

Exploring Predictive Uncertainty and Calibration in NLP: A Study on the Impact of Method & Data Scarcity

Dennis Ulmer¹ Jes Frellsen² Christian Hardmeier²

¹Department of Computer Science, IT University of Copenhagen

²Department of Applied Mathematics & Computer Science, Technical University of Denmark
dennis.ulmer@mailbox.org

Abstract

We investigate the problem of determining the predictive confidence (or, conversely, uncertainty) of a neural classifier through the lens of low-resource languages. By training models on sub-sampled datasets in three different languages, we assess the quality of estimates from a wide array of approaches and their dependence on the amount of available data. We find that while approaches based on pre-trained models and ensembles achieve the best results overall, the quality of uncertainty estimates can surprisingly *suffer* with more data. We also perform a qualitative analysis of uncertainties on sequences, discovering that a model’s total uncertainty seems to be influenced to a large degree by its data uncertainty, not model uncertainty. All model implementations are open-sourced in a software package.

1 Introduction

In 1877, Italian astronomer Giovanni Schiaparelli described the existence of “canals” on the surface of Mars, a finding that was described by a contemporary as a “very important and perplexing [problem]” (Young, 1895; p. 355). It later turned out that the structures, originally termed *canali* in Italian, were simply mistranslated, since the word can also refer to (natural) channels of water. By that point however, the possibility of irrigation on the red planet had already seeped into popular culture, and is still being referenced to this day. In the meantime, translation has become a task that is increasingly performed by neural networks, which — in the face of a word such as *canali* — might simply fall back on the most likely translation given the training data. And while the error above seems fairly innocuous, there are more safety-critical scenarios in which such ambiguities matter and can potentially have negative real-world consequences. Besides translation, there also exist other language-based problems in which the uncertainty surrounding a

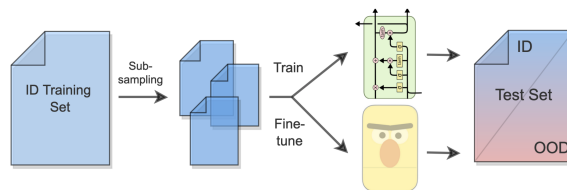


Figure 1: **Schematic of our experiments.** Training sets are sub-sampled and used to train LSTM-based models and fine-tune transformer-based ones, which are evaluated on in- and out-of-distribution test data.

model prediction can convey critical information, such as medical analyses (Esteva et al., 2019), legal case data (Frankenreiter and Livermore, 2020) or analyzing job applications (Zimmermann et al., 2016). Determining model confidence, or, conversely, uncertainty, consequently is an important mean to instill trust in end users and avert harm (Bhatt et al., 2021; Jacovi et al., 2021). While there exist many works on images (Lakshminarayanan et al., 2017; Snoek et al., 2019) and tabular data (Ruhe et al., 2019; Ulmer et al., 2020; Malinin et al., 2021), the quality of uncertainty estimates provided by neural networks remains underexplored in Natural Language Processing (NLP). In addition, as model underspecification due to insufficient data presents a risk (D’Amour et al., 2020), the increasing interest in less-researched languages with limited resources raises the question of how reliably uncertain predictions can be identified. This lets us pose the following research questions:

- RQ1** What are the best approaches in terms of uncertainty quality and calibration?
- RQ2** How are models impacted by the amount of available training data?
- RQ3** What are differences in how the different approaches estimate uncertainty?

Contributions ① We address these questions by conducting a comprehensive empirical study of eight different models for uncertainty estimation for classification and evaluate their effectiveness on three languages spanning distinct NLP tasks, involving sequence labeling and classification. ② We show that while approaches based on pre-trained models and ensembles achieve the best results overall, the quality of uncertainty estimates on OOD data can become worse using *more* data. ③ In a qualitative analysis, we also discover that a model’s total uncertainty seems to mostly consist of its data uncertainty. ④ We make our experimental code and model implementations available open-source in separate repositories, aiding future research in this direction.¹

2 Related Work

Notions of Uncertainty In the absence of additional information, the introductory example *canali* has two valid translations — *canals* and *channels*. This is an instance of *data* or *aleatoric* uncertainty, describing the irreducible ambiguity and noise in the data generating process. The other notion is *model* or *epistemic* uncertainty: Fitting parameters, there remains a degree of incertitude about the optimal values due to finite data. We can usually reduce this uncertainty by amassing more data,² for instance by supplying a translation system with other meanings of *canali*. These two concepts form the basis for uncertainty estimation in Machine Learning (Der Kiureghian and Ditlevsen, 2009; Hüllermeier and Waegeman, 2021).

Uncertainty in NLP Since uncertainty estimation literature is manifold on image data, we dedicate this part to related works in the realm of Natural Language Processing. There are several examples trying to incorporate uncertainty into models to either increase trustworthiness or performance, for instance in Machine Translation (Glushkova et al., 2021; Wei et al., 2020; Xiao et al., 2020), Summarization (Gidiotis and Tsoumakas, 2021), Information Retrieval (Penha and Hauff, 2021) and Active Learning (Siddhant and Lipton, 2018). To obtain uncertainties, Gan et al. (2017) use Stochastic-Gradient Langevin Dynamics (Welling and Teh,

2011) to obtain posterior weight samples for a LSTM. Shelmanov et al. (2021) apply MC Dropout with determinantal point processes to transformers for Natural Language Understanding. Several authors have also highlighted connections of multi-head attention to Bayesian inference (An et al., 2020; Hron et al., 2020). Shen et al. (2020) attempt to transfer the idea of prior networks (Malinin and Gales, 2018; Joo et al., 2020) onto recurrent neural networks. Another line of works investigates uncertainty properties themselves; For instance, Chen and Ji (2022) try to explain uncertainty estimates for BERT and RoBERTa. Another example is given by Xiao and Wang (2021), who use predictive uncertainty to explain hallucination in Language Generation. Xu et al. (2020) similarly use uncertainty as a tool to investigate challenges of neural summarization approaches. Lastly, due to the way that uncertainty estimates are evaluated, investigating distributional shift in NLP is also of interest, for instance through the work of Arora et al. (2021), Kamath et al. (2020), who focus on question answering and Tan et al. (2019) for text classification. The most similar work to ours is the text classification uncertainty benchmark by Van Landeghem et al. (2022), however they do not consider the impact of data or language, and test a different selection of models.

Calibration Calibration denotes the property of a model’s output to accurately reflect the true chance of a correct prediction — i.e. predicting a class with a confidence of 90% should yield the correct prediction for 90% of similar inputs, when repeated. There have been several studies testing this property in modern neural networks (Guo et al., 2017; Nixon et al., 2019; Minderer et al., 2021; Wang et al., 2021b) and proposing ways to improve it (Thulasidasan et al., 2019; Mukhoti et al., 2020; Karandikar et al., 2021; Zhao et al., 2021a; Tian et al., 2021). In NLP, calibration has been explored for pre-trained models (Desai and Durrett, 2020), including on out-of-distribution data (Dan and Roth, 2021), for neural machine translation (Wang et al., 2020) and for question-answering (Jiang et al., 2021). Likewise, authors have proposed several calibration schemes, for instance by focusing on classes of interest (Jagannatha and Yu, 2020), generating synthetic examples for regularization (Kong et al., 2020), using richer input representations (Zhang et al., 2021) and adapting prompts in a zero-shot setting (Zhao et al., 2021b).

¹The model zoo is available under <https://github.com/Kaleidophon/nlp-uncertainty-zoo>, with the code for the experiments available under <https://github.com/Kaleidophon/nlp-low-resource-uncertainty>.

²That is, unless the model class we chose is too restrictive.

3 Methodology

3.1 Models

We choose a variety of models that cover a range of different approaches based on the two most prominently used architectures in NLP: Long-Short Term Memory networks (LSTMs; Hochreiter and Schmidhuber, 1997) and transformers (Vaswani et al., 2017). Inside the first family, we use the Variational LSTM (Gal and Ghahramani, 2016b) based on MC Dropout (Gal and Ghahramani, 2016a), the Bayesian LSTM (Fortunato et al., 2017) implementing Bayes-by-backprop (Blundell et al., 2015) and the ST- τ LSTM (Wang et al., 2021a), modelling transitions in a finite-state automaton, as well as an ensemble (Lakshminarayanan et al., 2017). In the second family, we count the Variational Transformer (Xiao et al., 2020), also using MC Dropout, the SNGP Transformer (Liu et al., 2022), using a Gaussian Process output layer, and the Deep Deterministic Uncertainty transformer (DDU; Mukhoti et al., 2021), fitting a Gaussian mixture model on extracted features. We elaborate on implementation details in Appendix C.1.

3.2 Uncertainty Metrics

We employ the following metrics to quantify confidence or uncertainty — in all cases, lower values indicate lower confidence / certainty and conversely, higher values mean higher confidence / certainty. The following metrics were either chosen due to their frequent use in the literature, or because they are trying to capture uncertainty in a novel way.

Single prediction metrics We distinguish between metrics suitable for models using only a single prediction (or using the mean of multiple predictions, e.g. for an ensemble). The most straightforward of them is the maximum softmax probability by Hendrycks and Gimpel (2017). A variant of this is the softmax-gap, measuring the difference between the two largest predicted probabilities (Tagasovska and Lopez-Paz, 2019). Another common metric, predictive entropy, involves measuring the Shannon entropy of the output distribution, which is maximized for a uniform prediction:

$$-\sum_{k=1}^K p_{\theta}(y = k | \mathbf{x}) \log p_{\theta}(y = k | \mathbf{x})$$

Lastly, we consider the Dempster-Shafer metric (Sensoy et al., 2018), defined as $K/(K +$

$\sum_{k=1}^K \exp(z_k))$, where z_k denotes the logit corresponding to class k . It has been shown that probabilities for (ReLU) networks tend to saturate in the limit (Hein et al., 2019; Ulmer and Cinà, 2021), and since this metric considers logits, it might provide more informative estimates on OOD data.

Multiple prediction metrics For some of the included models, we can express uncertainty as some score based on a number of predicted distributions, e.g. from different ensemble members or forward passes for MC Dropout. Here we use the expectation with respect to the weight posterior to express the aggregation of multiple predictions, which will simply be evaluated using the mean of a number of Monte Carlo samples in practice. A simple uncertainty metric on this basis is the predictive variance between predictions for a class:

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E}_{q(\theta)} \left[\left(p_{\theta}(y = k | \mathbf{x}) - \mathbb{E}_{q(\theta)} [p_{\theta}(y = k | \mathbf{x})] \right)^2 \right],$$

where the expectation is evaluated over multiple sets of parameters, e.g. stemming from different dropout masks. Another possibility lies in using the mutual information between the label and model parameters given the data and input sample, which was introduced by Smith and Gal (2018):

$$\mathbb{H} \left[\mathbb{E}_{q(\theta)} [p_{\theta}(y | \mathbf{x})] \right] - \mathbb{E}_{q(\theta)} \left[\mathbb{H} [p_{\theta}(y | \mathbf{x})] \right] \quad (1)$$

where \mathbb{H} denotes the Shannon entropy as used for predictive entropy. The two terms of this equation can be identified as the total entropy and the aleatoric uncertainty, respectively. In theory, the remaining epistemic uncertainty of the model — in form of the mutual information — should be particularly high on OOD inputs.

Model-specific metrics Lastly, DDU by Mukhoti et al. (2021) uses the log-probability of the last layer network activation under a Gaussian Mixture Model fitted on the training set as an additional metric. Since all other models are trained or fine-tuned as classifiers, they are not able to assign log-probabilities to sequences.

Uncertainty for sequences Since some tasks require predictions for every time step of a sequence, we determine the uncertainty of a whole sequence in these cases by taking the mean over all step-wise uncertainties.³ A more principled approach for sequences is for instance provided by Malinin and

³We also just considered the *maximum* uncertainty over a sequence, with similar results.

Gales (2021), and we leave the extension and exploration of such methods for different uncertainty metrics, models and tasks to future work.

3.3 Dataset Selection & Creation

In-distribution training sets We choose three different languages, namely English (Cling Plus; Larson et al., 2019), Danish in the form of the Dan+ dataset (Plank et al., 2020) based on News texts from PAROLE-DK (Bilgram and Keson, 1998), Finnish (UD Treebank; Haverinen et al., 2014; Pyysalo et al., 2015; Kanerva and Ginter, 2022), corresponding to NLP tasks such as sequence classification, named entity recognition and part-of-speech tagging. An overview over the used data is given in Table 1. We do use standardized low-resource languages in the case of Finnish and Danish, and simulate a low-resource setting using English data.⁴ Starting with a sufficiently-sized training set and then sub-sampling allows us to create training sets of arbitrary sizes. By using languages from different families, we hope to be able draw conclusions that generalize across a single language. We employ a specific sampling scheme that tries to maintain the sequence length and class distribution of the original corpus, which we explain and verify in Appendix A.2.

Out-of-distribution Test Sets While it is possible to create OOD text by for instance withholding classes from the training set or appending text from a different source (Arora et al., 2021), we choose to pick entirely new OOD test sets that are qualitatively different: Out-of-scope voice commands by users in Larson et al. (2019),⁵ the Twitter split of the Dan+ dataset (Plank et al., 2020), and the Finnish OOD treebank (Kanerva and Ginter, 2022). In similar works for the image domain, OOD test sets are often chosen to be convincingly different from the training distribution, for instance MNIST versus Fashion-MNIST (Nalisnick et al., 2019; van

⁴The definition of low-resource actually differs greatly between works. One definition by Bird (2022) advocates the usage for (would-be) standardized languages with a large amount of speakers and a written tradition, but a lack of resources for language technologies. Another way is a task-dependent definition: For dependency parsing, Müller-Eberstein et al. (2021) define low-resource as providing less than 5000 annotated sentences in the Universal Dependencies Treebank. Hedderich et al. (2021); Lignos et al. (2022) lay out a task-dependent spectrum, from a several hundred to thousands of instances.

⁵Since all instances in this test set correspond to out-of-scope inputs and not to classes the model was trained on, we cannot evaluate certain metrics in Table 2.

Amersfoort et al., 2021). While there exist a variety of formalizations of types of distributional shift (Moreno-Torres et al., 2012; Wald et al., 2021; Arora et al., 2021; Federici et al., 2021), it is often hard to determine if and what kind of shift is taking place. Winkens et al. (2020) define *near OOD* as a scenario in which the inlier and outlier distribution are meaningfully related, and *far OOD* as a case in which they are unrelated. Unfortunately, this distinction is somewhat arbitrary and hard to apply in a language context, where OOD *could* be defined as anything ranging from a different language or dialect to a different demographic on an author or speaker or a new genre. Therefore, we use a similar methodology to the validation of the sub-sampled training sets to make an argument that the selected OOD splits are sufficiently different in nature from the training splits. The exact procedure along some more detailed results is described in Appendix A.3.

3.4 Model Training

Unfortunately, our datasets do not contain enough data to train transformer-based models from scratch. Therefore, we only fully train LSTM-based models, while using pre-trained transformers, namely BERT (English; Devlin et al., 2019), Danish BERT (Hvingelby et al., 2020), and FinBERT (Finnish; Virtanen et al., 2019), for the other approaches. The whole procedure is depicted in Figure 1. The way we optimize models is provided in Appendix C.3. We list training hardware, hyperparameter information in Appendix C.2, with the environmental impact described in Appendix C.5.

3.5 Evaluation

Apart from evaluating models on the task performance, we also evaluate the following calibration and uncertainty, painting a multi-faceted picture of the reliability of models. In all cases, we use the Almost Stochastic Order test (ASO; del Barrio et al., 2018; Dror et al., 2019) for significance testing, which is elaborated on in Appendix C.1.

Evaluation of Calibration First, we measure the calibration of models using the adaptive calibration error (ACE; Nixon et al., 2019), which is an extension of the expected calibration error (ECE; Naeini et al., 2015; Guo et al., 2017).⁶ Furthermore, we use the frequentist measure of coverage (Larry, 2004; Kompa et al., 2021). Coverage is based on

⁶See Appendix B for a short overview over the differences.

Language	Task	Dataset	OOD Test Set	# ID / OOD.	Sub-sampled Training Set Sizes
EN	Intent Classification	Clinic Plus (Larson et al., 2019)	Out-of-scope voice commands	15k / 1k	15k / 12.5k / 10k
DA	Named Entity Recognition	Dan+ News (Plank et al., 2020)	Tweets	4382 / 109	4k / 2k / 1k
FI	PoS Tagging	Finnish UD Treebank (Haverinen et al., 2014; Pyysalo et al., 2015; Kanerva and Ginter, 2022)	Hospital records, online forums, tweets, poetry	12217 / 2122	10k / 7.5k / 5k

Table 1: **Datasets.** The original and sub-sampled number of sequences for experiments are given on the right.

the prediction set $\hat{\mathbb{P}}(\mathbf{x})$ of a classifier given an input, which includes the most likely classes adding up to or surpassing $1 - \alpha$ probability mass. A well-tuned classifier should contain the correct class in this very set, and minimize its width. The extent to which this property holds can be determined by the *coverage percentage*, i.e. the number of times the correct class in indeed contain within the prediction set, and its cardinality, denoted simply as *width*.

Evaluation of Uncertainty We compare uncertainty scores on the ID and OOD test set and measure the area under the receiver-operator curve (AUROC) and under the precision-recall curve (AUPR), assuming that uncertainty will generally be higher on samples from the OOD test set.⁷ An ideal model should create very different distributions of confidence scores on ID and OOD data, thus maximizing AUROC and AUPR. However, we also want to find out to what extent uncertainty can give an indication of the correctness of the model, which is why we propose a new way to evaluate the *discrimination* property posed by Alaa and Van Der Schaar (2020) based on Leonard et al. (1992): A good model should be less certain for inputs that incur a higher loss. To measure this both on a token and sequence level, we utilize Kendall’s τ (Kendall, 1938), which, given two lists of measurements, determines the degree to which they are *concordant* — that is, to what extent the rankings of elements according to their measured values agree. This is expressed by a value between -1 and 1 , with the latter expressing complete concordance. In our case, these measurements correspond to the uncertainty estimate and the actual model loss, either for tokens (Token τ) or sequences (Sequence τ).

⁷We thus formulate a pseudo-binary classification task as common in the literature, using the model’s uncertainty score to try to distinguish the two test sets. Note that we do not advocate for actually using uncertainty for OOD detection, but only use it for evaluation purposes, since uncertainty on OOD examples should be high due to model uncertainty.

4 Experiments

4.1 RQ1: Uncertainty & Calibration

We present the results from our experiments using the largest training set sizes per dataset in Table 2.⁸

Task Performance Across datasets and models, we can identify several trends: some of the BERT-based models unsurprisingly perform better than LSTM based models, which can be explained with their pre-training procedure. We observe worse performance for some LSTM and BERT-variants, in particular the Variational, Bayesian and ST- τ LSTM, as well the SNGP BERT. In accordance with the ML literature (see e.g. Lakshminarayanan et al. (2017); Ovadia et al. (2019)), LSTM ensembles actually perform very strongly and on par or sometimes better than fine-tuned BERTs.

Calibration We also see BERT models to generally achieve lower calibration errors across all metrics measured, which is in line with previous works (Desai and Durrett, 2020; Dan and Roth, 2021). It is interesting to see that the correct prediction is almost always contained in the 0.95 confidence set across all models, however these number have to be interpreted in the context of the set’s width: It becomes apparent that for instance LSTMs achieve this coverage by spreading probability mass over many classes, while only BERT-based models, LSTM ensembles as well as the Bayesian LSTM (on Danish) and the Variational LSTM (on Finnish) are *confidently* correct.

Uncertainty Quality LSTM-based model seem to struggle to distinguish in- from out-of-distribution data based on predictive uncertainty. For Danish, only BERTs perform visibly above chance-level. For Finnish, the AUPR results suggest that although some OOD instances are quickly

⁸For English, some models were omitted due to convergence issues, which are discussed in Appendix C.4.

Model	Task (ID/OOD)		Calibration (ID/OOD)				Uncertainty (ID/OOD)				
	Acc.↑	F_1 ↑	ECE↓	ACE↓	%Cov.↑	\varnothing Width↓	AUROC↑	AUPR↑	Token τ ↑	Seq. τ ↑	
English	LSTM	.79 ±.00	.62 ±.01	77.95 ±.00	.49 ±.01	<u>1.00</u> ±.00	144.00 ±.00	.88 ⁺ ±.01	.60 ⁺ ±.01	—	.75 [○] ±.01
	Bayesian LSTM	.59 ±.06	.46 ±.05	77.66 ±.05	.22 ±.01	.88 ±.00	41.99 ±1.94	.86 [△] ±.01	.59 [×] ±.01	—	.66 [○] ±.02
	LSTM Ensemble	<u>.81</u> ±.00	<u>.64</u> ±.01	77.11 ±.00	.09 ±.00	.87 ±.00	4.27 ±.05	<u>.92⁺</u> ±.00	<u>.71⁺</u> ±.01	—	.73 [□] ±.01
	Variational BERT	.45 ±.16	.34 ±.13	77.92 ±.02	.22 ±.06	1.00 ±.00	115.11 ±11.38	.80 [×] ±.01	.53 [×] ±.01	—	.57 [○] ±.09
	DDU BERT	.79 ±.00	.64 ±.01	<u>77.02</u> ±.00	<u>.00</u> ±.00	.82 ±.00	<u>1.46</u> ±.04	.88 [○] ±.00	.62 [○] ±.01	—	<u>.87[○]</u> ±.00
Danish	LSTM	.93 / .92 ±.00 / ±.00	.26 / .19 ±.01 / ±.01	17.18 / 17.17 ±.00 / ±.00	.16 / .10 ±.01 / ±.01	<u>1.00</u> / <u>1.00</u> ±.00 / ±.00	19.00 / 19.00 ±.00 / ±.00	.50 [○] ±.02	.14 [○] ±.01	.50 [○] / .47 [○] ±.00 / ±.00	-.26 ⁺ / -.28 [○] ±.02 / ±.05
	Variational LSTM	.90 / .90 ±.02 / ±.02	.08 / .09 ±.02 / ±.02	16.74 / 16.72 ±.03 / ±.03	.26 / .17 ±.02 / ±.01	.99 / .98 ±.01 / ±.01	6.62 / 6.68 ±.37 / ±.33	.60 ⁺ ±.04	.21 ⁺ ±.02	.23 [○] / .23 [○] ±.06 / ±.05	-.04 ⁺ / -.02 [○] ±.02 / ±.05
	ST- τ LSTM	.92 / .92 ±.00 / ±.00	.12 / .09 ±.00 / ±.00	16.67 / 16.63 ±.00 / ±.01	.24 / .15 ±.01 / ±.01	1.00 / .99 ±.00 / ±.00	7.10 / 7.03 ±.07 / ±.08	.54 ⁺ ±.01	.15 ⁺ ±.01	.50 [○] / .48 [○] ±.00 / ±.00	-.05 ⁺ / -.01 [○] ±.03 / ±.05
	Bayesian LSTM	.93 / .93 ±.00 / ±.00	.07 / .07 ±.00 / ±.00	16.81 / 16.79 ±.00 / ±.00	.25 / .18 ±.01 / ±.01	1.00 / 1.00 ±.00 / ±.00	1.68 / 1.70 ±.04 / ±.05	.65 [○] ±.17	.31 [○] ±.30	.53 [○] / <u>.55[○]</u> ±.01 / ±.01	-.01 ⁺ / -.02 ⁺ ±.01 / ±.01
	LSTM Ensemble	<u>.95</u> / <u>.94</u> ±.00 / ±.00	<u>.33</u> / <u>.25</u> ±.01 / ±.01	16.37 / <u>16.35</u> ±.00 / ±.00	.18 / .13 ±.01 / ±.01	.98 / .97 ±.00 / ±.00	<u>1.62</u> / <u>1.58</u> ±.00 / ±.01	.60 [□] ±.02	.18 [□] ±.01	.44 [□] / .45 [□] ±.00 / ±.00	-.19 ⁺ / -.28 [□] ±.01 / ±.01
	SNGP BERT	.22 / .19 ±.35 / ±.34	.03 / .02 ±.03 / ±.02	17.19 / 17.18 ±.01 / ±.01	<u>.08</u> / <u>.06</u> ±.01 / ±.01	1.00 / 1.00 ±.00 / ±.00	18.84 / 18.83 ±.32 / ±.34	.86 [△] ±.06	.49 [△] ±.12	.17 [○] / .26 [○] ±.09 / ±.14	<u>.29⁺</u> / <u>.44[□]</u> ±.03 / ±.11
	Variational BERT	.94 / .89 ±.00 / ±.00	.29 / .17 ±.00 / ±.00	<u>16.36</u> / 16.43 ±.00 / ±.00	.20 / .22 ±.00 / ±.00	.99 / .98 ±.00 / ±.00	2.25 / 3.86 ±.01 / ±.08	.86 ⁺ ±.01	.46 ⁺ ±.02	.42 [○] / .17 [○] ±.01 / ±.01	-.35 [□] / -.41 [□] ±.01 / ±.01
DDU BERT	.92 / .89 ±.00 / ±.00	.25 / .17 ±.00 / ±.00	16.41 / 16.44 ±.00 / ±.00	.19 / .21 ±.01 / ±.01	.99 / .99 ±.00 / ±.00	3.48 / 4.04 ±.01 / ±.03	.86 [○] ±.01	.39 [○] ±.02	<u>.56[○]</u> / .25 [○] ±.00 / ±.01	-.24 [○] / -.38 [○] ±.01 / ±.03	
Finnish	LSTM	.75 / .69 ±.00 / ±.00	.57 / .53 ±.00 / ±.00	6.78 / 6.80 ±.00 / ±.00	.40 / .38 ±.01 / ±.01	1.00 / 1.00 ±.00 / ±.00	16.00 / 16.00 ±.00 / ±.00	.63 [△] ±.01	.69 ⁺ ±.01	.29 [○] / .19 [○] ±.00 / ±.01	-.28 ⁺ / -.27 ⁺ ±.02 / ±.02
	Variational LSTM	.27 / .26 ±.00 / ±.00	.03 / .03 ±.00 / ±.00	6.65 / 6.66 ±.01 / ±.01	.27 / .28 ±.01 / ±.01	.97 / .96 ±.00 / ±.00	1.35 / 1.37 ±.23 / ±.21	.51 ⁺ ±.01	.59 ⁺ ±.01	.00 [△] / .00 [○] ±.01 / ±.00	.01 [△] / .01 [○] ±.03 / ±.01
	ST- τ LSTM	.76 / .71 ±.00 / ±.00	.58 / .55 ±.00 / ±.00	6.18 / 6.21 ±.00 / ±.00	.20 / .22 ±.01 / ±.01	.97 / .96 ±.00 / ±.00	3.32 / 3.57 ±.01 / ±.01	.62 [△] ±.01	.69 ⁺ ±.01	.31 [○] / .21 [○] ±.01 / ±.01	-.14 ⁺ / -.12 [□] ±.01 / ±.01
	Bayesian LSTM	.27 / .26 ±.00 / ±.00	.03 / .03 ±.00 / ±.00	6.84 / 6.85 ±.00 / ±.00	<u>.11</u> / <u>.12</u> ±.00 / ±.00	1.00 / 1.00 ±.00 / ±.00	16.00 / 16.00 ±.00 / ±.00	.51 [○] ±.01	.60 ⁺ ±.01	.00 [○] / .00 [○] ±.00 / ±.00	.01 [○] / .04 ⁺ ±.01 / ±.00
	LSTM Ensemble	.81 / .75 ±.00 / ±.00	.62 / .57 ±.00 / ±.00	6.18 / 6.22 ±.00 / ±.00	.17 / .21 ±.01 / ±.00	.99 / .98 ±.00 / ±.00	3.46 / 3.80 ±.01 / ±.01	<u>.67⁺</u> ±.01	<u>.74⁺</u> ±.01	.29 [○] / .19 [○] ±.00 / ±.01	-.28 ⁺ / -.31 ⁺ ±.01 / ±.01
	Variational BERT	.87 / .81 ±.00 / ±.00	.74 / .70 ±.00 / ±.00	6.11 / 6.15 ±.00 / ±.00	.14 / .18 ±.00 / ±.01	.99 / .99 ±.00 / ±.00	4.68 / 5.19 ±.03 / ±.02	.64 [△] ±.01	.70 [○] ±.01	.14 [○] / .08 ⁺ ±.00 / ±.00	-.19 ⁺ / -.16 ⁺ ±.00 / ±.01
	SNGP BERT	.18 / .17 ±.10 / ±.10	.07 / .08 ±.02 / ±.02	6.82 / 6.83 ±.00 / ±.00	.16 / .15 ±.02 / ±.01	1.00 / .99 ±.00 / ±.01	15.00 / 15.00 ±.00 / ±.00	.54 [△] ±.05	.63 [△] ±.04	.15 [□] / .15 [□] ±.04 / ±.03	<u>.12[□]</u> / <u>.14[□]</u> ±.05 / ±.02
DDU BERT	.87 / .81 ±.00 / ±.00	.72 / .68 ±.03 / ±.03	<u>6.01</u> / <u>6.03</u> ±.00 / ±.00	.33 / .38 ±.02 / ±.02	.94 / .91 ±.00 / ±.00	<u>2.16</u> / <u>2.31</u> ±.06 / ±.06	.61 [○] ±.02	.69 [○] ±.02	<u>.39[○]</u> / <u>.26[○]</u> ±.04 / ±.03	-.07 [○] / -.16 [○] ±.05 / ±.04	

Table 2: **Results on the tested datasets.** Task performance is measured by macro F_1 and accuracy, calibration by different calibration errors, the coverage percentage the average prediction set width. For every result, and value on the ID and OOD test set is shown. For English, OOD scores are not available since the OOD set does not contain gold labels, and Token τ is missing due to CLINC being a sequence prediction task. Uncertainty quality is evaluated using its ability to discriminate between ID and OOD data, quantified by AUROC and AUPR. Furthermore, Kendall’s τ is measured between the uncertainty and losses on a sequence- and token-level. Displayed are mean and standard deviation over five random seeds, with bolding and underlining indicating almost stochastic dominance with $\epsilon_{\min} \leq 0.3$ over all other models. For last section, the best value over uncertainty metrics is given, with symbols indicating the type of metric achieving it: \circ Max. probability, \triangle Predictive entropy. \square Class variance. \diamond Softmax gap. $+$ Dempster-Shafer. \times Mutual information.

identified as uncertain, many other OOD remain undetected among in-distribution samples. For English, OOD samples are detected more effectively, which can be explained by them consisting of unknown voice commands, representing a potential instance of *semantic* shift, which has been shown to be easier to detect by classifiers (Arora et al., 2021). Furthermore, it is striking that uncertainty and loss on a token-level (Token τ) is only positive correlated for some models, using metrics such as the maximum probability score, softmax gap or the Dempster-Shafer metric, which are all entirely based on the categorical output distributions. On a

sequence-level (Sequence τ), the correlation is often *negative*, meaning that higher uncertainty goes hand in hand with a *higher* loss. Lastly, it should be noted that different uncertainty metrics yield diverse outcomes: There does not seem to be one superior metric across all experimental settings, as seen by the variety of markers shown in Table 2.

4.2 RQ2: Dependence on training data

After presenting the best results for the biggest training set sizes in Table 2, we now continue to analyze the difference between models and metrics in a more fine-grained way. In Figure 2, we show dif-

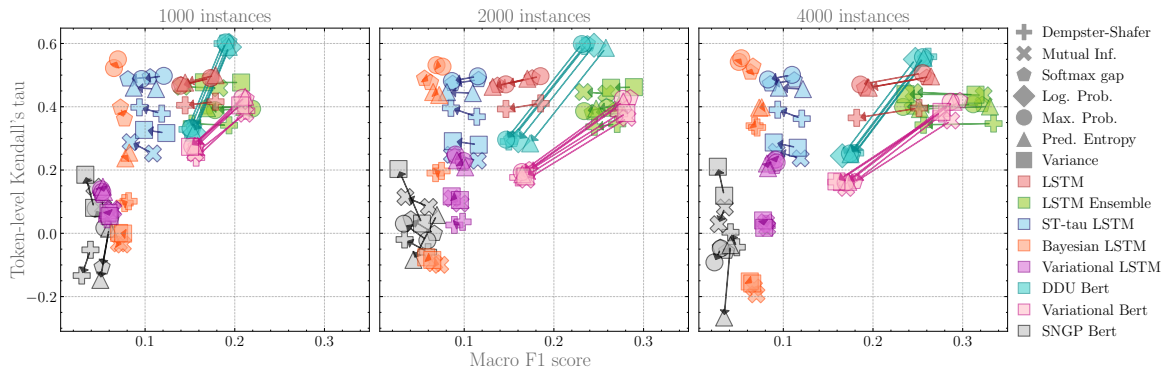


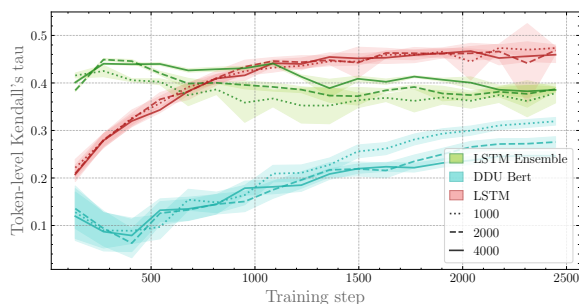
Figure 2: Scatter plot showing the difference between model performance (measured by macro F_1 and the quality of uncertainty estimates on a token-level (measured by Kendall’s τ). Shown are different models and uncertainty metrics and several training set sizes of the Dan+ dataset. Arrows indicate changes between the in-distribution and out-of-distribution test set. Best viewed electronically and in color.

ferences for the token-level correlation between a model’s loss and its uncertainty measured by Kendall’s τ , with arrows indicating the shift from measurements on the in- to the out-of-distribution test set. Here, we see the same trend of more training data having a larger influence on BERT models. Peculiarly, we also observe pre-trained models’ uncertainty to correlate less with their losses on the OOD data, while this property stays relative constant for LSTMs. We can recognize this trend also for the other datasets in Figure 2 and to a lesser degree on a sequence level Figure 14a in Appendix D.1, albeit with a *negative* correlation in general in the latter case. In Figures 11 and 12 in Appendix D.1, we show the AUROC and AUPR of different model-uncertainty metric combinations for all datasets and training set sizes. In both cases, we can notice that pre-trained models profit more from an increase in available training data than LSTM-based models that are trained from scratch. This improvement is observed both in task performance, as well as in the model’s ability to discern ID from OOD data using its uncertainty, but more so for the Danish than English or Finnish. Like in the previous section, we often see that uncertainty metrics of the same model perform quite similarly. These results outline a seeming paradox: Pre-trained and then fine-tuned models (often) perform better on the task at hand, and provide better uncertainty estimates, but only on in-distribution data. Models trained from scratch that have seen less data overall, however provide more reliable uncertainty estimates on OOD data, but are also worse calibrated (Section 4.1), with the exception of ensembles. This effect appears to largest on

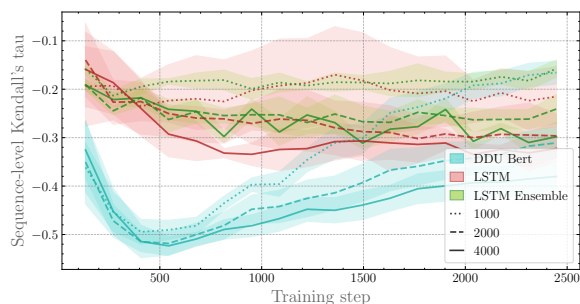
Danish, containing the least data.

Uncertainty quality over training Adding another facet to this issue, we plot the development of uncertainty estimate quality over the training for different models in Figure 3. We use LSTMs and DDU BERTs on Dan+ as representative for the observed differences between pre-trained transformers and models trained from scratch, with more examples given in Appendix D.2. On both a token and sequence-level, we can see that the correlation between uncertainty and loss dips for DDU BERT, before increasing again over the course of the training.⁹ Most curiously, the highest correlations are achieved with the models using the *least* training data. Such behavior is also present for LSTMs on a sequence level. We can also see that while the correlation is higher for DDU BERT on in-distribution data (see again Table 2), on OOD data, LSTMs actually more accurately reflect their knowledge using uncertainty. This again corroborates earlier insights from Section 4.1: Pre-trained models seem to provide better uncertainty estimates on in-distribution data, but yield worse results on OOD than LSTMs trained from scratch. Furthermore, the less training data is available, the more indicative predictive uncertainty seems to be of the correctness of a model. We see such behavior also to a lesser extent in the other, datasets (see Appendix D.2). Before we offer some potential explanations of this behavior, we try to gain an even more fine-grained understanding by analysing the differences in metrics and models on a token-level.

⁹Note that the OOD data used to create these results were not used for training.



(a) Development of token-level Kendall's τ .



(b) Development of sequence-level Kendall's τ .

Figure 3: **Development of correlation between uncertainty and loss**, shown on the Dan+ OOD test set over the training time using differently-sized training sets. Colored areas indicate the standard deviation over five runs.

4.3 RQ3: Qualitative Analysis

We investigate the development of uncertainty estimates over the course of a single sequence for different datasets, models, and uncertainty metrics. We showcase two examples in Figure 4, with more examples in Appendix D.3. By looking at the predictive entropy of models in Section 4.3, we can observe multiple things: First of all, we can observe some degree of agreement between models and their uncertainty: Processing sub-word tokens, uncertainty seems to increase, and the total uncertainty always appears to reduce considerably on punctuation. Interestingly, the highest uncertainty seems to be produced by the DDU and Variational BERT models as well as the ensembles. In Figure 4b, we compare the estimates for predictive entropy and mutual information, the latter of which is supposed to only express model uncertainty. Here, uncertainty is generally low, indicating a large part of the total uncertainty might actually be of an aleatoric nature (which is the gap between triangle and cross markers of the same color, due to Equation (1)). These insights indicate that while aleatoric uncertainty might be a constant factor for all models, epistemic uncertainty expectedly differs noticeably between them. We use all of these insights to discuss the choice of model next.

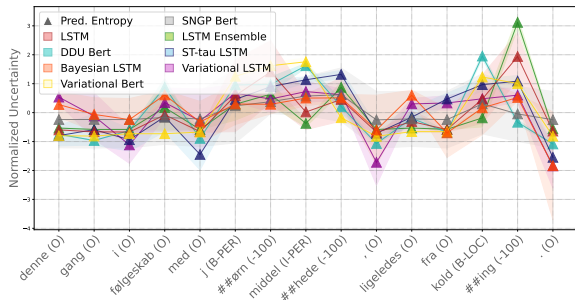
5 Discussion

Our experiments in Section 4 have uncovered interesting nuances about uncertainty estimation in Natural Language Processing. **With respect to RQ1**, we observe that fine-tuning BERTs and training LSTM ensembles on different languages produces high task scores with low calibration errors and high-quality uncertainty estimates, but only so on in-distribution data. On OOD data, uncer-

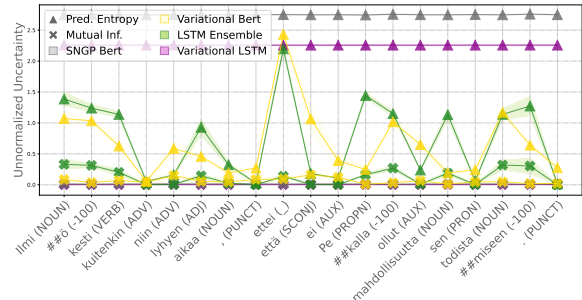
tainty estimates from fine-tuned models do actually become less indicative of potential model loss compared to LSTM-based models. We also find that among the variety of uncertainty metrics proposed, there does not appear to be a superior metric. Differences in Kendall's τ on a token and sequence level suggest that loss and uncertainties fluctuate over the course of sequence. **Answering RQ2**, it seems that paradoxically more training data seems to decrease the quality of uncertainty estimates on OOD data for pre-trained models. We speculate that fine-tuning models increasingly lets them forget relevant features that would produce higher uncertainty. This might explain why for LSTM-type models, this effect seems to be smaller. Lastly, **we conclude about RQ3** that all models' total uncertainty behave somewhat similarly, potentially due to the strong influence of aleatoric uncertainty. From these insights, we conclude that the approaches using pre-trained models overall give the best trade-off between task performance, uncertainty quality and calibrations, however their failure on OOD samples opens up further directions of research. Ensembles can provide an alternative here in data-scarce settings, when the task is sufficiently learnable without the need for pre-training.

6 Conclusion

In this work, we explore the current options for uncertainty estimation in NLP on three different languages and tasks, focusing on the impact of available data on the quality of uncertainty scores in a potential low-resource environment. We conclude the following: Fine-tuning pre-trained models produces the best results in terms of task performance, calibration and uncertainty quality, but only on in-distribution data. On out-of-distribution data, LSTM-based models produce more reliable



(a) Predictive entropy over the sentence "This time in company with Jørn Middelhede, also from Kolding".



(b) Predictive entropy and mutual information over the sentence "However, the phenomenon lasted for such a short time that Pekka did not have a chance to prove it".

Figure 4: **Uncertainty estimates on single sequences**, for predictive entropy of different models on Danish (Section 4.3) and predictive entropy and mutual information for multi-prediction models on Finnish (Figure 4b).

estimates, and could be preferred in cases pre-trained models might not be available, with LSTM-ensembles providing an especially attractive alternative. We discover that more training data seems to decrease quality of uncertainty on OOD, and show that the total uncertainty of models seems to often to be influenced by their aleatoric uncertainty.

Future Work We see our work as groundwork for future research: While uncertainty estimation is a thriving subject in Computer Vision, it remains understudied in NLP. Our experiments highlight that the model behavior on language data is not well-understood and open several lines for further investigation: One such line is the development of new methods for NLP that a) produce more faithful estimates on OOD data while retaining their ID performance and b) require less training data to so, in order to be applicable in low-resource settings. Additionally, our qualitative analyses along with existing works such as [Xiao and Wang \(2021\)](#); [Xu et al. \(2020\)](#) highlighted the potential to use uncertainty to understand model behavior.

Limitations

Even though the experiments test a large array of models and metrics, the here shown collection is by no means exhaustive, and thus only a selection of popular models or approaches from very different families were considered.

Another glaring shortcoming is the focus on only three European languages: By comparing members of the Uralic, North Germanic and West Germanic families, we only scratch the surface when it comes to the morphological diversity of human language. Further, we only focused on languages with a latin

writing systems, as well as specific text domains. This is due to resource constraints and the availability of suitable OOD test sets. We hope that follow-up works will refine our insights on a more representative sample of natural languages.

Lastly, we solely focused on sequence labelling and sequence predictions tasks. [Van Landeghem et al. \(2022\)](#) feature more sequence prediction tasks for English, however we are looking forward to similar studies on natural language generation and structured prediction tasks as well.

Ethics Statement

We do not foresee any immediate negative ethical consequences of our research.

Acknowledgements

We would like thank Mike Zhang and Joris Baan for their feedback on a draft, with a special thanks to the former for his input on the presentation of results in this work. Furthermore, we express our gratitude to Daniel Varab for verifying the translations of the Danish sentences shown in our works, and to Jenna Kanerva, Otto Tarkka, Antti Virtanen and the rest of the TurkuNLP group for the verification of the Finnish translations. The authors also acknowledge the IT University of Copenhagen's HPC resources made available for conducting the research reported in this paper.

JF was supported by the Novo Nordisk Foundation (NNF20OC0062606 and NNF20OC0065611), the Independent Research Fund Denmark (9131-00082B) and the Innovation Fund Denmark (0175-00014B).

References

- Ahmed Alaa and Mihaela Van Der Schaar. 2020. Discriminative jackknife: Quantifying uncertainty in deep learning via higher-order influence functions. In *International Conference on Machine Learning*, pages 165–174. PMLR. (Cited on p. 5)
- Bang An, Jie Lyu, Zhenyi Wang, Chunyuan Li, Changwei Hu, Fei Tan, Ruiyi Zhang, Yifan Hu, and Changyou Chen. 2020. [Repulsive attention: Rethinking multi-head attention as bayesian inference](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 236–255. Association for Computational Linguistics. (Cited on p. 2)
- Udit Arora, William Huang, and He He. 2021. [Types of out-of-distribution texts and how to detect them](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 10687–10701. Association for Computational Linguistics. (Cited on p. 2, 4, 6)
- James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(2). (Cited on p. 21)
- Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, et al. 2021. Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 401–413. (Cited on p. 1)
- Lukas Biewald. 2020. [Experiment tracking with weights and biases](#). Software available from wandb.com. (Cited on p. 20)
- Thomas Bilgram and Britt Keson. 1998. The construction of a tagged danish corpus. In *Proceedings of the 11th Nordic Conference of Computational Linguistics (NODALIDA 1998)*, pages 129–139. (Cited on p. 4)
- Steven Bird. 2022. Local languages, third spaces, and other high-resource scenarios. In *Proceedings of the 60th Annual Conference of the Association for Computational Linguistics*. Association for Computational Linguistics (ACL). (Cited on p. 4)
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. [Weight uncertainty in neural network](#). In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1613–1622. JMLR.org. (Cited on p. 3)
- Hanjie Chen and Yangfeng Ji. 2022. Explaining prediction uncertainty of pre-trained language models by detecting uncertain words in inputs. *arXiv preprint arXiv:2201.03742*. (Cited on p. 2)
- climeworks. 2022. climeworks. <https://climeworks.com/>. Accessed: 2022-06-22. (Cited on p. 23)
- Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. 2020. Underspecification presents challenges for credibility in modern machine learning. *arXiv preprint arXiv:2011.03395*. (Cited on p. 1)
- Soham Dan and Dan Roth. 2021. On the effects of transformer size on in-and out-of-domain calibration. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2096–2101. (Cited on p. 2, 5)
- Eustasio del Barrio, Juan A Cuesta-Albertos, and Carlos Matrán. 2018. An optimal transportation approach for assessing almost stochastic order. In *The Mathematics of the Uncertain*, pages 33–44. Springer. (Cited on p. 4, 21)
- Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? does it matter? *Structural safety*, 31(2):105–112. (Cited on p. 2)
- Shrey Desai and Greg Durrett. 2020. [Calibration of pre-trained transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 295–302. Association for Computational Linguistics. (Cited on p. 2, 5)
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics. (Cited on p. 4, 16)
- Rotem Dror, Segev Shlomov, and Roi Reichart. 2019. [Deep dominance - how to properly compare deep neural models](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2773–2785. Association for Computational Linguistics. (Cited on p. 4, 21)
- Piero Esposito. 2020. Blitz - bayesian layers in torch zoo (a bayesian deep learning library for torch). <https://github.com/piEsposito/blitz-bayesian-deep-learning/>. (Cited on p. 20)
- Andre Esteva, Alexandre Robicquet, Bharath Ramsundar, Volodymyr Kuleshov, Mark DePristo, Katherine Chou, Claire Cui, Greg Corrado, Sebastian Thrun, and Jeff Dean. 2019. A guide to deep learning in healthcare. *Nature medicine*, 25(1):24–29. (Cited on p. 1)

- Marco Federici, Ryota Tomioka, and Patrick Forré. 2021. An information-theoretic approach to distribution shifts. *Advances in Neural Information Processing Systems*, 34. (Cited on p. 4)
- Meire Fortunato, Charles Blundell, and Oriol Vinyals. 2017. Bayesian recurrent neural networks. *arXiv preprint arXiv:1704.02798*. (Cited on p. 3)
- Jens Frankenreiter and Michael A Livermore. 2020. Computational methods in legal analysis. *Annual Review of Law and Social Science*, 16:39–57. (Cited on p. 1)
- Yarin Gal and Zoubin Ghahramani. 2016a. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR. (Cited on p. 3)
- Yarin Gal and Zoubin Ghahramani. 2016b. [A theoretically grounded application of dropout in recurrent neural networks](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1019–1027. (Cited on p. 3, 22)
- Zhe Gan, Chunyuan Li, Changyou Chen, Yunchen Pu, Qinliang Su, and Lawrence Carin. 2017. [Scalable bayesian learning of recurrent neural networks for language modeling](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 321–331. Association for Computational Linguistics. (Cited on p. 2)
- Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wilson. 2018. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*. (Cited on p. 20)
- Alexios Gidiotis and Grigorios Tsoumakas. 2021. Uncertainty-aware abstractive summarization. *arXiv preprint arXiv:2105.10155*. (Cited on p. 2)
- Taisiya Glushkova, Chrysoula Zerva, Ricardo Rei, and André F. T. Martins. 2021. [Uncertainty-aware machine translation evaluation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 3920–3938. Association for Computational Linguistics. (Cited on p. 2)
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR. (Cited on p. 2, 4, 19)
- Katri Haverinen, Jenna Nyblom, Timo Viljanen, Veronika Laippala, Samuel Kohonen, Anna Missilä, Stina Ojala, Tapio Salakoski, and Filip Ginter. 2014. [Building the essential resources for Finnish: the Turku Dependency Treebank](#). *Language Resources and Evaluation*, 48:493–531. Open access. (Cited on p. 4, 5)
- Michael A. Hedderich, Lukas Lange, Heike Adel, Jan-nik Strötgen, and Dietrich Klakow. 2021. [A survey on recent approaches for natural language processing in low-resource scenarios](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 2545–2568. Association for Computational Linguistics. (Cited on p. 4)
- Matthias Hein, Maksym Andriushchenko, and Julian Bitterwolf. 2019. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 41–50. (Cited on p. 3)
- Dan Hendrycks and Kevin Gimpel. 2017. [A baseline for detecting misclassified and out-of-distribution examples in neural networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net. (Cited on p. 3)
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8):1735–1780. (Cited on p. 3)
- Jiri Hron, Yasaman Bahri, Jascha Sohl-Dickstein, and Roman Novak. 2020. Infinite attention: Nngp and ntk for deep attention networks. In *International Conference on Machine Learning*, pages 4376–4386. PMLR. (Cited on p. 2)
- Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3):457–506. (Cited on p. 2)
- Rasmus Hvingelby, Amalie Brogaard Pauli, Maria Barrett, Christina Rosted, Lasse Malm Lidgaard, and Anders Sjøgaard. 2020. [Dane: A named entity resource for danish](#). In *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*, pages 4597–4604. European Language Resources Association. (Cited on p. 4, 16)
- Alon Jacovi, Ana Marasović, Tim Miller, and Yoav Goldberg. 2021. Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in ai. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 624–635. (Cited on p. 1)
- Abhyuday Jagannatha and Hong Yu. 2020. Calibrating structured output predictors for natural language processing. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*,

- volume 2020, page 2078. NIH Public Access. (Cited on p. 2)
- Frederick Jelinek. 1980. Interpolated estimation of markov source parameters from sparse data. In *Proc. Workshop on Pattern Recognition in Practice, 1980*. (Cited on p. 16)
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977. (Cited on p. 2)
- Taejong Joo, Uijung Chung, and Min-Gwan Seo. 2020. Being bayesian about categorical probability. In *International Conference on Machine Learning*, pages 4950–4961. PMLR. (Cited on p. 2)
- Daniel Jurafsky and James H Martin. 2022. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, 3rd ed. draft*. (Cited on p. 16)
- Amita Kamath, Robin Jia, and Percy Liang. 2020. [Selective question answering under domain shift](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5684–5696. Association for Computational Linguistics. (Cited on p. 2)
- Jenna Kanerva and Filip Ginter. 2022. Out-of-domain evaluation of finnish dependency parsing. In *Proceedings of the 13th International Conference on Language Resources and Evaluation (LREC'22)*. (Cited on p. 4, 5)
- Archit Karandikar, Nicholas Cain, Dustin Tran, Balaji Lakshminarayanan, Jonathon Shlens, Michael C Mozer, and Rebecca Roelofs. 2021. Soft calibration objectives for neural networks. *Advances in Neural Information Processing Systems*, 34. (Cited on p. 2)
- Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93. (Cited on p. 5)
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. (Cited on p. 22)
- Benjamin Kompa, Jasper Snoek, and Andrew L Beam. 2021. Empirical frequentist coverage of deep learning uncertainty quantification procedures. *Entropy*, 23(12):1608. (Cited on p. 4)
- Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. 2020. [Calibrated language model fine-tuning for in- and out-of-distribution data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 1326–1340. Association for Computational Linguistics. (Cited on p. 2)
- Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. Quantifying the carbon emissions of machine learning. *Workshop on Tackling Climate Change with Machine Learning at NeurIPS 2019*. (Cited on p. 23)
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. [Simple and scalable predictive uncertainty estimation using deep ensembles](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6402–6413. (Cited on p. 1, 3, 5)
- Wassennan Larry. 2004. All of statistics: a concise course in statistical inference. (Cited on p. 4)
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 2019. [An evaluation dataset for intent classification and out-of-scope prediction](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1311–1316. Association for Computational Linguistics. (Cited on p. 4, 5)
- JA Leonard, Mark A Kramer, and LH Ungar. 1992. A neural network architecture that computes its own reliability. *Computers & chemical engineering*, 16(9):819–835. (Cited on p. 5)
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, et al. 2021. Datasets: A community library for natural language processing. *arXiv preprint arXiv:2109.02846*. (Cited on p. 20)
- Lisha Li, Kevin G. Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. 2017. [Hyperband: A novel bandit-based approach to hyperparameter optimization](#). *J. Mach. Learn. Res.*, 18:185:1–185:52. (Cited on p. 21)
- Constantine Lignos, Nolan Holley, Chester Palen-Michel, and Jonne Sälevä. 2022. [Toward more meaningful resources for lower-resourced languages](#). In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 523–532. Association for Computational Linguistics. (Cited on p. 4)
- Jeremiah Zhe Liu, Shreyas Padhy, Jie Ren, Zi Lin, Yeming Wen, Ghassen Jerfel, Zack Nado, Jasper Snoek, Dustin Tran, and Balaji Lakshminarayanan. 2022. A simple approach to improve single-model deep

- uncertainty via distance-awareness. *arXiv preprint arXiv:2205.00403*. (Cited on p. 3, 20, 22, 23)
- Kadan Lottick, Silvia Susai, Sorelle A. Friedler, and Jonathan P. Wilson. 2019. Energy usage reports: Environmental awareness as part of algorithmic accountability. *Workshop on Tackling Climate Change with Machine Learning at NeurIPS 2019*. (Cited on p. 23)
- Andrey Malinin, Neil Band, German Chesnokov, Yarin Gal, Mark JF Gales, Alexey Noskov, Andrey Ploskonosov, Liudmila Prokhorenkova, Ivan Provilkov, Vatsal Raina, et al. 2021. Shifts: A dataset of real distributional shift across multiple large-scale tasks. *arXiv preprint arXiv:2107.07455*. (Cited on p. 1)
- Andrey Malinin and Mark Gales. 2018. Predictive uncertainty estimation via prior networks. *Advances in neural information processing systems*, 31. (Cited on p. 2)
- Andrey Malinin and Mark J. F. Gales. 2021. [Uncertainty estimation in autoregressive structured prediction](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. (Cited on p. 3)
- Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. 2021. Revisiting the calibration of modern neural networks. *Advances in Neural Information Processing Systems*, 34. (Cited on p. 2)
- Jose G. Moreno-Torres, Troy Raeder, Rocío Alaíz-Rodríguez, Nitesh V. Chawla, and Francisco Herrera. 2012. [A unifying view on dataset shift in classification](#). *Pattern Recognit.*, 45(1):521–530. (Cited on p. 4)
- Jishnu Mukhoti, Andreas Kirsch, Joost van Amersfoort, Philip HS Torr, and Yarin Gal. 2021. Deterministic neural networks with appropriate inductive biases capture epistemic and aleatoric uncertainty. *arXiv preprint arXiv:2102.11582*. (Cited on p. 3)
- Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip Torr, and Puneet Dokania. 2020. Calibrating deep neural networks using focal loss. *Advances in Neural Information Processing Systems*, 33:15288–15299. (Cited on p. 2)
- Max Müller-Eberstein, Rob van der Goot, and Barbara Plank. 2021. [Genre as weak supervision for cross-lingual dependency parsing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 4786–4802. Association for Computational Linguistics. (Cited on p. 4)
- Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. 2015. [Obtaining well calibrated probabilities using bayesian binning](#). In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pages 2901–2907. AAAI Press. (Cited on p. 4, 19)
- Eric T. Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Görür, and Balaji Lakshminarayanan. 2019. [Do deep generative models know what they don't know?](#) In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. (Cited on p. 4)
- Hermann Ney, Ute Essen, and Reinhard Kneser. 1994. On structuring probabilistic dependences in stochastic language modelling. *Computer Speech & Language*, 8(1):1–38. (Cited on p. 16)
- Jeremy Nixon, Michael W. Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. 2019. [Measuring calibration in deep learning](#). In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 38–41. Computer Vision Foundation / IEEE. (Cited on p. 2, 4, 19, 20)
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. 2019. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32. (Cited on p. 5)
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037. (Cited on p. 20)
- Gustavo Penha and Claudia Hauff. 2021. [On the calibration and uncertainty of neural learning to rank models for conversational search](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 160–170. Association for Computational Linguistics. (Cited on p. 2)
- Barbara Plank, Kristian Nørgaard Jensen, and Rob van der Goot. 2020. Dan+: Danish nested named entities and lexical normalization. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6649–6662. (Cited on p. 4, 5)
- Sampo Pyysalo, Jenna Kanerva, Anna Missilä, Veronika Laippala, and Filip Ginter. 2015. [Universal Dependencies for Finnish](#). In *Proceedings of NoDaLiDa 2015*, pages 163–172. NEALT. (Cited on p. 4, 5)
- Alex Rogozhnikov. 2022. [Einops: Clear and reliable tensor manipulations with einstein-like notation](#). In *International Conference on Learning Representations*. (Cited on p. 20)

- David Ruhe, Giovanni Cina, Michele Tonutti, Daan de Bruin, and Paul Elbers. 2019. Bayesian modelling in practice: Using uncertainty to improve trustworthiness in medical applications. *arXiv preprint arXiv:1906.08619*. (Cited on p. 1)
- Victor Schmidt, Kamal Goyal, Aditya Joshi, Boris Feld, Liam Conell, Nikolas Laskaris, Doug Blank, Jonathan Wilson, Sorelle Friedler, and Sasha Lucioni. 2021. [CodeCarbon: Estimate and Track Carbon Emissions from Machine Learning Computing](#). (Cited on p. 20, 23)
- Murat Sensoy, Lance M. Kaplan, and Melih Kandemir. 2018. [Evidential deep learning to quantify classification uncertainty](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 3183–3193. (Cited on p. 3)
- Artem Shelmanov, Evgenii Tsymbalov, Dmitri Puzyrev, Kirill Fedyanin, Alexander Panchenko, and Maxim Panov. 2021. How certain is your transformer? In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1833–1840. (Cited on p. 2)
- Yilin Shen, Wenhua Chen, and Hongxia Jin. 2020. Modeling token-level uncertainty to learn unknown concepts in slu via calibrated dirichlet prior rnn. *arXiv preprint arXiv:2010.08101*. (Cited on p. 2)
- Aditya Siddhant and Zachary C. Lipton. 2018. [Deep bayesian active learning for natural language processing: Results of a large-scale empirical study](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2904–2909. Association for Computational Linguistics. (Cited on p. 2)
- Lewis Smith and Yarin Gal. 2018. [Understanding measures of uncertainty for adversarial example detection](#). In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, pages 560–569. AUAI Press. (Cited on p. 3)
- Jasper Snoek, Yaniv Ovadia, Emily Fertig, Balaji Lakshminarayanan, Sebastian Nowozin, D. Sculley, Joshua V. Dillon, Jie Ren, and Zachary Nado. 2019. [Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13969–13980. (Cited on p. 1)
- Andreas Stolcke. 2002. Srilm-an extensible language modeling toolkit. In *Seventh international conference on spoken language processing*. (Cited on p. 16)
- Natasa Tagasovska and David Lopez-Paz. 2019. Single-model uncertainties for deep learning. *Advances in Neural Information Processing Systems*, 32. (Cited on p. 3)
- Ming Tan, Yang Yu, Haoyu Wang, Dakuo Wang, Saloni Potdar, Shiyu Chang, and Mo Yu. 2019. [Out-of-domain detection for low-resource text classification tasks](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 3564–3570. Association for Computational Linguistics. (Cited on p. 2)
- Sunil Thulasidasan, Gopinath Chennupati, Jeff A. Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. 2019. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. *Advances in Neural Information Processing Systems*, 32. (Cited on p. 2)
- Junjiao Tian, Dylan Yung, Yen-Chang Hsu, and Zsolt Kira. 2021. A geometric perspective towards neural calibration via sensitivity decomposition. *Advances in Neural Information Processing Systems*, 34. (Cited on p. 2)
- Dennis Ulmer, Elisa Bagnani, Max Müller-Eberstein, Daniel Varab, Mike Zhang, Christian Hardmeier, and Barbara Plank. 2022a. Experimental standards for deep learning research: A natural language processing perspective. *arXiv preprint arXiv:2204.06251*. (Cited on p. 20)
- Dennis Ulmer and Giovanni Cinà. 2021. Know your limits: Uncertainty estimation with relu classifiers fails at reliable ood detection. In *Uncertainty in Artificial Intelligence*, pages 1766–1776. PMLR. (Cited on p. 3)
- Dennis Ulmer, Christian Hardmeier, and Jes Frellsen. 2022b. deep-significance-easy and meaningful statistical significance testing in the age of neural networks. *arXiv preprint arXiv:2204.06815*. (Cited on p. 21)
- Dennis Ulmer, Lotta Meijerink, and Giovanni Cinà. 2020. Trust issues: Uncertainty estimation does not enable reliable ood detection on medical tabular data. In *Machine Learning for Health*, pages 341–354. PMLR. (Cited on p. 1)
- Joost van Amersfoort, Lewis Smith, Andrew Jesson, Oscar Key, and Yarin Gal. 2021. On feature collapse and deep kernel learning for single forward pass uncertainty. *arXiv preprint arXiv:2102.11409*. (Cited on p. 4, 20)
- Jordy Van Landeghem, Matthew Blaschko, Bertrand Anckaert, and Marie-Francine Moens. 2022. Benchmarking scalable predictive uncertainty in text classification. *Ieee Access*. (Cited on p. 2, 9)

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008. (Cited on p. 3)
- Antti Virtanen, Jenna Kanerva, Rami Ilo, Jouni Luoma, Juhani Luotolahti, Tapio Salakoski, Filip Ginter, and Sampo Pyysalo. 2019. Multilingual is not enough: Bert for finnish. *arXiv preprint arXiv:1912.07076*. (Cited on p. 4, 16)
- Yoav Wald, Amir Feder, Daniel Greenfeld, and Uri Shalit. 2021. On calibration and out-of-domain generalization. *Advances in Neural Information Processing Systems*, 34. (Cited on p. 4)
- Cheng Wang, Carolin Lawrence, and Mathias Niepert. 2021a. [Uncertainty estimation and calibration with finite-state probabilistic rnns](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. (Cited on p. 3)
- Deng-Bao Wang, Lei Feng, and Min-Ling Zhang. 2021b. Rethinking calibration of deep neural networks: Do not be afraid of overconfidence. *Advances in Neural Information Processing Systems*, 34. (Cited on p. 2)
- Shuo Wang, Zhaopeng Tu, Shuming Shi, and Yang Liu. 2020. [On the inference calibration of neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3070–3079. Association for Computational Linguistics. (Cited on p. 2)
- Xiangpeng Wei, Heng Yu, Yue Hu, Rongxiang Weng, Luxi Xing, and Weihua Luo. 2020. [Uncertainty-aware semantic augmentation for neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2724–2735. Association for Computational Linguistics. (Cited on p. 2)
- Max Welling and Yee W Teh. 2011. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688. Cite-seer. (Cited on p. 2)
- Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, et al. 2020. Contrastive training for improved out-of-distribution detection. *arXiv preprint arXiv:2007.05566*. (Cited on p. 4)
- 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. (Cited on p. 20)
- Tim Z Xiao, Aidan N Gomez, and Yarin Gal. 2020. Wat zei je? detecting out-of-distribution translations with variational transformers. *arXiv preprint arXiv:2006.08344*. (Cited on p. 2, 3, 20)
- Yijun Xiao and William Yang Wang. 2021. [On hallucination and predictive uncertainty in conditional language generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 2734–2744. Association for Computational Linguistics. (Cited on p. 2, 9)
- Jiacheng Xu, Shrey Desai, and Greg Durrett. 2020. [Understanding neural abstractive summarization models via uncertainty](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6275–6281. Association for Computational Linguistics. (Cited on p. 2, 9)
- Charles Augustus Young. 1895. *Manual of Astronomy: A Text-book*. Ginn and Company. (Cited on p. 1)
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*. (Cited on p. 22)
- Shujian Zhang, Chengyue Gong, and Eunsol Choi. 2021. [Knowing more about questions can help: Improving calibration in question answering](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 1958–1970. Association for Computational Linguistics. (Cited on p. 2)
- Shengjia Zhao, Michael Kim, Roshni Sahoo, Tengyu Ma, and Stefano Ermon. 2021a. Calibrating predictions to decisions: A novel approach to multi-class calibration. *Advances in Neural Information Processing Systems*, 34. (Cited on p. 2)
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021b. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR. (Cited on p. 2)
- Tim Zimmermann, Leo Kotschenreuther, and Karsten Schmidt. 2016. Data-driven hr-r\`esum\`e analysis based on natural language processing and machine learning. *arXiv preprint arXiv:1606.05611*. (Cited on p. 1)

A Data

A.1 Pre-processing

Tokenization We use the corresponding BERT tokenizer for each language, including for LSTM-based models to ensure compatibility. For English, this corresponds to the original SentencePiece tokenizer used by Devlin et al. (2019), while we use the tokenizer of the Danish BERT (Hvingelby et al., 2020) and Finnish BERT (Virtanen et al., 2019) for those languages, respectively.

Tags for Sub-word Tokens For named entity recognition and part-of-speech tagging, we follow Jurafsky and Martin (2022), chapter 11.3.3 to deal with sub-word tokens: For every token that is split into sub-word tokens, we assign the tag only to the first sub-word token, and -100 for the rest, which ignores them for evaluation purposes.

A.2 Sub-sampling of Training Sets

Since we sub-sample some of the data splits in Table 1, this bears the dangers of producing unnatural samples of text. For that reason, we use this appendix to describe the sampling strategies in more detail.

Sub-sampling procedure The procedure for sub-sampling text is that sequences are first placed into buckets of the same label, then into sub-buckets of the same length. Then, the sampling procedure consists of first drawing a label based on the observed label frequencies, after which the draw of sequence length, proportional to the frequency of this length inside the bucket, determines the final bucket from which a sequence is again drawn uniformly. Lastly, the process for token classification involves the grouping into sequences by length at the highest level. Inside a bucket, a sequence is not drawn uniformly but with a probability according to the *alignment* of the sequence’s labels with the overall corpus label distribution. This alignment is calculated for each sequence by evaluating the expected log-probability of the sequence’s label distribution w.r.t to the label distribution of the corpus (i.e., the cross-entropy). The scores for all same-length sequences in a bucket are then normalized into a $[0, 1]$ interval in order to enable sampling, which is similar to the two-stage procedure used in the sequence classification case.

Validation of sub-sampled training sets We take multiple steps to validate the representative-

ness of our sub-sampled data splits. First, we plot the distributions of the 50 most frequent types in the original corpus in Figure 5, where we see that distributions converge with increasing sample size. Secondly, we plot sentence length distributions in Figure 6, where we also see increasing alignment with sample size.¹⁰ For Sequence and Token Classification tasks, we also plot the class distributions in Figure 7. Lastly, we train an interpolated trigram Kneser-Ney language model (Jelinek, 1980; Ney et al., 1994) with uniform interpolation weights trained on the original training set using SRILM (Stolcke, 2002) and sub-word tokens produced by the corresponding BERT tokenizer, sub-sample multiple splits and compare their perplexity scores to those of the original corpus in Table 3. While n -gram perplexities of sub-sampled training sets do lie over the ones of the original data, they are still upper-bounded by the in-distribution test-set perplexities. Furthermore, this verification was not aimed to give the most precise results, as also the scoring using an n -gram model can be rather crude. Thus, with all these results, we conclude that our sub-sampling procedure produces sufficiently representative samples of the original data for the different tasks discussed.

A.3 Selection of OOD Test Sets

In this appendix section, we present additional evidence that the OOD test splits shown in Table 1 are sufficiently different from the training data — meaning, out-of-distribution — to enable our chosen methodology. To that end, we re-use similar ideas as described in Appendix A.2, but with the opposite goal. In Figure 9, we plot the distribution of sequence lengths of the training set compared with the OOD test set, with the same done for the most frequent 25 types in Figure 10 and class labels in Figure 8. Lastly, we again use an interpolated Kneser-Ney trigram language model to compute the perplexity of the training compared to the OOD test set in Table 3. In all cases, OOD n -gram perplexities lie much over the training or sub-sampled data perplexities. Except for Finnish, they are also widely different from the test set perplexities. In that exceptional cases, an explanation could be given by the highly agglutinative nature of Finnish, increasing the sparsity of the language despite the subword tokenization.

¹⁰The distributions for Language Modelling are slightly distorted since we sample whole sets of sentences.

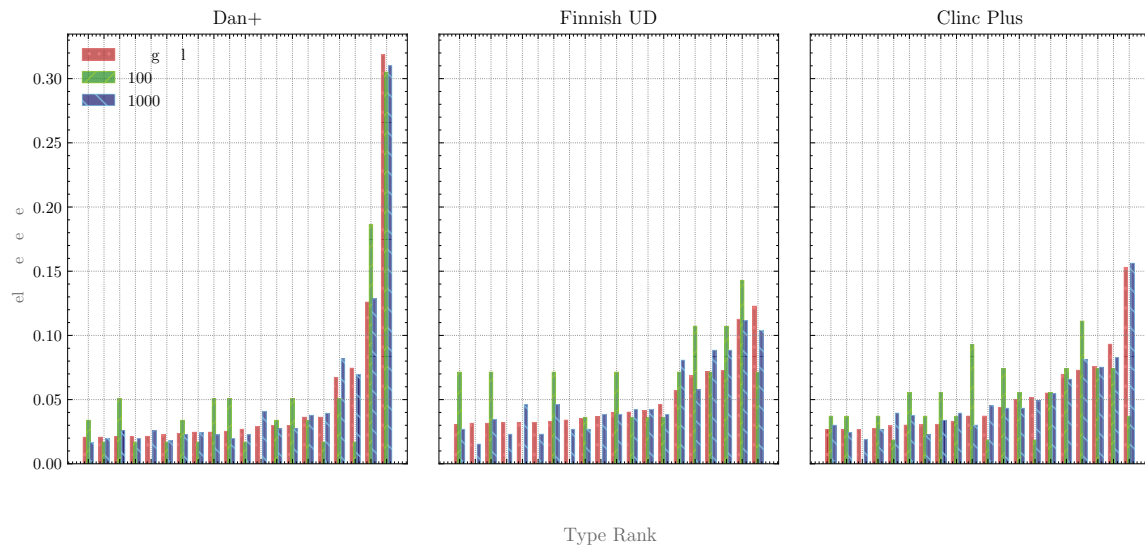


Figure 5: **Comparing the relative frequency of types in the original and sub-sampled training sets.** Shown are the top 20 types in the original training set, compared to sub-sampled training sets of 100 and 1000 sequences for Dan+, Finnish UD and Clinc Plus. It is shown that while the type frequencies differ noticeably for the small dataset, already 1000 sequences suffice to approximate the original frequencies. Numbers, stopwords and the most common punctuation were removed.

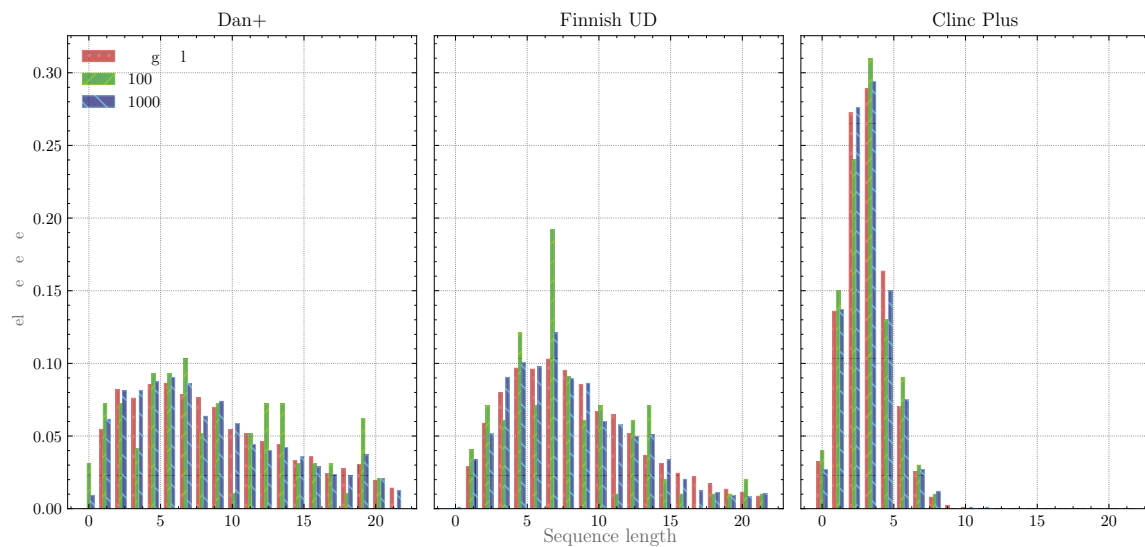


Figure 6: **Comparing the relative frequency of sequence lengths in the original and sub-sampled training sets.** Shown are sequence lengths between 0 and 25 in the original test, compared to OOD test sets for Dan+, Finnish UD, Clinc Plus. Not the whole distribution is shown in all cases, with many of the OOD sentences for Dan+ being very long. For Dan+ and Finnish UD, the sentence length distributions are noticeably different. For Clinc Plus, they are very similar.

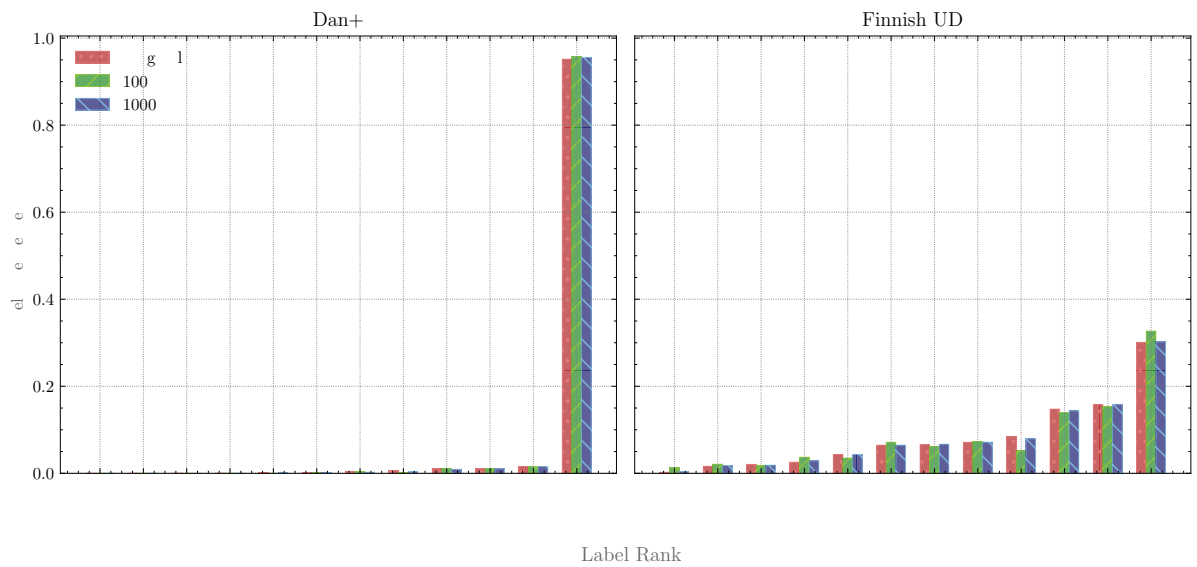


Figure 7: **Comparing the relative frequency of labels in the original training set, compared to sub-sampled training sets.** Shown are frequencies for 100 and 1000 sequences. For Danish, the most frequent label by far is the neutral label indicating that no named entity is present.

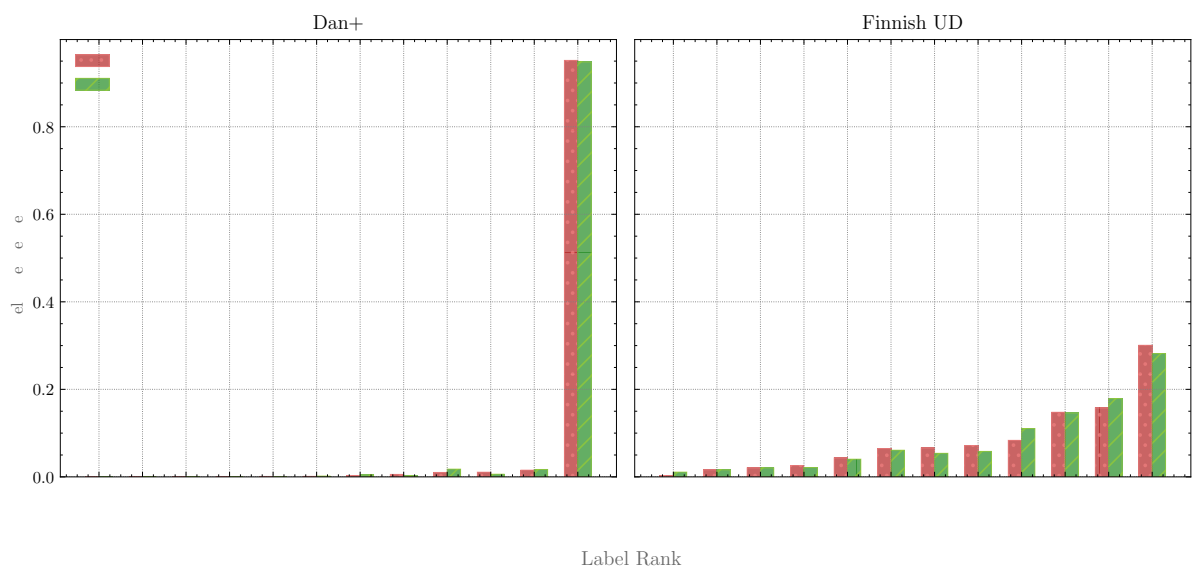


Figure 8: **Comparison of the relative class frequencies between original training set compared to the OOD test set.** The proportions stay largely the same for Danish, while different more for Finnish.

Language	Train ppl.↓	Sub-sampled Train ppl.↓			Test ppl.↓	OOD Test ppl.↓
		$n = 100$	$n = 500$	$n = 1000$		
English	31.54	43.97 ± 2.46	44.50 ± 0.68	44.9 ± 0.4	53.11	120.32
Danish	112.73	252.52 ± 13.25	247.09 ± 3.3	249.27 ± 3.15	418.71	524.32
Finnish	116.49	257.67 ± 10.96	257.66 ± 4.7	260.36 ± 5.36	1374.76	1284.82

Table 3: **Results of using an interpolated Kneser-Ney n -gram language model on selected datasets, including sub-sampled training splits and the OOD test set.** Scores of sub-sampled training sets were obtained over five different attempts.

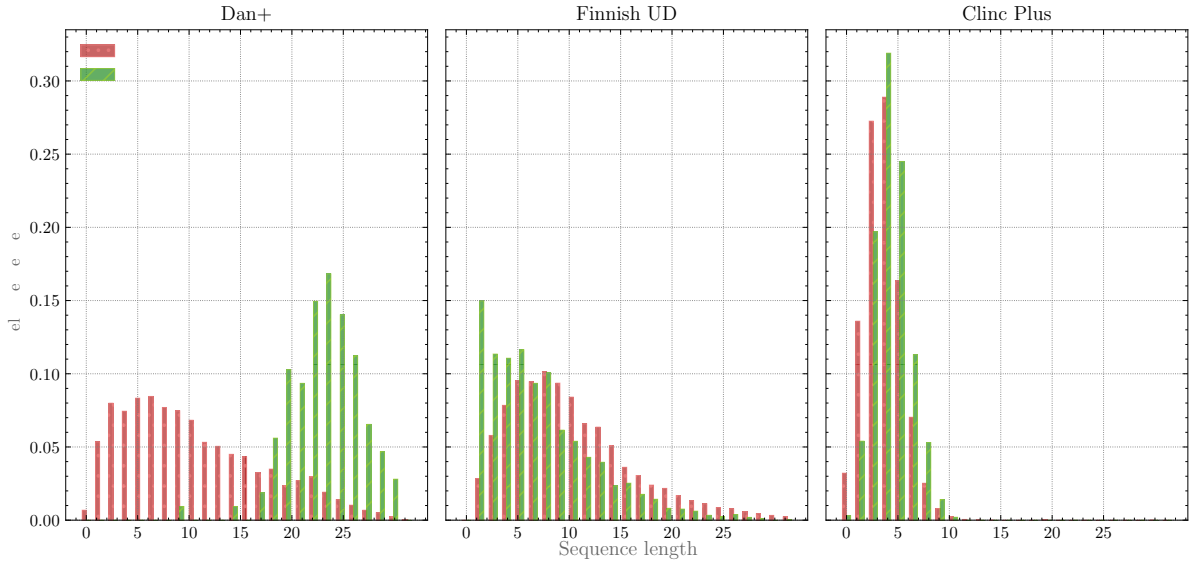


Figure 9: **Comparison of sequence length distribution between the original training set and the OOD test set.** For English, the distribution of lengths of voice assistant commands is quite similar, while the differences for Dan+ and Finnish UD are more pronounced.

B Calibration Metrics

Perfect calibration is defined as the the confidence of a neural network corresponding to the percentage of samples with that same predicted probability actually receiving the correct label by the network. Using a predicted label \hat{y} with probability \hat{p} , perfect calibration is defined as

$$P(\hat{y} = y | \hat{p} = p) = p, \quad \forall p \in [0, 1]$$

The expected calibration error (Naeini et al., 2015) quantifies the difference between the confidence and the calibration on a test set by collecting predictions into m bins:

$$\begin{aligned} \text{ECE} &= \mathbb{E}_{\hat{p}} \left[P(\hat{y} = y | \hat{p} = p) - p \right] \\ &\approx \sum_{m=1}^M \frac{|\mathbb{B}_m|}{N} \left| \text{acc}(\mathbb{B}_m) - \text{conf}(\mathbb{B}_m) \right| \\ &= \sum_{m=1}^M \frac{1}{N} \left| \sum_{b \in \mathbb{B}_m} \mathbb{1}(\hat{y}_b = y_b) - \hat{p}_b \right| \end{aligned} \quad (2)$$

where N is the number of data points and \mathbb{B}_m denotes the m -th bin.

The problem is that ECE is only defined for binary classification and depends highly on the number of bins chosen. For the former problem, Guo et al. (2017) present a naive extension to multi-class classification that only considers the most likely prediction. In order to consider all classes, Nixon et al. (2019) introduce the static calibration error

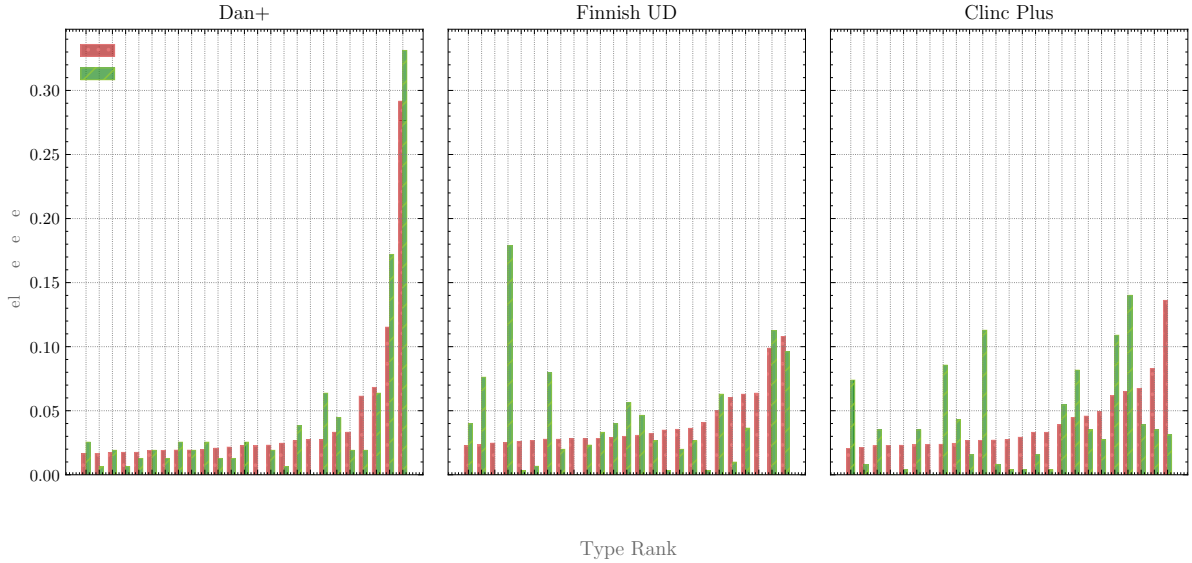


Figure 10: Comparison of the relative frequencies of the top 25 types in the original training set compared to the OOD test set. Even among the most frequent and therefore usually common tokens, the plots show differences between the in-distribution train and out-of-distribution test set. Numbers, stopwords and the most common punctuation were removed.

(SCE) as an extension to multi-class problems:

$$\text{SCE} = \frac{1}{K} \sum_{k=1}^K \sum_{m=1}^M \frac{N_{mk}}{N} \left| \text{acc}(\mathbb{B}_m, k) - \text{conf}(\mathbb{B}_m, k) \right|$$

Here, N_{mk} denotes the number of instances of class k in bin m , and $\text{acc}(\mathbb{B}_m, k)$, $\text{conf}(\mathbb{B}_m, k)$ the accuracies and confidences for class label k in bin m , respectively. However, we found this error not be very informative in our case, and therefore omitted corresponding results.

Secondly, Nixon et al. (2019) introduce the adaptive calibration error (ACE), which makes sure that every bin contains the same number of predictions. They define a calibration range by the $\lfloor N/R \rfloor$ -th index of sorted and thresholded predictions. Then, the error is defined as

$$\text{ACE} = \frac{1}{KR} \sum_{k=1}^K \sum_{r=1}^R \left| \text{acc}(\mathbb{B}_m, r) - \text{conf}(\mathbb{B}_m, r) \right|$$

C Implementation & Experiments

C.1 Implementation Details

Resources All models were implemented in PyTorch (Paszke et al., 2019). BERT models were implemented with the help of HuggingFace’s transformers library (Wolf et al., 2020). Linear algebra operations were often implemented using the EinOps package (Rogozhnikov, 2022).

The Bayesian LSTM was developed using the Blitz package (Esposito, 2020) for PyTorch and the SNGP transformer using gpytorch (Gardner et al., 2018). Huggingface’s datasets (Lhoest et al., 2021) were furthermore used for dataset creation and codecarbon (Schmidt et al., 2021) for carbon emissions tracking. Weights & Biases (Biewald, 2020) was used to track and manage hyperparameter searches and experiments. In general, we follow many of the experimental guidelines and suggestions laid out by Ulmer et al. (2022a).

Models For the DUE transformer, we used Principal Component Analysis on the latent representations for Clinc Plus to reduce the memory usage of the Gaussian Discriminant Analysis by reducing dimensionality to 64. We initially also experimented with the usage of the DUE transformer by (van Amersfoort et al., 2021), however found that it was not trivial to create the inducing points for the Gaussian process output layer in a sequential setting. For the Variational Transformer (Xiao et al., 2020), the authors do not specify exactly how MC Dropout is used. We use the existing dropout layers in the corresponding model, and use a number of forward passes with different dropout masks to make predictions. Since the number of classes is prohibitive for the original formulation of the SNGP transformer, we use the extension proposed by Liu et al. (2022) in Appendix A.1 and only store one $\hat{\Sigma}^{-1}$ matrix for all classes. Furthermore, we update the matrix

continuously during training and not just during the last epoch, in order to allow tracking of the predictive performance over the training time. Lastly, we also evaluate predictions using Monte-Carlo approximations instead of the mean-field approach, since this allows us to compute a wider variety of uncertainty metrics.

Evaluation When computing uncertainty estimates and losses for evaluation purposes, the measurements for a number of tokens were discarded. These include the ignore token with ID -100 , as well as the IDs corresponding to the [EOS], [SEP], [CLS] and [PAD] token, which might differ between tokenizers of different languages. For computing the ECE, we use 10 bins, and 10 value ranges for ACE.

Model Comparison We facilitate the comparison of models using the almost stochastic order test (ASO; del Barrio et al., 2018; Dror et al., 2019), as implemented by Ulmer et al. (2022b). One distribution is stochastically dominant over the other when its cumulative distribution function is equal or larger than its counterpart at all points. In an experimental setting, that implies that a model is producing higher scores than a baseline. The ASO test measures the deviation from the stochastic order using an approach rooted in optimal transport. We use the test with a confidence level $\alpha = 0.05$ and a decision threshold of $\tau = 0.3$.

C.2 Hyperparameters

We perform hyperparameter search using random sampling (Bergstra and Bengio, 2012) using hyperband scheduling (Li et al., 2017)¹¹ on the entire training set, even if models are trained on sub-sampled training sets later. This has the advantage of ensuring comparability between runs and eliminating suboptimal hyperparameter choices as a source of worse uncertainty estimation. We do 80 trials for LSTM-based models, and 30 for BERT-based models. Furthermore, the hyperparameters for the LSTM are identical for the LSTM ensemble (10 instances are used per ensemble). Hyperparameters were picked by best final validation loss over search trials.

Chosen Hyperparameters We summarize some common hyperparameters here and show the rest in Table 5. We commonly use a batch size of 32, and

¹¹Trials might be terminated using Hyperband after $10k$ steps.

Name	Tuned for	Search space
Learning rate	LSTM, LSTM Ensemble, Bayesian LSTM, ST- τ LSTM, Variational LSTM	$\mathcal{U}(0.1, 0.5)$
Learning rate	DDU BERT, SNGP BERT, Variational BERT	$\log\mathcal{U}(10^{-5}, 10^{-3})$
Spectral norm upper bound	DDU BERT, SNGP BERT	$\mathcal{U}(0.95, 0.99)$
Kernel amplitude	SNGP BERT	$\log\mathcal{U}(0.01, 0.5)$
β weight decay	SNGP BERT	$\log\mathcal{U}(10^{-3}, 0.5)$
Weight decay	LSTM, LSTM Ensemble, ST- τ LSTM, Variational BERT	$\mathcal{U}(0.1, 0.5)$
Layers	LSTM, LSTM Ensemble	{2, 3}
Dropout	LSTM, LSTM Ensemble, ST- τ LSTM, Variational BERT	$\mathcal{U}(0.1, 0.4)$
Layer Dropout	Variational LSTM	$\mathcal{U}(0.1, 0.4)$
Time Dropout	Variational LSTM	$\mathcal{U}(0.1, 0.4)$
Embedding Dropout	Variational LSTM	$\mathcal{U}(0.1, 0.4)$
Hidden size	LSTM, LSTM Ensemble	{350, 500, 650}
Prior σ_1	Bayesian LSTM	$\log\mathcal{U}(-0.8, 0.1)$
Prior σ_2	Bayesian LSTM	$\log\mathcal{U}(-0.8, 0.1)$
Prior π	Bayesian LSTM	$\log\mathcal{U}(0.1, 0.9)$
Posterior μ init	Bayesian LSTM	$\mathcal{U}(-0.6, 0.6)$
Posterior ρ init	Bayesian LSTM	$\mathcal{U}(-8, -2)$
Init weight	LSTM	$\mathcal{U}(0.1, 0.4)$
Number of centroids	ST- τ LSTM	{5, 10, 20, 30, 40}

Table 4: **List of searched hyperparameters.** LSTM Ensemble hyperparameters are not searched, but simply copied from the found LSTM hyperparameters.

sequence lengths of 35 for LSTM-based and 128 for BERT-based models. All LSTM-based models are trained using 2 layers, with the exception of the vanilla LSTM and the LSTM-ensemble on Clinc Plus with 3 layers. Their hidden size and embedding sizes are set to 650. For all models, gradient clipping is set to 10. For models using multiple predictions to compute uncertainty estimates, 10 predictions are used at a time.

C.3 Optimization

To make sure that all models are trained for the same number of steps regardless of the size of (sub-sampled) training set, we set the training duration to the number of steps corresponding to a number of epochs using the original training set size, and name it *epoch-equivalents* in the following. Due to the imbalance of classes in Finnish UD and Dan+, all models were trained using loss-weights that are inverse to the frequency of a label

Model	Hyperparameter	English	Danish	Finnish
LSTM	Weight decay	0.001337	0.001357	0.001204
	Learning rate	0.4712	0.4931	0.2205
	Init. weight	0.283	0.5848	0.5848
	Dropout	0.3379	0.2230	0.1392
Variational LSTM	Weight decay	–	10^{-7}	0.01953
	Learning rate	–	0.3031	0.7817
	Init. weight	–	0.1097	0.5848
	Embedding Dropout	–	0.1207	0.3566
	Layer Dropout	–	0.1594	0.3923
	Time Dropout	–	0.1281	0.1646
	Bayesian LSTM	Weight decay	0.001341	0.003016
Bayesian LSTM	Learning rate	0.1704	0.1114	0.1549
	Dropout	0.3410	0.3868	0.331
	Prior σ_1	0.9851	0.7664	0.3246
	Prior σ_2	0.5302	0.851	0.5601
	Prior π	1	1	0.1189
	Posterior μ init	–0.005537	–0.0425	0.4834
	Posterior ρ init	–7	–6	0.1124
	ST- τ LSTM	Weight decay	–	0.001189
ST- τ LSTM	Learning rate	–	0.01979	0.3601
	Dropout	–	0.1867	0.1737
	Num. centroids	–	5	30
	DDU Bert	Learning Rate	0.003077	0.00006168
DDU Bert	Spectral norm upper bound	0.9753	0.9211	0.941
	Weight decay	0.003	0.1868	0.09439
	Variational BERT	Learning Rate	0.0002981	0.00009742
Variational BERT	Weight decay	0.01591	0.02731	0.09927
	Dropout	0.2382	0.4362	0.4364
	SNGP Bert	Learning Rate	–	0.0002332
SNGP Bert	Spectral norm upper bound	–	0.99	0.96
	Beta Weight decay	–	0.001619	0.002438
	Beta length scale	–	2.467	2.254
	Kernel amplitude	–	0.3708	0.2466

Table 5: List of used model hyperparameters by dataset.

in the dataset.

Optimization of LSTMs We adopt different optimization schemes for transformer and LSTM-based models. For LSTMs, we choose stochastic gradient descent with a decaying learning rate schedule, decaying by 0.8695 after the equivalent of 14 epochs for every following epoch-equivalent for 55 epoch-equivalents in total. This corresponds to the setup in Gal and Ghahramani (2016b), modified from the setup in Zaremba et al. (2014).

Optimization of BERTs We fine-tune BERT models using the shorter duration of 20 epoch-equivalents, corresponding to the NLP experiments in Liu et al. (2022). Adam (Kingma and Ba, 2015) is used for optimization with default parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$ alongside a triangular learning rate, using the first 10% of the training duration as warm-up.

C.4 Convergence on Clinc Plus

Here, we briefly address the models missing from the English Clinc Plus experiments. For the ST- τ and Variational LSTM, we could not identify clear reasons on why models did not converge. Even after extensive hyperparameter searches and manual fine-tuning of hyperparameters (including different learning rate schedules and optimizers), we did not find a combination of options that resulted in convergence. We also observed strange behavior for the Bayesian LSTM, which, after reaching a validation accuracy of 0.5, would suddenly return to its initial training performance. This could potentially be explained by the model accidentally escaping a low-loss basin due to a learning rate that is still too high, and thus we changed the model to only be trained for 18 epoch-equivalents and initiate the learning rate decay after seven epoch-equivalents. The puzzling fact is that SNGP BERT did not converge on Clinc Plus, since the authors successfully used the dataset in their own work (Liu et al., 2022). We put forth the following explanations: First of all, we observed the model to generally possess a high variance, as demonstrated by the standard deviation on the Danish and Finnish data. Secondly, we make at least two changes to their implementation: Instead of using the mean-field approximation to the predictive distribution, we use the Monte Carlo approximation in order to compute metrics such as mutual information. Also, we update the covariance matrix $\hat{\Sigma}$ over the whole training time in order to track the predictive performance for our experiments, and not just during the last epoch.

C.5 Environmental Impact

The carbon efficiency was estimated to be 0.61 kgCO₂eq/kWh. 735 hours of computation were performed on a Tesla V100 GPU. This includes hyperparameter search, failed runs, debugging, and discarded runs. As a rough upper bound, we estimate the compute time for a single replication of all

experiments to take around 73 hours.¹² To lessen the environmental impact, all models and model predictions are published in the open-source repository. Total emissions are estimated to be 52.45 kgCO₂eq. We use direct air capture by climeworks to offset the emissions (climeworks, 2022). Estimations were conducted using the codecarbon package (Schmidt et al., 2021), a joint effort from authors of Lacoste et al. (2019) and Lottick et al. (2019).

D Additional Results

This section contains additional experimental results and plots that could not be added to the paper due to spatial constraints. We roughly follow the structure of Section 4.

D.1 Additional Scatter Plots

This section provides some additional scatter plots. For all plots presented here as well as Figure 2, some slight jitter sampled from $\mathcal{N}(0, 0.01)$ was added to x and y-coordinates to increase readability of overlapping points.

Clinc Plus In Figures 11a and 12a, we can see that the Variational Bert model actually *degrades* in performance as the more training data is added, both on a task and uncertainty dimensions, while other models stay relatively constant. The same trend can be detected using the sequence-level Kendall’s τ for ClinC Plus. We suspect that the smallest training size of 10k examples does already provide enough data for models to converge to similar solutions even after adding more data, and that the Variational Bert alone might be prone to overfitting in this case.

Dan+ Results for the Danish dataset are shown in Figures 11b and 12b. It is apparent that LSTM-based models stay mostly constant in their predictive performance, with the largest gains observed by the LSTM ensemble. We can also observe the DDU and Variational BERT to increase both in task performance and uncertainty quality with increasing training data. Interestingly, we can see for the SNGP BERT that uncertainty estimates become more indicative of OOD with more training samples, but mostly only using predictive entropy and

the maximum probability score. This might indicate that in these cases, the model actually achieves the desired distance-awareness posed by Liu et al. (2022). In Figure 14b, we can see a similar behavior of the SNGP-BERT and its metrics w.r.t. to the sequence-level correlation. Also, we see that the other BERT models and LSTM-Ensemble actually loose in uncertainty quality as more data is added.

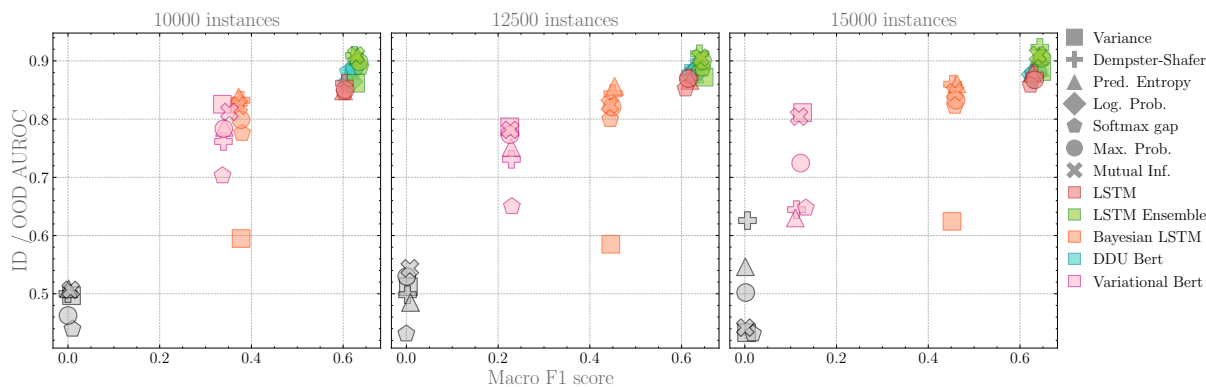
Finnish UD In Figures 11c and 12c, we see that the AUROC and AUPR scores of differnet models and metrics stay largely constant across dataset sizes, which could be explained with the larger amount of training data supplied compared to Dan+. On the token-level correlation between uncertainty and loss in Figure 13, we see the DDU BERT profiting most from more data. On a sequence-level, as depicted in Figure 14c, the correlation appears mostly static across training set sizes, with only small gaps between in-distribution and out-of-distribution data.

Overall, it seems that the range of dataset sizes for Dan+ show the most critical differences between models, while for the dataset sizes used for Finnish UD and ClinC Plus, enough data seems to be supplied for changes to be more miniscule. This result is particularly relevant for low-resource setting, although the dependency on the task can not be disentangled from these results.

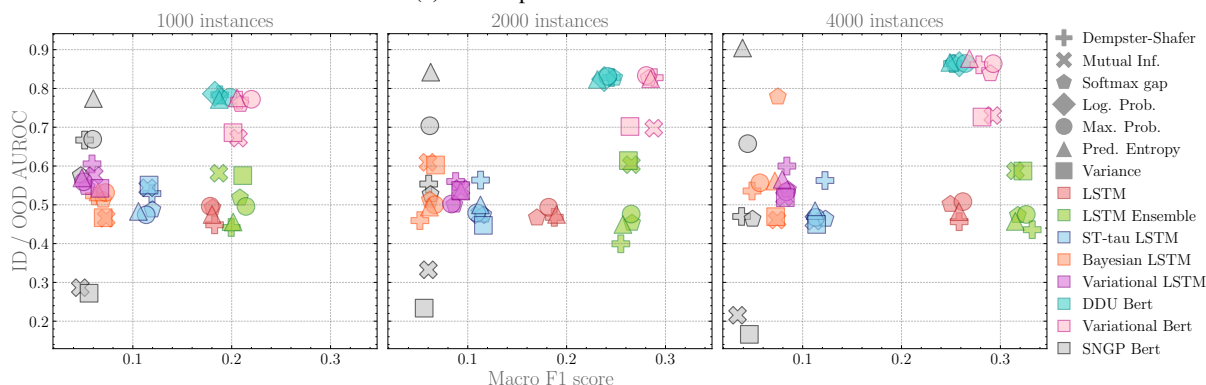
D.2 Additional Uncertainty over Training Plots

We extend the plots from Figure 3 for all tested models and datasets in Figure 15 and Figure 16, showing the correlation of predictive entropy on a token- and sequence-level, respectively. On the token-level, we see that token-level correlation is the highest for SNGP-BERT, although the correlation levels for training set sizes seems to be harder to differentiate between models and could also be due to variance between models runs. Secondly, on a sequence-level, we also see either similar correlation across training set sizes, or higher correlation for lower sizes. In all cases, we observe that some models start with a high correlation that decreases over the training time, as the model fits the in-distribution data better. That corroborates a trend described in Section 4.2, implying that uncertainty estimates become less reliable as the model tries to decrease the loss on the training data.

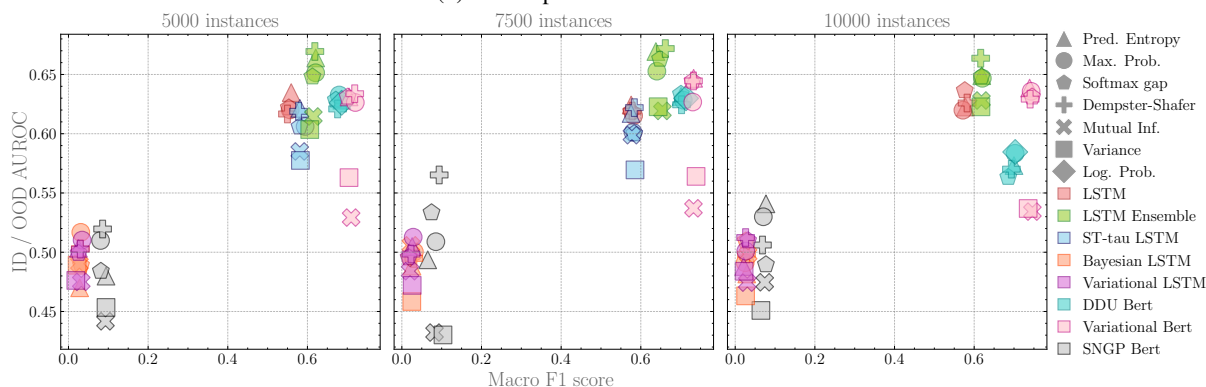
¹²Note that this number could be reduced further by using better hardware acceleration, larger batch sizes, and slightly reducing the training duration for some models. Most importantly, this number also includes compute used for hyperparameter search.



(a) Scatter plot for the Clin Plus dataset.



(b) Scatter plot for the Dan+ dataset.



(c) Scatter plot for the Finnish UD dataset.

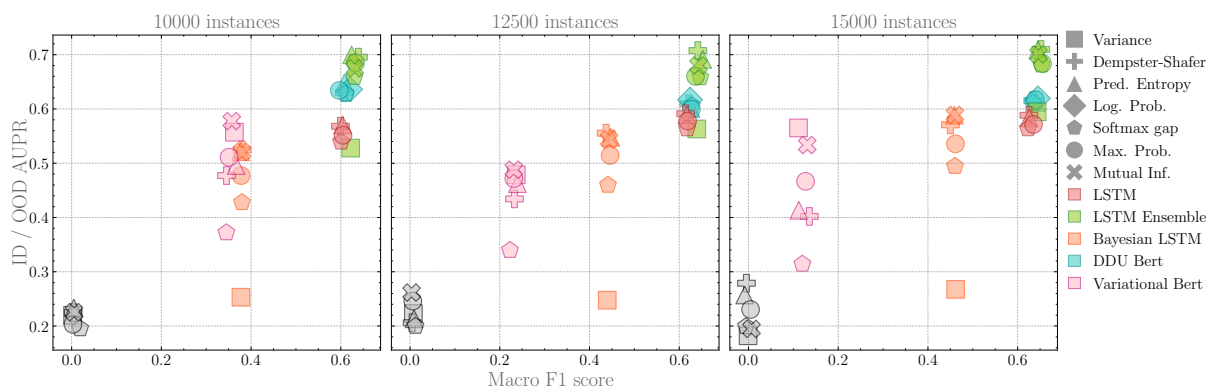
Figure 11: Scatter plot showing the difference between model performance (measured by macro F_1) and the quality of uncertainty estimates using AUROC. Shown are different models and uncertainty metrics and several training set sizes on the used datasets.

D.3 Qualitative Analysis

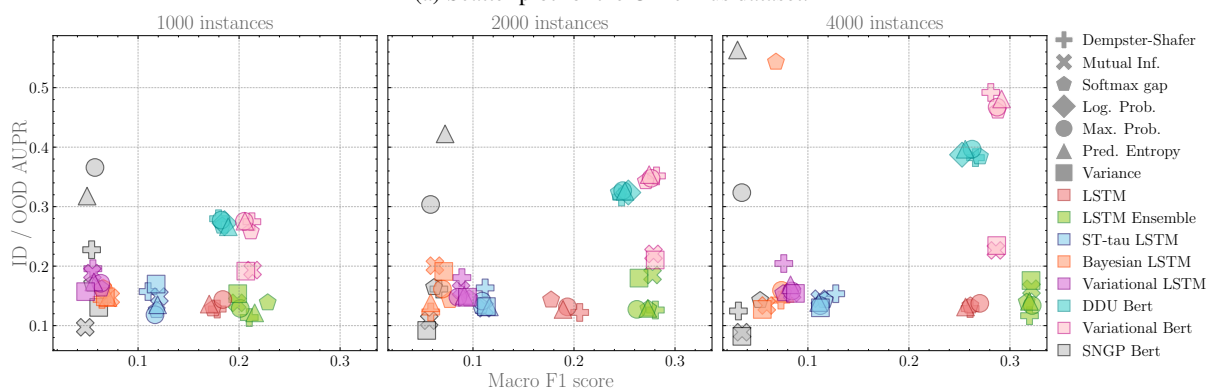
Dan+ We show more examples of the predictive entropies on samples from the Dan+ dataset in Figure 17, where uncertainty values were jointly normalized by subtracting the mean and dividing by the standard deviation over all models and time steps. We can make the following observations: Firstly, uncertainty seems to decrease on punctuation marks such as commas and full-stops. Secondly, uncertainty appears higher on sub-word tokens and some named entities. Thirdly, DDU

BERT and the LSTM ensemble produce the highest uncertainty values, which are also two of the best performing models on the task.

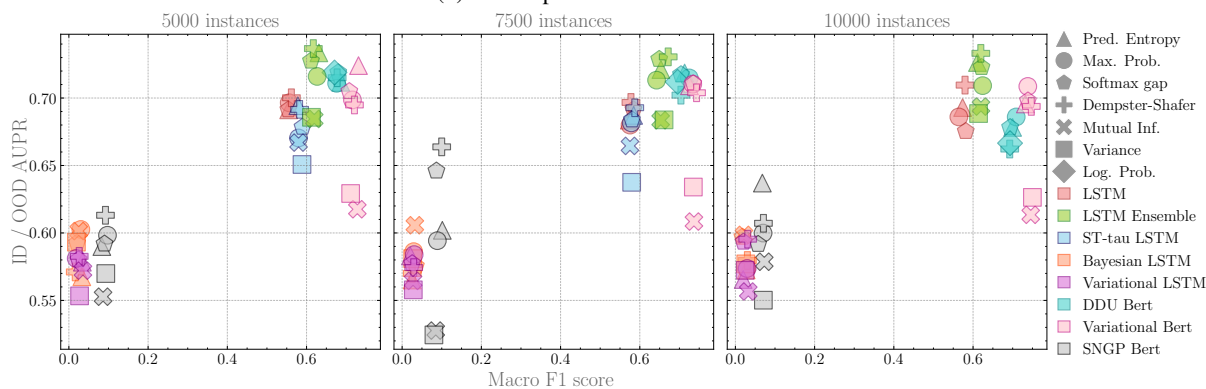
Finnish UD Here, we give more examples of the analysis on the Finnish UD dataset in Figure 18. First of all, we see that the Variational LSTM and SNGP BERT seem to produce almost constant uncertainty scores, which can be explained by their suboptimal performance in task, as shown by their results in Table 2. But even for the models that per-



(a) Scatter plot for the Clin Plus dataset.



(b) Scatter plot for the Dan+ dataset.



(c) Scatter plot for the Finnish UD dataset.

Figure 12: Scatter plot showing the difference between model performance (measured by macro F_1) and the quality of uncertainty estimates using AUPR. Shown are different models and uncertainty metrics and several training set sizes on the used datasets.

form better, such as the Variational BERT and the LSTM ensemble, the decomposition of predictive entropy into aleatoric and epistemic uncertainty reveals that model uncertainty generally remains low, and is overshadowed to a larger extent by the aleatoric uncertainty. We can observe that similar to Danish, uncertainty seems to be low on punctuation marks and high on subword tokens. Furthermore, aleatoric uncertainty seems to be higher on nouns and pronouns. This could be due to the sheer number of possible nouns and pronouns that could

fill such a gap in a sentence.

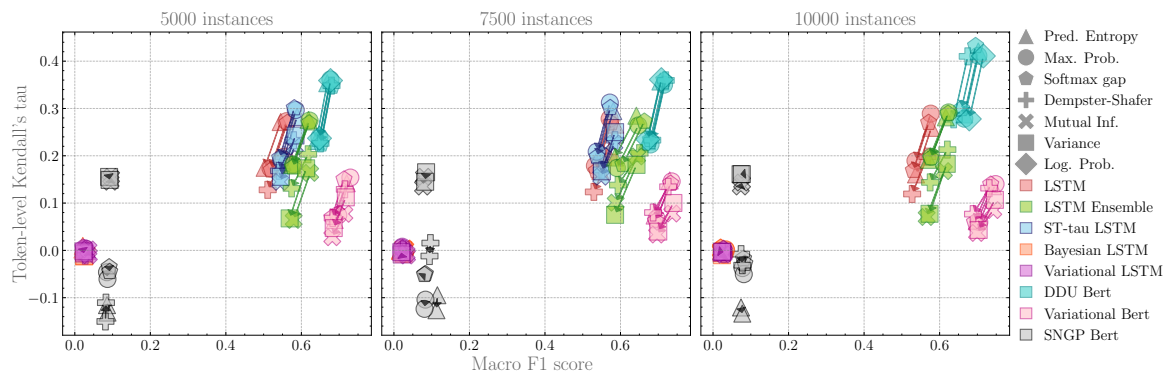
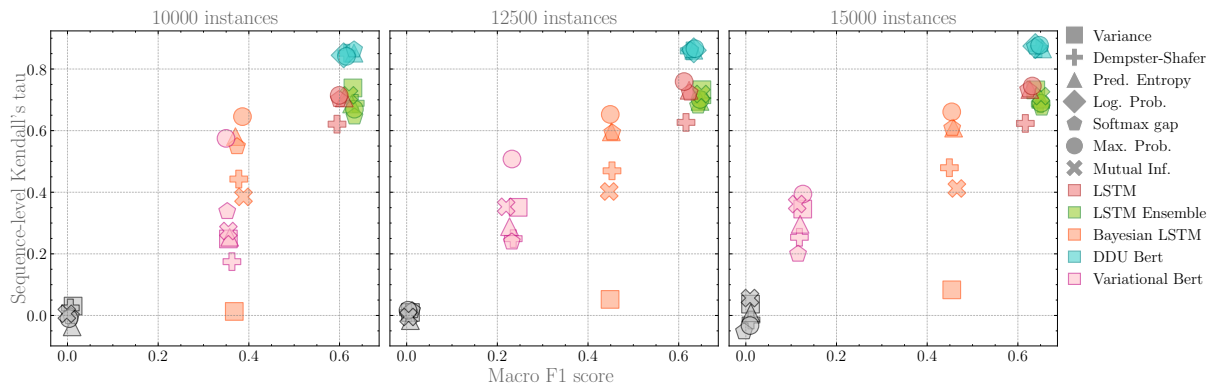
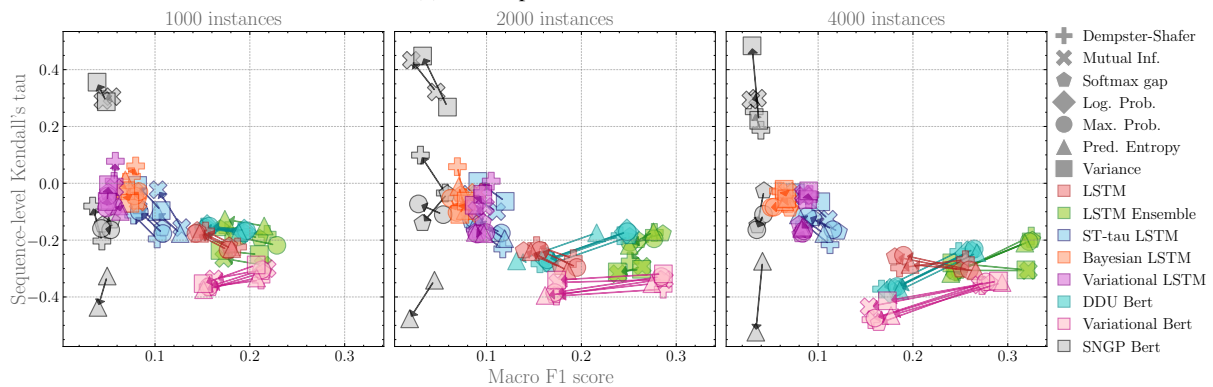


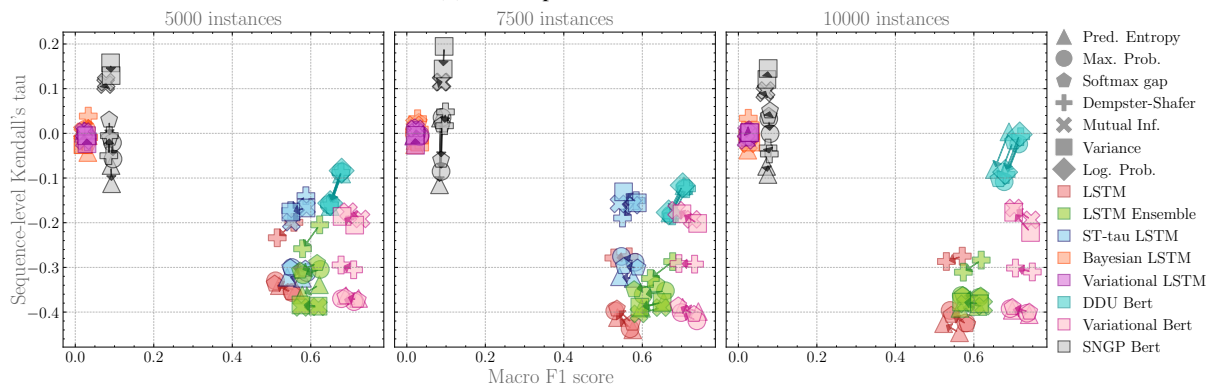
Figure 13: Scatter plot showing the difference between model performance (measured by macro F_1) and the quality of uncertainty estimates on a token-level (measured by Kendall's τ). Results are shown for different models and uncertainty metrics and several training set sizes on the Finnish UD dataset. Arrows indicate changes between the in-distribution and out-of-distribution test set.



(a) Scatter plot for the Clinc Plus dataset.

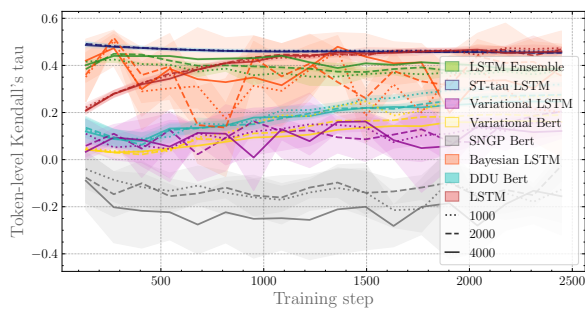


(b) Scatter plot for the Dan+ dataset.

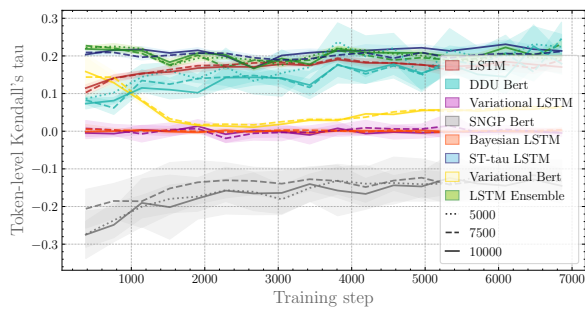


(c) Scatter plot for the Finnish UD dataset.

Figure 14: Scatter plot showing the difference between model performance (measured by macro F_1) and the quality of uncertainty estimates on a sequence-level (measured by Kendall's τ). Results are shown for different models and uncertainty metrics and several training set sizes on the Finnish UD and Clinc Plus dataset. Arrows indicate changes between the in-distribution and out-of-distribution test set.

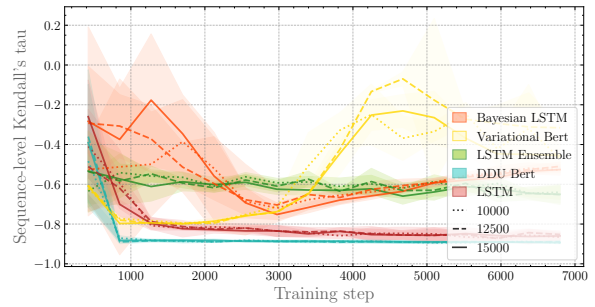


(a) Development of token-level Kendall's τ for Dan+.

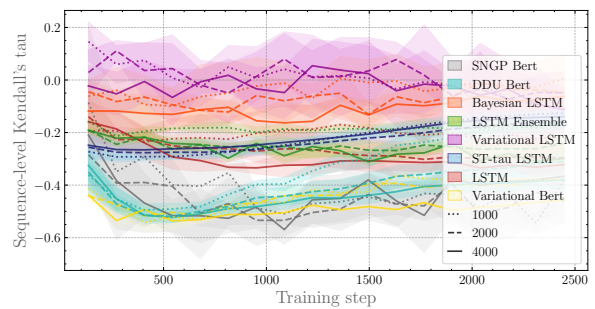


(b) Development of token-level Kendall's τ for Finnish UD.

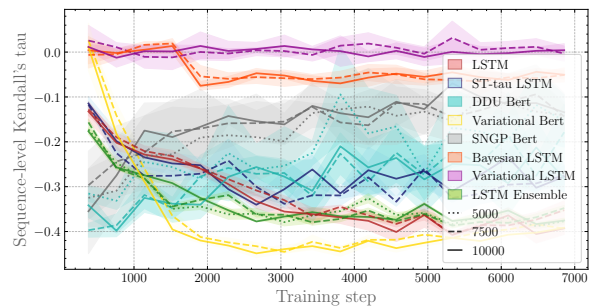
Figure 15: Development of correlation between token-level predictive entropy and loss on the Dan+ and Finnish UD OOD test set over the training time. Data is shown for several model types and using differently-sized training sets. Colored areas indicate the standard deviation over five runs.



(a) Development of sequence-level Kendall's τ for Clinc Plus.

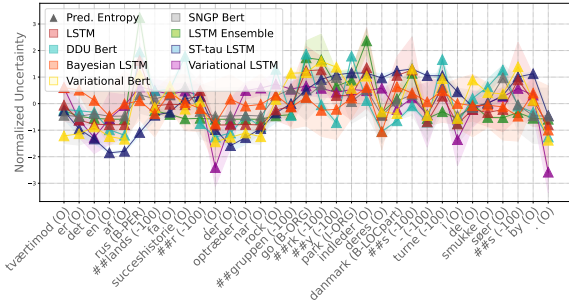


(b) Development of sequence-level Kendall's τ for Dan+.

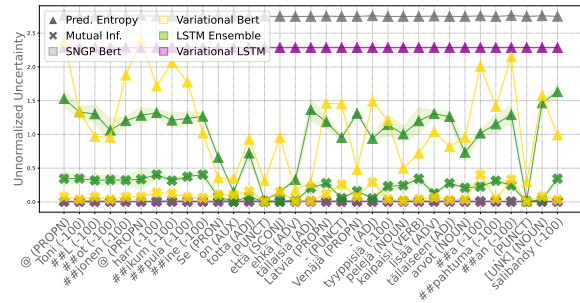


(c) Development of sequence-level Kendall's τ for Finnish UD.

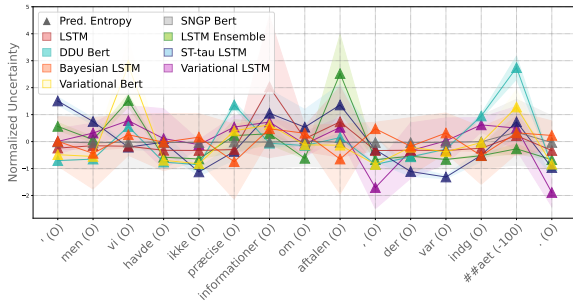
Figure 16: Development of correlation between sequence-level predictive entropy and loss on all OOD test sets over the training time. Data is shown for several model types and using differently-sized training sets. Colored areas indicate the standard deviation over five runs.



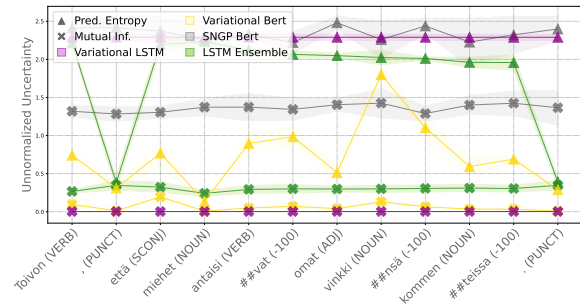
(a) Predictive entropy over the sentence "On the contrary, it is one of Russia's few success stories that performs when the rock group Gorky Park begins their Danish tour in the city of the beautiful lakes".



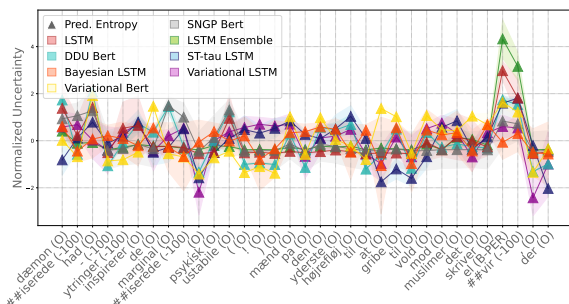
(a) Predictive entropy over the sentence "@ToniLotjonen @harrikumpulaine It is true that I'd maybe like to see more of such Latvia-Russia type games in these kinds of major sports events. #floorball".



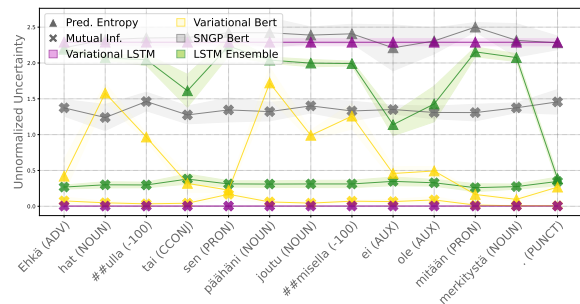
(b) Predictive entropy over the sentence "However, we did not have precise information about what was agreed upon".



(b) Predictive entropy over the sentence "I hope that the procedures done on the person in question stop and he gives his body (and mind) time to recover from that poisoning!".



(c) Predictive entropy over the sentence "Demonizing hate speech inspires the marginalized, PSYCHOLOGY UNSTABLE (!) Men on the far right to resort to violence against Muslims. This writes Elvir, who...".



(c) Predictive entropy over the sentence "Maybe the hat or how it got on my head doesn't matter".

Figure 17: Further examples for uncertainty estimates on single sequences. Taken from the Dan+ dataset.

Figure 18: Further examples for uncertainty estimates on single sequences. Taken from the Finnish UD dataset.