

PromptShots at the FinNLP-2022 ERAI Tasks: Pairwise Comparison and Unsupervised Ranking

Peratham Wiriyathamabhum
peratham.bkk@gmail.com

Abstract

This report describes our PromptShots submissions to a shared task on Evaluating the Rationales of Amateur Investors (ERAI). We participated in both pairwise comparison and unsupervised ranking tasks. For pairwise comparison, we employed instruction-based models based on T5-small and OpenAI InstructGPT language models. Surprisingly, we observed OpenAI InstructGPT language model few-shot trained on Chinese data works best in our submissions, ranking 3rd on the maximal loss (ML) pairwise accuracy. This model works better than training on the Google translated English data by a large margin, where the English few-shot trained InstructGPT model even performs worse than an instruction-based T5-small model finetuned on the English data. However, all instruction-based submissions do not perform well on the maximal potential profit (MPP) pairwise accuracy where there are more data and learning signals. The Chinese few-shot trained InstructGPT model still performs best in our setting. For unsupervised ranking, we utilized many language models, including many financial-specific ones, and Bayesian lexicons unsupervised-learned on both Chinese and English words using a method-of-moments estimator. All our submissions rank best in the MPP ranking, from 1st to 3rd. However, they all do not perform well for ML scoring. Therefore, both MPP and ML scores need different treatments since we treated MPP and ML using the same formula. Our only difference is the treatment of market sentiment lexicons.

1 Introduction

Evaluating the rationals of amateur investors (ERAI) (Chen et al., 2021a,b) is a shared task on evaluating social media opinions on the topic of investments and whether they are going to be useful or not. Mining high-quality opinions by inspecting their supporting rationales might utilize the wisdom of the crowd on social media. Previous work (Chen

et al., 2021a) proposes stylistic and semantic features to filter out noisy crowd opinions which may not be high-quality and profitable. There are two settings in this ERAI shared task, pairwise comparison and unsupervised ranking. These settings sort out the opinions based on two metrics, higher maximal potential profit (MPP) and lower maximal loss (ML). In pairwise comparison, two posts are given with a binary label whether the MPP and ML of the first post are more or less than the second post. In unsupervised ranking, the goal is to filter and keep the top 10% posts based on MPP and ML given a set of unranked posts.

For pairwise comparison, our best submission ranks 3rd on the maximal loss (ML) pairwise accuracy on the leaderboard¹. For unsupervised ranking, our best submission ranks 1st on the maximal potential profit (MPP) ranking. The codes for our systems are open-sourced and available at our GitHub repository².

2 Models

2.1 Pairwise Comparison

For pairwise comparison, we utilized instruction-based models based on T5-small (Raffel et al., 2020) and OpenAI InstructGPT language models (Ouyang et al., 2022) in a few-shot prompt-based setting (Brown et al., 2020b).

2.1.1 T5

T5 is an encoder-decoder language model which was trained by treating every text processing problem as a "text-to-text" problem to unify NLP tasks using only a single model, loss function, hyperparameter set, etc. The input texts will be encoded and the T5 decoder will decode them. Specifically, T5 was unsupervised-pretrained by denois-

¹<https://sites.google.com/nlg.csie.ntu.edu.tw/finnlp-2022-emnlp/era-shared-task>

²https://github.com/perathambkk/finnlp_era_shared_task_emnlp2022

ing masked inputs on the “Colossal Clean Crawled Corpus” (C4) dataset, Common Crawled from web scraping. Then, T5 can be further supervised fine-tuned using the “text-to-text” format and T5’s decoder will be in the teacher forcing mode where the decoder will be trained using input and a right-shifted target sequence. T5 architecture is pretty much the same as the vanilla transformer (Vaswani et al., 2017) except removing the LayerNorm bias term, placing the LayerNorm outside the residual path, and using a relative position embedding (Shaw et al., 2018).

In the T5 paper, the authors state that T5 can be specified which task it should perform by adding a task-specific textual prefix to the original input sequence before feeding into the model. Therefore, we take T5-small as an instruction-based model using the input prompt template, ‘post1 : %s post2 : %s </s>’, where the %s contains texts from the corresponding post, and the output prompt, ‘maximal potential profit (MPP) : %s maximal loss (ML) : %s </s>’, where the %s contains the MPP and the ML corresponding labels accordingly. This is similar to the baseline system in the FLUTE figurative language understanding dataset paper (Chakrabarty et al., 2022), however, in our case, the T5-small is expected to jointly predict both MPP and ML in one forward pass of preparing the probability tensor. We use top-p sampling for text generation (Holtzman et al., 2019).

2.1.2 OpenAI’s InstructGPT

The OpenAI API has many variants of InstructGPT language models based on the GPT-3 autoregressive language model to conveniently perform various NLP tasks with the prompt library. The InstructGPT was trained with a human-in-the-loop style and is claimed to be better at following instructions, more truthful, and less toxic than the GPT-3. In this shared task, we engineered the prompts for InstructGPT using a few-shot learning setting, as in the GPT-3 paper (Brown et al., 2020a), where few data instances were given from the target task/domain. Each data instance will become a prompt as ‘post1 : d[‘post1’] post2: d[‘post2’] > maximal potential profit (MPP)| %s# maximal loss (ML)| %s.’, where d[‘post1’] and d[‘post2’] are texts from the corresponding post. Then, we append the query we want to predict MPP and ML as just a truncated template, ‘post1 : d[‘post1’] post2: d[‘post2’] >’, and let the language model generate the rest.

We use the ‘text-davinci-002’ model and randomly construct those few-shot prompts where each prompt will be a length of around 4,000 because of the API token length limit. We use the same setting and the model pipeline for both of our submissions 2 and 3 where we use the Chinese posts as d[‘post1’] and d[‘post2’] for our submission 2 and the Google-translated English posts for our submission 3. By this we mean, for example, we use the same tokenizer for Chinese and English. Therefore, these systems are very simple and to-go prompt-based systems. We had done very minimal parameter tuning to the model, only prompt engineering. For a survey in prompt-based systems, please consider (Liu et al., 2021).

2.2 Unsupervised Ranking

For unsupervised ranking, we utilized many financial and general language models and Bayesian lexicons in both Chinese and English.

2.2.1 Base Model

Our first submission, our base model, consists of a stylistic length feature (Zong et al., 2020) derived from the opinion (sub)word lengths segmented using the ‘hfl/chinese-bert-wwm-ext’ tokenizer (Cui et al., 2021), prediction scores from FinBERT-FLS (Huang et al., 2020), a professional lexicon count from FinProLex (Chen et al., 2021a), and a market sentiment lexicon count from NTUSD-Fin (Chen et al., 2018).

In the measuring forecasting skill from text paper (Zong et al., 2020), the authors observe various linguistics phenomena indicating that skilled forecasters tend to write significantly longer justifications because of more rationale. For example, skilled forecasters also provide less readability, because of the usage of more complex languages, and less emotion, because of the usage of less emotional languages as neutral sentiments. Moreover, skilled forecasters tend to use more cardinal numbers, prepositions, and nouns. They tend to use fewer verbs and pronouns. Therefore, in this base model, we just stick with the lengths of justifications as our simplest skill indicator.

FinBERT (Huang et al., 2020) is essentially BERT (Devlin et al., 2019) customized for financial texts, pretrained on corporate filings, analyst reports, and earnings conference call transcripts, which differ from normal texts in both vocabulary and writing style. In the FinBERT paper, FinBERT outperforms all other methods, Loughran McDon-

al lexicon, and machine learning algorithms, especially in negative financial sentiment prediction, when finetuned for financial sentiment analysis (FinBERT-tone). FinBERT was finetuned in two additional tasks, labeling environment, social, and governance (ESG) discussions and labeling forward-looking statements (FLS), from firms' corporate social responsibility (CSR) reports and management discussion and analyses (MD&As) textual sentences. For this base model, we sum and normalize the prediction logit outputs from FinBERT-FLS classes, $\{FLS, NON_FLS, NOT_FLS\}$ on each sentence of the textual inputs as our scores.

FinProLex (Chen et al., 2021a) is a Chinese financial lexicon derived from Bloomberg Terminal and PTT Stock Taiwanese social media platform, containing 5,162 tokens from professional analysts' reports and social media posts paired with expertise scores. FinProLex uses Point-wise Mutual Information (PMI), as in (Turney, 2002; Li and Shah, 2017), to measure the association strengths between a word and either the positive or negative lexicon. The formula of the expert-like score (ELScore) of a given word w is as follows:

$$ELScore_w = PMI(w, analyst) - PMI(w, amateur), \quad (1)$$

$$ELScore_w = \log_2 \frac{p(w, analyst)}{p(w)p(analyst)} - \log_2 \frac{p(w, amateur)}{p(w)p(amateur)}, \quad (2)$$

where *analyst* and *amateur* are labels of whether a given word is from an analyst report or an amateur post. This is the difference value between the PMI scores measuring how much a term is associated with either analyst or amateur documents. This is similar to the term's sentiment score (S_{PMI}) (Li and Shah, 2017) which is

$$S_{PMI} = PMI(w, bullish) - PMI(w, bearish). \quad (3)$$

FinProLex tends to include hard words, complex semantics, noun phrase modifiers, content words, transition words, personal pronouns, and negative words as experts tend to use most of them, except personal pronouns and negative words which are used more by amateurs, based on the paper findings that can be summarized as experts tend to evaluate pricing and valuations while amateurs tend to predict the stock movements.

NTUSD-Fin is an English lexicon for market sentiment analysis from StockTwits, containing 8,331 words, 112 hashtags, and 115 emojis. We used their market sentiment scores which are also computed essentially from equation (3).

We aggregated the scores to predict MPP and ML using these heuristic functions (base-1),

$$MPP = len + FinProLex + |FinWord > 0| + (FLS + 0.5 NON_FLS - NOT_FLS), \quad (4)$$

$$ML = len + FinProLex + |FinWord < 0| + (FLS + 0.5 NON_FLS - NOT_FLS). \quad (5)$$

We simply used a weighted sum as our heuristic function. We grouped similar scores together. $len + FinProLex$ are stylistic features where we put an equal weight of 1 for each of them. We used *FinWord* as a switch feature for either MPP or ML that would behave differently because of the market sentiment based on our belief. MPP posts should be from a bullish market while ML posts should instead be from a bearish market. *FLS* has 3 different class scores so we weighted 1 for a positive class, 0.5 for a less positive one, and -1 for a negative class. The weights are just our rule-of-thumb (make-up numbers that we felt they made sense solely from our intuitions).

It is like trying to intuitively come up with a good feature weighting number for a Maximum Entropy (MaxEnt) model. From our heuristic functions, we just down-weighted some scores and specify some negative interactions. We did not normalize the weighting into probabilities but the ranking should be the same anyway. Most weightings are uniformly the same number.

- If a score should positively correlate with the target, we should give a high weight.
- If a score should weakly correlate with the target, we should give a low weight.
- If a score should negatively correlate with the target, we should give a high negative weight.
- For the rest that we are not certain of, they should retain a maximum entropy (uniformity).
- These heuristics can be estimated with intuitions and give an intuitive unsupervised aggregated scoring function.

However, we submitted the same function for MPP and ML to get a sense of using the same strategy for both bullish and bearish markets.

2.2.2 Bayesian Lexicons

Next, we added Bayesian lexicons (Eisenstein, 2017) (by fitting FinProLex and NTUSD-Fin), FinBERT-tone (Huang et al., 2020), (fitted) Loughran-McDonald financial sentiment lexicon (LM) (Loughran and McDonald, 2011) and Part-of-Speech (POS) features (Zong et al., 2020) into the score aggregators. We would like to note that these lexicons are not multi-word (only unigrams) so they are not expected to be able to handle negations except the creators of those lexicons had made them handle some kind of negations, like in the LM lexicon Fin-Pos list. The authors use bigram to quadgrams counts when that bigram to quadgram follows some negation patterns. Our second and third submissions differ in the normalization of scores and Bayesian lexicon variants.

Loughran-McDonald financial sentiment lexicon (LM) (Loughran and McDonald, 2011) was created because the Harvard Psychosociological Dictionary, specifically, the Harvard-IV-4 TagNeg (H4N) file, does not perform well in financial and accounting domains. Lots of Harvard dictionary negative words are not negative in finance.

Bayesian lexicon learns predictive weights for each word in a lexicon using a method-of-moments estimator from co-occurrence statistics without any labels as a special case of multinomial Naïve Bayes. For the second submission, we use the Dirichlet Compound Multinomial likelihood to reduce effective counts for repetitive words. For the third submission, we use the multinomial likelihood model. For example, when we fitted the LM lexicon using the pairwise comparison data, we gave 0.02626 to ‘good’, 0.00501 to ‘optimistic’, and 0.00278 to ‘highest’. For LM negative words, we gave 0.00243 to ‘decline’, 0.00234 to sharply, and 0.00186 to ‘difficult’.

Our POS features are motivated by the measuring forecasting skill from text paper. We simply counted cardinal numbers, nouns, and verbs from Chinese jieba segmented texts. Then, these counts were normalized into the range of $[0, 1]$.

For these submissions, we sum and normalize the prediction logit outputs from FinBERT-tone classes, $\{pos_tone, neg_tone\}$ on each sentence of the textual inputs as our scores.

For the second submission, we aggregated the

scores to predict MPP and ML using these heuristic functions (bayesdcm-2),

$$\begin{aligned} MPP = & len + FinProLex + |FinWord > 0| \\ & + (FLS + 0.5 \text{NON_FLS} - \text{NOT_FLS}) \\ & + (pos_tone - neg_tone + LM) \\ & + (nouns + cards - verbs), \end{aligned} \quad (6)$$

$$\begin{aligned} ML = & len + FinProLex + |FinWord < 0| \\ & + (FLS + 0.5 \text{NON_FLS} - \text{NOT_FLS}) \\ & + (pos_tone - neg_tone + LM) \\ & + (nouns + cards - verbs). \end{aligned} \quad (7)$$

For the third submission, we aggregated the scores to predict MPP and ML using these heuristic functions (multinomial-3),

$$\begin{aligned} MPP = & 0.5(len + FinProLex) \\ & + 0.33(FLS + 0.5 \text{NON_FLS} - \text{NOT_FLS}) \\ & + 0.33(pos_tone - neg_tone + LM) \\ & + 0.33(nouns + cards - verbs) \\ & + |FinWord > 0|, \end{aligned} \quad (8)$$

$$\begin{aligned} ML = & 0.5(len + FinProLex) \\ & + 0.33(FLS + 0.5 \text{NON_FLS} - \text{NOT_FLS}) \\ & + 0.33(pos_tone - neg_tone + LM) \\ & + 0.33(nouns + cards - verbs) \\ & + |FinWord < 0|. \end{aligned} \quad (9)$$

In these functions, we tried to group and reweigh the scores as normalization. If two or more scores mean the same thing, we might double count.

3 Experimental Results

In our experiments, most of our submissions (except T5-small) are intuition-based heuristics, and we did not even measure neither their training nor validation performance at all during the competition. We did not use any data augmentation techniques.

3.1 Pairwise Comparison

The experimental results in Table.1 show that the OpenAI InstructGPT language model few-shot trained on Chinese data works best in our submissions, ranking 3rd on the maximal loss (ML) pairwise accuracy, even better than instead training on the Google translated English data by a

Table 1: MPP and ML accuracies of our models in pairwise comparison test data. (The numbers in subscript are submission rankings on the leaderboard. The symbol † denotes a top-3 performance.)

Model	MPP acc.	ML acc.
T5-small	47.13 ₁₆	45.98 ₁₀
InstructGPT-zh	48.28 ₁₄	54.02 ₃ †
InstructGPT-en	47.13 ₁₆	41.38 ₁₃
FinNLP-22 best	62.07	59.77

Table 2: Average MPP and ML from top 10% posts of our models in unsupervised ranking test data. (The numbers in subscript are submission rankings on the leaderboard. The symbol † denotes a top-3 performance and the symbol ‡ denotes the score beats the baseline.)

Model	avg. MPP	avg. ML
Stylistic baseline	17.61%	-2.46%
base-1	22.53% ₃ † ‡	-7.80% ₁₁
bayesdcm-2	24.39 % ₁ † ‡	-13.04% ₁₆
multinomial-3	23.76% ₂ † ‡	-12.33% ₁₅
FinNLP-22 best	24.39% (ours)	-2.46%

large margin, where the English few-shot trained InstructGPT model even performs worse than an instruction-based T5-small model finetuned on the English data. However, all instruction-based submissions do not perform well on the maximal potential profit (MPP) pairwise accuracy where there are more data and learning signals, nonetheless, the Chinese few-shot trained InstructGPT model still performs best in our setting.

We additionally split the training data into a held-out train/val split and evaluated our methods on the val split in Table 3. The results are a bit different since the English version of the InstructGPT works better. However, we did not hope for an accurate cross-validation estimation given a small amount of data. Using leave-one-out validation (LOOCV) or k -fold cross-validation with a high value of k can produce a better estimation but they are costly. We might be able to generate more data pairs, but we decided to keep the same setting.

3.2 Unsupervised Ranking

For the unsupervised ranking task, we utilized many language models, including many financial-specific ones, and Bayesian lexicons unsupervisedly learned on both Chinese and English words. All of our submissions rank best in the MPP ranking, from 1st to 3rd in this task. However, they all do not perform well for the ML scoring. Therefore, both MPP and ML scores need different treatments

Table 3: Additional experiments on using our pairwise comparison methods on a held-out train/val split (ratio=0.3). The evaluation metric is accuracy.

Model	MPP acc.	ML acc.
T5-small	0.4833	0.6000
InstructGPT-zh	0.4667	0.4167
InstructGPT-en	0.6167	0.4667

Table 4: Additional experiments on using our unsupervised ranking methods to rank all posts of the pairwise data. The evaluation metrics are average MPP and average ML of the top 10% posts.

Model	avg. MPP	avg. ML
base-1	0.2083	-0.2108
bayesdcm-2	0.2085	-0.2104
multinomial-3	0.2085	-0.2104

substantially since we treated MPP and ML using the same formula. Our only difference is the treatment of market sentiment lexicons. We feel that the sentiment features, or mostly semantic features, might be negatively correlated, weakly correlated, or even uncorrelated with ML because the stylistic baseline performs best, and our base submission performs better than our Bayesian lexicon submissions.

We conducted additional experiments on unsupervised ranking by using the whole training set of the pairwise comparison data. We compared all posts using our scoring functions in Table 4. The results show not much difference among our methods. When we tried to evaluate using the pairwise comparison accuracy, the results show no difference (0.545 MPP comparison acc. and 0.525 ML comparison acc.) as our methods were not designed for that.

4 Conclusion

This report describes our systems for a shared task of evaluating the rationales of amateur investors at FinNLP-2022. From the experimental results in pairwise comparison, we conclude that few-shot prompted instruction-based language models can work reasonably well in low resource settings with minimal training efforts but might need quite accurate data from sources since using translated data seems not to perform well. From the experimental results in unsupervised ranking, financial language models perform well and Bayesian-fitting the lexicons helps improve the performance. Also, the heuristic function design needs to differ between MPP and ML.

Limitations

We only sampled a relatively small portion of models and draw conclusions. We also conducted experiments only on one dataset for evaluating the rationales of amateur investors. Besides, the dataset is in Chinese with English translation using Google Translate. Lots of our methods rely on the translated data.

Because we are limited to only three submissions, we don't know how each feature set contributes to the score. There were no ablations. However, the shared task organizers released the test data with ground truths in private.

The authors are self-affiliated and do not represent any entities. The authors also participated in the shared task under many severe unattended local personal criminal events in their home countries. There might be some unintentional errors and physical limitations based on these unlawful interruptions. Even at the time of drafting this report, the authors suffer from unknown toxin fumes spraying into their places. We want to participate in the shared task because it is fun and educational. We apologize for any errors in this report. We tried our best.

Ethics Statement

Scientific work published at EMNLP 2022 must comply with the [ACL Ethics Policy](#). We, the authors, hope the intended uses of our systems are for peace, well-being, and social good only. No harm.

Acknowledgments

We would like to thank anonymous reviewers for their constructive feedback and additional experiment suggestions.

References

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020a. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020b. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.

Tuhin Chakrabarty, Arkadiy Saakyan, Debanjan Ghosh, and Smaranda Muresan. 2022. [Flute: Figurative language understanding and textual explanations](#). *arXiv preprint arXiv:2205.12404*.

Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2018. [Nturd-fin: a market sentiment dictionary for financial social media data applications](#). In *Proceedings of the 1st Financial Narrative Processing Workshop (FNP 2018)*, pages 37–43.

Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2021a. [Evaluating the rationales of amateur investors](#). In *Proceedings of the Web Conference 2021*, pages 3987–3998.

Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2021b. [From opinion mining to financial argument mining](#). Springer Nature.

Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, and Ziqing Yang. 2021. [Pre-training with whole word masking for chinese bert](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3504–3514.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Jacob Eisenstein. 2017. [Unsupervised learning for lexicon-based classification](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. [The curious case of neural text degeneration](#). In *International Conference on Learning Representations*.

Allen Huang, Hui Wang, and Yi Yang. 2020. [Finbert—a deep learning approach to extracting textual information](#). *Available at SSRN 3910214*.

Quanzhi Li and Sameena Shah. 2017. [Learning stock market sentiment lexicon and sentiment-oriented word vector from StockTwits](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 301–310, Vancouver, Canada. Association for Computational Linguistics.

- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.
- Tim Loughran and Bill McDonald. 2011. When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *The Journal of finance*, 66(1):35–65.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana. Association for Computational Linguistics.
- Peter Turney. 2002. [Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 417–424, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Shi Zong, Alan Ritter, and Eduard Hovy. 2020. [Measuring forecasting skill from text](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5317–5331, Online. Association for Computational Linguistics.