

How Did This Get Funded?!

Automatically Identifying Quirky Scientific Achievements

Chen Shani*, Nadav Borenstein*, Dafna Shahaf

The Hebrew University of Jerusalem

{chenxshani, nadav.borenstein, dshahaf}@cs.huji.ac.il

Abstract

Humor is an important social phenomenon, serving complex social and psychological functions. However, despite being studied for millennia humor is computationally not well understood, often considered an AI-complete problem.

In this work, we introduce a novel setting in humor mining: automatically detecting funny and unusual scientific papers. We are inspired by the Ig Nobel prize, a satirical prize awarded annually to celebrate funny scientific achievements (example past winner: “Are cows more likely to lie down the longer they stand?”). This challenging task has unique characteristics that make it particularly suitable for automatic learning.

We construct a dataset containing thousands of funny papers and use it to learn classifiers, combining findings from psychology and linguistics with recent advances in NLP. We use our models to identify potentially funny papers in a large dataset of over 630,000 articles. The results demonstrate the potential of our methods, and more broadly the utility of integrating state-of-the-art NLP methods with insights from more traditional disciplines.

1 Introduction

Humor is an important aspect of the way we interact with each other, serving complex social functions (Martineau, 1972). Humor can function either as a lubricant or as an abrasive: it can be used as a key for improving interpersonal relations and building trust (Wanzer et al., 1996; Wen et al., 2015), or help us work through difficult topics. It can also aid in breaking taboos and holding power to account. Enhancing the humor capabilities of computers has tremendous potential to better understand interactions between people, as well as build more natural human-computer interfaces.

Nevertheless, computational humor remains a long-standing challenge in AI; It requires complex language understanding, manipulation capabilities, creativity, common sense, and empathy. Some even claim that computational humor is an AI-complete problem (Stock and Strapparava, 2002).

As humor is a broad phenomenon, most works on computational humor focus on *specific* humor types, such as knock-knock jokes or one-liners (Mihalcea and Strapparava, 2006; Taylor and Mazlack, 2004). In this work, we present a novel humor recognition task: identifying quirky, funny scientific contributions. We are inspired by the **Ig Nobel prize**¹, a satiric prize awarded annually to ten scientific achievements that “first make people laugh, and then think”. Past Ig Nobel winners include “Chickens prefer beautiful humans” and “Beauty is in the eye of the beer holder: People who think they are drunk also think they are attractive”.

Automatically identifying candidates for the Ig Nobel prize provides a unique perspective on humor. Unlike most humor recognition tasks, the humor involved is sophisticated, and requires common sense, as well as specialized knowledge and understanding of the scientific culture. On the other hand, this task has several characteristics rendering it attractive: the funniness of the paper can often be recognized from its *title* alone, which is short, with simple syntax and no complex narrative structure (as opposed to longer jokes). Thus, this is a relatively clean setting to explore our methods.

We believe humor in science is also particularly interesting to explore, as humor is strongly tied to creativity. Quirky contributions could sometimes indicate fresh perspectives and pioneering attempts to expand the frontiers of science. For example, Andre Geim won an Ig Nobel in 2000 for levitating a frog using magnets and a Nobel Prize in Physics

*Equal contribution

¹improbable.com/ig-about

in 2010. The Nobel committee explicitly attributed the win to his playfulness ([The Royal Swedish Academy of Science, 2010](#)).

Our contributions are:

- We formulate a novel humor recognition task in the scientific domain.
- We construct a dataset containing thousands of funny scientific papers.
- We develop multiple classifiers, combining findings from psychology and linguistics with recent NLP advances. We evaluate them both on our dataset and in a real-world setting, identifying potential Ig Nobel candidates in a large corpus of over 0.6M papers.
- We devise a rigorous, data-driven way to aggregate crowd workers’ annotations for *subjective* questions.
- We release data and code².

Beyond the tongue-in-cheek nature of our application, we more broadly wish to promote combining data-driven research with more-traditional works in areas such as psychology. We believe insights from such fields could complement machine learning models, improving performance as well as enriching our understanding of the problem.

2 Related Work

Humor in the Humanities. A large body of theoretical work on humor stems from linguistics and psychology. [Ruch \(1992\)](#) divided humor into three categories: incongruity, sexual, and nonsense (and created a three-dimensional humor test to account for them). Since our task is to detect humor in scientific contributions, we believe that the third category can be neglected under the assumption that no-nonsense article would (or at least, should) be published (notable exception: the Sokal hoax ([Sokal, 1996](#))).

The first category, incongruity, was first fully conceptualized by Kant in the eighteenth century ([Shaw, 2010](#)). The well-agreed extensions to incongruity theory are the linguistics incongruity resolution model and semantic script theory of humor ([Suls, 1972](#); [Raskin, 1985](#)). Both state that if a situation ended in a manner that contradicted our prediction (in our case, the title contains an unexpected term) and there exists a different, less likely rule to explain it – the result is a humorous experience. Simply put, the source of humor lies in

violation of expectations. Example Ig Nobel winners include: “Will humans swim faster or slower in syrup?” and “Coordination modes in the multi-segmental dynamics of hula hooping”.

The second category, sex-related humor is also common among Ig Nobel winning papers. Examples include: “Effect of different types of textiles on sexual activity. Experimental study” and “Magnetic resonance imaging of male and female genitals during coitus and female sexual arousal”.

Humor Detection in AI. Most computational humor detection work done in the context of AI relies on supervised or semi-supervised methods and focuses on specific, narrow, types of jokes or humor.

Humor detection is usually formulated as a binary text classification problem. Example domains include knock-knock jokes ([Taylor and Mazlack, 2004](#)), one-liners ([Miller et al., 2017](#); [Simpson et al., 2019](#); [Liu et al., 2018](#); [Mihalcea and Strapparava, 2005](#); [Blinov et al., 2019](#); [Mihalcea and Strapparava, 2006](#)), humorous tweets ([Maronikolakis et al., 2020](#); [Donahue et al., 2017](#); [Ortega-Bueno et al., 2018](#); [Zhang and Liu, 2014](#)), humorous product reviews ([Ziser et al., 2020](#); [Reyes and Rosso, 2012](#)), TV sitcoms ([Bertero and Fung, 2016](#)), short stories ([Wilmot and Keller, 2020](#)), cartoons captions ([Shahaf et al., 2015](#)), and even “That’s what she said” jokes ([Hossain et al., 2017](#); [Kiddon and Brun, 2011](#)). Related tasks such as irony, sarcasm and satire have also been explored in similarly narrow domains ([Davidov et al., 2010](#); [Reyes et al., 2012](#); [Ptáček et al., 2014](#)).

3 Problem Formulation and Dataset

Our goal in this paper is to automatically identify candidates for the Ig Nobel prize. More precisely, to automatically detect humor in scientific papers.

First, we consider the question of input to our algorithm. [Sagi and Yechiam \(2008\)](#) found a strong correlation between funny *title* and humorous *subject* in scientific papers. Motivated by this correlation, we manually inspected a subset of Ig Nobel winners. For the vast majority of them, reading the title was enough to determine whether it is funny; very rarely did we need to read the abstract, let alone the full paper. Typical past winners’ titles include “Why do old men have big ears?” and “If you drop it, should you eat it? Scientists weigh in on the 5-second rule”. An example of a non-informative title is “Pouring flows”, a paper calculating the optimal way to dunk a biscuit in a cup of tea.

²github.com/nadavborenstein/Iggy

Based on this observation, we decided to focus on the papers’ *titles*. More formally: Given a title t of an article, our goal is to learn a binary function $\varphi(t) \rightarrow \{0, 1\}$, reflecting whether the paper is humorous, or ‘Ig Nobel-worthy’. The main challenge, of course, lies in the construction of φ .

To take a data-driven approach to tackle this problem, we crafted a first-of-its-kind dataset containing titles of funny scientific papers². We started from the 211 Ig Nobel winners. Next, we manually collected humorous papers from online forums and blogs³, resulting in 1,707 papers. We manually verified all of these papers can be used as positive examples. In Section 6 we give more indication these papers are indeed useful for our task.

For negative examples, we randomly sampled 1,707 titles from Semantic Scholar⁴ (to obtain a balanced dataset). We then classify each paper into one of the following scientific fields: neuroscience, medicine, biology, or exact sciences⁵. We balanced the dataset in a per-field manner. While some of these randomly sampled papers could, in principle, be funny, the vast majority of scientific papers are not (we validated this assumption through sampling).

4 Humor-Theory Inspired Features

In deep learning, architecture engineering largely took the place of feature engineering. One of the goals of our work is to evaluate the value of features inspired by domain experts. In this section, we describe and formalize 127 features implementing insights from humor literature. To validate the predictive power of the features that require training, we divide our data to train and test sets (80%/20%). We now describe the four major feature families.

4.1 Unexpected Language

Research suggests that *surprise* is an important source of humor (Raskin, 1985; Suls, 1972). Indeed, we notice that titles of Ig Nobel winners often include an unexpected term or unusual language, e.g.: “On the rheology of *cats*”, “Effect of *coke* on sperm motility” and “*Pigeons*’ discrimination of *paintings* by Monet and Picasso”. To quantify unexpectedness, we create several different language-models (LMs):

³E.g., [reddit.com/r/ScienceHumour](https://www.reddit.com/r/ScienceHumour), popsci.com/read/funny-science-blog, goodsciencewriting.wordpress.com

⁴api.semanticscholar.org/corpus/

⁵Using scimagojr.com to map venues to fields.

N-gram Based LMs. We train simple N-gram LMs with $n \in \{1, 2, 3\}$ on two corpora – 630,000 titles from Semantic Scholar, and 231,600 one-line jokes (Moudgil, 2016).

Syntax-Based LMs. Here we test the hypothesis that humorous text has more surprising *grammatical structure* (Oaks, 1994). We replace each word in our Semantic Scholar corpus with its corresponding part-of-speech (POS) tag⁶. We then trained N-gram based LMs ($n \in \{1, 2, 3\}$) on this corpus.

Transformer-Based LMs. We use three different Transformers based (Vaswani et al., 2017) models: 1) BERT (Devlin et al., 2018) (pre-trained on Wikipedia and the BookCorpus), 2) SciBERT (Beltagy et al., 2019), a variant of BERT optimized on scientific text from Semantic Scholar, and 3) GPT-2 (Radford et al., 2019), a large Transformer-based LM, trained on a dataset of 8M web pages. We fine-tuned GPT-2 on our Semantic Scholar corpora (details in Appendix C.1).

Using the LMs. For each word in a title, we compute the word’s perplexity. For the N-gram LMs and GPT-2, we compute the probability to see the word given the previous words in the sentence ($n - 1$ previous words in the case of the N-gram models and all the previous words in the case of GPT-2). For the BERT-based models, we compute the masked loss of the word given the sentence. For each title, we computed the mean, maximum, and variance of the perplexity across all words in the title.

4.2 Simple Language

Inspired by previous findings (Ruch, 1992; Gultchin et al., 2019), we hypothesize that titles of funny papers tend to be simpler (e.g., the past Ig Nobel winners: “Chickens prefer beautiful humans” and “Walking with coffee: Why does it spill?”). We utilize several simplicity measures:

Length. Short titles and titles containing many short words tend to be simpler. We compute title length and word lengths (mean, maximum, and variance of word lengths in the title).

Readability. We use the automated readability index (Smith and Senter, 1967).

Age of Acquisition (AoA). A well-established measure for word’s difficulty in psychology (Brybaert and Biemiller, 2017), denoting word’s difficulty by the age a child acquires it. We compute mean, maximum and variance AoA.

⁶Obtained using NLTK (nltk.org)

AoA and Perplexity. Many basic words can be found in serious titles (e.g., ‘water’ in a hydraulics paper). Funny titles, however, contain simple words which are also *unexpected*. Thus, we combine AoA with perplexity. We compute word perplexity using the Semantic Scholar N-gram LMs and divide it by AoA. Higher values correspond to simpler and unexpected words. We compute the mean, maximum, minimum, and variance.

4.3 Crude Language

According to relief theory, crude and scatological connotations are often considered humorous (Shurcliff, 1968) (e.g., the Ig Nobel winners “Duration of urination does not change with body size”, “Acute management of the zipper-entrapped penis”).

We trained a Naive Bayes SVM (Wang and Manning, 2012) classifier over a dataset of toxic and rude Wikipedia comments (Zafar, 2018), and compute title probability to be crude. Similar to the AoA feature, we believe that crude words should also be unexpected to be considered funny. As before, we divide perplexity by the word’s probability of being benign. Higher values correspond to crude and unexpected words. We compute the mean, maximum, minimum, and variance.

4.4 Funny Language

Some words (e.g., nincompoop, razzmatazz) are inherently funnier than others (due to various reasons surveyed by Gultchin et al. (2019)). It is reasonable that the funniness of a title is correlated with the funniness of its words. We measure funniness using the model of Westbury and Hollis (2019), quantifying noun funniness based on humor theories and human ratings. We measure the funniness of each noun in a title. We also multiplied perplexity and funniness (for funny and unexpected) and use the mean, maximum, minimum, and variance.

4.5 Feature Importance

As a first reality check, we plotted the distribution of our features between funny and not-funny papers (see Appendix A.1 for representative examples). For example, we hypothesized that titles of funny papers might be linguistically similar to one-liners, and indeed we saw that the one-liner LM assigns lower perplexity to funny papers. Similarly, we saw a difference between the readability scores.

To measure the predictive power of our literature-inspired features, we use the Wilcoxon signed-rank

Feature	Wilcoxon value	P-value
Unexpected Language		
Avg. Semantic Scholar 2-gram LM	4850	3.6e-39
Avg. POS 2-gram LM	18926	3e-7
Avg. one-liners 2-gram LM	6919	9.2e-33
Avg. GPT-2 LM	7421	2.7e-31
Avg. BERT LM	17153	9e-10
Simple Language		
Readability	8931	8.6e-26
Title’s length	18493	2.4e-6
Avg. AoA values	16768	2.2e-10
Avg. AoA + 2-gram LM	4882	4.6e-39
Crude Language		
Crudeness classifier	17423	2.3e-9
Avg. crudeness + 2-gram LM	4755	1.8e-39
Funny Language		
Avg. funny nouns model	20101	1.7e-5
Avg. funny nouns + 2-gram LM	8886	3.2e-27

Table 1: Wilcoxon and p-values for representative features using our dataset (tested differentiating ability between funny and serious papers). Combining perplexity with other features seems particularly beneficial.

test⁷ (see Table 1). Interestingly, all feature families include useful features. Combining perplexity with other features (e.g., surprising and simple words) was especially prominent. In the next sections, we describe how we use those features to train models for detecting Ig Nobel worthy papers.

5 Models

We can now create models to automatically detect scientific humor. As mentioned in Section 4, one of our goals in this paper is to compare between the NLP SOTA huge-models approach and the literature-inspired approach. Thus, we trained a binary multi-layer perceptron (MLP) classifier using our dataset (described in Section 3, see reproducibility details in Appendix C.2), receiving as input the 127 features from Section 4. We named this classifier ‘**Iggy**’, after the Ig Nobel prize.

⁷A non-parametric paired difference test used to assess whether the mean ranks of two related samples differ.

As baselines representing the contemporary NLP approach (requiring huge compute and training data), we used **BERT** (Devlin et al., 2018) and **SciBERT** (Beltagy et al., 2019), which is a BERT variant optimized on scientific corpora, rendering it potentially more relevant for our task. We fine-tuned SciBERT and BERT for Ig Nobel classification using our dataset (see Appendix C.3 for implementation details).

We also experimented with two models combining **BERT/SciBERT with our features** (see Figure 6 in Appendix C.4), denoted as $BERT^f$ / $SciBERT^f$. In the spirit of the original BERT paper, we added two linear layers on top of the models and used a standard cross-entropy loss. The input to this final MLP is the concatenation of two vectors: our features’ embedding and the last hidden vector from BERT/SciBERT ([CLS]). See Appendix C.4 for implementation details.

For the sake of completeness, we note that we also conducted exploratory experiments with simple syntactic baselines (title length, maximal word length, title containing a question, title containing a colon) as well as BERT trained on sarcasm detection⁸. None of these baselines was strong enough on its own. We note that the colon-baseline tended to catch smart-aleck titles, but the topic was not necessarily funny. The sarcasm baseline achieved near guess-level accuracy (0.482), emphasizing the distinction between the two humor tasks.

6 Evaluation on the Dataset

We first evaluate the five models (Iggy, SciBERT, BERT, $SciBERT^f$ and $BERT^f$) on our labeled dataset in terms of general accuracy and Ig Nobel retrieval ability. As naive baselines, we added two bag of words (BoW) based classifiers: random forest (RF) and logistic regression (LR).

Accuracy. We randomly split the dataset to train, development, and test sets (80%–10%–10%), and used the development set to tune hyper-parameters (e.g., learning rate, number of training epochs). Table 2 summarizes the results. We note that all five models achieve very high accuracy scores and that the simple BoW models fall behind. This gives some indication about the inherent difficulty of the task. Both features-based Iggy and BERT-based models outperform simple baseline. $SciBERT^f$ outperforms the other models across all measures.

⁸kaggle.com/raghavkhemka/sarcasm-detection-using-bert-92-accuracy

Model	Accuracy	Precision	Recall
Iggy	0.897	0.901	0.893
SciBERT	0.910	0.911	0.911
$SciBERT^f$	0.922	0.919	0.926
BERT	0.904	0.906	0.893
$BERT^f$	0.900	0.899	0.902
RF	0.761	0.746	0.796
LR	0.781	0.754	0.837

Table 2: Accuracy of the different models on our dataset using cross validation with k=5. $SciBERT^f$ outperforms.

Model	Accuracy	Precision	Recall
Iggy	0.884	0.913	0.848
SciBERT	0.882	0.909	0.848
$SciBERT^f$	0.903	0.921	0.882
BERT	0.863	0.905	0.810
$BERT^f$	0.903	0.921	0.882
RF	0.713	0.708	0.725
LR	0.765	0.755	0.787

Table 3: Accuracy of the different models on our Ig-Nobel retrieval test set. The combination of SOTA pre-trained models and our features is superior.

Ig Nobel Winners Retrieval. Our positive examples consist of 211 Ig Nobel winners and additional 1,496 humorous papers found on the web. Thus, the portion of real Ig Nobel winning papers in our data is relatively small. We now measure whether our web-originated papers serve as a good proxy for Ig Nobel winners. Thus, we split the dataset differently: the test set consists of the 211 Ig Nobel winners, plus a random sample of 211 negative titles (slightly increasing the test set size to 12%). Train set consists of the remaining 2,992 papers. This experiment follows our initial inspiration of finding Ig Nobel-worthy papers, as we test our models’ ability to retrieve only the real winners.

Table 3 demonstrate that our web-based funny papers are indeed a good proxy for Ig Nobel winners. Similar to the previous experiment, the combination of SOTA pretrained models with literature based features is superior.

Based on both experiments, we conclude that our features are indeed informative for our Ig Nobel-worthy papers detection task.

7 Evaluation “in the Wild”

Our main motivation in this work is to recommend papers worthy of an Ig Nobel prize. In this section,

Title	Models
The kinematics of eating with a spoon: Bringing the food to the mouth, or the mouth to the food?	Iggy, BERT ^f , SciBERT ^f
Do bonobos say NO by shaking their head?	Iggy, BERT ^f , SciBERT ^f
Is Anakin Skywalker suffering from borderline personality disorder?	Iggy, BERT ^f , SciBERT ^f
Not eating like a pig: European wild boar wash their food	Iggy, BERT ^f
Why don't chimpanzees in Gabon crack nuts?	SciBERT ^f , BERT ^f
Why do people lie online? "Because everyone lies on the internet"	BERT ^f
Which type of alcohol is easier on the gut?	BERT ^f
Rainbow connection and forbidden subgraphs	BERT
A scandal of invisibility: making everyone count by counting everyone	SciBERT
Where do we look when we walk on stairs? Gaze behaviour on stairs, transitions, and handrails	SciBERT

Table 4: A sample of top rated papers found by our models.

we test our models in a more realistic setting; we run them on a large sample of scientific papers, ranking each paper according to their certainty in the label ('humorous'), and identifying promising candidates. We use the same dataset of 630k papers from Semantic Scholar used for training the LMs (Section 4). We compute funniness according to our models (excluding random forest and logistic regression, which performed poorly). Table 4 shows examples of top-rated titles. We use the Amazon Mechanical Turk (MTurk) crowdsourcing platform to assess models' performance.

In an exploratory study, we asked people to rate the funniness of titles on a Likert scale of 1-5. We noted that people tended to confuse funny *research topic* and funny *title*. For example, titles like "Are you certain about SIRT?" or "NASH may be trash" received high funniness scores, even though the research topic is not even clear from the title. To mitigate this problem, we redesigned the study to include *two* 5-point Likert scale questions: 1) whether the *title* is funny, and 2) whether the *research topic* is funny. This addition seems to indeed help workers understand the task better. Example papers rated as serious title, funny topic include "Hat-wearing patterns in spectators attending baseball games: a 10-year retrospective comparison". Funny title, serious topic include "Slicing the psychoanalytic pie: or, shall we bake a new one? Commentary on Greenberg". Unless stated otherwise, the evaluation in the remainder of the paper was done on the "funny topic" Likert scale.

We paid crowd workers \$0.04 per title. As this task is challenging, we created a qualification test with 4 titles (8 questions), allowing for one mis-

take. The code for task and test can be found in the repository². We also required workers to have completed at least 1,000 approved HITs with at least 97% success rate.

All algorithms classified and ranked (according to certainty) all 630k papers. However, in any reasonable use-case, only the top of the ranked list will ever be examined. There is a large body of work, both in academia and industry, studying how people interact with ranked lists (in particular, search result pages) (Kelly and Azzopardi, 2015; Beus, 2020). Many information retrieval algorithms assume the likelihood of the user examining a result to exponentially decrease with rank. The conventional wisdom is that users rarely venture into the second page of search results.

Thus, we posit that in our scenario of Ig Nobel recommendations, users will be willing to read only the several tens of results. We choose to evaluate the top-300 titles for each of our five models, to study (in addition to the performance at the top of the list) how performance decays. We also included a baseline of 300 randomly sampled titles from Semantic Scholar. Altogether we evaluated 1375 titles (due to overlap). Each title was rated by five crowd workers. Overall, 13 different workers passed our test. Seven workers annotated less than 300 titles, while four annotated above 1,300 each.

Decision rule. Each title was rated by five different crowd workers on a 1-5 scale. There are several reasonable ways to aggregate these five continuous scores to a binary decision. A commonly-used aggregation method is the majority vote. The majority vote should return the clear-cut humorous titles. However, we stress that humor is very subjective

Decision rule	Threshold	Expert corr.	Labeled data accuracy
Min. 1 annotator	3	0.7	0.84
	4	0.49	0.83
Min. 2 annotators	3	0.47	0.82
	4	0.19	0.73
Min. 3 annotators	3	0.15	0.78
	4	0.02	0.62

Table 5: Spearman correlation of MTurk annotators with our expert, along with accuracy of MTurk annotators on our labeled dataset for the various mapping methods of the form “minimum (min.) k annotators gave a score at least m (threshold)”.

(and in the case of scientific humor, quite subtle). Indeed, annotators had low agreement on the topic question (average pairwise Spearman $\rho = 0.27$).

Thus, we explored more aggregation methods⁹. Our hypothesis class is of the general form “at least k annotators gave a score at least m ”¹⁰. To pick the best rule, we conducted two exploratory experiments: In the first one, we recruited an expert scientist and thoroughly trained him on the problem. He then rated 90 titles and we measured the correlation of different aggregations with his ratings. Results are summarized in table 5: The highest-correlation aggregation is when at least one annotator crossed the 3 threshold (Spearman $\rho = 0.7$).

In the second experiment, we used the exact same experimental setup as the original task, but with *labeled* data. We used 100 Ig Nobel winners as positives and a random sample of 100 papers as negatives. The idea was to see how crowd workers rate papers that we know are funny (or not). Table 5 shows the accuracy of each aggregation method. Interestingly, the highest accuracy is achieved with the same rule as in the first experiment (at least one crossing 3). Thus, we chose this aggregation rule.

We believe the method outlined in this section could be more broadly applicable to aggregation of crowd sourced annotations for subjective questions.

Results. Figure 1 shows precision at k for the top-rated 300 titles according to each model. The random baseline is ~ 0.03 . Upon closer inspection, these seem to be false positives of the annotation.

We have argued that in our setting it is reasonable for users to read the first several tens of results.

⁹For completeness, see Figure 3 in Appendix A.2.

¹⁰There is a long-running debate about whether it is valid to average Likert scores. We believe we cannot treat the ratings in this study as interval data.

Model	Precision at k=50	Precision at k=300
Iggy	0.6	0.37
SciBERT	0.57	0.46
SciBERT ^f	0.53	0.41
BERT	0.44	0.41
BERT ^f	0.58	0.43

Table 6: Precision at k of our models on the Semantic Scholar corpus for $k=\{50, 300\}$. These relatively high scores suggest that our models are able to identify funny papers.

In this range, Iggy slightly outperforms the other four models (BERT is particularly bad, as it picks up on short, non-informative titles). For larger k values SciBERT and BERT^f take the lead. We note that even at $k = 300$, all models still achieve considerable (absolute) precision.

We obtain similar results using normalized discounted cumulative gain (nDCG), a common measure for ranking quality (see Table 6 for nDCG scores for the top 50 and the 300 papers). Overall, these relatively high scores suggest that our models are able to identify funny papers.

We stress that Iggy is a small and simple network ($\sim 33k$ parameters), compared to pretrained 110 million parameters BERT-based models. Yet despite its simplicity, Iggy’s performance is roughly comparable to BERT-based methods. We believe this demonstrates the power of implementing insights from domain experts. We hypothesize that if the fine-tuning dataset were larger, BERT^f and SciBERT^f would outperform the other models.

8 Analysis

8.1 Importance of Literature-based Features

Taking a closer look at the actual papers in the experiment of Section 7, the overlap between the three *feature-based* models is 26 – 56% (for $1 < k < 50$) and 39 – 62% (for $1 < k < 300$). BERT had very low overlaps with all other models (0% in top 50, 10% in all 300). SciBERT had almost no overlap in top 50 (maximum 2%), 10 – 40% in all 300 (see full details in Appendix A.3). We believe this implies that the features were indeed important and informative for both BERT^f and SciBERT^f.

8.2 Interpreting Iggy

We have seen Iggy performs surprisingly well, given its relative simplicity. In this section, we




Figure 1: Precision at k for our chosen decision rule. Iggy outperforms the other models for $0 < k < 50$. For larger k , SciBERT^F and BERT achieve better precision.

wish to better understand the reasons. We chose to analyze Iggy with Shapely additive explanations (SHAP) (Lundberg and Lee, 2017). SHAP is a feature attribution method to explain the output of any black-box model, shown to be superior to more traditional feature importance methods. Importantly, SHAP provides insights both globally and locally (i.e., for specific data points).

Global interpretability. We compute feature importance globally. Among top contributing features we see multiple features corresponding to incongruity (both alone and combined with funniness) and to word/sentence simplicity. Interestingly, features based on the one-liner jokes seem to play an important role (See Figure 4 in Appendix A.4).

Local interpretability. To understand how Iggy errs, we examined the SHAP decision plots for false positives and false negatives (See Figure 5 in Appendix A.4). These show the contribution of each feature to the final prediction for a given title, and thus can help “debugging” the model.

Looking at false negatives, it appears that various perplexity features misled Iggy, while funniness and joke LM steered it in the right direction. We see a contrary trend in false positives: perplexity helped, and joke LM confused the classifier.

We also observe that the model learned that a long title is an indication of a serious paper. We expected our rudeness classifier to play a bigger role in some of the titles (e.g., “Adaptive inter-population differences in blue tit life-history traits on Corsica”), but the signal was inconclusive, perhaps indicating our rudeness classifier is lacking.

8.3 Observations

We now take a more qualitative approach to understand the models. First, we set out to explore whether the models confuse funny titles and funny topics. Using the crowd sourced annotations from Section 7, we measure the portion of this mistake in the top-rated 300 titles of all five models. That is, we check in how many cases our models classify a title as “Ig Nobel-worthy” while the workers have classified it as “funny title and non-funny topic”. Iggy had the highest degree of such confusion (0.28). Similarly, BERT^F and SciBERT^F exhibit more confusion than the versions without features (0.24, 0.19 compared to 0.13, 0.08). Random baseline is 0.02. Examples of this kind of error include “A victim of the Occam’s razor.”, “While waiting to buy a Ferrari, do not leave your current car in the garage!”, and “Reinforcement learning: The good, the bad and the ugly?”. All were classified as Ig Nobel-worthy, although their topic is serious (or even unclear from the title).

Looking closer at the data, we observe that a high portion of these are *editorials* with catchy titles. As our dataset does not differentiate between editorials and real research contributions, filtering editorials is not straightforward. Interestingly, the portion of editorials is also greater in the lowest annotators’ agreement area, hinting that this confusion also occurs in humans.

In addition to editorials, we notice another category of papers causing the same type of confusion. There are papers dealing with disturbing or unfortunate topics (violence, death, sexual abuse), whose titles include literary devices used to lighten

the mood. Censored (for the readers’ own well-being) examples include “Licorice for hepatitis C: yum-yum or just ho-hum?”, “The song of the siren: Dealing with masochistic thoughts and behaviors”.

A note on scientific disciplines. Another observation we make concerns with the portion of Ig Nobel-worthiness across the different scientific disciplines. We notice that most papers classified by our models as funny belong to social sciences (“Dogs can discriminate human smiling faces from blank expressions”) or medicine (“What, if anything, can monkeys tell us about human amnesia when they can’t say anything at all?”), compared to exact sciences (“The kinematics of eating with a spoon: bringing the food to the mouth, or the mouth to the food?”). We believe this might be the case since, quite often, social sciences and medicine papers study topics that are more familiar to the layperson. We also note that although our models performed about the same across the different disciplines, they were slightly better in psychology.

9 Conclusions & Future Work

In this work, we presented a novel task in humor recognition – detecting funny and unusual scientific papers, which represents a subtle and sophisticated humor type. It has important characteristics (short, simple syntax, stand-alone) making it a (relatively) clean setting to explore computational humor.

We created a dataset of funny papers and constructed models, distilling humor literature into features as well as harnessing SOTA advances in NLP. We conducted experiments both on our dataset and in a real-world setting, identifying funny papers in a corpus of over 0.6M papers. All models were able to identify funny papers, achieving high nDCG scores. Interestingly, despite the simplicity of the literature-based Iggy, its performance was overall comparable to complex, BERT-based models.

Our dataset can be further used for various humor related tasks. For example, it is possible to use it to create an *aligned* corpus, pairing every funny paper title with a nearly identical but serious title, using methods similar to [West and Horvitz \(2019\)](#). This would allow us to understand why a paper is funny at a finer granularity, by identifying the exact words that make the difference. This technique will also allow exploring different types of “funny”.

Another possible use of our dataset is to collect additional meta-data about the papers (e.g., citations, author information) to explore questions

about whether funny science achieves disproportionate attention and engagement, who tends to produce it (and at which career stage), with implications to science of science and science communication.

Another interesting direction is to expand beyond paper titles and consider the paper abstract, or even full text. This could be useful in examples such as the Ig Nobel winner “Cure for a Headache”, which takes inspiration from woodpeckers to help cure headaches in humans.

Finally, we believe multi-task learning is a direction worth pursuing towards creating a more holistic and robust humor classifier. In multi-task learning, the learner is challenged to solve multiple problems at the same time, often resulting in better generalization and better performance on each individual task ([Ruder, 2017](#)). As multi-task learning enables unraveling cross-task similarities, we believe it might be particularly fruitful to apply to tasks highlighting different aspects of humor. We believe our dataset, combined with other task specific humor datasets, could assist in pursuing such a direction.

Despite the tongue-in-cheek nature of our task, we believe that computational humor has tremendous potential to create personable interactions, and can greatly contribute to a range of NLP applications, from chatbots to educational tutors. We also wish to promote complementing data-driven research with insights from more-traditional fields. We believe combining such insights could, in addition to improving performance, enrich our understanding of core aspects of being human.

Acknowledgments

We thank the reviewers for their insightful comments. We thank Omri Abend, Michael Doron and Meirav Segal, Ronen Tamari and Moran Mizrahi for their help, and Shuki Cohen for preliminary discussions. This work was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant no. 852686, SIAM), US National Science Foundation, US-Israel Binational Science Foundation (NSF-BSF) grant no. 2017741, and Amazon Research Awards.

References

Iz Beltagy, Arman Cohan, and Kyle Lo. 2019. Scibert: Pretrained contextualized embeddings for scientific

- text. *arXiv preprint arXiv:1903.10676*.
- Dario Bertero and Pascale Fung. 2016. A long short-term memory framework for predicting humor in dialogues. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 130–135.
- Johannes Beus. 2020. [Why \(almost\) everything you knew about google ctr is no longer valid](#).
- Vladislav Blinov, Valeria Bolotova-Baranova, and Pavel Braslavski. 2019. Large dataset and language model fun-tuning for humor recognition. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 4027–4032.
- Marc Brysbaert and Andrew Biemiller. 2017. Test-based age-of-acquisition norms for 44 thousand english word meanings. *Behavior research methods*, 49(4):1520–1523.
- Dmitry Davidov, Oren Tsur, and Ari Rappoport. 2010. Semi-supervised recognition of sarcastic sentences in twitter and amazon. In *Proceedings of the fourteenth conference on computational natural language learning*, pages 107–116. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- David Donahue, Alexey Romanov, and Anna Rumshisky. 2017. Humorhawk at semeval-2017 task 6: Mixing meaning and sound for humor recognition. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 98–102.
- Limor Gultchin, Genevieve Patterson, Nancy Baym, Nathaniel Swinger, and Adam Tauman Kalai. 2019. Humor in word embeddings: Cockamamie gobbledegook for nincompoops. *arXiv preprint arXiv:1902.02783*.
- Nabil Hossain, John Krumm, Lucy Vanderwende, Eric Horvitz, and Henry Kautz. 2017. Filling the blanks (hint: plural noun) for mad libs humor. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 638–647.
- Diane Kelly and Leif Azzopardi. 2015. How many results per page? a study of serp size, search behavior and user experience. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 183–192.
- Chloe Kiddon and Yuriy Brun. 2011. That’s what she said: double entendre identification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*, pages 89–94. Association for Computational Linguistics.
- Lizhen Liu, Donghai Zhang, and Wei Song. 2018. Modeling sentiment association in discourse for humor recognition. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 586–591.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Advances in neural information processing systems*, pages 4765–4774.
- Antonios Maronikolakis, Danae Sánchez Villegas, Daniel Preoțiuc-Pietro, and Nikolaos Aletras. 2020. Analyzing political parody in social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4373–4384.
- William H Martineau. 1972. A model of the social functions of humor. *The psychology of humor*, pages 101–125.
- Rada Mihalcea and Carlo Strapparava. 2005. Making computers laugh: Investigations in automatic humor recognition. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 531–538.
- Rada Mihalcea and Carlo Strapparava. 2006. Learning to laugh (automatically): Computational models for humor recognition. *Computational Intelligence*, 22(2):126–142.
- Tristan Miller, Christian F Hempelmann, and Iryna Gurevych. 2017. Semeval-2017 task 7: Detection and interpretation of english puns. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 58–68.
- Abhinav Moudgil. 2016. [short-jokes-dataset](#).
- Dallin D Oaks. 1994. Creating structural ambiguities in humor: getting english grammar to cooperate. *Humor*, 7(4):377–402.
- Reynier Ortega-Bueno, Carlos E Muniz-Cuza, José E Medina Pagola, and Paolo Rosso. 2018. Uo upv: Deep linguistic humor detection in spanish social media. In *Proceedings of the Third Workshop on Evaluation of Human Language Technologies for Iberian Languages (IberEval 2018) co-located with 34th Conference of the Spanish Society for Natural Language Processing (SEPLN 2018)*.
- Tomáš Ptáček, Ivan Habernal, and Jun Hong. 2014. Sarcasm detection on czech and english twitter. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 213–223.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Victor Raskin. 1985. Semantic theory. In *Semantic Mechanisms of Humor*, pages 59–98. Springer.
- Antonio Reyes and Paolo Rosso. 2012. Making objective decisions from subjective data: Detecting irony in customer reviews. *Decision support systems*, 53(4):754–760.
- Antonio Reyes, Paolo Rosso, and Davide Buscaldi. 2012. From humor recognition to irony detection: The figurative language of social media. *Data & Knowledge Engineering*, 74:1–12.
- Willibald Ruch. 1992. Assessment of appreciation of humor: Studies with the 3 wd humor test. *Advances in personality assessment*, 9:27–75.
- Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *ArXiv*, abs/1706.05098.
- Itay Sagi and Eldad Yechiam. 2008. Amusing titles in scientific journals and article citation. *Journal of Information Science*, 34(5):680–687.
- Nobel Prize Committee The Royal Swedish Academy of Science. 2010. Press release. <https://www.nobelprize.org/prizes/physics/2010/press-release/>. [Online; 5 October 2010].
- Dafna Shahaf, Eric Horvitz, and Robert Mankoff. 2015. Inside jokes: Identifying humorous cartoon captions. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1065–1074.
- Joshua Shaw. 2010. Philosophy of humor. *Philosophy Compass*, 5(2):112–126.
- Arthur Shurcliff. 1968. Judged humor, arousal, and the relief theory. *Journal of personality and social psychology*, 8(4p1):360.
- Edwin Simpson, Erik-Lân Do Dinh, Tristan Miller, and Iryna Gurevych. 2019. Predicting humorousness and metaphor novelty with gaussian process preference learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5716–5728.
- Edgar A Smith and RJ Senter. 1967. Automated readability index. *AMRL-TR. Aerospace Medical Research Laboratories (US)*, page 1.
- Alan D Sokal. 1996. Transgressing the boundaries: Toward a transformative hermeneutics of quantum gravity. *Social text*, (46/47):217–252.
- Oliviero Stock and Carlo Strapparava. 2002. Hacronym: Humorous agents for humorous acronyms. *Stock, Oliviero, Carlo Strapparava, and Anton Nijholt. Eds.*, pages 125–135.
- Jerry M Suls. 1972. A two-stage model for the appreciation of jokes and cartoons: An information-processing analysis. *The psychology of humor: Theoretical perspectives and empirical issues*, 1:81–100.
- Julia M Taylor and Lawrence J Mazlack. 2004. Computationally recognizing wordplay in jokes. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 26.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Sida I Wang and Christopher D 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.
- Melissa Bekelja Wanzer, Melanie Booth-Butterfield, and Steve Booth-Butterfield. 1996. Are funny people popular? an examination of humor orientation, loneliness, and social attraction. *Communication Quarterly*, 44(1):42–52.
- Miaomiao Wen, Nancy Baym, Omer Tamuz, Jaime Teevan, Susan T Dumais, and Adam Kalai. 2015. Omg ur funny! computer-aided humor with an application to chat. In *ICCC*, pages 86–93.
- Robert West and Eric Horvitz. 2019. Reverse-engineering satire, or” paper on computational humor accepted despite making serious advances”. *arXiv preprint arXiv:1901.03253*.
- Chris Westbury and Geoff Hollis. 2019. Wiggly, squiffy, lummoX, and boobs: What makes some words funny? *Journal of Experimental Psychology: General*, 148(1):97.
- David Wilmot and Frank Keller. 2020. **Modelling suspense in short stories as uncertainty reduction over neural representation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1763–1788, Online. Association for Computational Linguistics.
- Zafar. 2018. **Cleaned Toxic Comments**.
- Renxian Zhang and Naishi Liu. 2014. Recognizing humor on twitter. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 889–898. ACM.
- Yftah Ziser, Elad Kravi, and David Carmel. 2020. Humor detection in product question answering systems. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 519–528.

A Supplementary figures

A.1 Dataset Analysis

In Section 4 we presented 127 humor literature-based features. Here we present the distribution of two example features in funny vs. serious papers in our dataset (described in Section 3). These examples represent the general trend, as many features show predictive power (see Figure 2).

A.2 “In the Wild” Study Results

For the “in the wild” evaluation executed using Semantic Scholar data, we used crowdsourcing annotations (see Section 7). Each title was rated by five different crowd workers on a 1-5 scale, while our models provide binary decision. There are several reasonable ways to aggregate these five continuous scores to a binary decision. We choose a rule in a data-driven manner (see “Decision rule” in Section 7). For completeness, here we show the commonly-used aggregation method of *majority*. We show here the precision at k of our five models using the majority vote aggregation rule with a cutoff at 3 (see Figure 3). Iggy outperforms until $k = 30$, where SciBERT^f takes the lead afterwards.

A.3 Models’ Overlap

In Section 8.1 we discuss the importance of our literature-based features by showing that models who received them as input indeed found them useful. The overlap was measured on the top 50 and top 300 papers retrieved using our five models on the Semantic Scholar data (see Section 7 for the full experimental setup). The overlap between the 3 *features-based* models was found to be high (see Table 7). Both BERT and SciBERT had very low overlaps with all other models. We believe this implies that the features were indeed important for our SOTA based models, BERT^f and SciBERT^f.

A.4 SHAP Analysis

In Section 8.2 we analysed Iggy using SHAP (Lundberg and Lee, 2017). We compute feature importance globally (Figure 4). To understand how Iggy errs, we examined the SHAP decision plots for false positives and false negatives (Figure 5). Decision plots show the contribution of each feature to the final prediction for a given title. Thus, it can help “debugging” the model’s mistakes.

B Reproducibility

B.1 Code and Data Availability

Dataset, code, and data files can be found in our Github repository².

C Implementation details

C.1 Fine-Tuning GPT-2 LM

To fine-tune GPT-2 we used Huggingface’s Transformers package¹¹. We fine-tuned the model using learning rate = $5e-5$, one epoch, batch size of 4, weight decay = 0, max gradient norm = 1 and random seed = 42. Optimization was done using Adam with epsilon = $1e-8$. Model configurations were set to default.

C.2 Iggy Classifier

We used a simple MLP with a single hidden layer of 256 neurons. We trained the MLP until convergence, using Adam optimizer, a learning rate of 0.001 and an L2 penalty of 2.

C.3 Fine-tuning SciBERT & BERT

To fine-tune SciBERT & BERT we used Huggingface’s Transformers package. We fine-tuned both models with learning rate = $5e-5$ for 3 epochs with batch size of 32, maximal sequence length of 128 and random seed = 42. Optimization was done using Adam with warm-up = 0.1 and weight decay of 0.01. Model configurations were set to default.

C.4 SciBERT^f & BERT^f Models

As specified in Section 5, these models were constructed as follows (see Figure 6). Each model had two inputs – the raw text of the title, and a vector of our 127 features. The feature vector is fed to an MLP with a single hidden layer of 512 neurons and an output size of 512 neurons as well. The raw text is fed to a *frozen* SciBERT /BERT model. We collect the last hidden vector ([CLS]) from BERT /SciBERT. Next, we concatenate this vector to the output of the features-MLP network and pass the result to a second MLP with a single hidden layer of 1,024 neurons. The output of this MLP, then, is fed to a Softmax layer, which represents the final prediction of the model.

We train the model using a cross-entropy loss and the same parameters that were used to train the vanilla SciBERT /BERT model. Those parameters are described in Appendix C.3.

¹¹huggingface.co/transformers/

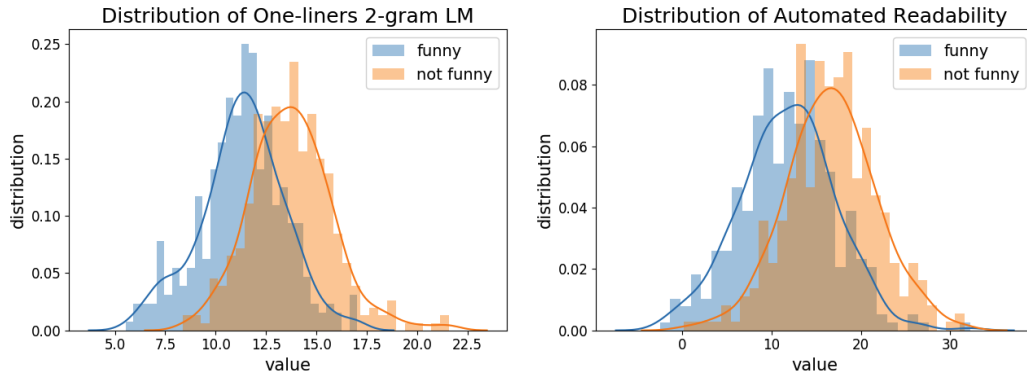


Figure 2: Distribution and Gaussian fit of two representative features: one-liners 2-gram LM mean perplexity (left) and automated readability index (right), indicating the predictive power of these features.




Figure 3: Precision at k for majority vote. Iggy outperforms until $k = 30$ and $BERT^f$ takes the lead afterwards.






Figure 4: Feature importance SHAP analysis done on the Iggy model. The top variables according to this plot contribute more than the bottom ones (have high predictive power). The analysis reveals that the highest contribution corresponds to short, funny, and simple words (where simplicity was measured using features such as AoA and readability). We also notice that features which are based on the one-liners LMs contributed much to the final prediction, meaning that there is indeed some similarity between funny titles and short jokes.

	Iggy	BERT ^f	SciBERT ^f	BERT
BERT ^f	0.56 0.62			
SciBERT ^f	0.26 0.39	0.36 0.48		
BERT	0 0	0 0.1	0 0.1	
SciBERT	0 0.3	0.02 0.4	0.02 0.2	0 0.1

Table 7: Models’ overlap for the top rated 50 and 300 (left number in a cell corresponds to the overlap in the top 50 and right number corresponds to the 300). The overlap between the 3 *features-based* models was found to be high compared with BERT and SciBERT. We believe this implies that the features were indeed important for our SOTA based models, BERT^f and SciBERT^f.



(a) SHAP decision plot for the 12 false negative of Iggy from our test set. Perplexity features misled Iggy, while funniness and joke LM ones provided informative input.



(b) SHAP decision plot for the 11 false positive of Iggy from our test set. Perplexity helped shifting the output towards the correct label, joke LM features confused the classifier

Figure 5: SHAP decision plot for Iggy’s false negatives and positives from our test set. Decision plots show the contribution of each feature to the final prediction for a given data point. Starting at the bottom of the plot, the prediction line shows how the SHAP values (i.e., the feature effects) accumulate to arrive at the model’s final score at the top of the plot. To get a better intuition, one can think of it in terms of a linear model where the sum of effects, plus an intercept, equals the prediction.




Figure 6: The flow of SciBERT^f / BERT^f. A 2-layers MLP receives an input the concatenation of two vectors: our features' embedding and the last hidden vector ([CLS]) from BERT /SciBERT.