

Semi-Supervised Exaggeration Detection of Health Science Press Releases

Dustin Wright and Isabelle Augenstein

Dept. of Computer Science

University of Copenhagen

Denmark

{dw|augenstein}@di.ku.dk

Abstract

Public trust in science depends on honest and factual communication of scientific papers. However, recent studies have demonstrated a tendency of news media to misrepresent scientific papers by exaggerating their findings. Given this, we present a formalization of and study into the problem of *exaggeration detection* in science communication. While there are an abundance of scientific papers and popular media articles written about them, very rarely do the articles include a direct link to the original paper, making data collection challenging. We address this by curating a set of labeled press release/abstract pairs from existing expert annotated studies on exaggeration in press releases of scientific papers suitable for benchmarking the performance of machine learning models on the task. Using limited data from this and previous studies on exaggeration detection in science, we introduce MT-PET, a multi-task version of Pattern Exploiting Training (PET), which leverages knowledge from complementary cloze-style QA tasks to improve few-shot learning. We demonstrate that MT-PET outperforms PET and supervised learning both when data is limited, as well as when there is an abundance of data for the main task.¹

1 Introduction

Factual and honest science communication is important for maintaining public trust in science (Nelkin, 1987; Moore, 2006), and the “dominant link between academia and the media” are press releases about scientific articles (Sumner et al., 2014). However, multiple studies have demonstrated that press releases have a significant tendency to sensationalize their associated scientific articles (Sumner et al., 2014; Bratton et al., 2019; Woloshin et al., 2009; Woloshin and

¹The code and data are available online at <https://github.com/copenlu/scientific-exaggeration-detection>




Figure 1: Scientific exaggeration detection is the problem of identifying when a news article reporting on a scientific finding has exaggerated the claims made in the original paper. In this work, we are concerned with predicting exaggeration of the main finding of a scientific abstract as reported by a press release.

Schwartz, 2002). In this paper, we explore how natural language processing can help identify exaggerations of scientific papers in press releases.

While Sumner et al. (2014) and Bratton et al. (2019) performed manual analyses to understand the prevalence of exaggeration in press releases of scientific papers from a variety of sources, recent work has attempted to expand this using methods from NLP (Yu et al., 2019, 2020; Li et al., 2017). These works focus on the problem of automatically detecting the difference in the strength of causal claims made in scientific articles and press releases. They accomplish this by first building datasets of main claims taken from PubMed abstracts and (unrelated) press releases from EurekAlert² labeled for their strength. With this, they train machine learning models to predict claim strength, and analyze unlabelled data using these models. This marks an important first step toward the goal of automatically identifying exaggerated scientific claims in science reporting.

However, existing work has only partially at-

²<https://www.eurekalert.org/>

tempted to address this task using NLP. Particularly, there exists no standard benchmark data for the exaggeration detection task with **paired** press releases and abstracts i.e. where the data consist of tuples of the form (press release, abstract) and the press release is written about the paired scientific paper. Collecting paired data labeled for exaggeration is critical for understanding how well any solution performs on the task, but is challenging and expensive as it requires domain expertise (Sumner et al., 2014). The focus of this work is then to curate a standard set of benchmark data for the task of scientific exaggeration detection, provide a more realistic task formulation of the problem, and develop methods effective for solving it using limited labeled data. To this end, we present MT-PET, a multi-task implementation of Pattern Exploiting Training (PET, Schick and Schütze (2020a,b)) for detecting exaggeration in health science press releases. We test our method by curating a benchmark test set of data from the expert annotated data of Sumner et al. (2014) and Bratton et al. (2019), which we release to help advance research on scientific exaggeration detection.

Contributions In sum, we introduce:

- A new, more realistic task formulation for scientific exaggeration detection.
- A curated set of benchmark data for testing methods for scientific exaggeration detection consisting of 563 press release/abstract pairs.
- MT-PET, a multi-task extension of PET which beats strong baselines on scientific exaggeration detection.

2 Problem Formulation

We first provide a formal definition of the problem of scientific exaggeration detection, which guides the approach described in §3. We start with a set of document pairs $\{(t, s) \in \mathcal{D}\}$, where s is a source document (e.g. a scientific paper abstract) and t is a document written about the source document s (e.g. a press release for the paper). The goal is to predict a label $l \in \{0, 1, 2\}$ for a given document pair (t, s) , where 0 implies the target document *undersells* source document, 1 implies the target document accurately reflects the source document, and 2 implies the target document *exaggerates* the source document.

Two realizations of this formulation are investigated in this work. The first (defined as **T1**) is an *inference* task consisting of labeled document pairs

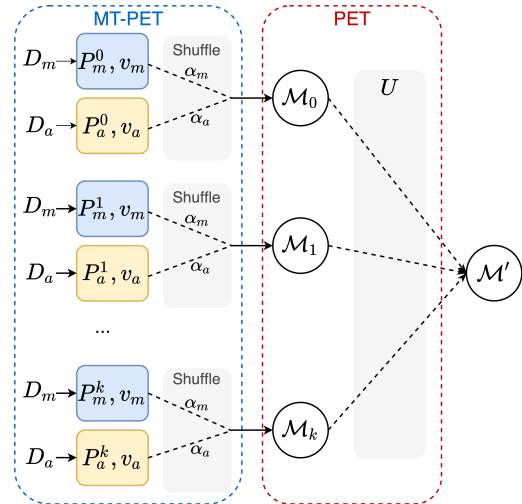


Figure 2: MT-PET design. We define pairs of complementary pattern-verbalizer pairs for a main task and auxiliary task. These PVPs are then used to train PET on data from both tasks.

used to learn to predict l directly. In other words, we are given training data of the form (t, s, l) and can directly train a model to predict l from both t and s . The second (defined as **T2**) is as a *classification* task consisting of a training set of documents $d \in \mathcal{D}'$ from **both** the source and the target domain, and a classifier is trained to predict the *claim strength* l' of sentences from these documents. In other words, we don't require **paired** documents (t, s) at train time. At test time, these classifiers are then applied to document pairs (t, s) and the predicted claim strengths (l'_s, l'_t) are compared to get the final label l . Previous work has used this formulation to estimate the prevalence of *correlation to causation* exaggeration in press releases (Yu et al., 2020), but have not evaluated this on paired labeled instances.

Following previous work (Yu et al., 2020), we simplify the problem by focusing on detecting when the *main finding* of a paper is exaggerated. The first step is then to identify the main finding from s , and the sentence describing the main finding in s from t . In our semi-supervised approach, we do this as an intermediate step to acquire unlabeled data, but for all labeled training and test data, we assume the sentences are already identified and evaluate on the sentence-level exaggeration detection task.

3 Approach

One of the primary challenges for scientific exaggeration detection is a lack of labeled training

data. Given this, we develop a semi-supervised approach for few-shot exaggeration detection based on pattern exploiting training (PET, Schick and Schütze (2020a,b)). Our method, multi-task PET (MT-PET, see Figure 2), improves on PET by using multiple complementary cloze-style QA tasks derived from different source tasks during training. We first describe PET, followed by MT-PET.

3.1 Pattern Exploiting Training (PET)

PET (Schick and Schütze, 2020a) uses the masked language modeling objective of pretrained language models to transform a task into one or more cloze-style question answering tasks. The two primary components of PET are *patterns* and *verbalizers*. *Patterns* are cloze-style sentences which mask a single token e.g. in sentiment classification with the sentence “We liked the dinner” a possible pattern is: “We liked the dinner. It was [MASK].” *Verbalizers* are single tokens which capture the meaning of the task’s labels in natural language, and which the model should predict to fill in the masked slots in the provided patterns (e.g. in the sentiment analysis example, the verbalizer could be `Good`).

Given a set of *pattern-verbalizer pairs* (PVPs), an ensemble of models is trained on a small labeled seed dataset to predict the appropriate verbalizations of the labels in the masked slots. These models are then applied on unlabeled data, and the raw logits are combined as a weighted average to provide soft-labels for the unlabeled data. A final classifier is then trained on the soft labeled data using a distillation loss based on KL-divergence.

3.2 Notation

We adopt the notation in the original PET paper (Schick and Schütze, 2020a) to describe MT-PET. In this, we have a masked language model \mathcal{M} with a vocabulary V and mask token $[MASK] \in V$. A pattern is defined as a function $P(x)$ which transforms a sequence of input sentences $\mathbf{x} = (s_0, \dots, s_{k-1})$, $s_i \in V^*$ to a phrase or sentence which contains exactly one mask token. Verbalizers $v(x)$ map a label in the task’s label space \mathcal{L} to a set of tokens in the vocabulary V which \mathcal{M} is trained to predict.

For a given sample $\mathbf{z} \in V^*$ containing exactly one mask token and $w \in V$ corresponding to a word in the language model’s vocabulary, $M(w|\mathbf{z})$ is defined as the unnormalized score that the language model gives to word w at the masked posi-

tion in \mathbf{z} . The score for a particular label as given in Schick and Schütze (2020a) is then

$$s_{\mathbf{p}}(l|\mathbf{x}) = M(v(l)|P(\mathbf{x})) \quad (1)$$

For a given sample, PET then assigns a score s for each label based on all of the verbalizations of that label. When applied to unlabeled data, this produces soft labels from which a final model \mathcal{M}' can be trained via distillation using KL-divergence.

3.3 MT-PET

In the original PET implementation, PVPs are defined for a single target task. MT-PET extends this by allowing for auxiliary PVPs from related tasks, adding complementary cloze-style QA tasks during training. The motivation for the multi-task approach is two-fold: 1) complementary cloze-style tasks can potentially help the model to learn different aspects of the main task; in our case, the similar tasks of exaggeration detection and claim strength prediction; 2) data on related tasks can be utilized during training, which is important in situations where data for the main task is limited.

Concretely, we start with a main task T_m with a small labeled dataset $(x_m, y_m) \in D_m$, where $y_m \in \mathcal{L}_m$ is a label for the instance, as well as an auxiliary task T_a with labeled data $(x_a, y_a) \in D_a$, $y_a \in \mathcal{L}_a$. Each pattern $P_m^i(x)$ for the main task has a corresponding complementary pattern $P_a^i(x)$ for the auxiliary task. Additionally, the labels in \mathcal{L}_a have their own verbalizers $v_a(x)$. Thus, with k patterns, the full set of PVP tuples is given as

$$\mathcal{P} = \{((P_m^i, v_m), (P_a^i, v_a)) | 0 \leq i < k\}$$

Finally, a large set of unlabeled data U for the *main task only* is available. MT-PET then trains the ensemble of k masked language models using the pairs defined for the main and auxiliary task. In other words, for each individual model both the main PVP (P_m, v_m) and auxiliary PVP (P_a, v_a) are used during training.

For a given model \mathcal{M}_i in the ensemble, on each batch we randomly select one task T_c , $c \in \{m, a\}$ on which to train. The PVP for that task is then selected as (P_c^i, v_c) . Inputs (x_c, y_c) from that dataset are passed through the model, producing raw scores for each label in the task’s label space.

$$s_{\mathcal{P}_c^i}(\cdot|\mathbf{x}_c) = \{\mathcal{M}_i(v_c(l)|P_c^i(\mathbf{x}_c)) | \forall l \in \mathcal{L}_c\} \quad (2)$$

The loss is calculated as the cross-entropy between the task label y_c and the softmax of the score s

Name	Pattern
$P_{T_1}^0(x)$	Scientists claim a . Reporters claim b .The reporters claims are [MASK]
$P_{T_1}^1(x)$	Academic literature claims a . Popular media claims b . The media claims are [MASK]
$P_{T_2}^0(x)$	[Reporters Scientists] say a . The claim strength is [MASK]
$P_{T_2}^1(x)$	[Academic literature Popular media] says a . The claim strength is [MASK]

Table 1: Patterns for both **T1** (exaggeration detection) and **T2** (claim strength prediction)

normalized over the scores for all label verbalizations (Schick and Schütze, 2020a), weighted by a term α_c .

$$q_{\mathbf{p}_c^i} = \frac{e^{s_{\mathbf{p}_c^i}(\cdot|\mathbf{x}_c)}}{\sum_{l \in \mathcal{L}_c} e^{s_{\mathbf{p}_c^i}(l|\mathbf{x}_c)}} \quad (3)$$

$$L_c = \alpha_c * \frac{1}{N} \sum_n H(y_c^{(n)}, q_{\mathbf{p}_c^i}^{(n)}) \quad (4)$$

where N is the batch size, n is a sample in the batch, H is the cross-entropy, and α_c is a hyperparameter weight given to task c .

MT-PET then proceeds in the same fashion as standard PET. Different models are trained for each PVP tuple in \mathcal{P} , and each model produces raw scores $s_{\mathbf{p}_m^i}$ for all samples in the unlabeled data. The final score for a sample is then a weighted combination of the scores of individual models.

$$s(l|\mathbf{x}_u^j) = \sum_i w_i * s_{\mathbf{p}_m^i}(l|\mathbf{x}_u^j) \quad (5)$$

where the weights w_i are calculated as the accuracy of model \mathcal{M}_i on the train set D_m before training. The final classification model is then trained using KL-divergence between the predictions of the model and the scores s as target logits.

3.4 MT-PET for Scientific Exaggeration

We use MT-PET to learn from data labeled for both of our formulations of the problem (**T1**, **T2**). In this, the first step is to define PVPs for exaggeration detection (**T1**) and claim strength prediction (**T2**).

To do this, we develop an initial set of PVPs and use PETAL (Schick et al., 2020) to automatically find verbalizers which adequately represent the labels for each task. We then update the patterns manually and re-run PETAL, iterating as such until we find a satisfactory combination of verbalizers and patterns which adequately reflect the task. Additionally, we ensure that the patterns between **T1** and **T2** are roughly equivalent. This yields 2

patterns for each task, provided in Table 1, and verbalizers given in Table 2. The verbalizers found by PETAL capture multiple aspects of the task labels, selecting words such as “mistaken,” “wrong,” and “artificial” for exaggeration, “preliminary” and “conditional” for downplaying, and multiple levels of strength for strength detection such as “estimated” (correlational), “cautious” (conditional causal), and “proven” (direct causal).

For unlabeled data, we start with unlabeled pairs of full text press releases and abstracts. As we are concerned with detecting exaggeration in the primary conclusions, we first train a classifier based on single task PET for conclusion detection using a set of seed data. The patterns and verbalizers we use for conclusion detection are given in Table 3 and Table 4. After training the conclusion detection model, we apply it to the press releases and abstracts, choosing the sentence from each with the maximum score $s_{\mathbf{p}}(1|\mathbf{x})$.

4 Data Collection

One of the main contributions of this work is a curated benchmark dataset for scientific exaggeration detection. Labeled datasets exist for the related task of claim strength detection in scientific abstracts and press releases (Yu et al., 2020, 2019), but these data are from press releases and abstracts which are unrelated (i.e. the given press releases are not written about the given abstracts), making them unsuitable for benchmarking exaggeration detection. Given this, we curate a dataset of paired sentences from abstracts and associated press releases, labeled by experts for exaggeration based on their claim strength. We then collect a large set of unlabeled press release/abstract pairs useful for semi-supervised learning.

4.1 Gold Data

The gold test data used in this work are from Sumner et al. (2014) and Bratton et al. (2019), who annotate scientific papers, their abstracts, and asso-

Pattern	Label	Verbalizers
$P_{T_1}^0$	Downplays	preliminary, competing, uncertainties
	Same	following, explicit
	Exaggerates	mistaken, wrong, hollow, naive, false, lies
$P_{T_1}^1$	Downplays	hypothetical, theoretical, conditional
	Same	identical
	Exaggerates	mistaken, wrong, premature, fantasy, noisy, artificial
$P_{T_2}^*$	NA	sufficient, enough, authentic, medium
	Correlational	inferred, estimated, calculated, borderline, approximately, variable, roughly
	Cond. Causal	cautious, premature, uncertain, conflicting, limited
	Causal	touted, proven, replicated, promoted, distorted

Table 2: Verbalizers for PVPs from both **T1** and **T2**. Verbalizers are obtained using PETAL (Schick et al., 2020), starting with the top 10 verbalizers per label and then manually filtering out words which do not make sense with the given labels.

ciated press releases along several dimensions to characterize how press releases exaggerate papers. The original data consists of 823 pairs of abstracts and press releases. The 462 pairs from Sumner et al. (2014) have been used in previous work to test claim strength prediction (Li et al., 2017), but the data, which contain press release and abstract conclusion sentences that are mostly paraphrases of the originals, are used as is.

We focus on the annotations provided for claim strength. The annotations consist of six labels which we map to the four labels defined in Li et al. (2017). The labels and their meaning are given in Table 5. This gives a claim strength label l_ρ for the press release and l_γ for the abstract. The final exaggeration label is then defined as follows:

$$l_e = \begin{cases} 0 & l_\rho < l_\gamma \\ 1 & l_\rho = l_\gamma \\ 2 & l_\rho > l_\gamma \end{cases}$$

As the original abstracts in the study are not provided, we automatically collect them using the Semantic Scholar API.³ We perform a manual inspection of abstracts to ensure the correct ones

³<https://api.semanticscholar.org/>

Name	Pattern
$P_0(x)$	[MASK]: a
$P_1(x)$	[MASK] - a
$P_2(x)$	“[MASK]” statement: a
$P_3(x)$	a ([MASK])
$P_4(x)$	([MASK]) a
$P_5(x)$	[Type: [MASK]] a

Table 3: Patterns for conclusion detection.

Label	Verbalizers
0	Text
1	Conclusion

Table 4: Verbalizers for PVPs for conclusion detection.

are collected, discarding missing and incorrect abstracts. Gold conclusion sentences are obtained by sentence tokenizing abstracts using SciSpaCy (Neumann et al., 2019) and finding the best matching sentence to the provided paraphrase in the data using ROUGE score (Lin, 2004). We then manually fix sentences which do not correspond to a single sentence from the abstract. Gold press release sentences are gathered in the same way from the provided press releases.

This results in a dataset of 663 press release/abstract pairs labeled for claim strength and exaggeration. The label distribution is given in Table 6. We randomly sample 100 of these instances as training data for few shot learning (**T1**), leaving 553 instances for testing. Additionally, we create a small training set of 1,138 sentences labeled for whether or not they are the main conclusion sentence of the press release or abstract. This data is used in the first step of MT-PET to identify conclusion sentences in the unlabeled pairs.

For **T2** we use the data from Yu et al. (2020, 2019). Yu et al. (2019) create a dataset of 3,061 conclusion sentences labeled for claim strength from structured PubMed abstracts of health observational studies with conclusion sections of 3 sentences or less. Yu et al. (2020) then annotate statements from press releases from EurekAlert. The selected data are from the title and first two sentences of the press releases, as Sumner et al. (2014) note that most press releases contain their main conclusion statements in these sentences, following an inverted pyramid structure common in journalism (Pöttker, 2003). Both studies use the

Sumner et al. (2014)	Description	Li et al. (2017)	Description
0	No relationship mentioned	-	-
1	Statement of no relationship	0	Statement of no relationship
2	Statements of correlation	1	Statement of correlation
3	Ambiguous statement of relationship		
4	Conditional statement of causation	2	Conditional statement of causation
5	Statement of “can”		
6	Statements of causation	3	Statement of causation

Table 5: Claim strength labels and their meaning from the original data in Sumner et al. (2014) and Bratton et al. (2019) and the mappings to the labels from Li et al. (2017). We use the labels from Li et al. (2017) in this study, including for deriving the exaggeration labels.

Label	Count
Downplays	113
Same	406
Exaggerates	144

Table 6: Number of labels per class for benchmark exaggeration detection data.

Method	P	R	F1
Supervised	28.06	33.10	29.05
PET	41.90	39.87	39.12
MT-PET	47.80	47.99	47.35

Table 7: Results for exaggeration detection with paired conclusion sentences from abstracts and press releases (**T1**). MT-PET uses 200 sentences for strength classification, 100 each from press releases and abstracts.

labeling scheme from Li et al. (2017) (see Table 5). The final data contains 2,076 labeled conclusion statements. From these two datasets, we select a random stratified sample of 4,500 instances for training in our full-data experiments, and subsample 200 for few-shot learning (100 from abstracts and 100 from press releases).

4.2 Unlabeled Data

We collect unlabeled data from ScienceDaily,⁴ a science reporting website which aggregates and re-releases press releases from a variety of sources. To do this, we crawl press releases from ScienceDaily via the Internet Archive Wayback Machine⁵ between January 1st 2016 and January 1st 2020 using Scrapy.⁶ We discard press releases without paper

⁴<https://www.sciencedaily.com/>

⁵<https://archive.org/web/>

⁶<https://scrapy.org/>

DOIs and then pair each press release with a paper abstract by querying for each DOI using the Semantic Scholar API. This results in an unlabeled set of 7,741 press release/abstract pairs. Additionally, we use only the title, lead sentence, and first three sentences of each press release.

5 Experiments

Our experiments are focused on the following primary research questions:

- **RQ1:** Does MT-PET improve over PET for scientific exaggeration detection?
- **RQ2:** Which formulation of the problem leads to the best performance?
- **RQ3:** Does few-shot learning performance approach the performance of models trained with many instances?
- **RQ4:** What are the challenges of scientific exaggeration prediction?

We experiment with the following model variants:

- **Supervised:** A fully supervised setting where only labeled data is used.
- **PET:** Standard single-task PET.
- **MT-PET:** We run MT-PET with data from one task formulation as the main task and the other formulation as the auxiliary task.

We perform two evaluations in this setup: one with **T1** as the main task and one with **T2**. For **T1**, we use the 100 expert annotated instances with paired press release and abstract sentences labeled for exaggeration (200 sentences total). For **T2**, we use 100 sentences from the press data from Yu et al. (2020) and 100 sentences from the abstract data in Yu et al. (2019) labeled for claim

Method	T2 , T1	P	R	F1	Press F1	Abstract F1
Supervised	200,0	49.28	51.07	49.03	54.78	59.41
PET	200,0	55.76	58.58	56.57	63.56	62.76
MT-PET	200,100	56.68	60.13	57.44	64.72	63.27
Supervised	4500,0	58.20	59.99	58.66	63.26	67.26
PET	4500,0	59.53	61.84	60.45	64.20	64.92
MT-PET	4500,100	60.09	62.68	61.11	63.93	64.69
PET+in domain MLM	200,100	57.18	60.12	58.06	64.29	62.69
PET+in domain MLM	4500,100	59.87	62.33	60.85	64.10	64.73

Table 8: Results on exaggeration detection via strength classification (**T2**) with varying numbers of instances. MT-PET uses 100 instances from paired press and abstract sentences (200 sentences total).

strength. We use RoBERTa base (Liu et al., 2019) from the HuggingFace Transformers library (Wolf et al., 2020) as the main model, and set α_m to be 1, and $\alpha_a = \min(2, \frac{|D_m|}{|D_a|})$. All methods are evaluated using macro-F1 score, and results are reported as the average performance over 5 random seeds.

5.1 Performance Evaluation

We first examine the performance with **T1** as the base task (see Table 7). In a purely supervised setting, the model struggles to learn and mostly predicts the majority class. Basic PET yields a substantial improvement of 10 F1 points, with MT-PET further improving upon this by another 8 F1 points. Accordingly, we conclude that training with auxiliary task data provides much benefit for scientific exaggeration detection in the **T1** formulation.

We next examine performance with **T2** (strength classification) as the main task in both few-shot and full data settings (see Table 8). In terms of base performance, the model can predict exaggeration better than **T1** in a purely supervised setting. For PET and MT-PET, we see a similar trend; with 200 instances for **T2**, PET improves by 7 F1 points over supervised learning, and MT-PET improves on this by a further 0.9 F1 points. Additionally, MT-PET improves performance on the individual tasks of predicting the claim strength of conclusions in press releases and scientific abstracts with 200 examples. While less dramatic, we still see gains in performance using PET and MT-PET when 4,500 instances from **T2** are used, despite the fact that there are still only 100 instances from **T1**. We also test if the improvement in performance is simply due to training on more in-domain data (“PET+in domain MLM” in Table 8). We observe gains for exaggeration detection using masked language modeling on data from **T1**, but MT-PET still per-

forms better at classifying the strength of claims in press releases and abstracts when 200 training instances from **T2** are used.

RQ1 Our results indicate that MT-PET does in fact improve over PET for both training setups. With **T1** as the main task and **T2** as the auxiliary task, we see that performance is substantially improved, demonstrating that learning claim strength prediction helps produce soft-labeled training data for *exaggeration detection*. Additionally, we find that the reverse holds with **T2** as main task and **T1** as auxiliary task. As performance can also be improved via masked language modeling on data from **T1**, this indicates that some of the performance improvement could be due to including data closer to the test domain. However, our error analysis in subsection 5.2 shows that these methods improve model performance on different types of data.

RQ2 We find that **T2** is better suited for scientific exaggeration detection in this setting, however, with a couple of caveats. First, the final exaggeration label is based on expert annotations for claim strength, so clearly claim strength prediction will be useful in this setup. Additionally, the task may be more forgiving here, as only the direction needs to be correct and not necessarily the final strength label (i.e. predicting ‘0’ for the abstract and any of ‘1,’ ‘2,’ or ‘3’ for the press release label will result in an exaggeration label of ‘exaggerates’).

RQ3 We next examine the learning dynamics of our few-shot models with different amounts of training data (see Figure 3), comparing them to MT-PET to understand how well it performs compared to settings with more data. MT-PET with only 200 samples is highly competitive with purely supervised learning on 4,500 samples (57.44 vs. 58.66).

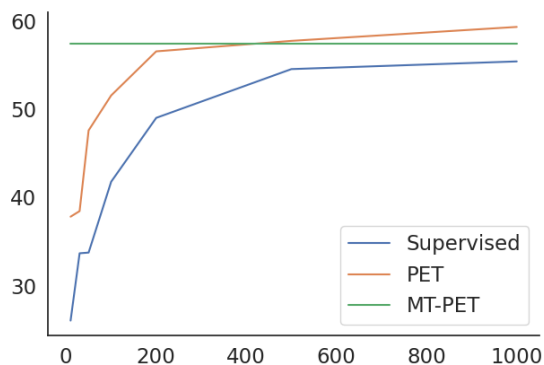


Figure 3: Learning curve for supervised learning and PET compared to performance of MT-PET using 200 instances from **T2** and 100 from **T1**.

Additionally, MT-PET performs at or above supervised performance up to 1000 input samples, and at or above PET up to 500 samples, again using only 200 samples from **T2** and 100 from **T1**.

5.2 Error Analysis

RQ4 Finally, we try to understand the difficulty of scientific exaggeration detection by observing where models succeed and fail (see Figure 4). The most difficult category of examples to predict involve direct causal claims, particularly exaggeration and downplaying when one document is a direct causal claim and the other an indirect causal claim (‘CON->CAU’, ‘CAU->CON’). Also, it is challenging to predict when both the press release and abstract conclusions are directly causal.

The models have the easiest time predicting when both statements involve correlational claims, and exaggerations involving correlational claims from abstracts. We also observe that MT-PET helps the most for the most difficult category: causal claims (see Figure 5 in Appendix A). The model is particularly better at differentiating when a causal claim in an abstract is *downplayed* by a press release. It is also better at identifying correlational claims than PET, where many claims involve association statements such as ‘linked to,’ ‘could predict,’ ‘more likely,’ and ‘suggestive of.’

The model trained with MLM on data from **T1** also benefits causal statement prediction, but mostly for when both statements are causal, whereas MT-PET sees more improvement for pairs where one causal statement is exaggerated or downplayed by another (see Figure 6 in Appendix A). This suggests that training with the patterns from

T1 helps the model to differentiate direct causal claims from weaker claims, while MLM training mostly helps the model to understand better how direct causal claims are written. We hypothesize that combining the two methods would lead to mutual gains.

6 Related Work

6.1 Scientific Misinformation Detection

Misinformation detection focuses on a variety of problems, including fact verification (Thorne et al., 2018; Augenstein et al., 2019), check-worthiness detection (Wright and Augenstein, 2020; Nakov et al., 2021), stance (Augenstein et al., 2016; Baly et al., 2018; Hardalov et al., 2021a) and clickbait detection (Potthast et al., 2018). While most work has focused on social media and general domain text, recent work has begun to explore different problems in detecting misinformation in scientific text such as SciFact (Wadden et al., 2020) and CiteWorth (Wright and Augenstein, 2021), as well as related tasks such as summarization (DeYoung et al., 2021; Dangovski et al., 2021).

Most work on scientific exaggeration detection has focused on flagging when the primary finding of a scientific paper has been exaggerated by a press release or news article (Sumner et al., 2014; Bratton et al., 2019; Yu et al., 2020, 2019; Li et al., 2017). Sumner et al. (2014) and Bratton et al. (2019) manually label pairs of press releases and scientific papers on a wide variety of metrics, finding that one third of press releases contain exaggerated claims, and 40% contain exaggerated advice. Li et al. (2017) is the first study into automatically predicting claim strength, using the data from Sumner et al. (2014) as a small labeled dataset. Yu et al. (2019) and Yu et al. (2020) extend this by building larger datasets for claim strength prediction, performing an analysis of a large set of unlabeled data to estimate the prevalence of claim exaggeration in press releases. Our work improves upon this by providing a more realistic task formulation of the problem, consisting of paired press releases and abstracts, as well as curating both labeled and unlabeled data to evaluate methods in this setting.

6.2 Learning from Task Descriptions

Using natural language to perform zero and few-shot learning has been demonstrated on a number of tasks, including question answering (Radford et al.), text classification (Puri and Catanzaro,

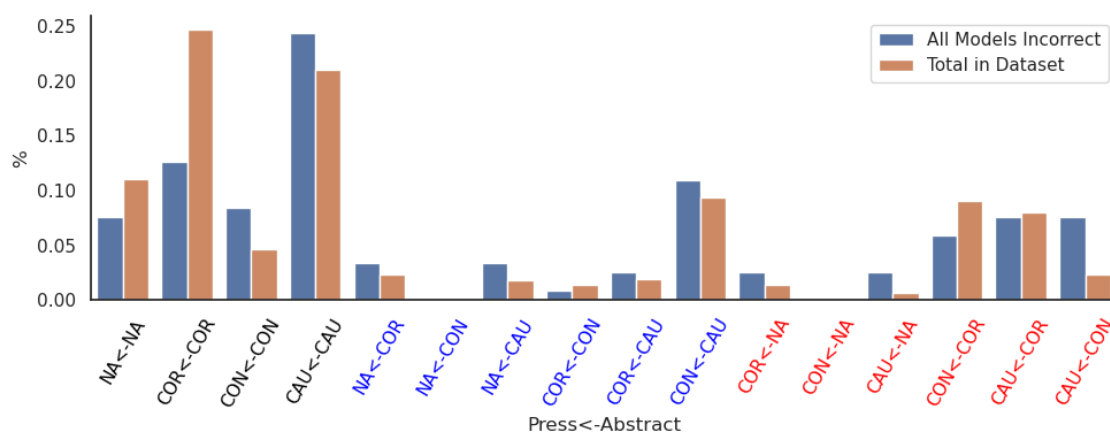


Figure 4: Proportion of examples by label which all models predict incorrectly.

2019), relation extraction (Bouraoui et al., 2020) and stance detection (Hardalov et al., 2021b,c). Methods of learning from task descriptions have been gaining more popularity since the creation of GPT-3 (Brown et al., 2020). Raffel et al. (2020) attempt to perform this with smaller language models by converting tasks into natural language and predicting tokens in the vocabulary. Schick and Schütze (2020a) propose PET, a method for few shot learning which converts tasks into cloze-style QA problems which can be solved by a pretrained language model in order to provide soft-labels for unlabeled data. We build on PET, showing that complementary cloze-style QA tasks can be trained on simultaneously to improve few-shot performance on scientific exaggeration detection.

7 Conclusion

In this work, we present a formalization of and investigation into the problem of scientific exaggeration detection. As data for this task is limited, we develop a gold test set for the problem and propose MT-PET, a semi-supervised approach based on PET, to solve it with limited training data. We find that MT-PET helps in the more difficult cases of identifying and differentiating direct causal claims from weaker claims, and that the most performant approach involves classifying and comparing the individual claim strength of statements from the source and target documents. The code and data for our experiments can be found online⁷. Future work should focus on building more resources e.g. datasets for exploring scientific exaggeration detec-

⁷<https://github.com/copenlu/scientific-exaggeration-detection>

tion, including data from multiple domains beyond health science. Finally, it would be interesting to explore how MT-PET works on combinations of more general NLP tasks, such as question answering and natural language inference or part-of-speech tagging and named entity recognition.

Acknowledgements



The research documented in this paper has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 801199.

Broader Impact Statement

Being able to automatically detect whether a press release exaggerates the findings of a scientific article could help journalists write press releases, which are more faithful to the scientific articles they are describing. We further believe it could benefit the research community working on fact checking and related tasks, as developing methods to detect subtle differences in a statement’s veracity is currently understudied.

On the other hand, as our paper shows, this is currently still a very challenging task, and thus, the resulting models should only be applied in practice with caution. Moreover, it should be noted that the predictive performance results reported in this paper are for press releases written by science journalists – one could expect worse results for press releases which more strongly simplify scientific articles.

References

- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. [MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims](#). 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 4685–4697, Hong Kong, China. Association for Computational Linguistics.
- Isabelle Augenstein, Tim Rocktäschel, Andreas Vlachos, and Kalina Bontcheva. 2016. [Stance Detection with Bidirectional Conditional Encoding](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 876–885, Austin, Texas. Association for Computational Linguistics.
- Ramy Baly, Mitra Mohtarami, James Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018. [Integrating stance detection and fact checking in a unified corpus](#). 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 21–27, New Orleans, Louisiana. Association for Computational Linguistics.
- Zied Bouraoui, José Camacho-Collados, and Steven Schockaert. 2020. [Inducing Relational Knowledge from BERT](#). 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 7456–7463. AAAI Press.
- Luke Bratton, Rachel C Adams, Aimée Challenger, Jacky Boivin, Lewis Bott, Christopher D Chambers, and Petroc Sumner. 2019. The Association Between Exaggeration in Health-Related Science News and Academic Press Releases: A Replication Study. *Wellcome open research*, 4.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Rumen Dangovski, Michelle Shen, Dawson Byrd, Li Jing, Desislava Tsvetkova, Preslav Nakova, and Marin Soljagic. 2021. We Can Explain Your Research in Layman’s Terms: Towards Automating Science Journalism at Scale. In *AAAI 2021*. AAAI Press.
- Jay DeYoung, Iz Beltagy, Madeleine van Zuylen, Bailey Kuehl, and Lucy Lu Wang. 2021. [MS2: Multi-Document Summarization of Medical Studies](#). *CoRR*, abs/2104.06486.
- Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2021a. [A Survey on Stance Detection for Mis- and Disinformation Identification](#).
- Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2021b. [Cross-Domain Label-Adaptive Stance Detection](#). In *Proceedings of EMNLP*. Association for Computational Linguistics.
- Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2021c. [Few-Shot Cross-Lingual Stance Detection with Sentiment-Based Pre-Training](#). *CoRR*.
- Yingya Li, Jieke Zhang, and Bei Yu. 2017. [An NLP Analysis of Exaggerated Claims in Science News](#). In *Proceedings of the 2017 Workshop: Natural Language Processing meets Journalism, NLPmJ@EMNLP, Copenhagen, Denmark, September 7, 2017*, pages 106–111. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A Package for Automatic Evaluation of Summaries. In *Text summarization branches out*, pages 74–81.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *CoRR*, abs/1907.11692.
- Andrew Moore. 2006. Bad Science in the Headlines: Who Takes Responsibility When Science is Distorted in the Mass Media? *EMBO reports*, 7(12):1193–1196.
- Preslav Nakov, Giovanni Da San Martino, Tamer Elsayed, Alberto Barrón-Cedeño, Rubén Míguez, Shaden Shaar, Firoj Alam, Fatima Haouari, Maram Hasanain, Nikolay Babulkov, Alex Nikolov, Gautam Kishore Shahi, Julia Maria Struß, and Thomas Mandl. 2021. [The CLEF-2021 CheckThat! Lab on Detecting Check-Worthy Claims, Previously Fact-Checked Claims, and Fake News](#). In *Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 - April 1, 2021, Proceedings, Part II*, volume 12657 of *Lecture Notes in Computer Science*, pages 639–649. Springer.
- Dorothy Nelkin. 1987. Selling Science: How the Press Covers Science and Technology.

- Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. [ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing](#). In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 319–327, Florence, Italy. Association for Computational Linguistics.
- Martin Potthast, Tim Gollub, Kristof Komlossy, Sebastian Schuster, Matti Wiegmann, Erika Patricia Garces Fernandez, Matthias Hagen, and Benno Stein. 2018. [Crowdsourcing a large corpus of clickbait on Twitter](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1498–1507, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Horst Pöttker. 2003. News and its Communicative Quality: The Inverted Pyramid—When and Why Did it Appear? *Journalism Studies*, 4(4):501–511.
- Raul Puri and Bryan Catanzaro. 2019. [Zero-shot Text Classification With Generative Language Models](#). *CoRR*, abs/1912.10165.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving Language Understanding by Generative Pre-Training. *Technical report*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Timo Schick, Helmut Schmid, and Hinrich Schütze. 2020. [Automatically Identifying Words That Can Serve as Labels for Few-Shot Text Classification](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 5569–5578. International Committee on Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2020a. [Exploiting Cloze Questions for Few-Shot Text Classification and Natural Language Inference](#). *Computing Research Repository*, arXiv:2001.07676.
- Timo Schick and Hinrich Schütze. 2020b. [It’s Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners](#). *Computing Research Repository*, arXiv:2009.07118.
- Petroc Sumner, Solveiga Vivian-Griffiths, Jacky Boivin, Andy Williams, Christos A Venetis, Aimée Davies, Jack Ogden, Leanne Whelan, Bethan Hughes, Bethan Dalton, et al. 2014. The Association Between Exaggeration in Health Related Science News and Academic Press Releases: Retrospective Observational Study. *BMJ*, 349.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a Large-scale Dataset for Fact Extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 809–819. Association for Computational Linguistics.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or Fiction: Verifying Scientific Claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 7534–7550. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 38–45. Association for Computational Linguistics.
- Steven Woloshin and Lisa M Schwartz. 2002. Press Releases: Translating Research Into News. *Jama*, 287(21):2856–2858.
- Steven Woloshin, Lisa M Schwartz, Samuel L Casella, Abigail T Kennedy, and Robin J Larson. 2009. Press Releases by Academic Medical Centers: Not So Academic? *Annals of Internal Medicine*, 150(9):613–618.
- Dustin Wright and Isabelle Augenstein. 2020. Claim Check-Worthiness Detection as Positive Unlabelled Learning. In *Findings of EMNLP*. Association for Computational Linguistics.
- Dustin Wright and Isabelle Augenstein. 2021. Cite-Worth: Cite-Worthiness Detection for Improved Scientific Document Understanding. In *Findings of ACL-IJCNLP*. Association for Computational Linguistics.
- Bei Yu, Yingya Li, and Jun Wang. 2019. Detecting Causal Language Use in Science Findings. In *EMNLP*, pages 4656–4666.
- Bei Yu, Jun Wang, Lu Guo, and Yingya Li. 2020. Measuring Correlation-to-Causation Exaggeration in Press Releases. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4860–4872.

A Error Analysis Plots

Extra plots from our error analysis are given in Figure 5 and Figure 6.

B Reproducibility

B.1 Computing Infrastructure

All experiments were run on a shared cluster. Requested jobs consisted of 16GB of RAM and 4 Intel Xeon Silver 4110 CPUs. We used a single NVIDIA Titan X GPU with 24GB of RAM.

B.2 Average Runtimes

The average runtime performance of each model is given in Table 9. Note that different runs may have been placed on different nodes within a shared cluster.

Setting	T1 , T2	Time
Supervised	100,0	00h01m28s
PET	100,0	00h11m14s
MT-PET	100,200	00h13m05s
Supervised	0,200	00h01m20s
PET	0,200	00h16m22s
MT-PET	100,200	00h18m43s
Supervised	0,4500	00h03m23s
PET	0,4500	00h40m23s
MT-PET	100,4500	00h31m48s

Table 9: Average runtimes for each model (runtimes are taken for the entire run of an experiment).

B.3 Number of Parameters per Model

We use RoBERTa, specifically the base model, for all experiments, which consists of 124,647,170 parameters.

B.4 Validation Performance

As we are testing a few shot setting, we do not use a validation set and only report the final test results.

B.5 Evaluation Metrics

The primary evaluation metric used was macro F1 score. We used the sklearn implementation of `precision_recall_fscore_support` for F1 score, which can be found here: https://scikit-learn.org/stable/modules/generated/sklearn.metrics.precision_recall_fscore_support.html. Briefly:

$$p = \frac{tp}{tp + fp}$$

$$r = \frac{tp}{tp + fn}$$
$$F1 = \frac{2 * p * r}{p + r}$$

where tp are true positives, fp are false positives, and fn are false negatives. Macro F1 is average F1 across all classes.

B.6 Hyperparameters

T1 Hyperparameters Supervised/PET training

We used the following hyperparameters for experiments with **T1** as the main task: Epochs: 10; Batch Size: 4; Learning Rate: 0.00005598; Warmup Steps: 50; Weight decay: 0.001. We also weigh the cross-entropy loss based on the label distribution. These hyperparameters are found by performing a hyperparameter search using 4-fold cross validation on the 100 training examples. The bounds are as follows: Learning rate: [0.000001, 0.0001; Warmup steps: {0, 10, 20, 30, 40, 50, 100}; Batch size: {4, 8}; Weight decay: {0.0, 0.0001, 0.001, 0.01, 0.1}; Epochs: [2, 10].

T2 Hyperparameters Supervised/PET training

Epochs: 10; Batch Size: 4; Learning Rate: 0.00003; Warmup Steps: 50; Weight Decay: 0.001. We also weigh the cross-entropy loss based on the label distribution.

Hyperparameters for Distillation We used the following hyperparameters for distillation (training the final classifier after PET) for both **T1** and **T2** as the main task: Epochs: 3; Batch Size: 4; Learning Rate: 0.00001; Warmup Steps: 200; Weight decay: 0.01; Temperature: 2.0. We also weigh the cross-entropy loss based on the label distribution.

B.7 Data

We build our benchmark test dataset from the studies of Sumner et al. (2014) and Bratton et al. (2019). The original data can be found at <https://figshare.com/articles/dataset/InSciOut/903704> and <https://osf.io/v8qhn/files/>. A link to the test data will be provided upon acceptance of the paper (and included in the supplemental material). Claim strength data from Yu et al. (2019) for abstracts can be found at https://github.com/junwang4/correlation-to-causation-exaggeration/blob/master/data/annotated_pubmed.csv.




Figure 5: Number of instances that each model predicted correctly which the supervised model predicted incorrectly.




Figure 6: Number of instances that each model predicted correctly which PET predicted incorrectly.

Claim strength data for press releases from Yu et al. (2020) can be found at https://github.com/junwang4/correlation-to-causation-exaggeration/blob/master/data/annotated_eureka.csv