

EVALUATION OF SEMANTIC ANSWER SIMILARITY METRICS

¹Farida Mustafazade and ²Peter F. Ebbinghaus

¹GAM Systematic

²Teufel Audio

ABSTRACT

There are several issues with the existing general machine translation or natural language generation evaluation metrics, and question-answering (QA) systems are indifferent in that context. To build robust QA systems, we need the ability to have equivalently robust evaluation systems to verify whether model predictions to questions are similar to ground-truth annotations. The ability to compare similarity based on semantics as opposed to pure string overlap is important to compare models fairly and to indicate more realistic acceptance criteria in real-life applications. We build upon the first to our knowledge paper that uses transformer-based model metrics to assess semantic answer similarity and achieve higher correlations to human judgement in the case of no lexical overlap. We propose cross-encoder augmented bi-encoder and BERTScore models for semantic answer similarity, trained on a new dataset consisting of name pairs of US-American public figures. As far as we are concerned, we provide the first dataset of co-referent name string pairs along with their similarities, which can be used both for training and as a benchmark.

KEYWORDS

Question-answering, semantic answer similarity, exact match, pre-trained language models, cross-encoder, bi-encoder, semantic textual similarity, automated data labelling

1. INTRODUCTION

Having reliable metrics for evaluation of language models in general, and models solving difficult question answering (QA) problems, is crucial in this rapidly developing field. These metrics are not only useful to identify issues with the current models, but they also influence the development of a new generation of models. In addition, it is preferable to have an automatic, simple metric as opposed to expensive, manual annotation or a highly configurable and parameterisable metric so that the development and the hyperparameter tuning do not add more layers of complexity. SAS, a cross-encoder-based metric for the estimation of semantic answer similarity [1], provides one such metric to compare answers based on semantic similarity.

The central objective of this research project is to analyse pairs of answers similar to the one in [Figure 1](#) and to evaluate evaluation errors across datasets and evaluation metrics.

The main hypotheses that we will aim to test thoroughly through experiments are twofold. Firstly, lexical-based metrics are not well suited for automated QA model evaluation as they lack a notion of context and semantics. Secondly, most metrics, specifically SAS and BERTScore, as described in [1], find some data types more difficult to assess for similarity than others.

After familiarising ourselves with the current state of research in the field in [Section 2](#), we describe the datasets provided in [1] and the new dataset of names that we purposefully tailor to our model in [Section 3](#). This is followed by [Section 4](#), introducing the four new semantic answer similarity approaches described in [1], our fine-tuned model as well as three lexical n-gram-based automated metrics. Then in [Section 5](#), we thoroughly analyse the evaluation datasets described in

Figure 1: Representative example of a question and all semantic answer similarity measurement results.

<p>Question: Who makes more money: NFL or Premier League? Ground-truth answer: National Football League Predicted Answer: the NFL EM: 0.00 F₁: 0.00 Top-1-Accuracy: 0.00 SAS: 0.9008 Human Judgment: 2 (definitely correct prediction) f_{BERT}: 0.4317 f'_{BERT}: 0.4446 Bi-Encoder: 0.5019</p>
--

the previous section and conduct an in-depth qualitative analysis of the errors. Finally, in [Section 6](#), we summarise our contributions.

2. RELATED WORK

We define semantic similarity as different descriptions for something that has the same meaning in a given context, following largely [2]’s definition of semantic and contextual synonyms. [3] noted that open-domain QA is inherently ambiguous because of the uncertainties in the language itself. The human annotators attach a label 2 to all predictions that are ”definitely correct”, 1 - ”possibly correct”, and 0 - ”definitely incorrect”. Automatic evaluation based on exact match (EM) fails to capture semantic similarity for definitely correct answers, where 60% of the predictions are semantically equivalent to the ground-truth answer. Just under a third of the predictions that do not match the ground-truth labels were nonetheless correct. They also mention other reasons for failure to spot equivalence, such as time-dependence of the answers or underlying ambiguity in the questions.

QA evaluation metrics in the context of SQuAD v1.0 [4] dataset are analysed in [5]. They thoroughly discuss the limitations of EM and F1 score from n-gram based metrics, as well as the importance of context including the relevance of questions to the interpretation of answers. A BERT matching metric (Bert Match) is proposed for answer equivalence prediction, which performs better when the questions are included alongside the two answers, but appending contexts didn’t improve results. Additionally, authors demonstrate better suitability of Bert Match in constructing top- k model’s predictions. In contrast, we will cover multilingual datasets, as well as more token-level equivalence measures, but limit our focus on similarity of answer pairs without accompanying questions or contexts.

Two out of four semantic textual similarity (STS) metrics that we analyse and the model that we eventually train depend on bi-encoder and BERTScore [6]. The bi-encoder approach model is based on the Sentence Transformer structure [7], which is a faster adaptation of BERT for the semantic search and clustering type of problems. BERTScore uses BERT to generate contextual embeddings, then match the tokens of the ground-truth answer and prediction, followed by creating a score from the maximum cosine similarity of the matched tokens. This metric is not one-size-fits-all. On top of choosing a suitable contextual embedding and model, there is an optional feature of importance weighting using inverse document frequency (idf). The idea is to limit the influence of common words. One of the findings is that most automated evaluation metrics demonstrate significantly better results on datasets without adversarial examples, even when these are introduced within the training dataset, while the performance of BERTScore suffers only slightly. [6] uses machine

translation (MT) and image captioning tasks in experiments and not QA. [8] apply BERT-based evaluation metrics for the first time in the context of QA. Even though they find that METEOR as an n-gram based evaluation metric proved to perform better than the BERT-based approaches, they encourage more research in the area of semantic text analysis for QA. Moreover, [5] uses only BERTScore base as one of the benchmarks, while we explore the larger model, as well as a finetuned variation of it.

Authors in [1] expand on this idea and further address the issues with existing general MT, natural language generation (NLG), which entails as well generative QA and extractive QA evaluation metrics. These include reliance on string-based methods, such as EM, F1-score, and top-n-accuracy. The problem is even more substantial for multi-way annotations. Here, multiple ground-truth answers exist in the document for the same question, but only one of them is annotated. The major contribution of the authors is the formulation and analysis of four semantic answer similarity approaches that aim to resolve to a large extent the issues mentioned above. They also release two three-way annotated datasets: a subset of the English SQuAD dataset [9], German GermanQuAD dataset [10], and NQ-open [3].

Looking into error categories (see Table 1 and Section 5) revealed problematic data types, where entities, particularly those involving names of any kind, turned out to be the leading category. [11] analyse Natural Questions (NQ) [12], TriviaQA [13] as well as SQuAD and address the issue that current QA benchmarks neglect the possibility of multiple correct answers. They focus on the variations of names, e.g. nicknames, and improve the evaluation of Open-domain QA models based on a higher EM score by augmenting ground-truth answers with aliases from Wikipedia and Freebase. In our work, we focus solely on the evaluations of answer evaluation metrics and generate a standalone names dataset from another dataset, described in greater detail in Section 3.

Our main assumption is that better metrics will have a higher correlation with human judgement, but the choice of a correlation metric is important. Pearson correlation is a commonly used metric in evaluating semantic text similarity (STS) for comparing the system output to human evaluation. [14] show that Pearson power-moment correlation can be misleading when it comes to intrinsic evaluation. They further go on to demonstrate that no single evaluation metric is well suited for all STS tasks, hence evaluation metrics should be chosen based on the specific task. In our case, most of the assumptions, such as normality of data and continuity of the variables behind Pearson correlation do not hold. Kendall's rank correlations are meant to be more robust and slightly more efficient in comparison to Spearman as demonstrated in [15].

Soon after Transformers took over the field, adversarial tests resulted in significantly lower performance figures, which increased the importance of adversarial attacks [16]. General shortcomings of language models and their benchmarks led to new approaches such as Dynabench [17]. Adversarial GLUE (AdvGLUE) [18] focuses on the added difficulty of maintaining the semantic meaning when applying a general attack framework for generating adversarial texts. There are other shortcomings of large language models, including environmental and financial costs [19]. Hence, analysing existing benchmarks is crucial in effectively supporting the pursuit of more robust models. We therefore carefully analyse state of the art benchmarking for semantic answer similarity metrics while keeping in mind the more general underlying shortcomings of large pre-trained language models.

3. DATA

We perform our analysis on three subsets of larger datasets annotated by three human raters and provided by [1]. Unless specified otherwise, these will be referred to by their associated dataset names.

Table 1: Category definitions and examples from annotated NQ-open dataset.

Category	Definition	Question	Gold label	Prediction
Acronym	An abbreviation formed from the initial letters of other words and pronounced as a word	what channel does the haves and have nots come on on directv	OWN	Oprah Winfrey Network
Alias	Indicate an additional name that a person sometimes uses	who is the man in black the dark tower	Randall Flagg	Walter Padick
Co-reference	Requires resolution of a relationship between two distinct words referring to the same entity	who is marconi in we built this city	the father of the radio	Italian inventor Guglielmo Marconi
Different levels of precision	When both answers are correct, but one is more precise	when does the sympathetic nervous system be activated	constantly	fight-or-flight response
Imprecise question	There can be more than one correct answers	b-25 bomber accidentally flew into the empire state building	Old John Feather Merchant	1945
Medical term	Language used to describe components and processes of the human body	what is the scientific name for the shoulder bone	shoulder blade	scapula
Multiple correct answers	There is no single definite answer	city belonging to mid west of united states	Des Moines	kansas city
Spatial	Requires an understanding of the concept of space, location, or proximity	where was the tv series pie in the sky filmed	Marlow in Buckinghamshire	bray studios
Synonyms	Gold label and prediction are synonymous	what is the purpose of a chip in a debit card	control access to a resource	security
Biological term	Of or relating to biology or life and living processes	where is the ground tissue located in plants	in regions of new growth	cortex
Wrong gold label	The ground-truth label is incorrect	how do you call a person who cannot speak	sign language	mute
Wrong label	The human judgement is incorrect	who wrote the words to the original pledge of allegiance	Captain George Thatcher Balch	Francis Julius Bellamy
Incomplete answer	The gold label answer contains only a subset of the full answer	what are your rights in the first amendment	religion	freedom of the press

3.1. Original datasets

SQuAD is an English-language dataset containing multi-way annotated questions with 4.8 answers per question on average. **GermanQuAD** is a three-way annotated German-language question/answer pairs dataset created by the deepset team which also wrote [1]. Based on the German counterpart of the English Wikipedia articles used in SQuAD, GermanQuAD is the SOTA dataset for German question answering models. To address a shortcoming of SQuAD that was mentioned in [20], GermanQuAD was created with the goal of preventing strong lexical overlap between questions and answers. Hence, more complex questions were encouraged, and questions were rephrased with synonyms and altered syntax. SQuAD and GermanQuAD contain a pair of answers and a hand-labelled annotation of 0 if answers are completely dissimilar, 1 if answers have a somewhat similar meaning, and 2 if the two answers express the same meaning. **NQ-open** is a five-way annotated open-domain adaption of [20]’s Natural Questions dataset. NQ-open is based on actual Google search engine queries. In case of NQ-open, the labels follow a different methodology as described in [3]. The assumption is that we only leave questions with a non-vague interpretation (see Table 1). Questions like *Who won the last FIFA World Cup?* received the label 1 because they have different correct answers without a precise answer at a point in time later than when the question was retrieved. There is yet another ambiguity with this question, which is whether it is discussing FIFA Women’s World Cup or FIFA Men’s World Cup. This way, the two answers can be correct without semantic similarity even though only one correct answer is expected.

The annotation of NQ-open indicates truthfulness of the predicted answer, whereas for SQuAD and GermanQuAD the annotation relates to the semantic similarity of both answers which can lead to differences in interpretation as well as evaluation. To keep the methodology consistent and improve NQ-open subset, vague questions with more than one ground-truth labels have been filtered out. We also manually re-label incorrect labels as well as filter out vague questions.

Table 2 describes the size and some lexical features for each of the three datasets. There were 2, 3 and 23 duplicates in each dataset respectively. Dropping these duplicates led to slight changes in the metric scores.

Table 2: Percentage distribution of the labels and statistics on the subsets of datasets used in the analyses. The average answer size column refers to the average of both the first and second answers as well as ground-truth answer and predicted answer (NQ-open only). $F_1 = 0$ indicates no string similarity, $F_1 \neq 0$ indicates some string similarity. Label distribution is given in percentages.

	SQuAD	GermanQuAD	NQ-open
Label 0	56.7	27.3	71.7
Label 1	30.7	51.5	16.6
Label 2	12.7	21.1	11.7
$F_1 = 0$	565	124	3030
$F_1 \neq 0$	374	299	529
Size	939	423	3559
Avg answer size	23	68	13

3.2. Augmented dataset

For NQ-open, the largest of the three datasets, names was the most challenging category to predict similarity. While names includes city and country names as well, we focus on the names of public figures in our work. To resolve this issue, we provide a new dataset that consists of $\sim 40,000$ (39,593) name pairs and employ the Augmented SBERT approach [21]: we use the cross-encoder

model to label a new dataset consisting of name pairs and then train a bi-encoder model on the resulting dataset. We discuss the deployed models in more detail in [Section 4](#).

The underlying dataset is created from an [open dbpedia-data dataset \[22\]](#) which includes the names of more than a million public figures that have a page on Wikipedia and DBpedia, including actors, politicians, scientists, sportsmen, and writers. Out of these we only use those with a U.S. nationality as the questions in NQ-open are on predominantly U.S. related topics. We then shuffle the list of 25,462 names and pair them randomly to get the name pairs that are then labelled by the cross-encoder model.

The dataset includes different ways of writing a person’s name including aliases. For example, *Gary A Labranche* and *Labranche Gary*, or aliases like *Lisa Marie Abato*’s stage name *Holly Ryder* as well as e.g. Chinese ways of writing such as *Rulan Chao Pian* and 卞趙如蘭. We filter out all examples where more than three different ways of writing a person’s name exist because in these cases these names don’t refer to the same person but were mistakenly included in the dataset. For example, names of various members of Tampa Bay Rays minor league who have one page for all members. Since most public figures in the dataset have a maximum of one variation of their name, we only leave out close to 800 other variations this way, and can add 14,131 additional pairs. These are labelled as 1 because they refer to the same person.

4. MODELS / METRICS

The focus of our research lies on different semantic similarity metrics and their underlying models. As a human baseline, [1] reports correlations between the labels by the first and the second annotator for subsets of SQuAD and GermanQuAD and omits these for the NQ-open subset since they are not publicly available. Maximum Kendall’s tau-b rank correlations are 0.64 for SQuAD and 0.57 for GermanQuAD. The baseline semantic similarity models considered are bi-encoder, BERTScore vanilla, and BERTScore trained, whereas the focus will be on cross-encoder (SAS) performance. [Table 3](#) outlines the exact configurations used for each model.

Table 3: Configuration details of each of the models used in evaluations. The architectures for the first two models and our model follow corresponding sequence classification. T-systems-onsite model, as well as our trained model, follow `XLMRobertaModel`, and the other two - `BertForMaskedLM` & `ElectraForPreTraining` architectures respectively. Most of the models use absolute position embedding.

	deepset/ gbert-large-sts	cross-encoder/ stsb-roberta-large	T-Systems-onsite/ cross-en-de-roberta -sentence-transformer	bert-base-uncased	deepset/ gelectra-base	Augmented cross-en-de-roberta -sentence-transformer
hidden_size	1,024	1,024	768	768	768	768
intermediate_size	4,096	4,096	3,072	3,072	3,072	3,072
max_position_embeddings	512	514	514	512	512	514
model_type	bert	roberta	xlm-roberta	bert	electra	xlm-roberta
num_attention_heads	16	16	12	12	12	12
num_hidden_layers	24	24	12	12	12	12
vocab_size	31,102	50,265	250,002	30,522	31,102	250,002
transformers_version	4.9.2	-	-	4.6.0.dev0	-	4.12.2

A cross-encoder architecture [23] concatenates two sentences with a special separator token and passes them to a network to apply multi-head attention over all input tokens in one pass. Pre-computation is not possible with the cross-encoder approach because it takes both input texts into account at the same time to calculate embeddings. A well-known language model that makes use of the cross-encoder architecture is BERT [24]. The resulting improved performance in terms of more accurate similarity scores for text pairs comes with the cost of higher time complexity, i.e. lower speed, of cross-encoders in comparison to bi-encoders. A bi-encoder calculates the

embeddings of the two input texts separately by mapping independently encoded sentences for comparison to a dense vector space which can then be compared using cosine similarity. The separate embeddings result in higher speed but result in reduced scoring quality due to treating the text pairs completely separate [25]. In our work, both cross- and bi-encoder architectures are based on Sentence Transformers [26].

The original bi-encoder applied in [1] uses the multi-lingual [T-Systems-onsite/cross-en-de-roberta-sentence-transformer](#) [27] that is based on [xlm-roberta-base](#) which was further trained on an unreleased multi-lingual paraphrase dataset resulting in the model [paraphrase-xlm-r-multilingual-v1](#). The latter then in turn was fine-tuned on an English-language STS benchmark dataset [28] and a machine-translated German [STS benchmark](#).

[1] used a separate English and German model for the cross-encoder because there is no multi-lingual cross-encoder implementation available yet. Similar to the bi-encoder approach, the English SAS cross-encoder model relies on [cross-encoder/stsb-roberta-large](#) which was trained on the same English STS benchmark. For German, a new cross-encoder model had to be trained, as there were no German cross-encoder models available. It is based on deepset's [gbert-large](#) [29] and trained on the same machine-translated German STS benchmark as the bi-encoder model, resulting in [gbert-large-sts](#).

BERTScore implementation from [6] is used for our evaluation, with minor changes to accommodate for missing key-value pairs for the [27] model type. For BERTScore trained, the last layer representations were used, while for vanilla type BERTScore, only the second layer. **BERTScore vanilla** is based on [bert-base-uncased](#) for English (SQuAD and NQ-open) and deepset's [gelectra-base](#) [29] for German (GermanQuAD), whereas **BERTScore trained** is based on the *multi-lingual* model that is used by the bi-encoder [27]. BERTScore trained outperforms SAS for answer-prediction pairs without lexical overlap, the largest group in NQ-open, but neither of the models perform well on names. We use our new name pairs dataset to train the Sentence Transformer with the same hyperparameters as were used to train [paraphrase-xlm-r-multilingual-v1](#) on the English-language STS benchmark dataset.

We did an automatic hyperparameter search [Table 6](#) for 5 trials with Optuna [30]. Note that cross-validation is an approximation of Bayesian optimization, so it is not necessary to use it with Optuna. The following set of hyperparameters was found to be the best: 'batch': 64, 'epochs': 2, 'warm': 0.45.

We have scanned all metrics from [Table 5](#) for time complexity on NQ-open as it is the largest evaluation dataset. Note that we haven't profiled training times as those are not defined for lexical-based metrics, but only measured CPU time for predicting answer pairs in NQ-open. N-gram based metrics are much faster as they don't have any encoding or decoding steps involved, and they take ~10s to generate similarity scores. The slowest is the cross-encoder as it requires concatenating answers first, followed by encoding, and it takes ~10 minutes. Concatenation grows on a quadratic scale with the input length due to self-attention mechanism. For the same dataset, bi-encoder takes ~2 minutes. BERTScore trained takes ~3 minutes, hence computational costs of BERTScore and bi-encoders are comparable. Additional complexity for all methods mentioned above except for SAS would be marginal when used during training on the validation set. Please note the following system description:

```
System name='Darwin', Release='20.6.0', Machine='x86_64',  
Total Memory=8.00GB, Total cores=4, Frequency=2700.00Mhz
```

5. ANALYSIS

To evaluate the shortcomings of lexical-based metrics in the context of QA, we compare BLEU, ROUGE-L, METEOR, F_1 and the semantic answer similarity metrics, i.e. Bi-Encoder, BERTScore vanilla, BERTScore trained, and Cross-Encoder (SAS) scores on evaluation datasets. To address the second hypothesis, we delve deeply into every single dataset and find differences between different types of answers.

5.1. Quantitative Analysis

As can be observed from Table 4 and Table 5, lexical-based metrics show considerably lower results than any of the semantic similarity approaches. BLEU lags behind all other metrics, followed by METEOR. Similarly, we found that ROUGE-L and F_1 achieve close results. In the absence of lexical overlap, METEOR gives superior results than the other n-gram-based metrics in the case of SQUAD, but ROUGE-L is closer to human judgement for the rest. The highest correlations are achieved in the case of BERTScore based trained models, followed closely by bi- and cross-encoder models. We found some inconsistencies regarding the performance of the cross-encoder based SAS metric. The superior performance of SAS doesn't hold up for the correlation metrics other than Pearson. We observed that SAS score underperformed when $F_1 = 0$ compared to all other semantic answer similarity metrics and overperformed when there is some lexical similarity.

NQ-open is not only by far the largest of the three datasets but also the most skewed one. We observe that the vast majority of answer-prediction pairs have a label 0 (see Table 2). In the majority of cases, the underlying QA model predicted the wrong answer.

All four semantic similarity metrics perform considerably worse on NQ-open than on SQUAD and GermanQuAD. In particular, answer-prediction pairs that have no lexical overlap ($F_1 = 0$) amount to 95 per cent of all pairs with the label 0 indicating incorrect predictions. Additionally, they perform only marginally better than METEOR or ROUGE-L.




Figure 2: Comparison of all (similarity) scores for the pairs in evaluation datasets. METEOR computations for GermanQuAD are omitted since it is not available for German.

Score distribution for SAS and BERTScore trained shows that SAS scores are heavily tilted towards 0 Figure 2.

In Figure 3, we analyse SQUAD subset dataset of answers and we observe a similar phenomenon as in [1] when there is no lexical overlap between the answer pairs: the higher in layers we go in case of BERTScore trained, the higher the correlation values with human labels are. Quite the opposite is observed in the case of BERTScore vanilla, where it is either not as sensitive to embedding representations in case of no lexical overlap or correlations decrease with higher embedding layers.

Table 4: Pearson, Spearman’s, and Kendall’s rank correlations of annotator labels and automated metrics on subsets of GermanQuAD. f_{BERT} is BERTScore vanilla and f'_{BERT} is BERTScore trained.

Metrics	GermanQuAD					
	$F_1 = 0$			$F_1 \neq 0$		
	r	ρ	τ	r	ρ	τ
BLEU	0.000	0.000	0.000	0.153	0.095	0.089
ROUGE-L	0.172	0.106	0.100	0.579	0.554	0.460
F_1 -score	0.000	0.000	0.000	0.560	0.534	0.443
Bi-Encoder	0.392	0.337	0.273	0.596	0.595	0.491
f_{BERT}	0.149	0.008	0.006	0.599	0.554	0.457
f'_{BERT}	0.410	0.349	0.284	0.606	0.592	0.489
SAS	0.488	0.432	0.349	0.713	0.690	0.574

Table 5: Pearson, Spearman’s, and Kendall’s rank correlations of annotator labels and automated metrics on subsets of SQuAD and NQ-open. f_{BERT} is BERTScore vanilla and f'_{BERT} is BERTScore trained, and \tilde{f}_{BERT} is the new BERTScore trained on names.

Metrics	SQuad						NQ-open					
	$F_1 = 0$			$F_1 \neq 0$			$F_1 = 0$			$F_1 \neq 0$		
	r	ρ	τ	r	ρ	τ	r	ρ	τ	r	ρ	τ
BLEU	0.000	0.000	0.000	0.182	0.168	0.159	0.000	0.000	0.000	0.052	0.054	0.051
ROUGE-L	0.100	0.043	0.041	0.556	0.537	0.455	0.220	0.163	0.159	0.450	0.458	0.377
METEOR	0.398	0.207	0.200	0.450	0.464	0.378	0.233	0.152	0.148	0.188	0.179	0.139
F1-score	0.000	0.000	0.000	0.594	0.579	0.497	0.000	0.000	0.000	0.394	0.407	0.337
Bi-Encoder	0.487	0.372	0.303	0.684	0.684	0.566	0.294	0.212	0.170	0.454	0.446	0.351
f_{BERT}	0.249	0.132	0.108	0.612	0.601	0.492	0.156	0.169	0.135	0.165	0.142	0.112
f'_{BERT}	0.516	0.391	0.318	0.698	0.688	0.571	0.319	0.225	0.181	0.452	0.449	0.354
SAS	0.561	0.359	0.291	0.743	0.735	0.613	0.422	0.196	0.158	0.662	0.647	0.512
New Bi-Encoder	0.501	0.391	0.318	0.694	0.690	0.572	0.338	0.252	0.203	0.501	0.501	0.392
\tilde{f}_{BERT}	0.519	0.399	0.324	0.707	0.698	0.581	0.351	0.257	0.208	0.498	0.507	0.398

5.2. Qualitative Analysis

This section is entirely dedicated to highlighting the major categories of problematic samples in each of the datasets.

5.2.1. SQuAD

In **SQuAD** there are only 16 cases where SAS completely diverges from human labels. In all seven cases where SAS score is above 0.5 and label is 0, we notice that the two answers have either a **common substring** or could be used often in the same context. In the other 9 extreme cases when the label is indicative of semantic similarity and SAS is giving scores below 0.25, there are three

Table 6: Experimental setup for hyperparameter tuning of cross-encoder augmented BERTScore.

Batch Size	{16, 32, 64, 128, 256}
Epochs	{1, 2, 3, 4}
warm	uniform(0.0, 0.5)




Figure 3: Pearson, Spearman’s, and Kendall’s rank correlations for different embedding extractions for when there is no lexical overlap ($F_1 = 0$), when there is some overlap ($F_1 \neq 0$) and aggregated for the SQuAD subset. f_{BERT} is BERTScore vanilla and f'_{BERT} is BERTScore trained.

spatial translations. There is an encoding-related example with 12 and 10 special characters each which seems to be a mislabelled example.

5.2.2. GermanQuAD

Overall, error analysis for GermanQuAD is limited to a few cases because it is the smallest dataset of the three and all language model based metrics perform comparably well - SAS in particular. SAS fails to identify semantic similarity in cases where the answers are **synonyms or translations** which also include technical terms that rely on Latin (e.g. *vis viva* and *living forces* (translated) (SAS score: 0.5), *Anorthotiko Komma Ergazomenou Laou* and *Progressive Party of the Working People* (translation) (0.04), *Nährgebiet* and *Akkumulationsgebiet* (0.45), *Zehrgebiet* and *Ablationsgebiet* (0.43)). This is likely the case because SAS does not use a multilingual model. Since multilingual models have not been implemented for cross-encoders yet, this remains an area for future research. Text-based **calculations and numbers** are also problematic: (translated) *46th day before Easter Sunday* and *Wednesday after the 7th Sunday before Easter* (0.41).

SAS also fails to recognise **aliases or descriptions of relations** that point to the same person or object: *Thayendanega* and *Joseph Brant* (0.028) are the same people. BERTScore vanilla and BERTScore trained both find some similarity (0.36, 0.22). *Goring House* and *Buckingham Haus* (0.29) refer to the same object but one is the official name, the other one a description of the same, again BERTScore vanilla and BERTScore trained identify more similarity (0.44, 0.37).

5.2.3. NQ-open

We also observe that similarity scores for answer-prediction pairs which include numbers, e.g. an amount, a date or a year, SAS, as well as BERTScore trained, diverge from labels. The only semantically similar entities to answers expected to contain a numeric value should be the exact value, not a unit more or less. Also, the position within the pairs seems to matter for digits and their string representation. For SAS that the pair of *11* and *eleven* has a score of 0.09 whereas the pair of *eleven* and *11* has a score of 0.89.

Figure 4 depicts the major error categories for when SAS scores range below 0.25 while human annotations indicate a label of 2. We observe that entities related to names, which includes spatial names as well as co-references and synonyms, form the largest group of scoring errors. After correcting for encoding errors and fixing the labels manually in the NQ-open subset, totalling 70 samples, the correlations have already improved by about a per cent for SAS. Correcting wrong labels in extreme cases where SAS score is below 0.25 and the label is 2 or when SAS is above




Figure 4: Subset of NQ-open test set, where SAS score < 0.01 and human label is 2, manually annotated for an explanation of discrepancies. Original questions and Google search has been used to assess the correctness of the gold labels.

0.5 and label is 0 improves results almost across the board for all models, but more so for SAS. After removal of duplicates, sample with imprecise questions, wrong gold label or multiple correct answers, we are left with 3559 ground-truth answer/prediction pairs compared to 3658 we started with.

An example for the better performance on names when applying our new bi-encoder and SBERT trained models can be seen in [Figure 5](#), where both models perform well in comparison to SAS and human judgement.

6. CONCLUSION

Existing evaluation metrics for QA models have various limitations. N-gram based metrics suffer from asymmetry, strictness, failure to capture multi-hop dependencies and penalise semantically-critical ordering, failure to account for relevant context or question, to name a few. We have found patterns in the mistakes that SAS was making. These include **spatial awareness**, **names**, **numbers**, **dates**, **context awareness**, **translations**, **acronyms**, **scientific terminology**, **historical events**, **conversions**, **encodings**.

The comparison to annotator labels is performed on answer pairs taken from subsets of SQuAD and GermanQuAD datasets, and for NQ-open we have a prediction and ground-truth answer pair. For cases with lexical overlap, ROUGE-L achieves comparative results to pre-trained semantic similarity evaluation models at a fraction of computation costs that the other models require. This holds for all GermanQuAD, SQuAD and NQ-open alike. We conclude that to further improve the semantic answer similarity evaluation in German, future work should focus on providing a larger dataset of answer pairs, since we could find only a few examples where SAS and the other metrics based on pre-trained language models didn't match with human annotator labels. Dataset size was one of the reasons why we focused more heavily on NQ-open dataset. In addition, focusing on the other two would mean less strong evidence on how the metric will perform when applied to model predictions behind a real-world application. Furthermore, all semantic similarity metrics failed to have a high correlation to human labels when there was no token-level overlap, which is arguably the most important use-case for a semantic answer similarity metric as opposed to, say, ROUGE-L. NQ-open happened to have the largest number of samples that satisfied this requirement. Removing duplicates and re-labelling led to significant improvements across the board. We have generated a

Figure 5: Representative example of a question and all semantic answer similarity measurement results.

Question: Who killed Natalie and Ann in Sharp Objects?
Ground-truth answer: Amma
Predicted Answer: Luke
EM: 0.00
F₁: 0.00
Top-1-Accuracy: 0.00
SAS: 0.0096
Human Judgment: 0
f_{BERT}: 0.226
f'_{BERT}: 0.145
Bi-Encoder: 0.208
\tilde{f}_{BERT}: 0.00
Bi-Encoder (new model): -0.034

names dataset, which was then used to fine-tune the bi-encoder and BERTScore model. The latter achieves and beats SOTA rank correlation figures when there is no lexical overlap for datasets with English as the core language. Bi-encoders outperformed cross-encoders on answer-prediction pairs without lexical overlap both in terms of correlation to human judgement and speed, which makes them more applicable in real-world scenarios. This, plus support for multilingual setups, could be essential for companies as well because models most probably won't understand the relationships between different employees and stakeholders mentioned in internal documents. A reason to have a preference towards BERTScore would be the ability to use BERTScore as a training objective to generate soft predictions, allowing the network to remain differentiable end-to-end.

An element of future research would be further improving the performance on names of public figures as well as spatial names like cities and countries. Knowledge-bases, such as Freebase or Wikipedia, as explored in [11], could be used to find an equivalent answer to named geographical entities. Numbers and dates which is the problematic data type in multi-lingual, as well as monolingual contexts, would be another dimension.

7. ACKNOWLEDGEMENTS

We would like to thank Ardhendu Singh, Julian Risch, Malte Pietsch and XCS224U course facilitators, Ankit Chadha in particular, as well as Christopher Potts for their constant support.

8. REFERENCES


- [1] J. Risch, T. Möller, J. Gutsch, and M. Pietsch, "Semantic answer similarity for evaluating question answering models," *arXiv preprint arXiv:2108.06130*, 2021.
- [2] X.-M. Zeng, "Semantic relationships between contextual synonyms," *US-China education review*, vol. 4, pp. 33–37, 2007.
- [3] S. Min, J. Boyd-Graber, C. Alberti, D. Chen, E. Choi, M. Collins, K. Guu, H. Hajishirzi, K. Lee, J. Palomaki, C. Raffel, A. Roberts, T. Kwiatkowski, P. Lewis, Y. Wu, H. Küttler, L. Liu, P. Minervini, P. Stenetorp, S. Riedel, S. Yang, M. Seo, G. Izacard, F. Petroni, L. Hosseini, N. D. Cao, E. Grave, I. Yamada, S. Shimaoka, M. Suzuki, S. Miyawaki, S. Sato, R. Takahashi, J. Suzuki, M. Fajcik, M. Docekal, K. Ondrej, P. Smrz, H. Cheng, Y. Shen, X. Liu, P. He, W. Chen, J. Gao, B. Oguz, X. Chen, V. Karpukhin, S. Peshterliev, D. Okhonko,

- M. Schlichtkrull, S. Gupta, Y. Mehdad, and W.-t. Yih, “Neurips 2020 efficientqa competition: Systems, analyses and lessons learned,” in *Proceedings of the NeurIPS 2020 Competition and Demonstration Track* (H. J. Escalante and K. Hofmann, eds.), vol. 133 of *Proceedings of Machine Learning Research*, pp. 86–111, PMLR, 06–12 Dec 2021.
- [4] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang, “Squad: 100,000+ questions for machine comprehension of text,” *arXiv:1606.05250*, 2016.
- [5] J. Bulian, C. Buck, W. Gajewski, B. Boerschinger, and T. Schuster, “Tomayto, tomahto. beyond token-level answer equivalence for question answering evaluation,” *arXiv preprint arXiv:2202.07654*, 2022.
- [6] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, “Bertscore: Evaluating text generation with bert,” *arXiv preprint arXiv:1904.09675*, 2019.
- [7] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” 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)*, (Hong Kong, China), pp. 3982–3992, Association for Computational Linguistics, Nov. 2019.
- [8] A. Chen, G. Stanovsky, S. Singh, and M. Gardner, “Evaluating question answering evaluation,” in *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, (Hong Kong, China), pp. 119–124, Association for Computational Linguistics, Nov. 2019.
- [9] P. Rajpurkar, R. Jia, and P. Liang, “Know what you don’t know: Unanswerable questions for squad,” *arXiv preprint arXiv:1806.03822*, 2018.
- [10] T. Möller, J. Risch, and M. Pietsch, “Germanquad and germandpr: Improving non-english question answering and passage retrieval,” *arXiv:2104.12741*, 2021.
- [11] C. Si, C. Zhao, and J. Boyd-Graber, “What’s in a name? answer equivalence for open-domain question answering,” *arXiv preprint arXiv:2109.05289*, 2021.
- [12] T. Kwiatkowski, J. Palomaki, O. Redfield, M. Collins, A. Parikh, C. Alberti, D. Epstein, I. Polosukhin, M. Kelcey, J. Devlin, K. Lee, K. N. Toutanova, L. Jones, M.-W. Chang, A. Dai, J. Uszkoreit, Q. Le, and S. Petrov, “Natural questions: a benchmark for question answering research,” *Transactions of the Association of Computational Linguistics*, 2019.
- [13] M. Joshi, E. Choi, D. S. Weld, and L. Zettlemoyer, “Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension,” *CoRR*, vol. abs/1705.03551, 2017.
- [14] N. Reimers, P. Beyer, and I. Gurevych, “Task-oriented intrinsic evaluation of semantic textual similarity,” in *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, (Osaka, Japan), pp. 87–96, The COLING 2016 Organizing Committee, Dec. 2016.
- [15] C. Croux and C. Dehon, “Influence functions of the spearman and kendall correlation measures,” *Stat Methods Appl (2010)* 19:497–515, 2010.
- [16] T. Niven and H.-Y. Kao, “Probing neural network comprehension of natural language arguments,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, (Florence, Italy), pp. 4658–4664, Association for Computational Linguistics, July 2019.
- [17] D. Kiela, M. Bartolo, Y. Nie, D. Kaushik, A. Geiger, Z. Wu, B. Vidgen, G. Prasad, A. Singh, P. Ringshia, *et al.*, “Dynabench: Rethinking benchmarking in nlp,” *arXiv preprint arXiv:2104.14337*, 2021.
- [18] B. Wang, C. Xu, S. Wang, Z. Gan, Y. Cheng, J. Gao, A. H. Awadallah, and B. Li, “Adversarial

- GLUE: A multi-task benchmark for robustness evaluation of language models,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [19] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, “On the dangers of stochastic parrots: Can language models be too big?,” in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, (New York, NY, USA), p. 610–623, Association for Computing Machinery, 2021.
- [20] T. Kwiatkowski, J. Palomaki, O. Redfield, M. Collins, A. Parikh, C. Alberti, D. Epstein, I. Polosukhin, J. Devlin, K. Lee, K. Toutanova, L. Jones, M. Kelcey, M.-W. Chang, A. M. Dai, J. Uszkoreit, Q. Le, and S. Petrov, “Natural questions: A benchmark for question answering research,” *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 452–466, Mar. 2019.
- [21] N. Thakur, N. Reimers, J. Daxenberger, and I. Gurevych, “Augmented sbert: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks,” *arXiv:2010.08240v2 [cs.CL]*, 2021.
- [22] C. Wagner, “Politicians on wikipedia and dbpedia,” 2017.
- [23] S. Humeau, K. Shuster, M.-A. Lachaux, and J. Weston, “Poly-encoders: Transformer architectures and pre-training strategies for fast and accurate multi-sentence scoring,” 2020.
- [24] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” *CoRR*, vol. abs/1810.04805, 2018.
- [25] J. Chen, L. Yang, K. Raman, M. Bendersky, J.-J. Yeh, Y. Zhou, M. Najork, D. Cai, and E. Emadzadeh, “DiPair: Fast and accurate distillation for trillion-scale text matching and pair modeling,” in *Findings of the Association for Computational Linguistics: EMNLP 2020*, (Online), pp. 2925–2937, Association for Computational Linguistics, Nov. 2020.
- [26] N. Reimers and I. Gurevych, “Sentence-bert: Sentence embeddings using siamese bert-networks,” *CoRR*, vol. abs/1908.10084, 2019.
- [27] P. May, “T-systems-onsite/cross-en-de-roberta-sentence-transformer,” *Hugging Face*, 2020.
- [28] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, and L. Specia, “SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, (Vancouver, Canada), pp. 1–14, Association for Computational Linguistics, Aug. 2017.
- [29] B. Chan, S. Schweter, and T. Möller, “German’s next language model,” in *Proceedings of the 28th International Conference on Computational Linguistics*, (Barcelona, Spain (Online)), pp. 6788–6796, International Committee on Computational Linguistics, Dec. 2020.
- [30] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, “Optuna: A next-generation hyperparameter optimization framework,” in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2623–2631, 2019.

9. AUTHORS

Peter F. Ebbinghaus is currently SEO Team Lead at Teufel. With a background in econometrics and copywriting, his work focuses on the quantitative analysis of e-commerce content. Based in Germany and Mexico, his research includes applied multi-lingual NLP in particular.



Farida Mustafazade is a London-based quantitative researcher in the GAM Systematic Cambridge investment team where she focuses on macro and sustainable macro trading strategies. Her area of expertise expands to applying machine learning, including NLP techniques, to financial as well as alternative data.

