge has a positive impact on its performance. However, it's important to note that while the primary focus of our work was not to propose a challenging model that performs better on multiple tasks and benchmarks, testing it on two languages and a single task could be considered a limitation. 2) **Different LLMs.** For our experiments, we primarily focused on models with "moderate" sizes, therefore more efficient and less computationally expensive. Furthermore, we selected T5 since we precisely knew its pre-training dataset. This choice was made in order to mitigate potential concerns about the incorporation of information during the intermediate tuning task that might have already been present in the model's training data. Nevertheless, we acknowledge the relevance of comparing various models in terms of architecture (e.g. encoder- or decoder-only models) and dimensions (e.g. > 1B parameters). 3) **Linguistic features.** Selecting the top 10 linguistic properties on the basis of the correlation with complexity scores is only one of the possible approaches to identifying the set of characteristics to be used for fine-tuning the model. Therefore, using a different selection method or a wider set of linguistic properties could provide more insights on how to complement the model's linguistic knowledge with additional information.

Acknowledgments

This work has been supported by FAIR - Future AI Research (PE00000013) and XAI-CARE - PNRR-MAD-2022-12376692 projects under the NRRP MUR program funded by the NextGenerationEU.

6. Bibliographical References

Jiangang Bai, Yujing Wang, Yiren Chen, Yaming Yang, Jing Bai, Jing Yu, and Yunhai Tong. 2021. [Syntax-BERT: Improving pre-trained transformers with syntax trees](#). In *Proceedings of the 16th*

- Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3011–3020, Online. Association for Computational Linguistics.
- Yonatan Belinkov and James Glass. 2019. [Analysis methods in neural language processing: A survey](#). *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Dominique Brunato, Andrea Cimino, Felice Dell’Orletta, Giulia Venturi, and Simonetta Montemagni. 2020. [Profiling-UD: a tool for linguistic profiling of texts](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 7145–7151, Marseille, France. European Language Resources Association.
- Ting-Yun Chang and Chi-Jen Lu. 2021. [Rethinking why intermediate-task fine-tuning works](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 706–713, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kevyn Collins-Thompson. 2015. Computational assessment of text readability: A survey of current and future research. *Recent Advances in Automatic Readability Assessment and Text Simplification. Special issue of International Journal of Applied Linguistics*, 165(2):97–135.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. [Universal Dependencies](#). *Computational Linguistics*, 47(2):255–308.
- Thomas Givón. 1991. Markedness in grammar: distributional, communicative and cognitive correlates of syntactic structure. *Studies in Language*, 15(2):335–370.
- Goran Glavaš and Ivan Vulić. 2021. [Is supervised syntactic parsing beneficial for language understanding tasks? an empirical investigation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3090–3104, Online. Association for Computational Linguistics.
- Rob van der Goot. 2023. [MaChAmp at SemEval-2023 tasks 2, 3, 4, 5, 7, 8, 9, 10, 11, and 12: On the effectiveness of intermediate training on an uncurated collection of datasets](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 230–245, Toronto, Canada. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). 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 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Fan Huang, Haewoon Kwak, and Jisun An. 2022. Chain of explanation: New prompting method to generate higher quality natural language explanation for implicit hate speech. *arXiv preprint arXiv:2209.04889*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *arXiv preprint arXiv:2205.11916*.
- Andrew K Lampinen, Ishita Dasgupta, Stephanie CY Chan, Kory Matthewson, Michael Henry Tessler, Antonia Creswell, James L McClelland, Jane X Wang, and Felix Hill. 2022. Can language models learn from explanations in context? *arXiv preprint arXiv:2204.02329*.
- Jiaoda Li, Ryan Cotterell, and Mrinmaya Sachan. 2022. [Probing via prompting](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1144–1157, Seattle, United States. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Alessio Miaschi, Dominique Brunato, Felice Dell’Orletta, and Giulia Venturi. 2020. [Linguistic profiling of a neural language model](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 745–756,

- Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Vishakh Padmakumar, Leonard Lausen, Miguel Ballesteros, Sheng Zha, He He, and George Karypis. 2022. [Exploring the role of task transferability in large-scale multi-task learning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2542–2550, Seattle, United States. Association for Computational Linguistics.
- Jason Phang, Thibault Févry, and Samuel R. Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4609–4622, Online. Association for Computational Linguistics.
- Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. [Intermediate-task transfer learning with pretrained language models: When and why does it work?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5231–5247, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Abhilasha Ravichander, Yonatan Belinkov, and Eduard Hovy. 2021. [Probing the probing paradigm: Does probing accuracy entail task relevance?](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3363–3377, Online. Association for Computational Linguistics.
- Gabriele Sarti, Dominique Brunato, and Felice Dell’Orletta. 2021. [That looks hard: Characterizing linguistic complexity in humans and language models](#). In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 48–60, Online. Association for Computational Linguistics.
- Gabriele Sarti and Malvina Nissim. 2022. It5: Large-scale text-to-text pretraining for italian language understanding and generation. *arXiv preprint arXiv:2203.03759*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Tu Vu, Tong Wang, Tsendsuren Munkhdalai, Alessandro Sordani, Adam Trischler, Andrew Mattarella-Micke, Subhransu Maji, and Mohit Iyyer. 2020. [Exploring and predicting transferability across NLP tasks](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7882–7926, Online. Association for Computational Linguistics.
- Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2019a. [Can you tell me how to get past sesame street? sentence-level pretraining beyond language modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4465–4476, Florence, Italy. Association for Computational Linguistics.
- Wei Wang, Bin Bi, Ming Yan, Chen Wu, Jiangnan Xia, Zuyi Bao, Liwei Peng, and Luo Si. 2019b. Structbert: Incorporating language structures into pre-training for deep language understanding. In *International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- Orion Weller, Kevin Seppi, and Matt Gardner. 2022. [When to use multi-task learning vs intermediate fine-tuning for pre-trained encoder transfer learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 272–282, Dublin, Ireland. Association for Computational Linguistics.

- 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. 2020. [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. Association for Computational Linguistics.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Shiwei Zhang and Xiuzhen Zhang. 2021. [Does QA-based intermediate training help fine-tuning language models for text classification?](#) In *Proceedings of the The 19th Annual Workshop of the Australasian Language Technology Association*, pages 158–162, Online. Australasian Language Technology Association.
- Junru Zhou, Zhuosheng Zhang, Hai Zhao, and Shuailiang Zhang. 2020. [LIMIT-BERT : Linguistics informed multi-task BERT](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4450–4461, Online. Association for Computational Linguistics.
- Sofia, Bulgaria. Association for Computational Linguistics.
- Maria Simi, Cristina Bosco, and Simonetta Montemagni. 2014. [Less is more? towards a reduced inventory of categories for training a parser for the Italian Stanford dependencies](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 83–90, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, and et al. 2019. [Universal dependencies 2.5](#). In *LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL)*.

7. Language Resource References

- Dominique Brunato, Lorenzo De Mattei, Felice Dell'Orletta, Benedetta Iavarone, and Giulia Venturi. 2018. [Is this sentence difficult? do you agree?](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2690–2699, Brussels, Belgium. Association for Computational Linguistics.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. [Universal Dependency annotation for multilingual parsing](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 92–97,

A. Intermediate Tasks Suffixes

As mentioned in Sec. 3, we define a set of specific suffixes to be postponed to each input sentence in the Italian/English datasets for the intermediate tasks. Tables 3 and 4 report the suffixes used in the experiments.

B. Model, Parameters and Training Details

We fine-tuned the T5 models with the following hyperparameters:

- Learning rate: 4e-5;
- Per device batch size: 4;
- Epochs: 20 epochs maximum for the complexity prediction task and 25 for the intermediate tasks.

We trained all the models on two NVIDIA GeForce RTX 4090 GPUs.

C. Linguistic Properties

The first step of our approach consists of enhancing the T5 models with multiple linguistic properties in an intermediate fine-tuning process. For the experiments conducted on the Italian datasets we considered the following properties which can be grouped into five main groups:

Raw text feature:

- *char_per_tok*: the **length of tokens** computed as the average number of characters per word in a sentence (excluding punctuation).

Distribution of morpho-syntactic categories:

- *ADJ*: percentage distribution of **adjectives** over the total amount of tokens in a sentence according to the UD parts-of-speech tagset¹⁰;
- *NUM*: percentage distribution of **numerals** over the total amount of tokens in a sentence according to the UD parts-of-speech tagset;
- *PUNCT*: percentage distribution of **punctuation marks** over the total amount of tokens in a sentence according to the UD parts-of-speech tagset;
- *lexical_density*: ratio of content parts-of-speech (verbs, nouns, adjectives and adverbs) over the total number of words in a sentence.

Distribution of verbs by inflectional morphology traits:

¹⁰<https://universaldependencies.org/u/pos/index.html>

- *aux_form_dist_Fin*: percentage distribution of auxiliary verbs (also including modal verbs) in **finite form** over the total number of auxiliary forms in a sentence according to the UD tagset¹¹;
- *aux_mood_dist_Ind*: percentage distribution of auxiliary verbs (also including modal verbs) in **indicative mood** over the total number of auxiliary moods in a sentence according to the UD tagset¹².

Distribution of syntactic dependency relations:

- *aux*: percentage distribution of **auxiliary verbs** (also including modal verbs) over the total amount of dependency relations in a sentence according to the UD tagset¹³;
- *mark*: percentage distribution of **markers**, i.e. words introducing a finite clause subordinate to another clause, over the total amount of dependency relations in a sentence according to the UD tagset.

Property referring to local syntactic tree structure:

- *obj_post*: percentage distribution of direct objects that occur in a post-verbal position in a sentence.

Property referring to the use of subordination:

- *subord_prop_dist*: distribution of subordinate clauses as defined in the UD scheme¹⁴.

For the experiments conducted on the English datasets, we considered the following sets of properties:

Distribution of morpho-syntactic categories:

- *AUX*: percentage distribution of **auxiliary verbs** over the total amount of tokens in a sentence according to the UD parts-of-speech tagset;
- *DET*: percentage distribution of **determiners**, i.e. word, also including definite and indefinite articles, that modify nouns or noun phrases, over the total amount of tokens in a sentence according to the UD parts-of-speech tagset;
- *NUM*: percentage distribution of **numerals** over the total amount of tokens in a sentence according to the UD parts-of-speech tagset;

¹¹<https://universaldependencies.org/u/feat/VerbForm.html>

¹²<https://universaldependencies.org/u/feat/Mood.html>

¹³<https://universaldependencies.org/u/dep/index.html>

¹⁴<https://universaldependencies.org/u/overview/complex-syntax.html#subordination>

Properties	Suffixes
char_per_tok	Il numero medio di caratteri per token nella frase è uguale a <extra_id_0> (transl. The average number of characters per token in the sentence is ...).
upos_dist_ADJ	La distribuzione di aggettivi nella frase è uguale a <extra_id_0> (transl. The distribution of adjectives in the sentence is ...).
upos_dist_NUM	La distribuzione dei numerali nella frase è uguale a <extra_id_0> (transl. The distribution of numbers in the sentence is ...).
lexical_density	Il rapporto fra parole piene e tutte le parole della frase è uguale a <extra_id_0> (transl. The ratio of content words over all words in the sentence is ...).
dep_dist_aux	La distribuzione dei verbi ausiliari nella frase è uguale a <extra_id_0> (transl. The distribution of auxiliary verbs in the sentence is ...).
dep_dist_mark	La distribuzione di marcatori che introducono una clausola subordinata ad un'altra clausola nella frase è uguale a <extra_id_0>. (transl. The distribution of markers introducing a subordinative clause to another clause in the sentence is ...)
aux_mood_dist_Ind	La distribuzione di verbi ausiliari all'indicativo nella frase è uguale a <extra_id_0> (transl. The distribution of indicative mood auxiliary verbs in the sentence is ...).
obj_post	La distribuzione dei verbi in posizione postverbale nella frase è uguale a <extra_id_0> (transl. The distribution of direct objects in a post-verbal position in the sentence is ...).
upos_dist_PUNCT	La distribuzione della punteggiatura nella frase è uguale a <extra_id_0> (transl. The distribution of punctuation in the sentence is ...).
subord_prop_dist	La distribuzione delle subordinate nella frase è uguale a <extra_id_0>. (transl. The distribution of subordinates in the sentence is ...)

Table 3: The suffixes used in the multi-task intermediate fine-tuning experiments for the Italian sentences.

Properties	Suffixes
aux_form_dist_Fin	The distribution of finite form auxiliary verbs in the sentence is <extra_id_0>.
aux_mood_dist_Ind	The distribution of indicative mood auxiliary verbs in the sentence is <extra_id_0>.
dep_dist_compound	The distribution of compounds in the sentence is <extra_id_0>.
dep_dist_nummod	The distribution of numerical modifiers in the sentence is <extra_id_0>.
upos_dist_AUX	The distribution of auxiliary verbs in the sentence is <extra_id_0>.
upos_dist_DET	The distribution of determiners in the sentence is <extra_id_0>.
upos_dist_NUM	The distribution of numerals in the sentence is <extra_id_0>.
upos_dist_PRON	The distribution of pronouns in the sentence is <extra_id_0>.
upos_dist_SYM	The distribution of symbols in the sentence is <extra_id_0>.
subord_prop_dist	The distribution of subordinates in the sentence is <extra_id_0>.

Table 4: The suffixes used in the multi-task intermediate fine-tuning experiments for the English sentences.

- *PRON*: percentage distribution of **pronouns** over the total amount of tokens in a sentence according to the UD parts-of-speech tagset;
- *SYM*: percentage distribution of **symbols**, i.e. word-like entity that differs from ordinary words by form, function, or both, over the total amount of tokens in a sentence according to the UD parts-of-speech tagset.

Distribution of verbs by inflectional morphology traits:

- *aux_form_dist_Fin*: percentage distribution of auxiliary verbs (also including modal verbs) in **finite form** over the total number of auxiliary forms in a sentence according to the UD tagset;
- *aux_mood_dist_Ind*: percentage distribution of auxiliary verbs (also including modal verbs) in **indicative mood** over the total number of auxiliary moods in a sentence according to the UD tagset.

Distribution of syntactic dependency relations:

- *compound*: percentage distribution of **compounds**, including all categories of multi-word expressions, over the total amount of dependency relations in a sentence according to the UD tagset;
- *nummod*: percentage distribution of **numeric modifiers**, i.e. number phrases that serve to modify the meaning of the noun with a quantity, over the total amount of dependency relations in a sentence according to the UD tagset.


Property referring to the use of subordination:

- *subord_prop_dist*: distribution of subordinate clauses as defined in the UD scheme.

D. Multilingual Models Results

We report in Figure 7 and 8 the results (in terms of Spearman correlation scores) for the target task

Intermediate Tasks = Italian



Intermediate Tasks = English

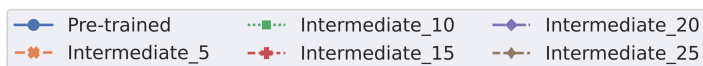



Figure 7: Spearman correlation coefficients for the target task on Italian sentences obtained with the multilingual models using pre-trained and LiT5 models.

obtained with the multilingual models on the Italian (Figure 7) and English (Figure 8) datasets. Table 5 reports instead the performance differences between *Pre-trained* and top-performing multilingual LiT5 models (trained on Italian/English sentences) in both scenarios of limited (1/5) and full (5/5) data availability.

Model	LiT5		LiT5	
	1/5 Training	Diff	5/5 Training	Diff
Intermediate = Italian, Target = Italian				
mt5-small	25 epochs	0.49	10 epochs	0.37
mt5-base	10 epochs	0.22	5 epochs	0.11
Intermediate = English, Target = Italian				
mt5-small	25 epochs	0.31	5 epochs	0.37
mt5-base	5 epochs	0.13	10 epochs	0.08
Intermediate = English, Target = English				
mt5-small	5 epochs	0.46	5 epochs	0.15
mt5-base	5 epochs	0.17	10 epochs	0.01
Intermediate = Italian, Target = English				
mt5-small	15 epochs	0.52	10 epochs	0.16
mt5-base	10 epochs	0.25	10 epochs	0.02

Table 5: Performance differences between *Pre-trained* and top-performing LiT5 models (trained on Italian/English sentences) in both scenarios of limited (1/5) and full (5/5) data availability.

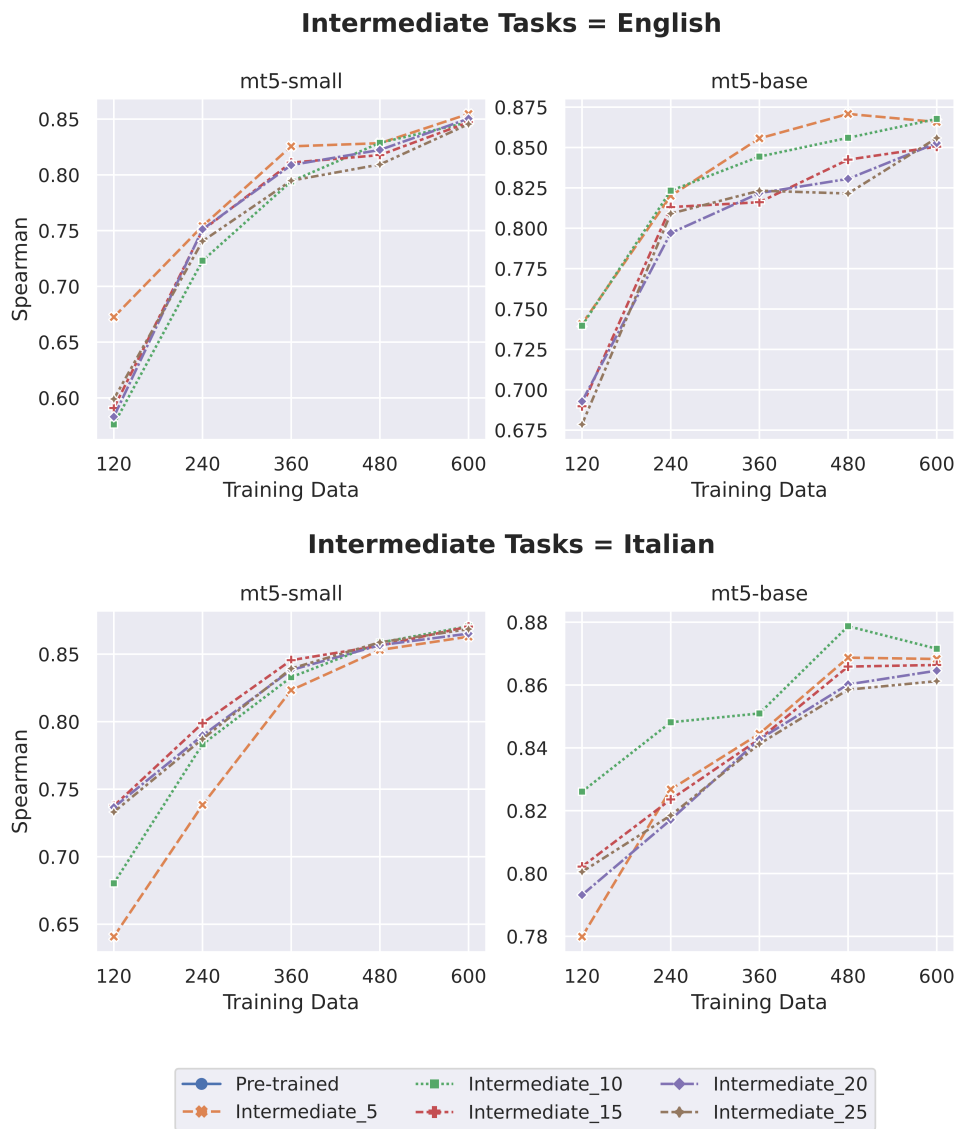


Figure 8: Spearman correlation coefficients for the target task on English sentences obtained with the multilingual models using pre-trained and LiT5 models.