

Switching Contexts: Transportability Measures for NLP

Guy Marshall^{*†}, Mokanarangan Thayaparan^{*†}, Philip Osborne[†], André Freitas^{†‡}

Department of Computer Science, University of Manchester, United Kingdom[†]

Idiap Research Institute, Switzerland[‡]

{guy.marshall, philip.osborne}@postgrad.manchester.ac.uk

{mokanarangan.thayaparan, andre.freitas}@manchester.ac.uk

Abstract

This paper explores the topic of transportability, as a sub-area of generalisability. By proposing the utilisation of metrics based on well-established statistics, we are able to estimate the change in performance of NLP models in new contexts. Defining a new measure for transportability may allow for better estimation of NLP system performance in new domains, and is crucial when assessing the performance of NLP systems in new tasks and domains. Through several instances of increasing complexity, we demonstrate how lightweight domain similarity measures can be used as estimators for the transportability in NLP applications. The proposed transportability measures are evaluated in the context of Named Entity Recognition and Natural Language Inference tasks.

1 Introduction

The empirical evaluation of the quality of NLP models under a specific task is a fundamental part of the scientific method of the NLP community. However, commonly, many proposed models are found to perform well in the specific context in which they are evaluated and state-of-the-art claims are usually found not transportable to similar but different settings. The current evaluation metrics may only indicate which algorithm or setup performs best: they are unable to estimate performance in a new context, to demonstrate internal validity, or to verify causality. To offset this, statistical significance testing is sometimes applied in conjunction with performance measures (e.g. F1-score, BLEU) to attempt to establish validity. However, statistical significance testing has been shown to be lacking. Dror et al. (2018) reviewed NLP papers from ACL17 and TACL17 and found that only a third of these papers use significance

testing. Further, many papers did not specify the type of test used, and some even employed an inappropriate statistical test.

Performance is measured in NLP tasks primarily through F1 score or task-specific metrics such as BLEU. The limited scope of these as performance evaluation techniques has been shown to have issues. Søggaard et al. (2014) highlights the data selection bias in NLP system performance. Gorman and Bedrick (2019) show issues of using standard splits, as opposed to random splits. We support their statement that “practitioners who wish to firmly establish that a new system is truly state-of-the-art augment their evaluations with Bonferroni-corrected random split hypothesis testing”. In an NLI task, using SNLI and MultiNLI datasets with a set of different models, it has been shown that permutations of training data leads to substantial changes in performance (Schluter and Varab, 2018).

Further, the lack of transportability for NLP tasks has been raised by specialists in applied domains. For example, healthcare experts have expressed their frustration in the limitations of algorithms built in research settings for practical applications (Demner-Fushman and Elhadad, 2016) and the reduction of performance “outside of their development frame” (Maddox and Matheny, 2015). More generally, “machine learning researchers have noted current systems lack the ability to recognize or react to new circumstances they have not been specifically programmed or trained for” (Pearl, 2019).

The advantages of “more transportable” approaches, such as BERT, in terms of their performance in multiple different domains, is currently not expressed (other than the prevalence of such architectures across a range of state-of-the-art tasks and domains). To support analysis and investigation into the insight that could be gained by examination of these properties, we suggest metrics

* equal contribution

and a method for measuring the transportability of models to new domains. This has immediate relevance for domain experts, wishing to implement existing solutions on novel datasets, as well as for NLP researchers wishing to assemble new dataset, design new models, or evaluate approaches.

To support this, we propose feature gradient, and show it to have promise as a way to gain lexical or semantic insight into factors influencing the performance of different architectures in new domains. This differs from data complexity, being a comparative measure between two datasets. We aim to start a conversation about evaluation of systems in a broader setting, and to encourage the creation and utilisation of new datasets.

This paper focuses on the design and evaluation of a lightweight transportability measure in the context of the empirical evaluation of NLP models. A further aim is to provide a category of measures which can be used to estimate the stability of the performance of a system across different domains. An initial transportability measure is built by formalising properties of performance stability and variation under a statistical framework. The proposed model is evaluated in the context of Named Entity Recognition tasks (NER) and Natural Language Inference (NLI) tasks across different domains.

Our contribution is to present a measure that evaluates the transportability and robustness of an NLP model, to evaluate domain similarity measures to understand and anticipate the transportability of an NLP model, and to compare state of the art models across different datasets for NER and NLI.

2 Relevant background and related work

2.1 Terminology

To quote [Campbell and Stanley \(2015\)](#), “External validity asks the question of generalizability: To what populations, settings, treatment variables, and measurement variables can this effect be generalized?”. For [Pearl and Bareinboim \(2014\)](#), transportability is how generalisable an experimentally identified causal effect is to a new population where only observational studies can be conducted. “However, there is an important difference, not often distinguished, between what might be called the potential (or generic) transferability of a study and its actual (or specific) transferability to another policy or practice decision context at another time and place.” ([Walker et al., 2010](#))

[Bareinboim and Pearl \(2013\)](#) explore transfer of causal information, culminating in an algorithm for identifying transportable relations. Transportability in this sense does not permit retraining in the new population, and guides our choices in this paper. Other definitions of transfer learning allow for training of the model in the new context ([Pan and Yang, 2010](#)), or highlight the distinction between evidential knowledge and causal assumptions ([Singleton et al., 2014](#)).

2.2 Transportability: Models evaluated across different datasets

[Rezaeinia et al. \(2019\)](#) consider improving transportability by demonstrating word embeddings’ accuracy degrades over different datasets, and propose an algorithmic method for improved word embeddings by using word2vec, adding gloVe when missing, and filling any further missing values with random entries. In a medical tagging task, [Ferrández et al. \(2012\)](#) used different train/test datasets, and compared precision and recall with self-trained vs transported-trained, finding that some tag-categories performed better than others. They postulate that degradation differences were due to the differing prevalence of entities in the transported training data. Another term from this domain is “portability”, in the sense that a model could be successfully used with consideration of implementation issues such as different data formats and target NLP vocabularies ([Carroll et al., 2012](#)). [Blitzer et al. \(2007\)](#) created a multi-domain dataset for sentiment analysis, and propose a measure of domain similarity for sentiment analysis based on the distance between the probability distributions in terms of characteristic functions of linear classifiers.

In image processing, domain transfer is an active area of research. [Pan et al. \(2010\)](#) propose transfer component analysis as a method to learn subspaces which have similar data properties and data distributions in different domains. They state that domain adaptation is “a special setting of transfer learning which aims at transferring shared knowledge across different but related tasks or domains”. In computer vision, [Peng et al. \(2019\)](#) combine multiple datasets into a larger dataset DomainNet, and consider multi-source domain adaptation, formalising for binary classification. They demonstrate multi-source training improves model accuracy, and publish baselines for state of the art methods.

2.3 Generalisability

The language used in literature is not consistent. Bareinboim and Pearl (2013) highlights that generalisability goes under different “rubrics” such as external validity, meta-analysis, overgeneralisation, quasi-experiments and heterogeneity.

Boulenger et al. (2005) disambiguate terms in the context of healthcare economics (such as generalisability, external validity, and transferability), and created a self-reporting checklist to attempt to quantify transferability. They define generalisability as “the degree to which the results of a study hold true in other settings”, and “the data, methods and results of a given study are transferable if (a) potential users can assess their applicability to their setting and (b) they are applicable to that setting”. They advocate a user-centric view of transferability, considering specific usability aspects such as explicit currency conversion rates.

Antonanzas et al. (2009) create a transferability index at general, specific and global levels. Their “general index” is comprised of “critical factors”, which utilise Boulenger et al.’s factors, adding subjective dimensions.

3 Transportability in NLP

3.1 Definitions

To support a rigorous discussion, notational conventions are introduced. Extending the choices of Pearl and Bareinboim (2011), we denote a domain \mathcal{D} with population Π , governed by feature probability distribution P , which is data taken from a particular domain. We denote the *source* with a 0 subscript.

Definition 1. *Generalisability:* A system Ψ has performance p for solving task T_0 in domain \mathcal{D}_0 . Generalisability is how the system Ψ performs for solving task T_i in domain \mathcal{D}_j , relative to the original task, without retraining.

Special cases, such as transportability or transferance, have some $i, j = 0$ in the definition above.

Definition 2. *Transportability:* A system Ψ has performance p for solving task T_0 in domain \mathcal{D}_0 . Transportability is the performance of system Ψ for solving task T_0 in a new domain \mathcal{D}_i , relative to the original task, without retraining.

Across multiple \mathcal{D}_i , we have relative performance $\tau_p(\mathcal{D}_0, \mathcal{D}_i)$, from which we can establish statistical measures for transportability performance and variation.

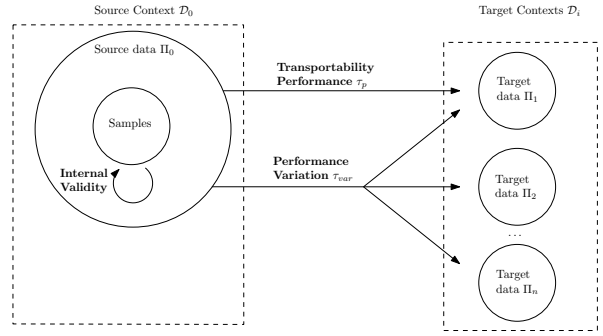


Figure 1: Schematic representation of the definitions

Transfer learning is a specific way of achieving transportability (between populations or domains) or generalisability (including between tasks). Singleton et al. (2014) state that “transport encompasses transfer learning in attempting to use statistical evidence from a source on a target, but differs by incorporating causal assumptions derived from a combination of empirical source data and outside domain knowledge.”. Note that this is different to *generalisation* in the Machine Learning sense, which is akin to internal validity (Marsland, 2011). Figure 1 shows the definitions associated with transportability discussed in this paper.

Table 1 summarises terminology, of how the target differs from source ($\Psi_0, T_0, \mathcal{D}_0(\Pi_0)$).

Term	Ψ	\mathbf{T}	\mathcal{D}	Π
Cross-validation	0	0	0	i
New modeling	i	0	0	0
Transportability	0	0	i	i
Transferability	0	0,i	i	i
Generalisability	0	0,i	0,i	0,i

Table 1: Terminology through variation from a source. Table body is subscripts.

Chance, bias and confounding are the three broad categories of “threat to validity”. Broadly, chance and bias can be assessed by cross-validity, as it applies a model to the same task in the same domain on different data population. Confounding, error in interpretation of what is being measured, is more difficult to assess. Transportability is concerned with the transfer of learned information, with particular advances in the transport of causal knowledge.

Generalisability is the catch-all term for how externally valid a result or model is. Any combination of task, domain and data can be used.

3.2 Transportability performance

We define transportability performance τ_p as the gradient of the change in the performance metric’s score from one domain to another. This measure does not take into account the underlying probability distributions, only the change in resulting performance measure.

$$\tau_p(\mathcal{D}_0, \mathcal{D}_i) = \frac{p(\Psi, T, \mathcal{D}_i)}{p(\Psi, T, \mathcal{D}_0)} \quad (1)$$

The measure uses a ratio in order to allow comparison between different systems. To generalise this measure across different settings, we can take an average to give Equation 2. Note that this is the average percentage change in performance, not an aggregated performance measure.

$$\tau_p(\mathcal{D}_0) = \frac{1}{n} \sum_{i=1}^n \frac{p(\Psi, T, \mathcal{D}_i)}{p(\Psi, T, \mathcal{D}_0)} \quad (2)$$

An analogous definition holds for different tasks over the same domain, $\tau_p(T)$.

3.3 Performance variation

Performance variation reflects how stable performance is across different contexts and can include, for example, to what extent the sampling method from the source data effects the performance metric of the algorithm. Part of this is data representativeness, the extent to which the source data representation also represents the target data.

More formally, performance variation $\tau_s(\Psi, T, \mathcal{D})$ is the change in performance of (Ψ, T, \mathcal{D}) across different contexts. This is useful in order to gain specific insight into external validity and generalisability. Indeed, we can assess the change in performance between source context \mathcal{D}_0 and target context \mathcal{D}_i . The source context has a privileged position, in that it is this space which the “learning” takes place, and the proposed metric for performance variation to multiple different domains is based on τ_p to reflect this. Through repeated measurement in different contexts, we can go further.

Definition 3. *Performance Variation:* For a model trained on domain \mathcal{D}_0 and applied on n new domains \mathcal{D}_i , we define the performance variation as the coefficient of variation of performance across

this set of domains so that:

$$\tau_{var}(\mathcal{D}_0) = \left(1 + \frac{1}{4n}\right) \frac{\sqrt{\frac{\sum_{i=1}^n (\tau_p(\mathcal{D}_0, \mathcal{D}_i) - \tau_p(\mathcal{D}_0))^2}{n-1}}}{\tau_p(\mathcal{D}_0)} \quad (3)$$

The $1 + \frac{1}{4n}$ term corrects for bias. In order to be meaningful, the target contexts must to have a good coverage of different domains. Enumerating these would be a task of ontological proportions, but can be pragmatically approximated by using the available Gold Standard datasets.

We can also assess ability to generalise not just over different domains, but also different tasks, provided they can be meaningfully assessed by the same performance measure. We can consider n different domain-task combinations, and with $\bar{\tau}_p = \sum_{i,j=0}^n \tau_p(\Psi, T_i, \mathcal{D}_j)/n$, this gives a more general form for Equation 3, with n large:

$$\tau_{var} = \frac{\sqrt{\frac{\sum_{i \geq 0, j \geq 0} (\tau_p(\Psi, T_i, \mathcal{D}_j) - \bar{\tau}_p)^2}{n-1}}}{\bar{\tau}_p} \quad (4)$$

In the case where different tasks cannot be assessed by the same measure, we are still able to compare different systems by looking at how the respective measures change.

3.3.1 Performance variation properties

For a purely random system, the transportability should be related to how similar the distributions of “answers” in the test dataset are. A random system should really be transportable by our measures. Similarly, we can consider trivial systems, such as identity and constant functions, which are necessarily entirely transportable. That is, for a system that is an identity function $\Psi = I$, $\tau_p = f(P)$, and $\tau_{var}(I, T, \mathcal{D}_i) = \tau_{var}(I, T, \mathcal{D}_j) = 0, \forall i, j$. Note that we would not expect the same performance of these functions on different tasks.

A stable system will have $\tau_{var}(\Psi, T, \mathcal{D}_0) \approx \tau_{var}(\Psi, T, \mathcal{D}_i) \forall i$, reflecting that it is resilient to the domain on which it is trained.

3.3.2 Factors influencing performance variation

Through repeated measurement, we can quantify how F_1 -score changes with respect to different measures A (e.g. dataset complexity), $\frac{\partial F_1}{\partial A}$, with other properties held constant.

NLP system performance is dependent on A . This list may include gold standard feature distribution (in terms of representativeness of the semantic

or linguistic phenomena), and task difficulty or sensitivity.

Users of NLP systems would benefit from being able to estimate the performance of an existing NLP system on a new domain, without performing the full implementation. Important for the performance of an NLP system, especially for few or zero shot learning, is having a common set of features (or phenomena) across domains. We proceed to propose three measures of increasing complexity, in order to attempt to understand how “similar” two domains are.

Lexical feature difference: A measure grounded on lexical features (i.e. bag of words). The intuition behind this measure is for treating the set of lexical features as a representation. Linguistic space is observed as materialised tokens, which in turn are in some higher-dimensional semantic space, which enable interpretation. The measure considers the overlap of these linguistic spaces, and indeed the extent to which the linguistic space is covered by the data. Due to the simplicity of this measure, correlation between this and actual transportability performance is likely to be weaker than other measures but is simpler to calculate.

$$\text{Lexical Feature Difference} = 1 - \frac{|\mathcal{D}_i \cap \mathcal{D}_0|}{|\mathcal{D}_i|}, i > 0 \quad (5)$$

Where $|\mathcal{D}_i|$ is the number of features in the target domain \mathcal{D}_i , and $|\mathcal{D}_i \cap \mathcal{D}_0|$ is the number of features overlapping. This measure is then the proportion of unseen features in the new dataset. If all features of \mathcal{D}_i are found in \mathcal{D}_0 , then the feature difference is 0. If no features of \mathcal{D}_i are found in \mathcal{D}_0 , then the feature difference is 1. The feature overlap is task specific, and therefore appropriate to consider for transportability, but not generalisability.

In the simplest case, the transported performance of a bag of words model should be precisely the lexical feature difference combined with distributions of the source and target domains. The feature set can range from binary lexical features to latent vector spaces. For different models, which target different aspects of semantic phenomena, different semantic and syntactic features will matter more. For this reason, considering a set of measures for domain complexity is warranted. In the context of this work, two measures are used over more complex feature spaces.

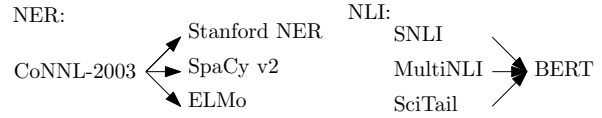


Figure 2: Overview of the experiments undertaken, indicating the models being applied to each dataset

Cosine distance: Specifically, we use Doc2Vec (Le and Mikolov, 2014) to embed the documents from each domain in a 300-dimensional feature-vector space, normalise, and calculate cosine distance to compare source and target domains.

Kullback–Leibler divergence: Considering each domain as a distribution of features, we can use relative entropy to understand the difference between the source and target domains. Similar to cosine distance, we convert the corpus to a vector using Doc2Vec and normalize. We treat these values as discrete probability distributions to calculate the KL divergence.

The usefulness of any of these domain similarity measures depends on the semantic phenomena and supporting corpora underlying the system, for example if the system requires a large training dataset, it may be more appropriate to use a measure which considers the underlying probability distributions in each feature. In this case, we can restrict to the case of the same task in order to keep the essential features reasonably consistent across domains. This makes this a measure of transportability (rather than generalisability).

There are additional dimensions of transportability potentially worthy of further investigation and quantification: (i) domain similarity (e.g. missing features), (ii) data efficiency (redundant/repeated features), (iii) data preparation (initial setup and formatting) and (iv) data manipulation required (data pipeline).

4 Experiments

4.1 Setup

The experiments aim to evaluate the consistency of the proposed transportability measures in the context of two standard tasks: named entity recognition and natural language inference. For reproducibility purposes the code and supporting data are available online¹.

We calculated the F1 score of multiple models on

¹<https://github.com/ai-systems/transportability>

Dataset		Model		
		Stanford	SpaCy	ELMo
CoNLL-2003	Train	98.69	99.32	99.97
	Dev	93.22	81.56	98.17
	Test	88.78	88.11	93.79
Wiki		66.31	52.14	79.4
WNUT	Train	51.63	27.03	36.3
	Dev	53.59	32.23	48.8
	Test	47.11	26.28	58.1

Table 2: NER F1 scores for different models trained on CoNLL dataset transported across different corpora

multiple datasets (Figure 2). Note that in general the applicability of the proposed transportability measures are not limited to the use of F1 score, but this is simpler as the same measure applies for both tasks. All models and datasets are standard. For NER, the datasets were chosen as they have the consistent tags: Location, Person and Organisation. Stanford NER (Finkel et al., 2005) is a CRF classifier, SpaCy v2 is a CNN, ELMo (Peters et al., 2018) is a vector embedding model which outperforms GloVe and word2vec. Each of the three models used are trained on the CoNLL-2003 dataset (Sang and De Meulder, 2003). We evaluated these models on CoNLL-2003, Wikipedia NER (Ghaddar and Langlais, 2017) (Wiki) and WNUT datasets (Baldwin et al., 2015) for NER in twitter microposts.

For NLI, we chose to use standard datasets. SNLI (Bowman et al., 2015) is well established with a limited range of NLI statements, MultiNLI (Williams et al., 2018) is multigenre with a more diverse range of texts, and SciTail (Khot et al., 2018) is based on scientific exam questions. We applied BERT (Devlin et al., 2018), a state of the art embedding model, to these datasets.

4.2 Results

NER: Table 2 shows results for the NER task, trained on CoNLL. Unsurprisingly, all models performed better when the target was in the CoNLL domain. The reduced performance on Wiki was more extreme than expected, particularly for ELMo, which was expected to be resilient to domain change (i.e. transportable). Table 6 and Table 4 illustrate the transportability and domain similarity scores for different NER models respectively.

NLI: Table 3 shows results for the NLI task, using BERT. We find that, despite the vast training data, BERT’s performance is substantially higher when it has been trained on data from that domain. BERT trained on SciTail performs poorly when transported to SNLI or MultiNLI. Table 7 and Table 5 illustrates the transportability and domain similarity scores for different NLI corpora.

4.3 Analysis

Every model had $\tau_p \ll 1$, meaning they performed worse on the new domain. This is as expected, though this would not be true in general.

NER: Examining the F1 scores (88.11 vs. 88.78) of SpaCy and Stanford they appear almost comparable. However, the latter transports much more effectively, with τ_p score difference (0.671 Vs 0.524 when transporting to Wiki) (refer Table 6).

ELMo is one of the state of the art approaches for NER, as evidenced by the high F1 scores for the source corpus. However, Stanford NER transports equally well, and when transported outperforms ELMo for twitter domain. While the absolute F1 score difference between them is 5, the τ_p scores are almost identical, with a difference of 0.003. In terms of transportability, it is notable that an approach that employs CRF tagger with linguistic features outperforms significantly the CNN-based SpaCy approach and stands in comparison to a computationally expensive model like ELMo.

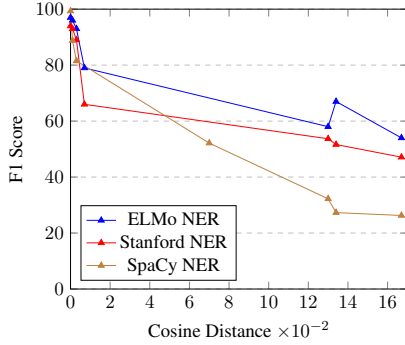
Stanford NER also has the lowest τ_{var} . This indicates this to be the most robust model out of the three. This conclusion was facilitated by the τ_p and τ_{var} measures.

NER for English is assumed to be an accomplished task as supported by the traditional F1 scores. By using τ_p we argue that there is a need for more robust models, with better transportability performance.

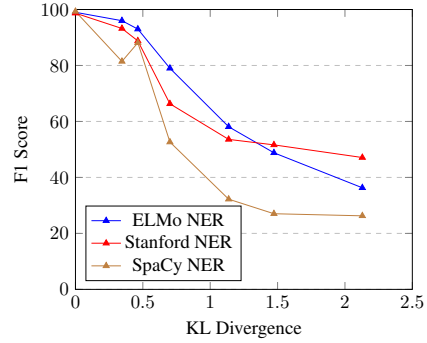
Figure 3a and Figure 3b illustrates the decrease in F1 scores as cosine distance and KL divergence increase. A simple 3 parameter non-linear regression model on KL Divergence and Cosine distance is able to predict the F1 score with an mean error of 3.33 and 2.66 respectively. Considering the lexical difference has similar results (Table 4). This implies that by using these measures we may be able to anticipate the accuracy of a model in a new domain based on easy to compute domain similarity, which is straightforward to compute.

Source Dataset	Target Dataset								
	SNLI Dataset			MultiNLI Dataset		SciTail Dataset			
	Train	Dev	Test	Train	Dev	Train	Dev	Test	
SNLI (Train)	96.81	90.83	90.40	72.51	72.29	54.04	61.34	52.72	
Multi NLI (Train)	77.13	79.05	79.31	97.78	83.50	66.52	67.79	67.26	
SciTail (Train)	42.68	44.36	44.20	47.49	44.49	99.88	94.78	93.08	

Table 3: NLI accuracy scores for BERT model trained on one dataset transported to a different dataset



(a) NER F1 scores Vs Doc2Vec cosine distance from training (CoNLL) corpus



(b) NER F1 scores Vs KL Divergence from training (CoNLL) corpus

Figure 3: NER F1 score plotted against different measures of corpus similarity

Dataset		Lexical	Cosine	KL Divergence
CoNLL	Train	0.000	0.000	0.000
	Dev	0.121	0.001	0.345
	Test	0.197	0.003	0.463
Wiki		0.290	0.007	0.701
WNUT	Train	0.421	0.134	2.129
	Dev	0.511	0.167	1.473
	Test	0.481	0.130	1.137

Table 4: Domain similarity scores between the training corpus (CoNLL-2003) across other NER datasets

NLI: Applying BERT to different domains was not as resilient to domain transport as we expected. The average τ_p is 0.612 over transported domains, despite these being standard corpora from the domains. We found MultiNLI(Train) to be more transportable than the others, since its performance in new domains is not much worse than new data from the same domain. This is as expected, since MultiNLI has been built to have good domain coverage. Specifically, MultiNLI has $\tau_p = 0.744$ and $\tau_{var} = 8.582$, whilst SNLI has $\tau_p = 0.646$ and $\tau_{var} = 15.22$ and SciTail has $\tau_p = 0.446$ and

$\tau_{var} = 3.921$. SciTail transports poorly, and does so reliably! SNLI transports in between, but variably, being quite “hit or miss” with different samples of SciTail. These results suggest a threshold for τ_p of perhaps 0.8 as being “appropriate” for transportability performance. A threshold for τ_{var} is more difficult to establish and would benefit from further investigation. Clearly, these measures depend on the domains chosen.

As with NER, we found lexical difference indicative of transported performance, and that for NLI, accuracy scores decrease with increasing lexical difference, cosine distance and KL divergence (Tables 3 and 5, and Figures 4a and 4b). A simple 3 parameter non-linear regression model on KL Divergence and Cosine distance is able to predict the accuracy score with a mean error of 3.98 and 1.95 respectively.

4.4 Discussion

τ_p and τ_{var} as complementary to traditional measures. We are not breaking new ground in terms of evaluation methodology, but the experiments demonstrate that traditional F1 and accuracy measures do not capture a complete picture. Transportability measure are not only simple enough to calculate and convey but also evaluates a model

Dataset	Measurement	SNLI			MultiNLI		SciTail		
		Train	Dev	Test	Train	Dev	Train	Dev	Test
SNLI (Train)	Lexical	0.000	0.003	0.003	0.086	0.088	0.136	0.115	0.119
	Cosine	0.000	0.002	0.002	0.008	0.007	0.233	0.242	0.242
	KL Divergence	0.000	3.277	4.283	6.489	8.982	16.02	17.50	18.20
MultiNLI (Train)	Lexical	0.008	0.008	0.008	0.000	0.008	0.063	0.063	0.047
	Cosine	0.008	0.018	0.016	0.000	0.002	0.298	0.307	0.306
	KL Divergence	11.07	7.613	6.333	0.000	3.342	33.10	35.27	34.69
SciTail (Train)	Lexical	0.282	0.282	0.282	0.277	0.278	0.000	0.028	0.025
	Cosine	0.233	0.230	0.231	0.262	0.298	0.000	0.001	0.002
	KL Divergence	11.17	7.04	7.492	5.220	6.682	0.000	1.097	1.424

Table 5: Domain similarity scores between the source training corpus and target corpora

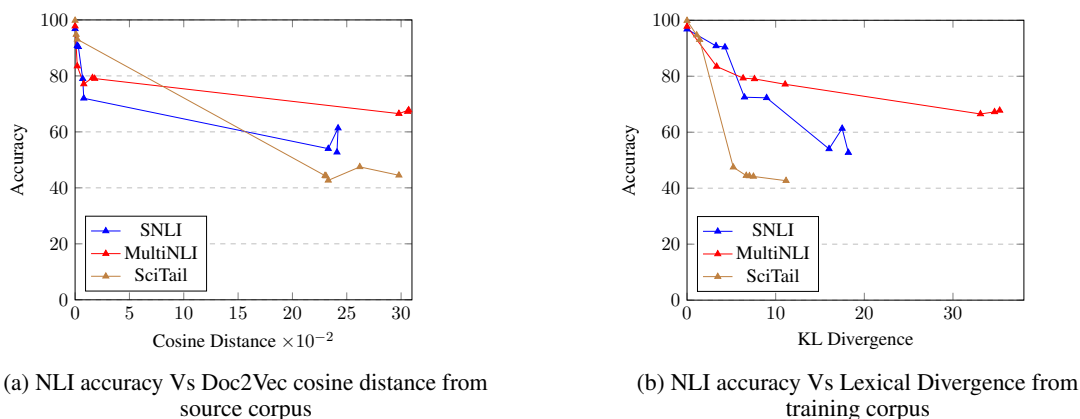


Figure 4: NLI accuracy score plotted against different measures of corpus similarity

	Stanford	SpaCy	ELMo
$\tau_p(wiki)$	0.671	0.524	0.794
$\tau_p(wnut)$	0.514	0.287	0.477
$\tau_p(wnut \& wiki)$	0.553	0.346	0.556
τ_{var}	15.051	35.171	32.666

Table 6: Transportability measures for NER models

with regards to generalisability and robustness.

Low cost ways of anticipating performance for a new task or domain. Most of the state of the art models are computationally expensive. With the transportability and domain similarity measures we are able to predict performance in a new domain with reasonable accuracy. These similarity measures are relatively simpler to run.

5 Conclusion

We have presented a model of transportability for NLP tasks, together with metrics to allow for the

	SNLI	MultiNLI	SciTail
τ_p	0.646	0.744	0.446
τ_{var}	15.22	8.582	3.921

Table 7: Transportability measures for NLI corpora

quantification in the change in performance. We have shown that the proposed transportability measure allows for direct comparison of NLP systems' performance in new contexts. Further, we demonstrated domain similarity as a measure to model corpus and domain complexity, and predict NLP system performance in unseen domains. This paper lays the foundations for further work in more complex transportability measures and estimation of NLP system performance in new contexts.

References

Fernando Antonanzas, Roberto Rodríguez-Ibeas, Carmelo Juárez, Florencia Hutter, Reyes Lorente, and Mariola Pinillos. 2009. Transferability indices

- for health economic evaluations: methods and applications. *Health economics*, 18(6):629–643.
- Timothy Baldwin, Marie-Catherine de Marneffe, Bo Han, Young-Bum Kim, Alan Ritter, and Wei Xu. 2015. Shared tasks of the 2015 workshop on noisy user-generated text: Twitter lexical normalization and named entity recognition. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 126–135.
- Elias Bareinboim and Judea Pearl. 2013. A general algorithm for deciding transportability of experimental results. *Journal of causal Inference*, 1(1):107–134.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, pages 440–447.
- Stephanie Boulenger, John Nixon, Michael Drummond, Philippe Ulmann, Stephen Rice, and Gerard de Pourville. 2005. Can economic evaluations be made more transferable? *The European Journal of Health Economics*, 6(4):334–346.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Donald T Campbell and Julian C Stanley. 2015. *Experimental and quasi-experimental designs for research*. Ravenio Books.
- Robert J Carroll, Will K Thompson, Anne E Eyler, Arthur M Mandelin, Tianxi Cai, Raquel M Zink, Jennifer A Pacheco, Chad S Boomershine, Thomas A Lasko, Hua Xu, et al. 2012. Portability of an algorithm to identify rheumatoid arthritis in electronic health records. *Journal of the American Medical Informatics Association*, 19(e1):e162–e169.
- D Demner-Fushman and Noemie Elhadad. 2016. Aspiring to unintended consequences of natural language processing: a review of recent developments in clinical and consumer-generated text processing. *Yearbook of medical informatics*, 25(01):224–233.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. The hitchhiker’s guide to testing statistical significance in natural language processing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392.
- Óscar Ferrández, Brett R South, Shuying Shen, F Jeff Friedlin, Matthew H Samore, and Stéphane M Meystre. 2012. Generalizability and comparison of automatic clinical text de-identification methods and resources. In *AMIA Annual Symposium Proceedings*, volume 2012, page 199. American Medical Informatics Association.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 363–370. Association for Computational Linguistics.
- Abbas Ghaddar and Phillippe Langlais. 2017. Winer: A wikipedia annotated corpus for named entity recognition. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 413–422.
- Kyle Gorman and Steven Bedrick. 2019. We need to talk about standard splits. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 2786–2791.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. Scitail: A textual entailment dataset from science question answering. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196.
- Thomas M Maddox and Michael A Matheny. 2015. Natural language processing and the promise of big data: Small step forward, but many miles to go. *Circulation. Cardiovascular quality and outcomes*, 8(5):463.
- Stephen Marsland. 2011. *Machine learning: an algorithmic perspective*. Chapman and Hall/CRC.
- Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. 2010. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, 22(2):199–210.
- Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- Judea Pearl. 2019. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3):54–60.
- Judea Pearl and Elias Bareinboim. 2011. Transportability of causal and statistical relations: A formal approach. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- Judea Pearl and Elias Bareinboim. 2014. External validity: From do-calculus to transportability across populations. *Statistical Science*, 29(4):579–595.

- Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. 2019. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*.
- Seyed Mahdi Rezaeinia, Rouhollah Rahmani, Ali Ghodsi, and Hadi Veisi. 2019. Sentiment analysis based on improved pre-trained word embeddings. *Expert Systems with Applications*, 117:139–147.
- Erik F Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*.
- Natalie Schluter and Daniel Varab. 2018. When data permutations are pathological: the case of neural natural language inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4935–4939.
- Kyle W Singleton, Alex AT Bui, and William Hsu. 2014. Transfer and transport: incorporating causal methods for improving predictive models. *Journal of the American Medical Informatics Association*, 21(e2):e374–e375.
- Anders Søgaard, Anders Johannsen, Barbara Plank, Dirk Hovy, and Héctor Martínez Alonso. 2014. What’s in a p-value in nlp? In *Proceedings of the eighteenth conference on computational natural language learning*, pages 1–10.
- Damian G Walker, Yot Teerawattananon, Rob Anderson, and Gerry Richardson. 2010. Generalisability, transferability, complexity and relevance. In *Evidence-Based Decisions and Economics*, pages 56–66. Wiley and Sons.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.