Writing Strategies for Science Communication: Data and Computational Analysis

Tal August[†]   **Lauren Kim**[†]   **Katharina Reinecke**[†]   **Noah A. Smith**[†◇]

[†]Paul G. Allen School of Computer Science & Engineering,
University of Washington, Seattle, WA, USA
[◇]Allen Institute for Artificial Intelligence, Seattle, WA, USA
`{taugust, lkim143, reinecke, nasmith}@cs.washington.edu`

Abstract

Communicating complex scientific ideas without misleading or overwhelming the public is challenging. While science communication guides exist, they rarely offer empirical evidence for how their strategies are used in practice. Writing strategies that can be automatically recognized could greatly support science communication efforts by enabling tools to detect and suggest strategies for writers. We compile a set of writing strategies drawn from a wide range of prescriptive sources and develop an annotation scheme allowing humans to recognize them. We collect a corpus of 128K science writing documents in English and annotate a subset of this corpus.[1] We use the annotations to train transformer-based classifiers and measure the strategies' use in the larger corpus. We find that the use of strategies, such as storytelling and emphasizing the most important findings, varies significantly across publications with different reader audiences.

1 Introduction

Communicating scientific discoveries to a general audience of readers is difficult. A researcher or writer interested in doing so is faced with the challenging task of translating complex scientific ideas in an engaging manner without misleading or overwhelming their audience. There are many guides to science communication (e.g., Blum et al., 2006), but they rarely offer empirical evidence for how their advice is used, or proven effective, in practice. The potential science communicator is then confronted with the additional hurdle of understanding how to implement these guidelines in their writing.

Effective science communication requires understanding the unique needs and expectations of different audiences and stakeholders in science (Nisbet and Scheufele, 2009). We envision natural language processing technologies that help science writers communicate more effectively. These technologies might automatically classify common strategies in a writer's own text, support writers to adapt language to specific readers, or guide readers through personalized article recommendations.

As a first step, we compile a set of strategies from a wide range of prescriptive science writing sources in English and develop an annotation scheme allowing humans to recognize these strategies in texts about science. We introduce a new corpus of 128K university press releases, science blogs, and science magazines and annotate a subset of 337 texts. We use the annotations to train transformer-based classifiers to explore the communicative goals of science writing by analyzing variations in the strategies' use across several scientific communication forums.

Our paper is the first computational analysis of writing strategies driven by science communication theory. We find that most strategies are prevalent throughout our corpus and that publication venues with varying audiences use strategies differently. For example, press releases emphasize the impacts of science more than magazine articles, which instead tell more stories about the science. We also find that higher quality newspaper articles, as rated by expert journalists, use more storytelling and analogies than lower quality articles.

2 Defining Science Communication Writing Strategies

The goal of general science communication is to increase public awareness, enjoyment, interest, and understanding about science (Burns et al., 2003). Based on the idea of compositionality in discourse theory (Bender and Lascarides, 2019), we can think of the communicative intent of science writing as being made up of smaller communication goals rep-

---

[1]Available at `https://github.com/talaugust/scientific-writing-strategies`.

resented in particular passages of an article (Grosz and Sidner, 1986; Louis and Nenkova, 2013a). Our computational approach builds on this theoretical assumption by annotating sentences and letting an article inherit the attributes we find in its sentences.

Past work on science communication has taken a similar view (Louis and Nenkova, 2013a) by using syntactic relations to characterize an article's communicative goals, allowing them to emerge inductively rather than from a theory of science communication. Our complementary approach starts with science communication guides to construct theory-driven communicative goals (referred to as "writing strategies" and consisting of lexical to multi-sentence features), and explores their use in a diverse range of science communication text.

To define our writing strategies, we categorized and grouped advice from style guides for science communicators. These guides were a mix of online resources, books, and academic articles (see Table 5 in the appendix). We selected the guides based on discussions with three expert science communicators at a large research university's press department and through online searches. We stopped adding resources when we reached saturation (Holton, 2007), meaning that each new resource had fewer new strategies and suggesting that our resources provided good coverage.

Two authors open-coded (Holton, 2007) the suggestions from each guide by assigning each piece of advice in a resource a code that represented its high-level strategy (such as "avoid jargon"). The authors then looked at other resources to see whether the same advice appeared there. Each new piece of advice was added with a new code. After coding all resources, the authors grouped the codes into a set of 10 suggested writing strategies. Appendix A.1 provides additional details on the coding and categorization. The strategies are as follows (examples of each are given in Table 9 in the appendix):

LEDE A few sentences at the beginning of an article, called a lede (spelled "lede" for easier differentiation with its homograph "lead"), that draws a reader in and makes them want to read more.

MAIN Sentences describing the main findings being reported by the original paper in order to not overwhelm the reader with details.

IMPACT Writing about the real world impact of the science or findings being reported in order to excite readers. This can include future technolo-

gies, breakthroughs the findings might enable, or their societal implications.

EXPLANATION Explanations about scientific subjects to improve reader understanding. This could be explaining a certain topic or word, or what researchers did in a study and what the findings mean.

ANALOGY Analogies or metaphors used as a way to explain concepts or make ideas in the article more relatable.

STORY Stories to engage readers and make the reported science more interesting. This can include short story snippets, or coming back to an underlying story throughout an article.

PERSONAL Including personal details about researchers in order to make them more approachable and add depth to the story.

JARGON Avoiding specialized terminology or jargon as much as possible as it can overwhelm readers.

ACTIVE Using the active voice to make the writing more lively and engaging.

PRESENT Similar to ACTIVE, using present tense verbs also to make the writing more lively and engaging.

Some of these strategies are specific to science writing, such as emphasizing the real world impact of the findings (IMPACT), while others are often thought of as general rules for good writing, such as using the active voice (ACTIVE). Both types of strategies were commonly referenced in the resources we analyzed, which suggests that engaging science writing shares traits with engaging writing in other disciplines while also containing its own set of unique strategies.

3 Dataset

In order to study the use of these strategies by science writers and build classifiers for automatic identification, we collected a corpus of 128K documents from a variety of science communication sources. We focused on four major types of U.S.-based venues, representing a broad spectrum of science communication for different audiences: blog sites, popular science magazines, university press releases, and scientific journal magazines.

Table 1: Sites and number of articles in corpus after filtering.

Venue	Site	#Articles
Blogs	*Sciencedaily.com*	38,191
	Phys.org	74,732
Magazines	*The Atlantic*	2,771
	Scientific American	5,174
Press releases	*Harvard*	752
	Stanford	599
	Rochester	219
	Northwestern	197
Journals	*Science*	4,173
	Nature	1,445
Total		128,253

Past work has shown that blog sites usually write to scientifically literate and engaged readers (Ranger and Bultitude, 2016), while university press releases often write to other science journalists (Bratton et al., 2019; Sumner et al., 2014). We selected popular science magazines since they target a more general audience, and scientific journal magazines as they often write to those involved in research, though not necessarily in the same domain (Nielsen and Schmidt Kjærgaard, 2011). The choices of website or publication we collected from each venue category were based either on previous work covering those categories (e.g., blog posts; Vadapalli et al., 2018) or as a convenience sample based on what was widely available. One note is that while past work has used the blogs sites we selected as sources for high quality science blogs (*sciencedaily.com* and *phys.org*), these sites also source a large portion of their content from press releases, often only changing headlines and lede sentences.

We scraped articles from each of these sources for all of 2016–2019 using the Wayback Machine,[2] resulting in 137,828 articles. Appendix A.2 provides more details on site selection.

3.1 Filtering

To focus on science communication, specifically, we removed articles matching U.S.-centric political keywords such as *Trump*, *democrats*, and *Senate*. Appendix A.3 lists all filter keywords. We also removed all articles over 15,000 or under 1,500 characters, since these represented ei-

ther multiple articles on the same page, article previews, or scraper errors. After filtering we had a total of 128,253 articles. In total 7% of documents were filtered (3.5% removed for political keywords and 3.5% for length). Table 1 details the sites for each venue and the number of articles after filtering. URLs for all scraped articles are available at https://github.com/talaugust/scientific-writing-strategies.

3.2 Annotation

Recall that our goal is to measure the use of strategies from Section 2 in our corpus. We sample 337 articles stratified across sites to gather a spread of articles and balance the articles across venues. Each article was given to two annotators who were trained on the writing strategies and instructed to annotate the use of strategies at the sentence level. Concretely, each "annotation" corresponds to a contiguous chunk of one or more sentences labeled with one of the seven strategies. A sentence can be labeled with multiple strategies. Three of the strategies, JARGON, ACTIVE, and PRESENT, were not annotated because we believe they can be reliably detected using existing methods based on lexical and syntactic features; see Section 5.1. Figure 5 (in the appendix) presents an example of an annotated excerpt and the task interface.[3]

We conducted annotation in sets of 50 articles. After each set, one author measured agreement and manually evaluated a subset of annotations by both annotators. This author then acted as a coordinator for the annotators, providing suggestions or revisions to annotation guidelines. Additionally, annotators were able to look at the other's annotations after completing an article. Figure 4 in the appendix plots Krippendorff's α after each batch of 50 articles.

3.3 Annotator Agreement

While our strategies emerged from prescriptive advice in guides, our annotators had to interpret these strategies in the context of real-world scientific writing. This, to our knowledge, has not been done systematically before, though we suspect editors do it frequently. Because writing strategies are somewhat subjective and we wanted our categorizations to be flexible to the different good-faith

---

[2] https://archive.org/web/

[3] Code for the interface is available at https://github.com/talaugust/scientific_article_annotation.

Table 2: Simple agreement rates for each writing strategy at a sentence, paragraph, and document level. See Table 6 (in the appendix) for α statistics.

Strategy	Sentence	Paragraph	Document
LEDE	0.97	0.95	0.65
MAIN	0.89	0.85	0.96
IMPACT	0.92	0.90	0.68
EXPLANATION	0.69	0.70	0.84
ANALOGY	0.95	0.92	0.68
STORY	0.84	0.86	0.71
PERSONAL	0.90	0.86	0.66

Table 3: Number of sentences and average number of words per annotation span for each strategy across articles in our annotated dataset (337 articles).

Strategy	# Sentences	Avg. words
LEDE	595	34.1±20.6
MAIN	1,596	35.9±19.1
IMPACT	1,102	40.0±19.1
EXPLANATION	4,677	55.0±28.1
ANALOGY	410	33.6±17.3
STORY	1,736	74.4±60.9
PERSONAL	727	50.5±33.1
Total	10,843	48.7±32.7

interpretations of each strategy, we were not aiming to obtain perfect agreement on each strategy. Table 2 reports on simple agreement rates for the annotated strategies.

Previous work annotating spans of text for communicative goals, such as framing (Card et al., 2015), propaganda techniques (Da San Martino et al., 2019), hate speech (Sap et al., 2020), statement strength (Tan and Lee, 2014), and agency (Sap et al., 2017), have shown that reaching high agreement is difficult. While agreement measures differ across these annotations tasks due to differences in how spans were annotated (e.g., preselected sentences or open selection), past work has reported Krippendorff's α agreement levels ranging between $0.3 < \alpha < 0.7$, in which we fall along the lower to moderate end ($0.3 < \alpha < 0.5$). We discuss the use of α in Appendix A.4 and there Table 6 reports α agreement for each strategy.

The annotators identified a total of 10,843 sentences (316,263 tokens) with one of the seven strategies in 337 articles. Table 3 details the number of sentences and average number of words for each strategy annotation span.

3.4 Abandoned Strategies

Two categories that achieved low agreement (as measured by α) were EXPLANATION and PERSONAL, which we drop from further analyses for the following reasons. We found that annotator 1 (a_1) annotated many more EXPLANATION strategies than a_2 (3,402 vs. 1,275). While the majority of a_2's EXPLANATION annotations agreed with a_1's (886 out of 1,275), a_2 highlighted fewer in general, suggesting that this lower agreement was due to the annotators having different thresholds for the EXPLANATION strategy.

For PERSONAL, our discussions with both annotators revealed that a_2 had assumed PERSONAL strategies were any reference to the researchers in an article (e.g., "Professor X, head of the mutation lab at the Academy") while a_1 focused on personal aspects of a researcher (e.g., "Professor X, who has been in a wheelchair since birth"). While we tried to resolve this difference early on, annotators still had difficulty reaching agreement on PERSONAL strategies. This discrepancy lowered agreement but highlighted an interesting nuance in the PERSONAL strategy.

Because both strategies are still important for science communication, we report on their use in the corpus in Figures 6 and 7 in the appendix.

4 Hypotheses

Our annotated corpus allowed us to begin to explore how strategies relate to the communicative goals of different science communication venues. To do this, we introduce hypotheses informed by existing literature.

Hypotheses **H1–H4** are based on our expectations for how strategies can differentiate science writing venues in our corpus based on their underlying goals. For these hypotheses, we evaluate strategy use across our corpus.

Hypotheses **H5** and **H6** focus on how strategies might relate to other important issues in science communication. They build on past research in science communication exploring article quality (Louis and Nenkova, 2013b) and sensationalism (Sumner et al., 2014). These issues are introduced with their own annotated datasets, and since we have no strategy annotations for these datasets, we report only on the aggregated predictions of our classifiers.

H1: LEDE is used once or twice within an article, but consistently across our entire corpus. Because the LEDE strategy is well adopted as a common strategy in general journalism (Pöttker, 2003), and LEDE sentences are only used at the beginning of an article, we expect low but consistent use of LEDE across the corpus.

H2: Press releases use higher IMPACT than other venues. One goal of press releases is to encourage other science writers to pick up a story. Because a key component to selecting stories is impactful findings (Hayden et al., 2013), we expect that press releases will emphasize this more.

H3: Magazines use lower JARGON, higher ACTIVE and PRESENT and higher STORY than other venues. Magazines target a broader readership compared to other venues, making it likely they use these strategies that are common in prescriptive guides for general good writing (Strunk, 2007) to relate to a wider audience.

H4: Blog sites use higher JARGON and MAIN, and lower IMPACT compared to other sites. Blog sites predominantly focus on other science educated or interested readers (e.g., *phys.org* reports a readership of "5 million scientists, researchers, and engineers every month"[4]). This suggests that blogs' focus is less on attracting a broad readership (higher JARGON) or encouraging news uptake (lower IMPACT) and more on communicating the main point of a journal article (higher MAIN).

H5: Higher quality science news articles use more STORY and ANALOGY than lower quality articles. Past work on science news quality has suggested that features related to storytelling and figurative language (e.g., coherence and descriptive language) are associated with higher quality articles (Louis and Nenkova, 2013b).

H6: Press releases that sensationalize the claims of the original journal paper use higher IMPACT and MAIN than press releases that do not. While emphasizing the impact of findings is a useful tool in engaging readers, work has shown that press releases will sometime sensationalize the claims of the original journal article (Sumner et al., 2014). Sumner et al. (2014) categorized sensationalism into three categories: exaggerated advice (suggesting actions the original paper did not), exaggerated causal claims (making causal claims the

---

[4] https://sciencex.com/help/about-us/

paper did not), and inference to humans from animal research. These categories most nearly relate to our strategies on the findings of a paper: IMPACT and MAIN, and we hypothesize these strategies' overuse is related to sensationalism.

5 Strategy Classification

We used our collected corpus and annotations to automate recognition of writing strategies and to evaluate our hypotheses. We describe our methods for classifying strategies with rules and with human annotations (Sections 5.1 and 5.2). We then discuss methods for using these classifiers to estimate the use of strategies in our corpus and overall classifier performance (Section 5.3).

5.1 Rule-Based Strategies

As discussed earlier, three of the strategies could be reasonably identified using rules, and were not annotated.

JARGON We used common science jargon word lists drawn from Rakedzon et al. (2017) and Gardner and Davies (2013) to detect jargon use. The word list from Rakedzon et al. (2017) consists of 2,949 words common in scientific journal abstracts and articles while rare in common usage. We augmented this list with the core Academic Vocabulary List (AVL, Gardner and Davies, 2013). The AVL is a list of the top 3,000 word lemmas based on 120 million words of academic texts from the Corpus of Contemporary American English (Davies, 2009). High JARGON means higher use of these specialized terms, which is negatively associated with the strategy (i.e., since the recommended strategy is to avoid specialized terms).

ACTIVE We identified active and passive clauses by counting the 'nsubj' and 'nsubj:pass' words from a parse of each article using Stanford NLP's dependency labels in the Stanford NLP Pipeline.[5] We normalized all active clauses by the number of verbs in an article.

PRESENT For measuring present tense, we normalized the number of present tense verbs using Stanford NLP's Universal Features (similar to POS tags and part of the same Stanford NLP Pipeline; Manning et al., 2014) over all verbs in an article.

Table 4: Mean and standard deviation of strategy classifier performance, including calibration error, on held-out test set based on 5 random seeds. Baseline accuracy for most frequent class (MFC), which always predicts the negative class, is included.

Strategy	Prec.	Recall	F_1	Calibr. Err.	Acc.	MFC Acc.
LEDE	$.31_{.02}$	$.56_{.03}$	$.40_{.02}$	$.05_{.004}$	$.95_{.002}$	.95
MAIN	$.38_{.04}$	$.51_{.03}$	$.43_{.03}$	$.11_{.01}$	$.86_{.005}$	.86
IMPACT	$.40_{.03}$	$.55_{.03}$	$.46_{.02}$	$.09_{.01}$	$.91_{.004}$	.90
ANALOGY	$.52_{.04}$	$.60_{.02}$	$.55_{.03}$	$.04_{.01}$	$.96_{.001}$	.95
STORY	$.38_{.03}$	$.49_{.04}$	$.43_{.02}$	$.12_{.01}$	$.84_{.009}$	.84

5.2 Sentence Classifiers

For the remainder of our strategies, we trained classifiers based on the annotations collected to estimate the prevalence of each strategy in our corpus. Each classifier takes a single sentence as input and provides a binary label (present or absent) for a given strategy. Apart from pretraining, the classifiers were trained separately for each strategy. We base our classifiers off RoBERTa (Liu et al., 2019) as it is a high-performing contextual word representation learner that has achieved state-of-the-art results on multiple NLP benchmarks, and which comes pretrained. We use Huggingface's RoBERTa implementation.[6]

We start by continuing to pretrain RoBERTa on additional in-domain text to tailor the model more closely to our task. This additional pretraining followed two phases as in Gururangan et al. (2020): pretraining on 11.90M general news articles from REALNEWS (Zellers et al., 2019) for 12.5K steps (domain-adaptive pretraining), and then pretraining on a held-out subset of 100k documents from the unannotated portion of our own corpus for 10 epochs (task-adaptive pretraining). Appendix A.5 includes details for both pretraining steps.

Finally, we finetune our pretrained RoBERTa model on each sentence-level classification task separately, making 5 binary classifiers (LEDE, MAIN, IMPACT, ANALOGY, STORY). Our pretrained RoBERTa models were finetuned using a 80%, 10%, 10% train, validation, test setup using all annotated articles. Articles were randomly split across the sets, meaning no two sentences from the same article could occur across sets. Appendix A.6 includes additional finetuning details.

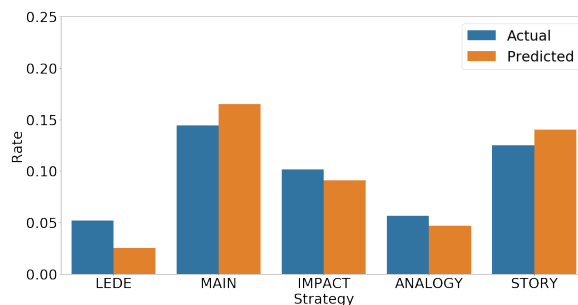Using classifiers optimized for individual classi-

Figure 1: The rate of occurrence of predicted versus actual strategies on our human-annotated test set based on PCC and the classifiers.

fication can lead to biases when estimating category proportions (Hopkins and King, 2010). Past work has suggested that using a **well-calibrated** classifier leads to better proportion estimation in large unlabeled corpora (Card and Smith, 2018). Calibration refers to the long-run accuracy of predicted probabilities, where a well-calibrated probabilistic classifier at the level μ is one that predicts class k with probability μ when the proportion of instances correctly assigned to k is also μ.

Following Card and Smith (2018), we perform model selection based on calibration error on held-out data during hyperparameter tuning. We estimate calibration error using the adaptive binning procedure from Nguyen and O'Connor (2015). After picking our most well-calibrated classifiers, we measure the rate of each strategy across a collection of documents by averaging the classifiers' predicted posterior probabilities of a positive label. This is referred to as **Probabilistic Classify and Count** (PCC; Bella et al., 2010) and is a standard method for predicting label distributions in a corpus using a probabilistic classifier (Card and Smith, 2018).

5.3 Evaluation

We evaluate our classifiers in two ways: on individual examples (i.e., for reporting F_1), and in aggregate on a held-out annotated test set.

Our goal for the classifiers is to estimate aggregated proportions in our corpus, not to achieve perfect performance. For this reason, we report on classifier F_1 performance only to establish that the classifiers are reasonably able to detect strategies. Table 4 details the precision, recall, F_1 scores, average calibration error and accuracy of the trained classifiers, and baseline accuracies for the most frequent class predictions for each strategy. Our classifiers achieved F_1 scores between .40 and .55, which is comparable to other classifications of communicative goals, such as propaganda technique detection (e.g., F_1 scores between 0.39 and 0.61 in Da San Martino et al., 2019).

Because we are most interested in estimated proportions, we also compared the classifiers' predicted strategy rates in our held-out test set with the actual rates of the annotated strategies. Actual strategy rate is calculated as the number of sentences containing a strategy divided by the total number of sentences in the test set. Figure 1 illustrates this comparison. While we do see some discrepancies between actual rates and our predicted rates, these differences are small ($<$ .05 rate difference, or less than 5% of sentences) and the trend of each strategy remains the same (e.g., STORY and MAIN are the most common, LEDE and ANALOGY are the least), suggesting that the classifiers estimate strategy rates with sufficiently high accuracy to begin comparing rates across strategies. Figure 6 in the appendix compares the predicted rate of strategies in the full dataset compared to actual rates of strategies in the test set broken down by site.

We additionally evaluated how accurate our automatic measures for JARGON, ACTIVE, and PRESENT were by randomly sampling 5 sentences from the top and bottom 10% of articles containing JARGON, ACTIVE, and PRESENT (as measured by our rule-based approaches) and manually inspecting them for correct word classifications. The rules for each measure are in line with our intuitions about JARGON, ACTIVE, and PRESENT with a large majority of words (Over 80% in the 30 sentences evaluated) being identified correctly as jargon, active voice, or present tense. Table 10 in the appendix provides examples of this evaluation.

6 Evaluating Strategy Applications

Evaluating our classifier output against gold-standard human annotations, as reported in Section 5.3, establishes the validity of our classifiers. We next turn to our hypotheses introduced in Section 4 to illustrate how we can use the strategies, classifiers, and corpus to explore the communicative goals of science writing. We introduce each hypothesis and report on its results separately.

H1: LEDE is used once or twice within an article, but consistently across our entire corpus. Figure 2a plots the estimated number of LEDE sentences per article across each site. Supporting **H1**, the majority of sites peak at either 0 or 1 LEDE sentences, with all sites tapering off quickly after that. *theatlantic.com* does have a higher number of predicted LEDE sentences (with 20% of articles containing more than 2 sentences). This might be due to *theatlantic.com* articles being longer (since they are full magazine articles) and therefore using more text to entice readers.

H2: Press releases use higher IMPACT than other venues. We find support for **H2**: press release sites like *news.harvard.edu*, *rochester.edu*, and *news.stanford.edu* have larger modes than other sites for IMPACT sentences in Figure 2c. For example, close to 15% of articles in *rochester.edu* have over 5 estimated IMPACT sentences, compared to 7 or 8% of articles in *scientificamerican.com* or *theatlantic.com* having that same number. This is especially striking because *scientificamerican.com* and *theatlantic.com* generally have much longer texts, since they are full magazines, compared to press releases.

H3: Magazines use lower JARGON, higher ACTIVE and PRESENT and higher STORY than other venues. Texts from *theatlantic.com* and *scientificamerican.com*, the two magazine sites, had the lowest and third lowest use of JARGON in the corpus with average rates below 0.2 (i.e., less than 20% of words), macro-averaged across articles (see Figure 6 in the appendix). Magazines also had the highest use of ACTIVE and some of the highest PRESENT. Additionally, *theatlantic.com* was the only site to have close to 5% of articles estimated to have more than an 15 STORY sentences (Figure 2e).

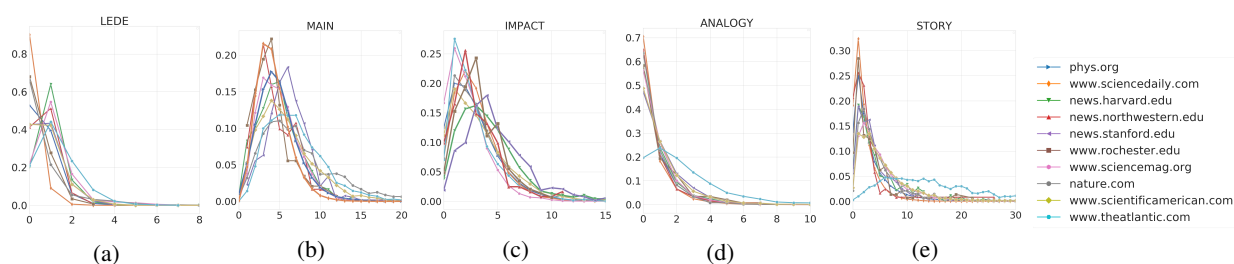**H4: Blog sites use higher JARGON and MAIN, and lower IMPACT compared to other sites.**

Figure 2: Histogram of number of sentences per article (x-axis) estimated to use each strategy (proportions on y-axis). For an enlarged figure, see Figure 7 in the appendix. Figure 6 plots the estimated proportion of sentences using each strategy for all sites.

We find mixed support for **H4**: blog sites do use higher JARGON and MAIN, but not lower IMPACT compared to other sites. The two blog sites, *sciencedaily.com* and *phys.org*, used the highest and third highest amount of JARGON (both above 20%), and showed high levels of MAIN compared to other sites, especially *phys.org*, which had one of the highest rates of MAIN (close to 0.20, or 20% of sentences, see Figure 6 in the appendix). However, we do not find that these blog sites used lower IMPACT; in fact, we see the opposite. Blog sites use almost the same level of IMPACT as press releases. This might be due to some blog posts focusing on breaking science news, similar to press releases (Ranger and Bultitude, 2016), or due to rehosting press releases.

Delineating venues Based on our results for **H1–H4**, we see that strategies delineate different venues well. Blogs often focus on scientific terms and the main findings of a paper, press releases emphasize the impact of the findings, and magazines avoid complex scientific terms, instead telling stories and using active, engaging, writing.

We visualize these differences by representing each site as a vector of strategy rates (e.g., *phys.org* would be a vector of length eight) and calculate a single principal component from these vectors using principal component analysis.[7] Figure 3 plots each site along this axis, showing that the four venues we explore (blogs, press releases, magazines, and science journal magazines) group together clearly. The one overlap is science journal magazines, which fall between magazines and press releases. This is especially interesting because the goal of journal magazines is to both advertise research published in the journals (i.e., *na-*

ture.com reports *Nature* findings) while also being closer in length to a magazine, making them a mixture of both press releases and magazines.

H5: Higher quality science news articles use more STORY and ANALOGY than lower quality articles. To evaluate this hypothesis, we use the corpus of *New York Times* science articles introduced by Louis and Nenkova (2013a).[8] The corpus consists of three labels of article quality: TYPICAL, VERY GOOD, and GREAT. These labels were drawn from whether the article appeared in that year's "Best American Science Writing" anthology (GREAT), was written by an author whose work had appeared in the year's anthology (VERY GOOD), or was neither (TYPICAL). The articles were drawn from the *New York Times* annotated corpus (Sandhaus, 2008) and filtered for science-related keywords (e.g., biology, biologist).

For a clear differentiation of article quality, we apply our strategy classifiers to only the GREAT and TYPICAL articles in the dataset. We select science articles from the years 2001 to 2007 for a total of 55 GREAT articles (6,211 sentences) and 15,532 TYPICAL articles (1,079,768 sentences).[9] To test for significance we perform χ^2 tests of independence and augment these with the ϕ coefficient, which is similar to Cohen's d as an effect size calculation for categorical variables (Fleiss, 1994).

Results: Sentences in GREAT articles use STORY and ANALOGY slightly but significantly more often than TYPICAL articles (STORY: 0.38 vs. 0.33, $p < 0.001$, $\phi = 0.01$, ANALOGY: 0.05 vs. 0.03, $p < 0.001$, $\phi = 0.01$), supporting **H5**. GREAT articles also used more ACTIVE and PRESENT than TYPICAL articles (ACTIVE: 0.73 vs. 0.67,

---

[7]We also considered calculating multiple components, but a single component covers the majority (65%) of the variance across the full vector.

[8]Available at `http://www.cis.upenn.edu/~nlp/corpora/scinewscorpus.html`

[9]We obtained similar results when uniformly sampling 55 TYPICAL articles.

Figure 3: All sites plotted along a their single principal component, which accounts for 65% of the total variance in strategies. Sites from left to right are: *sciencedaily.com*, *phys.org*, *news.northwestern.edu*, *nature.com*, *news.stanford.edu*, *news.harvard.edu*, *rochester.edu*, *scientificamerican.com*, *sciencemag.org*, and *theatlantic.com*.

$p < 0.001$, $\phi = 0.01$; PRESENT: 0.23 vs. 0.21, $p < 0.001$, $\phi = 0.004$).

H6: Press releases that sensationalize the claims of the original journal paper use higher IMPACT and MAIN than press releases that do not. Sumner et al. (2014) introduced a dataset of 462 press releases annotated for the three categories of sensationalism and their associated journal articles from 20 prominent U.K. universities.[10] We split press releases into 'sensationalized' and 'not sensationalized' for each area of sensationalism.

Results: **H6** is not supported. The difference in IMPACT and MAIN usage is small and not significant for all types of sensationalism. Partly this is an encouraging result, as it suggests that using the strategies does not risk sensationalizing the science. Future work might explore strategies that help writers avoid sensationalism.

7 Related Work

We highlight additional areas of research relevant to our work beyond those already discussed.

Science communication. Over the past twenty years, science communication has shifted from improving scientific literacy to fostering participation in science (Hetland, 2014). A growing body of research shows that scientific literacy is only one of many factors that influence public decision making and cannot be divorced from cultural values (Nisbet and Scheufele, 2009; Bubela et al., 2009).

Scientific writing. There is a wealth of work exploring writing in scientific journals (i.e., when scientists communicate within their discipline). Because of the natural structure of scientific journal papers, much work has looked at ways of automatically identifying content in these papers (Liakata et al., 2010; Guo et al., 2010; Liakata et al., 2012). Kröll et al. (2014) examined the use of

guidelines for science journal papers. Our work instead focuses on general science writing.

8 Future Work

We annotated writing strategies at the sentence level, but some strategies, such as STORY and ANALOGY, might be better annotated at the fragment level to account for longer or shorter use of strategies. Future work can explore more fine-grained analysis of these strategies (e.g., with metaphor detectors; Gao et al., 2018). We also hope to build on these findings by exploring how effective the strategies are at engaging different readers.

9 Conclusion

In this paper we compile writing strategies from theory and practical advice, collect a large corpus and annotate a subset of it to measure strategies' use. We observe how strategies covary with intended audience. For example, blog sites, which target researchers, use more jargon and focus on the main findings of a paper, while magazine articles, which target a much broader audience of readers, tell more stories and use more active voice. Our findings also suggest that science newspaper articles judged by experts to have higher quality use more metaphorical language and tell more stories. We expect that our strategy formulations, classifiers, annotations and dataset will enable NLP-powered tools to support effective science communication for different audiences.

---

[10]While these press releases were written in the U.K., we expect science writing to be more invariant to regional dialects than other genres of writing.

5335

References

American Association for the Advancement of Science AAAS. Communication fundamentals.

Antonio Bella, Cesar Ferri, José Hernández-Orallo, and Maria Jose Ramirez-Quintana. 2010. Quantification via probability estimators. In *2010 IEEE International Conference on Data Mining*, pages 737–742. IEEE.

Emily M Bender and Alex Lascarides. 2019. Linguistic fundamentals for natural language processing ii: 100 essentials from semantics and pragmatics. *Synthesis Lectures on Human Language Technologies*.

Deborah Blum, Mary Knudson, Robin Marantz Henig, et al. 2006. *A field guide for science writers*. Oxford University Press, USA.

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. *BMJ*, 4.

Belinda Bray, Bev France, and John K Gilbert. 2012. Identifying the essential elements of effective science communication: What do the experts say? *International Journal of Science Education, Part B*, 2(1):23–41.

Anita Bruzzese. 2018. How to explain technical information to non-techies.

Tania Bubela, Matthew C Nisbet, Rick Borchelt, Fern Brunger, Cristine Critchley, Edna Einsiedel, Gail Geller, Anil Gupta, Jürgen Hampel, Robyn Hyde-Lay, et al. 2009. Science communication reconsidered. *Nature biotechnology*, 27(6):514–518.

Katie Burke. 2018. 12 tips for scientists writing for the general public.

Terry W. Burns, D. John O'Connor, and Susan M. Stocklmayer. 2003. Science communication: a contemporary definition. *Public Understanding of Science*, 12(2):183–202.

Dallas Card, Amber E. Boydstun, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In *ACL*, pages 438–444, Beijing, China.

Dallas Card and Noah A Smith. 2018. The importance of calibration for estimating proportions from annotations. In *NAACL*, pages 1636–1646.

Curtis Crawford. *Writing for a General Audience: Science Journalism*, page 251–258.

Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news article. In *EMNLP-IJCNLP*, pages 5636–5646, Hong Kong, China.

Mark Davies. 2009. The 385+ million word corpus of contemporary american english (1990–2008+): Design, architecture, and linguistic insights. *International journal of corpus linguistics*, 14(2):159–190.

Joseph L Fleiss. 1994. Measures of effect size for categorical data.

Kristin Gagnier and Kelly Fisher. 2017. Communicating science to non-scientific audiences.

Ge Gao, Eunsol Choi, Yejin Choi, and Luke Zettlemoyer. 2018. Neural metaphor detection in context. In *EMNLP*, pages 607–613, Brussels, Belgium.

Dee Gardner and Mark Davies. 2013. A New Academic Vocabulary List. *Applied Linguistics*, 35(3):305–327. Downloaded from: https://www.academicvocabulary.info/.

Barbara J. Grosz and Candace L. Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.

Yufan Guo, Anna Korhonen, Maria Liakata, Ilona Silins Karolinska, Lin Sun, and Ulla Stenius. 2010. Identifying the information structure of scientific abstracts: An investigation of three different schemes. In *BioNLP*, page 99–107, USA. Association for Computational Linguistics.

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. In *ACL*.

Thomas Hayden, Michelle Nijhuis, et al. 2013. *The Science Writers' Handbook: Everything You Need to Know to Pitch, Publish, and Prosper in the Digital Age*. Da Capo Lifelong Books.

Per Hetland. 2014. Models in science communication: formatting public engagement and expertise. *Nordic Journal of Science and Technology Studies*, 2(2):5–17.

Judith A Holton. 2007. The coding process and its challenges. *The Sage handbook of grounded theory*, (Part III):265–89.

Daniel J Hopkins and Gary King. 2010. A method of automated nonparametric content analysis for social science. *American Journal of Political Science*, 54(1):229–247.

Mark Kröll, Gunnar Schulze, and Roman Kern. 2014. A study of scientific writing: Comparing theoretical guidelines with practical implementation. In *COLING Workshop on Synchronic and Diachronic Approaches to Analyzing Technical Language*, pages 48–55.

Christoph Kueffer and Brendon MH Larson. 2014. Responsible use of language in scientific writing and science communication. *BioScience*, 64(8):719–724.

Maria Liakata, Shyamasree Saha, Simon Dobnik, Colin Batchelor, and Dietrich Rebholz-Schuhmann. 2012. Automatic recognition of conceptualization zones in scientific articles and two life science applications. *Bioinformatics*, 28(7):991–1000.

Maria Liakata, Simone Teufel, Advaith Siddharthan, and Colin Batchelor. 2010. Corpora for the conceptualisation and zoning of scientific papers. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).

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. *arXiv preprint arXiv:1907.11692*.

Annie Louis and Ani Nenkova. 2013a. A corpus of science journalism for analyzing writing quality. *Dialogue & Discourse*, 4(2):87–117.

Annie Louis and Ani Nenkova. 2013b. What makes writing great? first experiments on article quality prediction in the science journalism domain. *Transactions of the Association for Computational Linguistics*, 1:341–352.

Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *ACL: system demonstrations*, pages 55–60.

Andrew Maynard and Dietram A. Scheufele. 2020. What does research say about how to effectively communicate about science?

Khanh Nguyen and Brendan O'Connor. 2015. Posterior calibration and exploratory analysis for natural language processing models. In *EMNLP*, pages 1587–1598, Lisbon, Portugal. Association for Computational Linguistics.

Kristian Hvidtfelt Nielsen and Rikke Schmidt Kjærgaard. 2011. News coverage of climate change in nature news and science now during 2007. *Environmental Communication*, 5(1):25–44.

Matthew C Nisbet and Dietram A Scheufele. 2009. What's next for science communication? promising directions and lingering distractions. *American journal of botany*, 96(10):1767–1778.

Horst Pöttker. 2003. News and its communicative quality: The inverted pyramid—when and why did it appear? *Journalism Studies*, 4(4):501–511.

Tzipora Rakedzon, Elad Segev, Noam Chapnik, Roy Yosef, and Ayelet Baram-Tsabari. 2017. Automatic jargon identifier for scientists engaging with the public and science communication educators. *PloS one*, 12(8). Downloaded from: http://scienceandpublic.com/.

Mathieu Ranger and Karen Bultitude. 2016. 'the kind of mildly curious sort of science interested person like me': Science bloggers' practices relating to audience recruitment. *Public Understanding of Science*, 25(3):361–378.

Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12).

Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. 2020. Social bias frames: Reasoning about social and power implications of language. In *ACL*.

Maarten Sap, Marcella Cindy Prasetio, Ariel Holtzman, Hannah Rashkin, and Yejin Choi. 2017. Connotation frames of power and agency in modern films. In *EMNLP*.

Scientifica. Tips for communicating your scientific research to non-experts.

William Strunk. 2007. *The elements of style*. Penguin.

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*.

Chenhao Tan and Lillian Lee. 2014. A corpus of sentence-level revisions in academic writing: A step towards understanding statement strength in communication. In *ACL*, pages 403–408, Baltimore, Maryland. Association for Computational Linguistics.

Raghuram Vadapalli, Bakhtiyar Syed, Nishant Prabhu, Balaji Vasan Srinivasan, and Vasudeva Varma. 2018. Sci-blogger: A step towards automated science journalism. In *CIKM*, page 1787–1790. ACM.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. In *Advances in Neural Information Processing Systems*, pages 9051–9062.

Table 5: Resources used to identify scientific communication writing strategies.

Title	Type
12 Tips for Scientists Writing for the General Public (Burke, 2018)	Online article
Communication Fundamentals (AAAS)	Online article
Communicating with the Public from AAAS (Gagnier and Fisher, 2017)	Online article
Explaining Tech to Non-Techies (Bruzzese, 2018)	Online article
Tips for Communicating Scientific Research to Non Experts (Scientifica)	Online article
What Does Research Say about Effective Communication about Science? (Maynard and Scheufele, 2020)	Online article
Identifying Essentials of Scientific Communication (Bray et al., 2012)	Journal article
Responsible Use of Language in Scientific Writing and Science Communication (Kueffer and Larson, 2014)	Journal article
'The Kind of Mildly Curious Sort of Science Interested Person Like Me' (Ranger and Bultitude, 2016)	Journal article
Science Journalism (Writing for a General Audience) (Crawford)	Book chapter
A Field Guide for Science Writers (Blum et al., 2006)	Book
The Science Writer's Handbook (Hayden et al., 2013)	Book

A Appendix

A.1 Open-Coding Details

Using the selected style guides, two authors open-coded (Holton, 2007) the guidelines from each guide and grouped these guidelines into suggested writing strategies. Some resources had lists of guidelines (e.g., "12 ways to..."), for which the authors coded each listed guideline as a separate strategy. For resources in prose (e.g., books and academic articles), the authors highlighted all guidance on writing strategies for science communication (e.g., "Have an engaging, new, first sentence."). Because the eventual goal was to identify these strategies in a document, the authors focused on document-specific strategies rather than process-specific strategies (e.g., "make sure to have a friend read the draft before sending it in."). Table 5 lists all resources used.

A.2 Corpus Collection

We selected universities based on the Carnegie Classification of Institutions of Higher Education[11] for large 4-year universities with doctoral programs and very high research activity (i.e., R1 institutions) in the US. We additionally filtered for STEM dominant research institutions. We randomly sampled 10 university websites from this filtered set of universities; however, many universities either did not have a single unified press department (e.g., each

school handled press separately), or the majority of press was unrelated to research output. As Table 1 shows, a majority of articles came from blog sites, while few came from press releases. This is due to the fact that press releases focus on research coming from that particular institution, substantially limiting the number of articles produced by these sites.

A.3 Cleaning Keywords

We selected the following keywords for filtering based on inspection of politicized articles from the sites we scraped between 2016 and 2019. All keywords are lower cased: 'trump', 'president', 'republican', 'refugee', 'congress', 'country', 'obama', 'senate', 'white house', 'democrat', 'political', 'epa', 'attorney', 'politics'. An article was considered political if the title contained any of the keywords and the body contained at least 4 of the keywords. We inspected all articles selected for annotation (337) and found none that were political.

A.4 Annotation Agreement

Table 6 reports Krippendorff's α at the sentence level for each strategy. Because most strategies do not occur often (i.e., usually less than 10% of sentences), simple agreement rate skews high due to a majority of negative examples. α corrects for this skew by taking into account random chance of overlapping annotations.
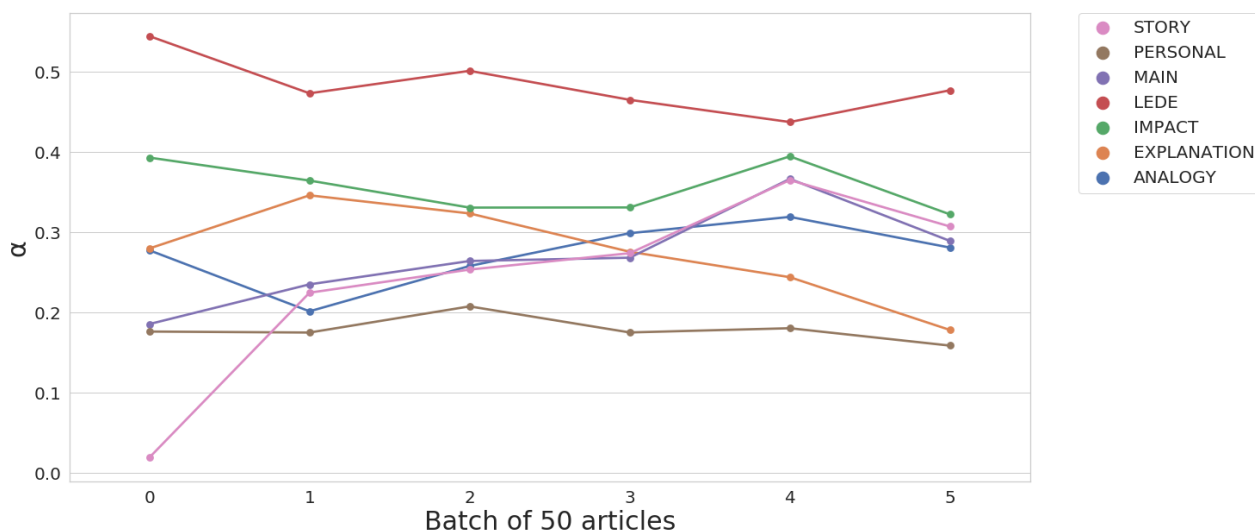
---

[11] https://carnegieclassifications.iu.edu/

Figure 4: Sentence-level α agreement over time.

Figure 5: Interface for annotation and example annotations.

Table 6: Sentence level Krippendorff's α for each writing strategy.

Strategy	Sentence level α
LEDE	0.47
MAIN	0.29
IMPACT	0.31
EXPLANATION	0.17
ANALOGY	0.25
STORY	0.29
PERSONAL	0.16

A.5 Pretraining Details

We followed the pretraining recommendations of Gururangan et al. (2020) and pretrain RoBERTa in two additional steps: domain- and task-adaptive pretraining. Both steps are to tailor the model to domain and task specific language. Domain adaptive pretraining was done on 11.90M articles from REALNEWS (Zellers et al., 2019) for 12.5k training steps and task adaptive pretraining was done on 100k articles from a held out portion of our corpus for 10 epochs. Hyperparameters for pretraining are in Table 7.

A.6 Finetuning Details

Articles were broken down into sentences for classification. We employed random search for hyperparameter tuning with 5-fold cross validation on the training set of the annotated articles. We ran a total of 10 search trials. Table 8 details the final

Table 7: Hyperparameters for domain- and task-adaptive pretraining. Based on pretraining in (Gururangan et al., 2020).

Hyperparameter	Assignment
number of epochs	10 (Task) or 12.5K (Domain)
batch size	256 (Task) or 2058 (Domain)
learning rate	0.0001 or 0.0005
learning rate optimizer	Adam
Adam epsilon	1e-6
Adam beta weights	0.98
Weight decay	0.01
warmup proportion	0.06
Learning rate decay	linear

Table 8: Final hyperparameters for finetuning the science strategy classifiers and bounds for hyperparameter tuning random search.

Hyperparameter	Assignment	Bounds
Number of epochs	3	[3, 5, 10]
Batch size	32	[16, 32]
Learning rate	2e-5	[1e-5, 2e-5, 3e-5]
Warmup proportion	0	[0, 0.06, 0.1]
Weight decay	.001	[.001, .01, .02]
Max sequence length	128	[64, 128, 256]

hyperparameters for our classifiers. Table 4 reports the precision, recall, and accuracy, calibration error and F_1 scores of the finetuned classifiers on the held-out test set.
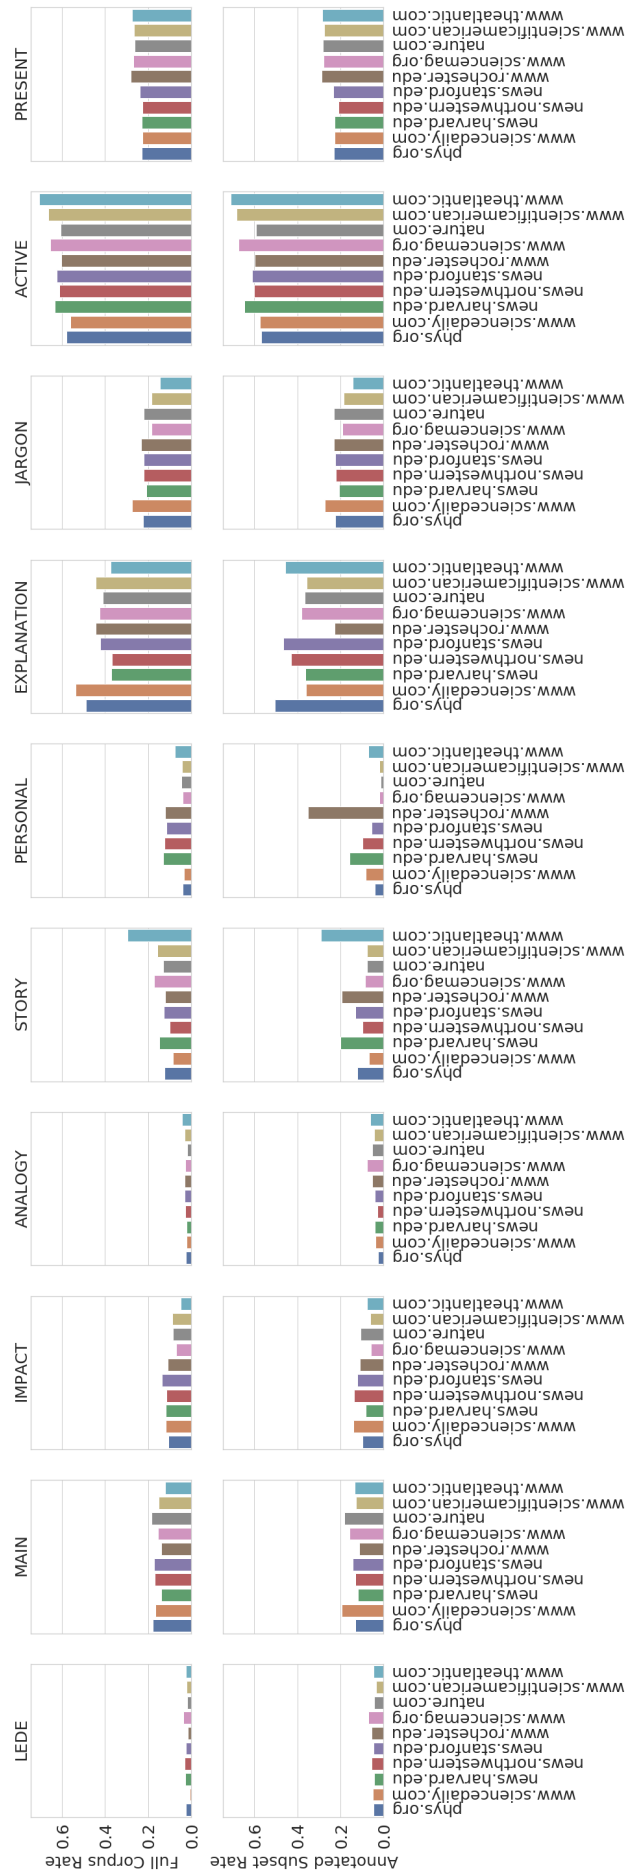
Figure 6: Strategy rate in full corpus based on classifier predictions (top) and in annotated subset (bottom).

Table 9: Strategies, examples of their descriptions in guidelines, and examples of their use.

Strategy	Example guideline	Example sentence
LEDE	Have an engaging, new, first sentence	On Wednesday, astronomers released what they said were the most detailed images ever taken of the surface of our sun.
MAIN	Give biggest, most important findings only	In their study, published in the journal Science Advances, the researchers describe a newly identified biomarker for detection of liver metastases.
IMPACT	Remember, "Why should I care?" for the reader	As date-palm growers adapt to climate change and battle pests and diseases, they might want to tap into the pool of ancient genes hidden in archaeological archives.
EXPLANATION	No matter how complicated a topic, the audience should be able to get the big idea	This idea suggests that as humans increasingly relied on peaceable social interactions to flourish, our ancestors began selecting mates with less aggressive features for facial appearance and other traits.
ANALOGY	Relate complex topics to simple ones (e.g., use metaphors)	The male climbed onto a platform and changed positions like a swimsuit model posing for a photograph
STORY	Tell stories for your reader	Ms. Moser was 23. It had taken her months to convince the clinic at NewYork-Presbyterian Hospital/Columbia University Medical Center in Manhattan that she wanted, at such a young age, to find out whether she carried the gene for Huntington's disease.
PERSONAL	Give readers a personal picture of scientists	But Dobson, bounding ahead in khaki hiking pants with her blond ponytail swinging, appears unfazed.
JARGON	Write in English (don't use jargon)	So if this black hole is, at least in astronomical terms, right there, how has it eluded astronomers for so long?
ACTIVE	Use the active voice	At night, hippos wander into grasslands to graze. During the day, they return to rivers to keep cool and protect themselves from sunburn.
PRESENT	Use the present tense	"Life continues but I don't think Dominica will ever be the same again," John says.

Table 10: Examples of highlights for JARGON, ACTIVE, and PRESENT. Excerpts were randomly sampled from articles automatically classified as having high (top 10% of articles) and low (bottom 10%) of each measure. **Black words** are those highlighted by our automatic measures. Red is incorrect highlights by the measures and blue are words our measures did not highlight that we believe should have been.

| Label	Sentences
High ACTIVE	The **challenges** associated with news writing, meanwhile, are...well, **they** 're challenging.
	All **four** regularly write about policy in popular news outlets — particularly prolific are Frakt and Carroll , **who** write for The New York Times.'
	For example, **they** are more likely to be immigrants.
	According to a team of scientists led by Nenad Sestan at Yale School of Medicine, this **process** might play out over a much longer time frame , and perhaps isn't as inevitable or irreparable as commonly believed.
	Observations from these scopes could reveal the planet's rotation rate, the composition and thickness of its atmosphere, and whether **it** has clouds.
Low ACTIVE	These secondary sediments were later eroded in the delta, exposing an inverted relief of the structure that is observed today.
	According to the World Health Organization, most significant **constituents** of air pollution include particulate matter (PM), ozone, nitrogen dioxide, and sulfur dioxide.
	Ke, working together with his graduate student Pengfei Wang, was instrumental in advancing the technology to its new version.
	Being able to touch, explore the shape, feel the weight and even smell the replica of an artefact has the potential to transform cultural heritage experiences.
	Some **deployments** might seem unusual .
High PRESENT	A stubborn myth **persists** that when policymakers **manage** recreational fishing they **'re managing** a food source.
	Professor Tanja Kallio and doctoral candidate Sami Tuomi **consider** the realisation of this goal entirely possible.
	"However, scientifically we **are** in the dark about the consequences of rewilding, and we **worry** about the general lack of critical thinking surrounding these often very expensive attempts at conservation.
	They also **suggest** that angler organizations should be more involved in promoting more responsible management processes and monitoring.
	First, the carbon nanotube and a solvent **are** combined in one vessel, while a nitrogen-containing compound and a solvent go into another.
Low PRESENT	The archaeologists identified the remains of Captain Matthew Flinders by the lead plate placed on top of his coffin.
	His team found a way to reengineer inhibitory interneurons to improve their function.
	"We were very lucky that Captain Flinders had a breastplate made of lead, meaning it would not have corroded."
	Near-infrared observations conducted with SPHERE allowed the astronomers to decompose the observed continuum emission into four components : young stellar population (about 120 million years old), hot dust (with a temperature of around 800 K), scattered light from the hidden Seyfert 1 nucleus and a very hot stellar background.
	These chemicals **are** potentially found in a huge variety of everyday products, including disinfectants, pesticides and toiletries.
High JARGON	The **researchers** found that **sustained** and unprecedented rise in **infant mortality** in England from 2014 to 2017 was not **experienced** evenly across the **population**.
	Exploiting this **reduction** of **complexity** and **degree** of **control** the team was able to **monitor** the **microscopic processes** in their **quantum** many body **system** and to **identify** ways to **enhance** and **manipulate** the **magnetic** order in their **system**.
	Often **patients** have to stop taking **medication** because of **adverse** side **effects** and wait for their bodies to recover before they can begin again, Shimada said.
	The next step for Fang and his **research** team is to **develop** computer **stimulations** to understand the **effects** of nanoparticle **shapes** sizes and surface modifiers.
	Exposure to **potentially harmful chemicals** is a reality of life.
Low JARGON	We are looking for an **alternative location** outside of Amsterdam, the plan says.
	These days unlicensed, recognizable **portrayals** of guns in games look from the outside the same as they did in the days of marketing deals: the guns look real and shoot well.
	Dora Linda Nishihara was driving in San Antonio one dark evening in early December when she suddenly disappeared from sight .
	The others didn't **respond** to requests for comment.
	"We were told if we become the first couple to do this **experiment** we'll become famous and HBO already tried to reach me", Yevgenievna says.
	She has been deaf since birth.
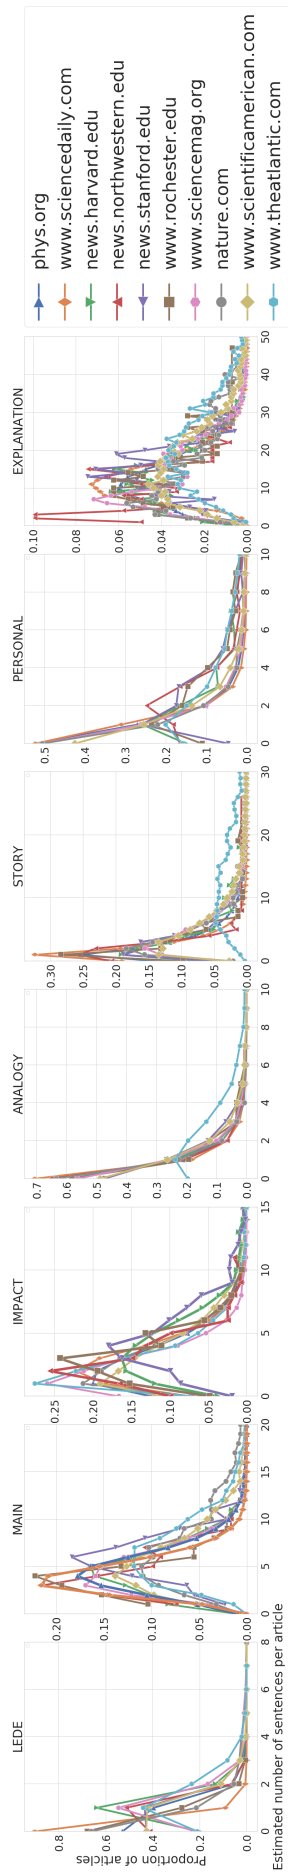
Figure 7: Histogram of number of sentences per article (x-axis) estimated to use each strategy (proportions on y-axis).