

# EXPLAINABOARD: An Explainable Leaderboard for NLP

Pengfei Liu<sup>1†</sup>, Jinlan Fu<sup>2</sup>, Yang Xiao<sup>2</sup>, Weizhe Yuan<sup>1</sup>, Shuaichen Chang<sup>3</sup>,  
Junqi Dai<sup>2</sup>, Yixin Liu<sup>1</sup>, Zihuiwen Ye<sup>1</sup>, Graham Neubig<sup>1‡</sup>

<sup>1</sup>Carnegie Mellon University, <sup>2</sup>Fudan University, <sup>3</sup>The Ohio State University,  
<sup>†</sup>pliu3@cs.cmu.edu, <sup>‡</sup>gneubig@cs.cmu.edu

## Abstract

With the rapid development of NLP research, leaderboards have emerged as one tool to track the performance of various systems on various NLP tasks. They are effective in this goal to some extent, but generally present a rather simplistic one-dimensional view of the submitted systems, communicated only through holistic accuracy numbers. In this paper, we present a new conceptualization and implementation of NLP evaluation: the EXPLAINABOARD, which in addition to inheriting the functionality of the standard leaderboard, also allows researchers to (i) diagnose strengths and weaknesses of a single system (e.g. what is the best-performing system bad at?) (ii) interpret relationships between multiple systems. (e.g. where does system A outperform system B? What if we combine systems A, B and C?) and (iii) examine prediction results closely (e.g. what are common errors made by multiple systems or in what contexts do particular errors occur?). So far, EXPLAINABOARD covers more than 400 systems, 50 datasets, 40 languages, and 12 tasks.<sup>1</sup> We not only released an online **platform** at the website<sup>2</sup> but also make our evaluation tool an **API** with *MIT Licence* at Github<sup>3</sup> and PyPi<sup>4</sup> that allows users to conveniently assess their models offline. We additionally release all **output files** from systems that we have run or collected to motivate “output-driven” research in the future.

## 1 Introduction

Natural language processing (NLP) research has been and is making astounding strides forward.

<sup>1</sup>EXPLAINABOARD keeps updated and is recently upgraded by supporting (1) **multilingual multi-task benchmark** and (2) more complicated task: **machine translation**, which reviewers also suggested.

<sup>2</sup><http://explainaboard.nlpedia.ai/>

<sup>3</sup><https://github.com/neulab/explainaboard>

<sup>4</sup><https://pypi.org/project/interpret-eval/>

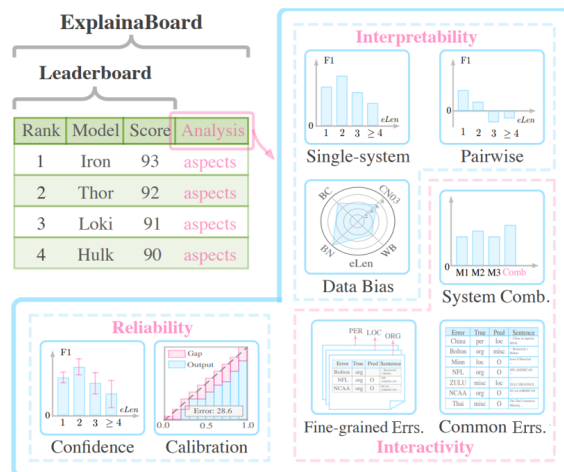


Figure 1: Illustration of the EXPLAINABOARD concept. Compared to vanilla leaderboards, EXPLAINABOARD allows users to perform *interpretable* (single-system, pairwise analysis, data bias), *interactive* (system combination, fine-grained/common error analysis), and *reliable* analysis (confidence interval, calibration) on systems in which they are interested. “Comb.” denotes “combination” and “Errs” represents “errors”. “PER, LOC, ORG” refer to different labels.

This is true both for classical tasks such as machine translation (Sutskever et al., 2014; Wu et al., 2016), as well as for new tasks (Lu et al., 2016; Rajpurkar et al., 2016), domains (Beltagy et al., 2019), and languages (Conneau and Lample, 2019). One way this progress is quantified is through *leaderboards*, which report and update performance numbers of state-of-the-art systems on one or more tasks. Some prototypical leaderboards include GLUE and SuperGLUE for natural language understanding (Wang et al., 2018, 2019), XTREME and XGLUE (Hu et al., 2020; Liang et al., 2020) for multilingual understanding, the WMT shared tasks (Barrault et al., 2020) for machine translation, and GEM and GENIE for natural language generation (Gehrmann et al., 2021; Khashabi et al., 2021), among many others.

These leaderboards serve an important purpose:

they provide a standardized evaluation setup, often on multiple tasks, that eases reproducible model comparison across organizations. However, at the same time, due to the prestige imbued by a top, or high, place on a leaderboard they also can result in a singular focus on raising evaluation numbers at the cost of deeper scientific understanding of model properties (Ethayarajh and Jurafsky, 2020). In particular, we argue that, among others, the following are three major limitations of the existing leaderboard paradigm:

- **Interpretability:** Most existing leaderboards commonly use a single number to summarize system performance holistically. This is conducive to system ranking but at the same time, the results are opaque, making the strengths and weaknesses of systems less interpretable.
- **Interactivity:** Existing leaderboards are static and non-interactive, which limits the ability of users to dig deeper into the results. Thus, (1) they usually do not flexibly support more complex evaluation settings (e.g. multi-dataset, multi-metric, multi-language) (2) users may miss opportunities to understand the relationships between different systems. For example, where does model A outperform model B? Would the performance be further improved if we combine the Top-3 models?
- **Reliability:** Given the increasing role that leaderboards have taken in guiding NLP research, it is important that information expressed in them is reliable, especially on datasets with small sample sizes, but most current leaderboards do not give an idea of the reliability of system rankings.

In this paper, we describe EXPLAINABOARD (see Fig. 1), a software package and hosted leaderboard that satisfies all of the above desiderata. It also serves as a prototype implementation of some desirable features that may be included in future leaderboards, even independent of the provided software itself. We have deployed EXPLAINABOARD for 9 different tasks and 41 different datasets, and demonstrate how it can be easily adapted to new tasks of interest.

We expect that EXPLAINABOARD will benefit different steps of the research process:

(i) **System Development:** EXPLAINABOARD provides more detailed information regarding the submitted systems (e.g. fine-grained results, confidence intervals), allowing system developers to

diagnose successes and failures of their own systems, or compare their systems with baselines and understand where improvements of their proposed methods come from. This better understanding can lead to more efficient and effective system improvements. Additionally, EXPLAINABOARD can help system developers uncover their systems’ advantages over others, even when these systems have not achieved state-of-the-art performance holistically. (ii) **Leaderboard Organization:** The EXPLAINABOARD software both provides a ready-made platform for easy leaderboard development over different NLP tasks, and helps upgrade traditional leaderboards to allow for more fine-grained analysis. For example, we have already established an ExplainaBoard<sup>5</sup> for the existing XTREME benchmark.<sup>6</sup> (iii) **Broad Analysis and Understanding:** Because EXPLAINABOARD encourages system developers to provide their system outputs in an easy-to-analyze format, these will also help researchers, particularly those just starting in a particular NLP sub-field, get a broad sense of what current state-of-the-art models can and cannot do. This not only helps them quickly track the progress of different areas, but also can allow them to understand the relative advantages of diverse systems, suggesting insightful ideas for what’s left and what’s next.<sup>7</sup>

## 2 ExplainaBoard

As stated above, EXPLAINABOARD extends existing leaderboards, improving their interpretability, interactivity, and reliability. It does so by providing a number of functionalities that are applicable to a wide variety of NLP tasks (as illustrated in Tab. 1). Many of these functionalities are grounded in existing research on evaluation and fine-grained diagnostics.

### 2.1 Interpretability

*Interpretable evaluation* (Popović and Ney, 2011; Strymne, 2011; Neubig et al., 2019; Fu et al., 2020a), is a research area that considers methods that break down the holistic performance of each system into different interpretable groups. For example, in a

<sup>5</sup><http://explainaboard.nlpedia.ai/leaderboard/xtreme/>

<sup>6</sup><https://sites.research.google/xtreme/>

<sup>7</sup>Since the first release of EXPLAINABOARD, we have received invitations from multiple companies, startups, and researchers to collaborate, and we are working together to make it better for the community.

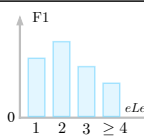
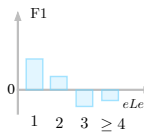
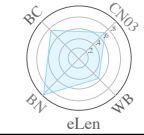
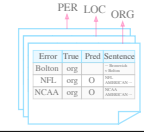
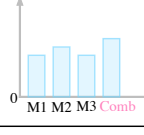
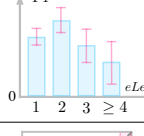
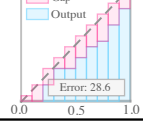
Aspect	Functionality	Input	Output
Interpretability	Single-system Analysis	One model	 <p><b>Performance Histogram:</b> the input model is good at dealing with short entities, while achieving lower performance on long entities.</p>
	Pairwise Analysis	Two models (M1,M2)	 <p><b>Performance Gap Histogram (M1–M2):</b> M1 is better at dealing with short entities, while M2 is better at dealing with long entities.</p>
Interactivity	Data Bias Analysis	Multi-dataset	 <p><b>Data Bias Chart:</b> For the <code>entity length</code> attribute, the average entity length (We average the length of all test entities on a given data set.) of these datasets order by descending is <math>BN &gt; BC &gt; CN03 &gt; WB</math>.</p>
	Fine-grained Error Analysis	Single- or Pairwise-system diagnostic results	 <p><b>Error Table:</b> Error analysis allows the user to print out the entities that are incorrectly predicted by the given model, as well as the true label of the entity, the mispredicted label, and the sentence where the entity is located.</p>
Reliability	System Combination	Multi-models (M1,M2,M3)	 <p><b>Ensemble Chart:</b> The combined result of model M1, M2, and M3 is shown by the histogram with x-label value <code>comb</code>. The combined result is better than the single models.</p>
	Confidence	One model	 <p><b>Error Bars:</b> the error bars represent 95% confidence intervals of the performance on the specific bucket.</p>
	Calibration	One model	 <p><b>Reliability Diagram:</b> Confidence histograms (red) and reliability diagrams (blue). that indicate the accuracy of model probability estimates</p>

Table 1: A graphical breakdown of the functionality of EXPLAINBOARD, with examples from an NER task.

Named Entity Recognition (NER) task, we may examine the accuracy along different dimensions of a concerned entity (such as “entity frequency,” telling us how well the model does on entities that appear in the training data a certain number of times) or sentences (such as “sentence length,” telling us how well the model does on entities that appear in longer or shorter sentences) (Fu et al., 2020a). This makes it possible to understand where models do well and poorly, leading to further insights beyond those that can be gleaned by holistic evaluation numbers. Applying this to a new task involves the following steps: (i) *Attribute definition*: define attributes by which we can partition the test set into different groups. (ii) *Bucketing*: partition into different buckets based on defined attributes and calculate performance w.r.t each bucket.

Generally, previous work on interpretable evaluation has been performed over single tasks, while EXPLAINBOARD allows for comprehensive eval-

uation of different types of tasks in a single software package. We concretely show several ways interpretable evaluation can be defined within EXPLAINBOARD below:

**F1<sup>8</sup>: Single-system Analysis: What is a system good or bad at?** For an individual system as input, generate a *performance histogram* that highlights the buckets where it performs well or poorly. For example, in Tab. 1 we demonstrate an example from NER where the input system does worse in dealing with longer entities ( $eLen \geq 4$ ).

**F2: Pairwise Analysis: Where is one system better (worse) than another?** Given a pair of systems, interpret where the performance gap occurs. Researchers could flexibly choose two systems they are interested in (e.g. selecting two rows from the leaderboard), and EXPLAINBOARD will output a *performance gap histogram* to describe how the

<sup>8</sup>“F” represents “Functionality”.

**performance differences** change over different buckets of different attributes. Tab. 1 demonstrates how we can see one system is better than the other at longer or shorter entities.

**F3: Data Bias Analysis: What are the characteristics of different evaluated datasets?** The defined attributes do not only help us interpret system performance, but also make it possible for users to take a closer look at characteristics of diverse datasets. For example, from Fig. 1 shows an example of analyzing differences in average entity length across several datasets.

## 2.2 Interactivity

EXPLAINBOARD also allows users to dig deeper, interacting with the results in more complex ways.

**F4: Fine-grained Error Analysis: What are common mistakes that most systems make and where do they occur?** EXPLAINBOARD provides flexible fine-grained error analyses based on the above-described performance evaluation:

1. Users can choose multiple systems and see their *common error cases*, which can be useful to identify *challenging samples* or even *annotation errors*.
2. In single-system analysis, users can choose particular buckets in the performance histogram<sup>9</sup> and see corresponding error samples in that bucket (e.g. which long entities does the current system mispredict?).
3. In pairwise analysis, users can select a bucket, and the unique errors (e.g. system A succeeds while B fails and vice versa) of two models will be displayed.

**F5: System Combination: Is there potential complementarity between different systems?**

System combination (Ting and Witten, 1997; González-Rubio et al., 2011; Duh et al., 2011) is a technique to improve performance by combining the output from multiple existing systems. In EXPLAINBOARD, users can choose multiple systems and obtain combined results calculated by voting over multiple base systems.<sup>10</sup> In practice, for NER task, we use the recently proposed SPANNER (Fu

<sup>9</sup>Each bin of the performance histogram is clickable, returning an error case table.

<sup>10</sup>With the system combination button of Explainboard, we observed the-state-of-the-art performance of some tasks (e.g., NER, Chunking) can be further improved.

et al., 2021) as a combiner, and for text summarization we employed REFACTOR, a state-of-the-art ensemble approach (Liu et al., 2021). Regarding the other tasks, we adopt the majority voting method for system combination.

## 2.3 Reliability

The experimental conclusions obtained from the evaluation metrics are not necessarily statistically reliable, especially when the experimental results can be affected by many factors. EXPLAINBOARD also makes a step towards more reliable interpretable evaluation.

**F6: Confidence Analysis: To what extent can we trust the results of our system?**

EXPLAINBOARD can perform confidence analysis over both holistic and fine-grained performance metrics. As shown in Tab. 1, for each bucket, there is an error bar whose width reflects how reliable the performance value is. We claim this is an important feature for fine-grained analysis since the numbers of test samples in each bucket are imbalanced, and with the confidence interval, one could know how much uncertainty there is. In practice, we use bootstrapping method (Efron, 1992; Ye et al., 2021) to calculate the confidence interval.

**F7: Calibration Analysis: How well is the confidence of prediction calibrated with its correctness?**

One commonly-cited issue with modern neural predictors is that their probability estimates are not accurate (i.e. they are poorly *calibrated*), often being over-confident in the correctness of their predictions (Guo et al., 2017). We also incorporate this feature into EXPLAINBOARD, allowing users to evaluate how well-calibrated their systems of interest are.

## 3 Tasks, Datasets and Systems

We have already added to EXPLAINBOARD 12 NLP tasks, 50 datasets, and 400 models,<sup>11</sup> which cover many or most of top-scoring systems on these tasks. We briefly describe them below, and show high-level statistics in Tab. 2.

**Text Classification** Prediction of one or multiple pre-defined label(s) for a given input text. The current interface includes datasets for sentiment

<sup>11</sup>265 of these models are implemented by us, as unfortunately it is currently not standard in NLP to release the system outputs that EXPLAINBOARD needs.

classification (Pang et al., 2002), topic identification (Wang and Manning, 2012), and intention detection (Chen et al., 2013).

**Text-Span Classification** Prediction of a pre-defined class from the input of a text and a span, such as aspect-based sentiment classification task (Pappas and Popescu-Belis, 2014). We collect top-perform system outputs from (Dai et al., 2021).

**Text Pair Classification** Prediction of a class given two texts, such as the natural language inference task (Bowman et al., 2015).

**Sequence Labeling** Prediction of a label for each token in a sequence. The EXPLAINBOARD currently includes four concrete tasks: named entity recognition (Tjong Kim Sang and De Meulder, 2003), part-of-speech tagging (Toutanova et al., 2003), text chunking (Ando and Zhang, 2005), and Chinese word segmentation (Chen et al., 2015).

**Structure Prediction** Prediction of a syntactic or semantic structure from text, where EXPLAINBOARD currently covers semantic parsing tasks (Berant et al., 2013; Yu et al., 2018).

**Text Generation** EXPLAINBOARD also considers text generation tasks, and currently mainly focuses on conditional text generation, for example, text summarization (Rush et al., 2015; Liu and Lapata, 2019) and machine translation. System outputs on text summarization are expanded based on the previous work’s collection (Bhandari et al., 2020) as well as recently state-of-the-art systems (Liu and Liu, 2021) while outputs from machine translation are collected from the WMT20.<sup>12</sup>

## 4 Case Study

Here, we briefly showcase the actual EXPLAINBOARD interface through a case study on analyzing state-of-the-art NER systems.

### 4.1 Experimental Setup

**Attribute Definition** We define attributes following Fu et al. (2020a) and three of them are used below: entity length, sentence length and label of entity.

**Collection of Systems Outputs** Currently, we collect system outputs by either implementing them by ourselves or collecting from other researchers

<sup>12</sup><http://www.statmt.org/wmt20/metrics-task.html>

	Task	Data	Model	Attr.
Text Classification	Sentiment	8	40	2
	Topic	4	18	2
	Intention	1	3	2
Text-Span Classification	Aspect Sentiment	4	20	4
Text Pair Classification	NLI	2	6	7
Sequence Labeling	NER	3	74	9
	POS	3	14	4
	Chunking	3	14	9
	CWS	7	64	7
Structure Pred.	Semantic Parsing	4	12	4
Text Generation	Summarization	2	36	7
	Translation	4	60	9

Table 2: Brief descriptions of tasks, datasets and systems that EXPLAINBOARD currently supports. “Attr.” denotes Attribute. “Pred.” denotes “Prediction”.

(Fu et al., 2020b; Schweter and Akbik, 2020; Yamada et al., 2020). Using these methods, we have gathered 74 models on six NER datasets with system output information.

### 4.2 Analysis using ExplainaBoard

Fig. 2 illustrates different types of results driven by four functionality buttons<sup>13</sup> over the top-3 NER systems: LUKE (Yamada et al., 2020), FLERT (Schweter and Akbik, 2020) and FLAIR (Akbik et al., 2019).

**Box A** breaks down the performance of the top-1 system over different attributes.<sup>14</sup> We can intuitively observe that even the state-of-the-art system does worse on longer entities. Users can further print error cases in the longer entity bucket by clicking the corresponding bin.

**Box B** shows the 1st system’s (LUKE) performance minus the 2nd system’s (FLERT) performance. We can see that although LUKE surpasses FLERT holistically, it performs worse when dealing with PERSON entities.

**Box C** identifies samples that all systems mispredict. Further analysis of these samples uncovers challenging patterns or annotation errors.

**Box D** examines potential complementarity among these top-3 systems. The result shows that, by a simple voting ensemble strategy, a new state-

<sup>13</sup>As it is relatively challenging to define calibration in structure prediction tasks, this feature is currently only provided for classification tasks. We will explore more in the future.

<sup>14</sup>Due to the page limitation, we only show three.

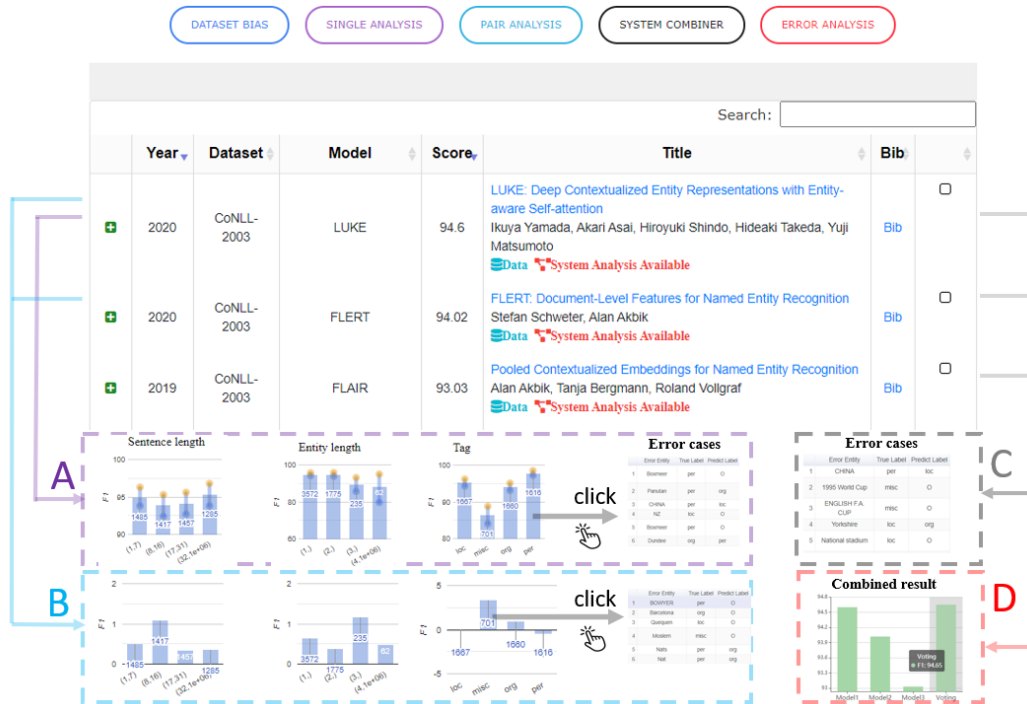


Figure 2: An example of the actual EXPLAINBOARD interface for NER over three top-performing systems on the CoNLL-2003 dataset. **Box A** shows the single-system analysis results obtained when users select the top-1 system and click the “Single Analysis” button. **Box B** shows the pairwise analysis results when top-2 systems are chosen and “Pair Analysis” is clicked. Users can click any bin of the histogram, which results in a fine-grained error case table. **Box C** represents a table with common errors of these top-3 systems. **Box D** illustrates the combined result of the top-3 systems.

of-the-art (94.65 F1) has been achieved on the CoNLL-2003 dataset.

## 5 Usage

**Example Use-cases** To show the practical utility of EXPLAINBOARD, we first present examples of how it has been used as an analysis tool in existing published research papers. Fu et al. (2020b) (Tab.4) utilize *single-system* analysis with the attribute of label consistency for NER task while Zhong et al. (2019) (Tab.4-5) use it for text summarization with attributes of density and compression. Fig.4 and Tab.3 in Fu et al. (2020a) leverage the *data bias* analysis and *pairwise system* diagnostics to interpret top-performing NER systems while Tab.4 in Fu et al. (2020c) use *single and pairwise* system analysis to investigate what’s next for the Chinese Word Segmentation task. Liu et al. (2021) use *system combiner* functionality to make ensemble analysis of summarization systems and Fig.1 in Ye et al. (2021) use *reliability analysis* functionality to observe how confidence intervals change in different buckets of a performance histogram.

**Using ExplainaBoard** Researchers can use EXPLAINBOARD in different ways: (i) We maintain a website where each task-specific EXPLAINBOARD allows researchers to interact with it, interpreting systems and datasets that they are interested in from different perspectives. (ii) We also release our back-end code for different NLP tasks so that researchers could flexibly use them to process their own system outputs, which can assist their research projects.

**Contributing to ExplainaBoard** The community can contribute to EXPLAINBOARD in several ways: (i) Submit system outputs of their implemented models. (ii) Add more informative attributes for different NLP tasks. (iii) Add new datasets or benchmarks for existing or new tasks.

## 6 Implications and Roadmap

EXPLAINBOARD presents a new paradigm in leaderboard development for NLP. This is just the beginning of its development, and there are many future directions.

**Research Revolving on System Outputs**<sup>15</sup> Due to the ability to analyze, contrast, or combine results from many systems EXPLAINBOARD incentivizes researchers to submit their results to explain-BOARD to better understand them and showcase their systems’ strengths. At the same time, EXPLAINBOARD will serve as a central repository for system outputs across many tasks, allowing for future avenues of research into cross-system analysis or system combination.

**Enriching ExplainaBoard with Glass-box Analysis** EXPLAINBOARD currently performs black-box analysis, solely analyzing system outputs without accessing model internals. On the other hand, there are many other glass-box interpretability tools that look at model internals, such as the AllenNLP Interpret (Wallace et al., 2019) and Language Interpretability Tool (Tenney et al., 2020). Expanding leaderboards to glass-box analysis methods (see Lipton (2018); Belinkov and Glass (2019) for a survey) is an interesting future work.

In the future, we aim to improve the applicability and usefulness by following action items: (1) Collaborate with more leaderboard organizers of diverse tasks and set up corresponding EXPLAIN-BOARDS for them. (2) Cover more tasks, datasets, models, as well as functionalities.

## Acknowledgements

We thank all reviewers for their valuable comments and authors who share their system outputs with us: Ikuya Yamada, Stefan Schweter, Colin Raffel, Yang Liu, Li Dong. We also thank Vijay Viswanathan, Yiran Chen, Hiroaki Hayashi for useful discussion and feedback about EXPLAINBOARD.

The work was supported in part by the Air Force Research Laboratory under agreement number FA8750-19-2-0200. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government.

<sup>15</sup>We released system outputs of EXPLAINBOARD: <http://explainboard.nlpedia.ai/download.html>

## References

- Alan Akbik, Tanja Bergmann, and Roland Vollgraf. 2019. [Pooled contextualized embeddings for named entity recognition](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 724–728, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rie Ando and Tong Zhang. 2005. [A high-performance semi-supervised learning method for text chunking](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 1–9, Ann Arbor, Michigan. Association for Computational Linguistics.
- Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(WMT20\)](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55, Online. Association for Computational Linguistics.
- Yonatan Belinkov and James Glass. 2019. [Analysis methods in neural language processing: A survey](#). *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Manik Bhandari, Pranav Narayan Gour, Atabak Ashfaq, Pengfei Liu, and Graham Neubig. 2020. [Re-evaluating evaluation in text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9347–9359, Online. Association for Computational Linguistics.
- 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*, pages

- 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Xinchi Chen, Xipeng Qiu, Chenxi Zhu, Pengfei Liu, and Xuanjing Huang. 2015. [Long short-term memory neural networks for Chinese word segmentation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1197–1206, Lisbon, Portugal. Association for Computational Linguistics.
- Zhiyuan Chen, Bing Liu, Meichun Hsu, Malu Castellanos, and Riddhiman Ghosh. 2013. [Identifying intention posts in discussion forums](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1041–1050, Atlanta, Georgia. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). 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 7057–7067.
- Junqi Dai, Hang Yan, Tianxiang Sun, Pengfei Liu, and Xipeng Qiu. 2021. Does syntax matter? a strong baseline for aspect-based sentiment analysis with roberta. In *The 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Kevin Duh, Katsuhito Sudoh, Xianchao Wu, Hajime Tsukada, and Masaaki Nagata. 2011. Generalized minimum bayes risk system combination. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 1356–1360.
- Bradley Efron. 1992. Bootstrap methods: another look at the jackknife. In *Breakthroughs in statistics*, pages 569–593. Springer.
- Kawin Ethayarajh and Dan Jurafsky. 2020. [Utility is in the eye of the user: A critique of NLP leaderboards](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4846–4853, Online. Association for Computational Linguistics.
- Jinlan Fu, Xuanjing Huang, and Pengfei Liu. 2021. [Spanner: Named entity re-/recognition as span prediction](#). In *Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Conference on Natural Language Processing (ACL-IJCNLP)*, Virtual.
- Jinlan Fu, Pengfei Liu, and Graham Neubig. 2020a. [Interpretable multi-dataset evaluation for named entity recognition](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6058–6069, Online. Association for Computational Linguistics.
- Jinlan Fu, Pengfei Liu, and Qi Zhang. 2020b. [Rethinking generalization of neural models: A named entity recognition case study](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7732–7739. AAAI Press.
- Jinlan Fu, Pengfei Liu, Qi Zhang, and Xuanjing Huang. 2020c. [RethinkCWS: Is Chinese word segmentation a solved task?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5676–5686, Online. Association for Computational Linguistics.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Aremu Anuoluwapo, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna Clinciu, Dipanjan Das, Kaustubh D Dhole, et al. 2021. The gem benchmark: Natural language generation, its evaluation and metrics. *arXiv preprint arXiv:2102.01672*.
- Jesús González-Rubio, Alfons Juan, and Francisco Casacuberta. 2011. [Minimum Bayes-risk system combination](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1268–1277, Portland, Oregon, USA. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. [On calibration of modern neural networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Jungo Kasai, Yejin Choi, Noah A Smith, and Daniel S Weld. 2021. Genie: A leaderboard for human-in-the-loop evaluation of text generation. *arXiv preprint arXiv:2101.06561*.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fengei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. 2020. Xglue: A new benchmark dataset for cross-lingual pre-training, understanding and generation. *arXiv preprint arXiv:2004.01401*.
- Zachary C Lipton. 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57.



- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Yixin Liu and Pengfei Liu. 2021. [Simcls: A simple framework for contrastive learning of abstractive summarization](#). In *Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, Virtual.
- Yixin Liu, Dou Ziyi, and Pengfei Liu. 2021. [Refsum: Refactoring neural summarization](#). In *The 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. [Hierarchical question-image co-attention for visual question answering](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 289–297.
- Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, and Xinyi Wang. 2019. [compare-mt: A tool for holistic comparison of language generation systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 35–41, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. [Thumbs up? sentiment classification using machine learning techniques](#). In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, pages 79–86. Association for Computational Linguistics.
- Nikolaos Pappas and Andrei Popescu-Belis. 2014. [Explaining the stars: Weighted multiple-instance learning for aspect-based sentiment analysis](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 455–466, Doha, Qatar. Association for Computational Linguistics.
- Maja Popović and Hermann Ney. 2011. [Towards automatic error analysis of machine translation output](#). *Computational Linguistics*, 37(4):657–688.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. [A neural attention model for abstractive sentence summarization](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389, Lisbon, Portugal. Association for Computational Linguistics.
- Stefan Schweter and Alan Akbik. 2020. [Flert: Document-level features for named entity recognition](#). *arXiv preprint arXiv:2011.06993*.
- Sara Stymne. 2011. [Blast: A tool for error analysis of machine translation output](#). In *Proceedings of the ACL-HLT 2011 System Demonstrations*, pages 56–61, Portland, Oregon. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- Ian Tenney, James Wexler, Jasmijn Bastings, Tolga Bolukbasi, Andy Coenen, Sebastian Gehrmann, Ellen Jiang, Mahima Pushkarna, Carey Radebaugh, Emily Reif, et al. 2020. [The language interpretability tool: Extensible, interactive visualizations and analysis for nlp models](#). *arXiv preprint arXiv:2008.05122*.
- Kai Ming Ting and Ian H. Witten. 1997. [Stacked generalizations: When does it work?](#) In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97, Nagoya, Japan, August 23-29, 1997, 2 Volumes*, pages 866–873. Morgan Kaufmann.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Kristina Toutanova, Dan Klein, Christopher D. Manning, and Yoram Singer. 2003. [Feature-rich part-of-speech tagging with a cyclic dependency network](#). In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 252–259.
- Eric Wallace, Jens Tuyls, Junlin Wang, Sanjay Subramanian, Matt Gardner, and Sameer Singh. 2019. [AllenNLP interpret: A framework for explaining predictions of NLP models](#). 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): System Demonstrations*, pages 7–12, Hong Kong, China. Association for Computational Linguistics.

- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. Super-glue: A stickier benchmark for general-purpose language understanding systems. *arXiv preprint arXiv:1905.00537*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Sida Wang and Christopher Manning. 2012. [Baselines and bigrams: Simple, good sentiment and topic classification](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 90–94, Jeju Island, Korea. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. [LUKE: Deep contextualized entity representations with entity-aware self-attention](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454, Online. Association for Computational Linguistics.
- Zihuiwen Ye, Pengfei Liu, Jinlan Fu, and Graham Neubig. 2021. [Towards more fine-grained and reliable NLP performance prediction](#). In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, Online.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.
- Ming Zhong, Danqing Wang, Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2019. [A closer look at data bias in neural extractive summarization models](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 80–89, Hong Kong, China. Association for Computational Linguistics.