

Evidence > Intuition: Transferability Estimation for Encoder Selection

Elisa Bagnina[⊖] Max Müller-Eberstein[⊖] Mike Zhang[⊖] Barbara Plank^{⊖,▲,♣}

[⊖]Department of Computer Science, IT University of Copenhagen, Denmark

[▲]Center for Information and Language Processing (CIS), LMU Munich, Germany

[♣]Munich Center for Machine Learning (MCML), Munich, Germany

{elba, mamy, mikz}@itu.dk b.plank@lmu.de

Abstract

With the increase in availability of large pre-trained language models (LMs) in Natural Language Processing (NLP), it becomes critical to assess their fit for a specific target task a priori—as fine-tuning the entire space of available LMs is computationally prohibitive and unsustainable. However, *encoder transferability estimation* has received little to no attention in NLP. In this paper, we propose to generate quantitative evidence to predict which LM, out of a pool of models, will perform best on a target task *without* having to fine-tune all candidates. We provide a comprehensive study on LM ranking for 10 NLP tasks spanning the two fundamental problem types of classification and structured prediction. We adopt the state-of-the-art Logarithm of Maximum Evidence (LogME) measure from Computer Vision (CV) and find that it positively correlates with final LM performance in 94% of the setups. In the first study of its kind, we further compare transferability measures with the de facto standard of human practitioner ranking, finding that evidence from quantitative metrics is more robust than pure intuition and can help identify unexpected LM candidates.

1 Introduction

Advances in Deep Learning-based NLP and CV build on expressive representations from encoder models pre-trained on massive corpora. Downstream models make use of latent information in these representations to extract relevant features for the task at hand. Within this paradigm, deciding which pre-trained encoder to use in any task-specific architecture is crucial, however training a model using each encoder candidate is infeasible. In absence of prior heuristics (e.g., via related work), the choice of encoder has therefore prevalently been based on practitioner intuition rather than quantitative evidence.

In NLP, prior work has examined the different yet related task of performance prediction (Xia et al., 2020a; Ye et al., 2021), surveyed and categorized LMs (Xia et al., 2020b), and used probing to predict LM performance specifically for dependency parsing (Müller-Eberstein et al., 2022b), but has yet to extensively investigate how to rank the increasingly large number of pre-trained LM encoders across various tasks and domains. Preliminary work by You et al. (2021) shows that the LogME estimator holds promise, including the first steps for encoder selection in NLP. With their main focus being on CV, however, they evaluate only a limited set of tasks and models for NLP and use self-reported benchmark scores instead of running controlled experiments which should include, e.g., the variance across initializations, domains, and fine-tuning strategies (Section 2). As such, we seek to answer: *How well can we estimate the transferability of pre-trained LMs to specific NLP tasks?* To do so, we contribute:

- The broadest encoder selection study in NLP to date, on 10 domain-diverse classification and structured prediction tasks (Section 3);
- An extensive evaluation and analysis across multiple dimensions of variation, including seven general vs. domain-specific LMs, [CLS] vs. mean representations, and head vs. full model fine-tuning (Section 4);
- A study with NLP experts, comparing the prevailing ranking of LMs by human intuition with LogME’s empirical evidence (Section 5);
- Guidelines for applying and interpreting transferability measures to NLP (Section 6), and an open-source toolkit for efficient, task-adaptive LM pre-selection.¹

[⊖] The authors contributed equally to this work.

¹<https://github.com/mainlp/logme-nlp>

	DATASET	TASK	TRAIN / DEV	$ \mathcal{Y} $	METRIC
CLASSIFICATION	AGNews (Zhang et al., 2015)	Topic Classification	84K / 12K	4	micro-F1
	Airline (Crowdfunder, 2020)	Sentiment Analysis	10K / 1.5K	3	micro-F1
	SciERC (Luan et al., 2018)	Relation Classification	1.9K / 275	7	macro-F1
	MNLI (Williams et al., 2018)	Natural Language Inference	393K / 20K	3	micro-F1
	QNLI (Rajpurkar et al., 2016)	Q&A/Natural Language Inference	105K / 5.4K	2	micro-F1
	RTE (Giampiccolo et al., 2007)	Natural Language Inference	2.5K / 3K	3	micro-F1
STR. PRED.	EWT (Silveira et al., 2014)	Dependency Labeling	12.5k / 2k	36	micro-F1
	CrossNER (Liu et al., 2021)	Named Entity Recognition	15K / 3.5K	4	span-F1
	CrossNER (Liu et al., 2021)	Named Entity Recognition	200 / 450	17	span-F1
	JobStack (Jensen et al., 2021)	De-identification	18K / 2K	11	span-F1

Table 1: **Datasets.** Indicated are the 10 datasets used in this study, distinguished between the two NLP problem types C and SP for a wide variety of tasks and domains. C tasks cover AGNews (news articles), Twitter Airline Sentiment (Airline; Twitter feedback), SciERC (AI proceedings), MNLI (speech, (non-)fiction, government), QNLI (Wikipedia) and RTE (Wikipedia, news). Within the SP tasks, we experiment on the English Web Treebank (EWT; social media, reviews, emails), CrossNER (news, scientific Wikipedia) and JobStack (Stack Overflow job ads). For each task, we report their TRAIN/DEV split, label space, and task-specific performance metric.

2 Transferability Estimation

Transferability estimation aims to quantify the ability of a model to transfer knowledge learned from one task to another (Eaton et al., 2008; Sinapov et al., 2015). Formally, given a pool of L pre-trained LMs $\{\phi_l\}_{l=1}^L$ and a dataset \mathcal{D} , we calculate a predictive score $S_l(\mathcal{D})$ for each ϕ_l which ideally correlates with the model’s final performance $P_l(\mathcal{D})$. $S_l(\mathcal{D})$ is computed without fine-tuning ϕ_l on \mathcal{D} such that the optimal ϕ_l^* can be chosen from a large model pool at a low computational cost.

The CV community has begun to explore methods for encoder pre-selection and ranking through metrics such as LogME and the Log Expected Empirical Prediction (LEEP; Nguyen et al., 2020). These are widely-used state-of-the-art methods in CV. Recent work introduced the Gaussian Bhattacharyya Coefficient (GBC; Pándy et al., 2021) and Optimal Transport based Conditional Entropy (OTCE; Tan et al., 2021), the exploration of which we leave for future work. However, in the NLP field, related work focus on choosing a task and *not* an LM encoder for transferability (Vu et al., 2020; Padmakumar et al., 2022), leaving the ranking of encoders an unexplored question.

LogME LogME measures the suitability of all encoded dataset features $F \in \mathbb{R}^{|D| \times h}$ (e.g., embeddings with dimensionality h) to predict all scalar labels $y \in \mathbb{R}^{|D|}$ via the probability density $p(y|F)$. As this density is intractable, it is estimated by mapping $F \rightarrow y$ using a linear transformation w ; this is akin to training a linear probe with optimal param-

eters w^* and using the likelihood $p(y|F, w^*)$ as a proxy for feature suitability. Because a simple linear model will overfit on the training data, it would be beneficial to obtain the marginal likelihood, or evidence, by integrating over all possible values of w : $p(y|F) = \int p(y|F, w)p(w)dw$. To once again make this computation tractable, You et al. (2021) reformulate it as an efficient, iterative evidence maximization problem where both w as well as y are drawn from lightly parametrized, isotropic Gaussian distributions. The normalized logarithm of the maximized evidence (LogME) can then be used as $S_l(\mathcal{D})$ to rank encoder models directly.

NLP Setting LogME has shown promise for CV, and an initial study on the GLUE benchmark (Wang et al., 2018) indicate the same for NLP (You et al., 2021). However, for NLP, there are notable differences in setups across tasks. We adapt and apply LogME extensively to a wide range of NLP settings to identify empirically grounded guidelines.

In particular, we investigate variations concerning the task, instance granularity, domain, and tuning strategy. First, compared to most image classification tasks, NLP tasks are subject to differences in granularity, i.e., **classification** (C) and **structured prediction** (SP). Furthermore, there is less clarity than for individual images as to which representation best captures the full language input (Mosbach et al., 2020). Therefore, for C setups we experiment with two representations: i.e., using [CLS]/<s> versus mean over sequence/subwords.

Second, depending on differences in the data

domain, NLP practitioners are often faced with a pool of domain-adapted LMs in addition to more general-purpose encoders—the correct choice of which may not be immediately apparent.

Finally, the best performance in NLP is often achieved using full fine-tuning, while CV models usually do not fine-tune the encoder (Peters et al., 2019). It will therefore be crucial to investigate whether the predictive performance of $S_l(\mathcal{D})$ holds when it is computed based on untuned F while $P_l(\mathcal{D})$ is based on fully fine-tuned representations.

3 Experimental Setup

Applying seven architecturally and domain-diverse pre-trained LMs with up to four configurations each to 10 datasets and a wide variety of tasks, we investigate LogME’s predictive power for transferability estimation in NLP—for a total of 280 setups. We refer to Table 1 for our detailed set of tasks.

Language Models We pick seven pre-trained LMs with a wide domain and architectural variety from the Transformers library’s model hub (Wolf et al., 2020). Three are “general-purpose” models, namely BERT_{base} (Devlin et al., 2018), RoBERTa_{base} (Liu et al., 2019), and DistilBERT_{base} (Sanh et al., 2019). Four models are pre-trained on domain-specific corpora, namely Clinical-BioBERT (Alsentzer et al., 2019), BioBERT (Lee et al., 2020), Twitter-RoBERTa_{base} (Barbieri et al., 2020), and SciBERT_{base} (Beltagy et al., 2019). Note that for BioBERT variants domain-adaptive pre-training has been applied (Gururangan et al., 2020).

Model Setups The model setup follows the same structure for each task: A pre-trained LM encoder and a 3-layer perceptron head, following Tenney et al. (2019). The input to the latter is either the [CLS] token or mean over sequence subwords for C tasks or mean over token subwords for SP tasks. While it is common in CV to keep the encoder frozen and only fine-tune the task-specific head, we also evaluate the practice of full model fine-tuning, as is more common in NLP (Peters et al., 2019). Considering these variations (frozen vs. fine-tuning, and [CLS] vs. mean), we obtain up to four setups per C task and two setups per SP task. Each experiment is run with five random seeds. Details for reproducibility can be found in Appendix A.

Evaluation Following You et al. (2021), we evaluate LogME’s predictive power for ranking

LMs according to their final performance by using the two correlation coefficients Pearson’s ρ and weighted Kendall’s τ_w (Vigna, 2015), both in $[-1, 1]$. Kendall’s τ_w further allows us to estimate the probability of a higher-ranked LM actually performing better by computing $\frac{\tau_w+1}{2}$.

4 Analysis of Results

Our results across all setups are consolidated in Figure 1 and Figure 2 (C: blue, SP: beige).^{2,3} The left of each figure plots the performance using frozen LM embeddings (✳) against LogME scores, while on the right, full LM fine-tuning is applied (⚡).⁴

Figure 1 shows the results of using mean-pooled embeddings in both C/SP settings. For ✳, we obtain $\rho > 0.8$ on 8/10 tasks and $\tau_w > 0.7$ on 6/10 tasks, indicating a strong relationship between model performance and LogME. After fine-tuning (⚡), we observe a general reduction in ρ and τ_w (most on CrossNER, EN-EWT), however overall correlations remain positive to a significant degree.

For C setups using the alternative [CLS]/<s> representations (Figure 2), LogME correlates highly at $\rho > 0.95$ on 5/6 tasks and $\tau_w > 0.7$ on 4/6 tasks when using head-only tuning (✳). After full fine-tuning (⚡), SciERC, RTE and AGNews have lower correlations, particularly with the high-variance RoBERTa model. However, the remaining tasks maintain a stable correlation, with $\rho > 0.6$ and $\tau_w > 0.3$ across 5/6 tasks.

Overall, LogME has a positive correlation with final performance in 30/32 cases. In more detail, LogME has a $\tau_w > 0.41$ in 20/32 setups, meaning that selecting a higher ranked model is the better choice 71% of the time. LogME both identifies intuitive, domain-specific scenarios (e.g., Twitter-RoBERTa performing well on Airline Twitter), but also finds cases that may be unintuitive, such as DistilBERT’s occasionally high performance for CrossNER and JobStack. This finding holds across C, SP, domains as well as different input representations. For the latter, we note that, surprisingly, even the untuned representation of [CLS]/<s> seems to contain useful information with comparable performance to mean pooling.

Comparing ✳ versus ⚡, we notice that, as expected, model performance improves, but in

²Exact results can be found in Appendix B.

³We successfully reproduce the results reported in You et al. (2021) for MNL, QNLI and RTE.

⁴Note that LogME is only computed on frozen embeddings and does not differ between ✳ and ⚡.




Figure 1: **Results of Mean Pooling (μ)**. We plot the model’s LogME scores against their task-specific performances on each dataset based on mean pooling the token embeddings (**left**: Frozen embeddings ($*$), **right**: Full model fine-tuning (⚡)). Task-types are indicated in specific colors: Lightblue for C and beige for SP. Further reported are the Pearson correlation coefficient (ρ) and weighted Kendall’s tau (τ).

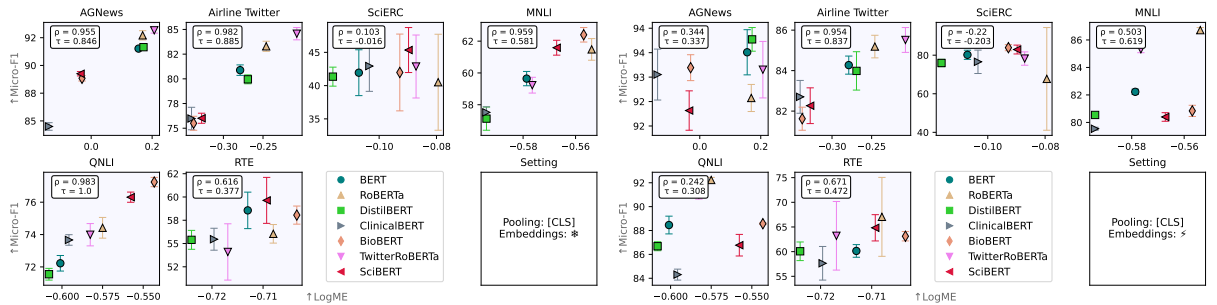


Figure 2: **Results of [CLS]**. We plot the model’s LogME scores against their task-specific performances on each dataset based on the vector representation of the [CLS] token (**left**: Frozen embeddings ($*$), **right**: Full model fine-tuning (⚡)). Reported are the Pearson correlation coefficient (ρ) and weighted Kendall’s tau (τ).

general, LogME’s predictive power decreases. The fully fine-tuned model makes predictions on updated representations such that decreases in predictive performance are inevitable unless the initial LM already represents a local optimum for the task at hand. This fact is crucial for NLP practitioners where full fine-tuning is the standard practice. Taking these factors into account, LogME’s efficiency is especially beneficial, as it offers an $86\times$ speedup over full model fine-tuning (You et al., 2021), and its positive correlation in 94% of our evaluated setups indicates that it is an effective score for transferability estimation in NLP.

5 Human Performance

Given the lack of prior work examining transferability estimation of pre-trained LM encoders, the most common method for encoder selection employed today is practitioner intuition. As such, we conduct a small-scale study with 12 NLP

practitioners and ask them to perform the same ranking as in Section 3. Despite having access to model details and associated papers, this task is difficult even for experts. While for LogME, the range of τ_w is in $[-0.20; 1.00]$, human rankings fall into a wider range of $[-0.54; 1.00]$, indicating higher uncertainty. Similarly, we observe that human correlations are negative thrice as often as for LogME. Additionally, LogME provides a continuous scale for comparing models, while human rankings offer no indication of relative performance differences. At the same time, they are more inaccurate for tasks without an associated domain-specific model (e.g., news, mixture of genres in EWT). Moreover, even when domains are clear (e.g., Twitter, science), LogME tends to be more accurate than the predictions of most human participants. Finally, the high variance between practitioners and the fact that no single person was an expert in all setups further reinforces the necessity of quantitative transferability scores.

6 Conclusion

We show the value of transferability estimation for selecting high-performing LMs before full model fine-tuning in experiments, covering the two fundamental NLP settings of classification and structured prediction. By adopting the state-of-the-art LogME scoring method, we are able to rank LMs on a continuous scale which correlates with final performance—with the better encoder being chosen in 71% of cases. Additionally, we identify NLP-specific guidelines for transferability estimation: In particular, predicting the best LM for tasks/domains which greatly deviate from an encoder’s pre-training setup and require large amounts of full fine-tuning may require larger pre-selections of LMs due to the higher uncertainty of the scoring methods. Finally, our human study showed that practitioners frequently misconstrue the performance of LMs even on domain-specific tasks. As such transferability quantification methods provide valuable evidence over intuition.

Limitations

A key limitation that practitioners should consider is that, while LogME is viable for the quantitative transferability estimation of LM encoders, there is a noticeable drop in predictive accuracy after full model fine-tuning. We attribute this to the misalignment between the frozen representations of the encoder, which LogME is applied to, and the representations after fine tuning. As stated in Section 4, unless the untuned LM already constitutes a local optimum for the task at hand, task-specific shifts in its parameters and representations are inevitable.

This similarly applies to cases where the untuned representations differ substantially from what a fully fine-tuned model uses during training. Specifically, for the relation classification task of SciERC, it is important to note that the input given to the model is augmented with special tokens delimiting the entities involved in the relation (Baldini Soares et al., 2019) which are unknown to the untuned model and thus the representations that LogME is computed on. Furthermore, for EN-EWT we suspect that dependency labeling is a more fundamental task solvable with high accuracy by most LMs, especially after fine-tuning as reflected in micro-F1 scores between 93–95. This is mirrored by work on probing untuned LMs which identifies high levels of inherent dependency information (Tenney et al., 2019; Müller-Eberstein et al., 2022a).

Such sensitivity to representational shifts is not exclusive to LogME: In preliminary experiments, we examined LEEP (Nguyen et al., 2020) as an alternative predictive score $S_l(\mathcal{D})$. Its original use was to rank the transferability of a classifier trained on one dataset, to a new task—leaving the ranking of pre-trained LMs for future work. LEEP has so far only been applied to CV tasks, but we apply it to LM ranking on the collection of NLP tasks above. Our initial experiments achieved low and unintuitive correlations between LEEP’s $S_l(\mathcal{D})$ and $P_l(\mathcal{D})$. We speculate that this is due to the absence of a normalizing factor over the number of source classes, i.e., the high number of embedding dimensions in our case (see Equation 2 in Nguyen et al., 2020). While it would further be valuable to investigate methods beyond LEEP and LogME, as mentioned in Section 2, we leave their evaluation on NLP to future work. At the time of writing, the former two were the most extensively explored in CV, in addition to the original LogME work containing an initial study showing promise for NLP.

Finally, our human ranking study in Section 5 was limited by the number of practitioners with a publication record which we could contact confidentially. However, the group still constituted a diverse set over seniority, gender, and cultural background. A larger group would cover a broader range of backgrounds and may produce different rankings. However, as the surveyed group already displayed high variance, overall predictive performance is unlikely to be significantly higher.

Keeping these limitations in mind, correlations do remain mostly positive for LogME and scores are well suited to be applied to high-dimensional embedding spaces, such that it offers a predictive and efficient measure for quantifying transferability compared to human practitioner intuition.

Ethics Statement

It is difficult to foresee ethical issues for this work due to the broad applicability of LM encoder pre-selection. To the best of our knowledge, in the CV community from which our evaluated scoring methods originate, there have been no harmful applications thus far. In fact, as fine-tuning the entire space of available language models is unsustainable and unethical in terms of climate sustainability, efficient encoder pre-selection methods such as LogME provide a positive first step towards tackling this problem.

Acknowledgements

We would like to wholeheartedly thank the members NLPnorth group at the IT University of Copenhagen, The Center for Information and Language Processing at the Ludwig Maximilian University of Munich, and the Dialogue Modeling Group at the University of Amsterdam for comments on earlier versions of this work and participating in the human practitioner ranking study.

MZ and BP are supported by the Independent Research Fund Denmark (DFF) grant 9131-00019B. EB, MME, and BP are supported by the Independent Research Fund Denmark (DFF) Sapere Aude grant 9063-00077B. BP is supported by the ERC Consolidator Grant DIALECT 101043235.

References

- Emily Alsentzer, John Murphy, William Boag, Weihung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. [Matching the blanks: Distributional similarity for relation learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2895–2905, Florence, Italy. Association for Computational Linguistics.
- Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. [TweetEval: Unified benchmark and comparative evaluation for tweet classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Crowdfunder. 2020. Twitter us airline sentiment. <https://www.kaggle.com/crowdfunder/twitter-airline-sentiment>. Accessed: 2022-02-20.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Eric Eaton, Marie Desjardins, and Terran Lane. 2008. Modeling transfer relationships between learning tasks for improved inductive transfer. In *Joint european conference on machine learning and knowledge discovery in databases*, pages 317–332. Springer.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360.
- Kristian Nørgaard Jensen, Mike Zhang, and Barbara Plank. 2021. [De-identification of privacy-related entities in job postings](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 210–221, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- 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](#). *CoRR*, abs/1907.11692.
- Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. 2021. Crossner: Evaluating cross-domain named entity recognition. In *AAAI*.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- Marius Mosbach, Anna Khokhlova, Michael A. Hedderich, and Dietrich Klakow. 2020. [On the Interplay Between Fine-tuning and Sentence-level Probing for Linguistic Knowledge in Pre-trained Transformers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2502–2516, Online. Association for Computational Linguistics.
- Max Müller-Eberstein, Rob van der Goot, and Barbara Plank. 2022a. [Probing for labeled dependency trees](#). In *Proceedings of the 60th Annual Meeting of the*

- Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7711–7726, Dublin, Ireland. Association for Computational Linguistics.
- Max Müller-Eberstein, Rob van der Goot, and Barbara Plank. 2022b. [Sort by structure: Language model ranking as dependency probing](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1296–1307, Seattle, United States. Association for Computational Linguistics.
- Cuong Nguyen, Tal Hassner, Matthias Seeger, and Cedric Archambeau. 2020. Leep: A new measure to evaluate transferability of learned representations. In *International Conference on Machine Learning*, pages 7294–7305. PMLR.
- 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. *arXiv preprint arXiv:2204.11117*.
- Michal Pándy, Andrea Agostinelli, Jasper Uijlings, Vittorio Ferrari, and Thomas Mensink. 2021. Transferability estimation using bhattacharyya class separability. *arXiv preprint arXiv:2111.12780*.
- Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. [To tune or not to tune? adapting pretrained representations to diverse tasks](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 7–14, Florence, Italy. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108.
- Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.
- Jivko Sinapov, Sanmit Narvekar, Matteo Leonetti, and Peter Stone. 2015. Learning inter-task transferability in the absence of target task samples. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 725–733.
- Yang Tan, Yang Li, and Shao-Lun Huang. 2021. Otce: A transferability metric for cross-domain cross-task representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15779–15788.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Sebastiano Vigna. 2015. A weighted correlation index for rankings with ties. In *Proceedings of the 24th international conference on World Wide Web*, pages 1166–1176.
- 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, 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.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 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.
- Mengzhou Xia, Antonios Anastasopoulos, Ruochen Xu, Yiming Yang, and Graham Neubig. 2020a. [Predicting performance for natural language processing tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8625–8646, Online. Association for Computational Linguistics.
- Patrick Xia, Shijie Wu, and Benjamin Van Durme. 2020b. [Which *BERT? A survey organizing contextualized encoders](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7516–7533, Online. Association for Computational Linguistics.
- Zihuiwen Ye, Pengfei Liu, Jinlan Fu, and Graham Neubig. 2021. [Towards more fine-grained and reliable NLP performance prediction](#). In *Proceedings of the*

16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, pages 3703–3714, Online. Association for Computational Linguistics.

Kaichao You, Yong Liu, Jianmin Wang, and Mingsheng Long. 2021. Logme: Practical assessment of pre-trained models for transfer learning. In *International Conference on Machine Learning*, pages 12133–12143. PMLR.

Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.

Appendix

A Reproducibility

Each model is trained on an NVIDIA A100 GPU with 40GBs of VRAM and an AMD Epyc 7662 CPU. The seed numbers the models are initialized with are 4012, 5060, 9908, 8857, 8823. We run the models for 30 epochs with a patience of 3 on each respective dev. data. We use a batch size of 16, 32, or 64 depending on the size of the dataset. When keeping the language model weights frozen, we use a learning rate of $1e-3$. For full model fine-tuning, the learning rate is set at $5e-5$. On GLUE, JobStack, and CrossNER (News), we observed training instability and set the learning rate to $5e-7$. The evaluated LMs have between 66M parameters (DistilBERT_{base}) and 125M parameters (RoBERTa_{base}), taking between 10 minutes (SciERC) and 3 days (e.g., AGNews, MNLI) to fully fine-tune. Keeping the LM frozen and only fine-tuning the task-specific head is around 70% more time-efficient. Computing LogME requires one forward pass to embed the dataset instances, before completing the score calculation in under 1 minute.

B Exact Results

In [Table 2](#) and [Table 3](#), we present the exact performance numbers shown in [Figure 1](#) and [Figure 2](#). The results here are separated by task.

DATASET	LANGUAGE MODEL	μ			[CLS]		
		LOGME	FROZEN	TUNED	LOGME	FROZEN	TUNED
AGNews	bert-base-uncased	0.0822	92.62±0.13	93.51±0.23	0.1555	91.52±0.10	93.51±0.46
	roberta-base	0.1628	93.30±0.17	91.70±1.81	0.1689	92.71±0.42	92.57±0.28
	distilbert-base-uncased	0.1786	92.26±0.37	93.85±0.11	0.1716	91.65±0.22	93.77±0.24
	emilyalsentzer/Bio_ClinicalBERT	-0.1801	87.52±0.29	92.62±0.62	-0.1384	84.50±0.34	93.05±0.52
	dmis-lab/biobert-v1.1	-0.0548	90.05±0.29	93.16±0.23	-0.0300	88.84±0.25	93.19±0.26
	cardiffnlp/twitter-roberta-base	0.1768	92.55±0.21	93.34±0.55	0.2070	93.16±0.25	93.15±0.58
	allenai/scibert_scivocab_uncased	-0.0527	90.06±0.17	92.40±0.47	-0.0348	89.25±0.20	92.32±0.40
	ρ, τ_w		0.954, 0.330	0.240, 0.559		0.955, 0.846	0.344, 0.337
Airline	bert-base-uncased	-0.2484	82.58±0.19	84.03±1.10	-0.2789	80.88±0.53	84.27±0.45
	roberta-base	-0.2407	84.10±0.52	85.43±0.77	-0.2460	83.29±0.51	85.19±0.56
	distilbert-base-uncased	-0.2612	81.71±0.39	83.89±1.21	-0.2691	79.95±0.45	83.99±0.95
	emilyalsentzer/Bio_ClinicalBERT	-0.3205	78.46±0.57	83.17±0.62	-0.3402	75.98±1.14	82.70±0.81
	dmis-lab/biobert-v1.1	-0.3295	76.67±0.93	82.55±0.96	-0.3376	75.50±0.71	81.62±0.57
	cardiffnlp/twitter-roberta-base	-0.2094	84.89±0.21	86.05±0.90	-0.2074	84.57±0.62	85.51±0.61
	allenai/scibert_scivocab_uncased	-0.3122	77.58±0.51	82.05±0.58	-0.3275	76.01±0.51	82.27±0.88
	ρ, τ_w		0.982, 0.953	0.922, 0.912		0.982, 0.885	0.954, 0.837
SciERC	bert-base-uncased	-0.0663	49.56±3.55	75.84±3.21	-0.1071	41.94±3.45	80.20±2.37
	roberta-base	-0.0752	51.07±5.34	78.80±3.34	-0.0794	40.51±7.22	67.71±26.6
	distilbert-base-uncased	-0.0816	45.98±5.17	73.13±3.20	-0.1161	41.35±1.43	75.95±1.93
	emilyalsentzer/Bio_ClinicalBERT	-0.0669	48.64±2.26	73.61±2.44	-0.1034	42.94±3.80	76.57±6.06
	dmis-lab/biobert-v1.1	-0.0546	56.64±4.85	81.60±4.13	-0.0928	41.98±5.77	83.89±1.58
	cardiffnlp/twitter-roberta-base	-0.0825	46.75±3.35	76.65±1.58	-0.0871	42.87±4.73	78.25±3.46
	allenai/scibert_scivocab_uncased	-0.0377	58.83±1.61	80.12±4.68	-0.0897	45.35±3.38	82.93±2.30
	ρ, τ_w		0.930, 0.825	0.631, 0.521		0.103, -0.016	-0.220, -0.203
MNLI	bert-base-uncased	-0.5818	59.18±0.39	81.85±0.31	-0.5786	59.64±0.45	82.23±0.19
	roberta-base	-0.5815	64.18±0.19	86.57±0.24	-0.5539	61.48±0.68	86.71±0.19
	distilbert-base-uncased	-0.5938	58.13±0.58	79.64±0.39	-0.5940	57.13±0.73	80.54±0.09
	emilyalsentzer/Bio_ClinicalBERT	-0.6154	56.53±0.35	79.21±2.44	-0.5940	57.52±0.45	79.54±0.11
	dmis-lab/biobert-v1.1	-0.5841	60.12±0.33	80.89±4.13	-0.5569	62.40±0.51	80.84±0.41
	cardiffnlp/twitter-roberta-base	-0.5826	61.77±0.36	85.41±1.58	-0.5765	59.23±0.13	85.32±0.25
	allenai/scibert_scivocab_uncased	-0.5787	59.57±0.44	80.41±4.68	-0.5672	61.59±0.28	80.40±0.31
	ρ, τ_w		0.698, 0.429	0.532, 0.384		0.959, 0.581	0.503, 0.619
QNLI	bert-base-uncased	-0.5823	75.75±0.11	88.17±0.19	-0.6008	72.23±0.48	88.46±0.74
	roberta-base	-0.5557	78.09±0.39	92.17±0.26	-0.5749	74.42±0.64	92.23±0.22
	distilbert-base-uncased	-0.5881	74.25±0.44	86.26±0.33	-0.6079	71.55±0.36	86.68±0.38
	emilyalsentzer/Bio_ClinicalBERT	-0.5908	74.69±0.38	84.13±0.27	-0.5957	73.67±0.33	84.31±0.46
	dmis-lab/biobert-v1.1	-0.5502	78.21±0.26	88.19±0.42	-0.5432	77.25±0.29	88.57±0.07
	cardiffnlp/twitter-roberta-base	-0.5728	77.49±0.20	91.22±0.41	-0.5826	73.99±0.69	91.03±0.39
	allenai/scibert_scivocab_uncased	-0.5737	76.84±0.30	87.24±0.26	-0.5577	76.31±0.32	86.77±0.90
	ρ, τ_w		0.933, 0.960	0.663, 0.621		0.983, 1.000	0.242, 0.308
RTE	bert-base-uncased	-0.7160	56.26±1.28	62.09±1.21	-0.7131	58.56±1.96	60.14±1.28
	roberta-base	-0.7133	58.35±8.00	68.99±1.58	-0.7081	56.04±1.00	67.05±8.00
	distilbert-base-uncased	-0.7220	53.96±1.89	57.63±3.00	-0.7237	55.40±1.02	60.07±1.89
	emilyalsentzer/Bio_ClinicalBERT	-0.7155	58.13±0.94	58.99±1.86	-0.7196	55.46±1.18	57.64±3.41
	dmis-lab/biobert-v1.1	-0.7160	56.97±1.74	59.98±4.95	-0.7034	58.05±0.97	63.12±0.95
	cardiffnlp/twitter-roberta-base	-0.7176	54.02±0.53	66.63±0.34	-0.7169	54.10±3.03	63.21±6.95
	allenai/scibert_scivocab_uncased	-0.7121	55.46±1.79	64.65±1.51	-0.7093	59.64±2.47	64.83±2.66
	ρ, τ_w		0.616, 0.450	0.597, 0.324		0.616, 0.377	0.671, 0.472

Table 2: **Exact Results of Classification Tasks.** We indicate the LOGME score of each model (LANGUAGE MODEL) and its performance on a wide variety of datasets (DATASET) in different settings (FROZEN, TUNED) by either taking the representations of the tokens and apply mean pooling (μ) or the representation of the [CLS] token. Given the LogME scores and the performance metrics, we can calculate the Pearson correlation coefficient (ρ) and the weighted Kendall’s tau (τ_w).

DATASET	LANGUAGE MODEL	μ		
		LOGME	FROZEN	TUNED
EN-EWT	bert-base-uncased	1.2367	85.04±0.12	94.16±0.10
	roberta-base	1.2681	86.10±0.13	94.85±0.18
	distilbert-base-uncased	1.2864	86.98±0.12	93.36±0.09
	emilyalsentzer/Bio_ClinicalBERT	1.2617	85.05±0.20	93.10±0.06
	dmis-lab/biobert-v1.1	1.2583	85.95±0.25	93.16±0.19
	cardiffnlp/twitter-roberta-base	1.2826	86.50±0.07	94.82±0.15
	allenai/scibert_scivocab_uncased	1.2837	87.54±0.26	93.29±0.09
	ρ, τ_w		0.858, 0.760	-0.022, 0.013
CrossNER (News)	bert-base-uncased	0.8397	87.66±0.33	92.53±0.17
	roberta-base	0.8290	88.08±0.65	94.59±0.17
	distilbert-base-uncased	0.8867	88.41±0.79	91.21±0.64
	emilyalsentzer/Bio_ClinicalBERT	0.6527	69.86±1.26	78.01±0.47
	dmis-lab/biobert-v1.1	0.7666	81.48±0.92	89.63±0.35
	cardiffnlp/twitter-roberta-base	0.8460	88.55±0.53	94.23±0.13
	allenai/scibert_scivocab_uncased	0.7897	82.38±0.39	88.16±0.18
	ρ, τ_w		0.974, 0.732	0.897, 0.257
CrossNER (Sci.)	bert-base-uncased	1.4339	43.22±1.51	38.68±17.3
	roberta-base	1.4297	47.00±0.90	62.27±4.02
	distilbert-base-uncased	1.4444	45.96±2.85	37.97±18.6
	emilyalsentzer/Bio_ClinicalBERT	1.3772	32.89±1.66	20.96±13.8
	dmis-lab/biobert-v1.1	1.4166	43.24±1.81	47.73±5.17
	cardiffnlp/twitter-roberta-base	1.4207	45.51±0.94	54.05±4.61
	allenai/scibert_scivocab_uncased	1.4205	43.98±1.24	53.44±4.13
	ρ, τ_w		0.906, 0.471	0.537, 0.010
JobStack	bert-base-uncased	1.7750	73.64±1.30	78.49±1.06
	roberta-base	1.7827	74.06±1.96	81.51±1.02
	distilbert-base-uncased	1.7998	74.96±2.03	77.02±0.34
	emilyalsentzer/Bio_ClinicalBERT	1.7056	61.13±0.99	67.07±0.60
	dmis-lab/biobert-v1.1	1.7508	68.32±2.58	74.65±0.49
	cardiffnlp/twitter-roberta-base	1.7793	73.72±3.00	79.99±1.06
	allenai/scibert_scivocab_uncased	1.7621	71.66±1.93	78.72±1.54
	ρ, τ_w		0.981, 1.000	0.863, 0.409

Table 3: **Exact Results of Structured Prediction Tasks.** We indicate the LOGME score of each model (LANGUAGE MODEL) and its performance on a wide variety of datasets (DATASET) in different settings (FROZEN, TUNED) by taking the representations of the tokens and apply mean pooling (μ). Here we do not take the representation of the [CLS] token as this has no meaning for the structured prediction task. Given the LogME scores and the performance metrics, we can calculate the Pearson correlation coefficient (ρ) and the weighted Kendall’s tau (τ_w).