

# Making Science Simple: Corpora for the Lay Summarisation of Scientific Literature

Tomas Goldsack<sup>1</sup>, Zhihao Zhang<sup>2</sup>, Chenghua Lin<sup>1\*</sup>, Carolina Scarton<sup>1</sup>

<sup>1</sup>Department of Computer Science, University of Sheffield, UK

<sup>2</sup>School of Economics and Management, Beihang University, China

{tgold sack1, c.lin, c.scarton}@sheffield.ac.uk

zhzhang@buaa.edu.cn

## Abstract

Lay summarisation aims to jointly summarise and simplify a given text, thus making its content more comprehensible to non-experts. Automatic approaches for lay summarisation can provide significant value in broadening access to scientific literature, enabling a greater degree of both interdisciplinary knowledge sharing and public understanding when it comes to research findings. However, current corpora for this task are limited in their size and scope, hindering the development of broadly applicable data-driven approaches. Aiming to rectify these issues, we present two novel lay summarisation datasets, PLOS (large-scale) and eLife (medium-scale), each of which contains biomedical journal articles alongside expert-written lay summaries. We provide a thorough characterisation of our lay summaries, highlighting differing levels of readability and abstractiveness between datasets that can be leveraged to support the needs of different applications. Finally, we benchmark our datasets using mainstream summarisation approaches and perform a manual evaluation with domain experts, demonstrating their utility and casting light on the key challenges of this task. Our code and datasets are available at [https://github.com/TGoldsack1/Corpora\\_for\\_Lay\\_Summarisation](https://github.com/TGoldsack1/Corpora_for_Lay_Summarisation).

## 1 Introduction

Scientific publications contain information that is essential for the preservation and progression of our understanding across all scientific disciplines. Typically being highly technical in nature, such articles tend to assume a degree of background knowledge and make use of domain-specific language, making them difficult to comprehend for one lacking the required expertise (i.e., a *lay person*). These factors often limit the impact of research to only its direct community (Albert et al., 2015, 2022) and, more

<sup>1\*</sup>Corresponding author.

## Technical Abstract

The virus SARS-CoV-2 can exploit biological vulnerabilities (e.g. host proteins) in susceptible hosts that predispose to the development of severe COVID-19. To identify host proteins that may contribute to the risk of severe COVID-19, we undertook proteome-wide genetic colocalisation tests, and polygenic (pan) and cis-Mendelian randomisation analyses leveraging publicly available protein and COVID-19 datasets...

## Lay Summary

Individuals who become infected with the virus that causes COVID-19 can experience a wide variety of symptoms. These can range from no symptoms or minor symptoms to severe illness and death. Key demographic factors, such as age, gender and race, are known to affect how susceptible an individual is to infection. However, molecular factors, such as unique gene mutations and gene expression levels can also have a major impact on patient responses by affecting the levels of proteins in the body...

Figure 1: The first few sentences of the abstract and lay summary of an eLife article, illustrating differences in the language and focus on background information.

dangerously, can cause readers (members of the public, journalists, etc.) to misinterpret research findings (Kuehne and Olden, 2015).

This latter point is especially important for biomedical research which, in addition to having particularly dynamic and confusing terminology (Smith, 2006; Peng et al., 2021), has the potential to directly impact people’s decision-making regarding health-related issues, with a pertinent example of this being the widespread misinformation seen during the COVID-19 pandemic (Islam et al., 2020). Aiming to address these challenges, some academic journals choose to publish *lay summaries* that clearly and concisely explain the context and significance of an article using non-specialist language. Figure 1 illustrates how simplifying jar-

gon (e.g., “SARS-CoV-2” → “the virus that causes COVID-19”) and focusing on background information allows a reader to better understand a complex scientific topic. However, in addition to placing an extra burden on authors, lay summaries are not yet ubiquitous and focus only on newly published articles.

Automatic text summarisation can provide significant value in the generation of scientific lay summaries. Although previous use of summarisation techniques for scientific articles has largely focused on generating a technical summary (e.g., the abstract), only a few have addressed the task of lay summarisation and introduced datasets to facilitate its study (Chandrasekaran et al., 2020; Guo et al., 2021; Zaman et al., 2020). However, compared to datasets ordinarily used for training supervised summarisation models, these resources are relatively small (ranging from 572 to 6,695 articles), presenting a significant barrier to the deployment of large data-driven approaches that require training on large amounts of parallel data. Furthermore, these resources are somewhat fragmented in terms of their framing of the task, making use of article and summary formats that limit their applicability to broader biomedical literature. These factors hinder the progression of the field and the development of usable models that can be used to make scientific content accessible to a wider audience.

To help alleviate these issues, we introduce **two new datasets** derived from different academic journals **within the biomedical domain** - PLOS and eLife (§3). Both datasets use the full journal article as the source, enabling the training of models which can be broadly applied to wider literature. PLOS is significantly larger than currently available datasets and makes use of short author-written lay summaries (150-200 words), whereas eLife’s summaries are approximately twice as long and written by expert editors who are well-practiced in the simplification of scientific content. Given these differences in authorship and length, we expect the lay summaries of eLife to simplify content to a greater extent, meaning our datasets are able to cater to different audiences and applications (e.g., personalised lay summarisation). We confirm this via an in-depth characterisation of the lay summaries within each dataset, quantifying ways in which they differ from the technical abstract and from each other (§4). Finally, we benchmark our datasets with popular summarisation approaches

using automatic metrics and conduct an expert-based manual evaluation, highlighting the utility of our datasets and key challenges for the task of lay summarisation (§5). This paper also presents a literature review (§2), conclusions (§6), and a discussion on its limitations (§7).

## 2 Related Work

Past attempts to automatically summarise scientific content in layman’s terms have been scarce, with the most prominent example being the LaySumm subtask of the CL-SciSumm 2020 shared task series (Chandrasekaran et al., 2020) which attracted a total of 8 submissions. Alongside the task, a training corpus of 572 articles and author-generated lay summaries from a multi-disciplinary collection of Elsevier-published scientific journals was provided, with submissions being evaluated on a blind test set of 37 articles. It was noted by the task organisers that the data provided was insufficient for training a model to produce a realistic lay summary.

Guo et al. (2021) also make use of a single publication source to retrieve lay summaries: The Cochrane Database of Systematic Reviews (CDSR). Their dataset contains the abstracts of 6,695 systematic reviews paired with their respective plain-language summaries, covering various healthcare domains. Although larger than other available datasets for lay summarisation, CDSR is constrained in that it only uses the abstracts of systematic reviews as source documents, and thus models trained using CDSR will be unlikely to generalise well to inputs that are longer than an abstract or the abstracts of other types of publication.

Alternatively, Zaman et al. (2020) introduce a dataset derived from the ‘Eureka-Alert’ science news website for the combined tasks of simplification and summarisation. Summaries consist of news articles (average length > 600 words) that aim to describe the content of a scientific publication to the non-expert. However, the extensive size of reference summaries is likely to present additional challenges in model training and their news-based format limits their applicability (e.g., in automating lay summarisation for journals).

Compared to previous resources, our datasets contain articles and lay summaries of a format that we consider to be more broadly applicable to wider literature. Additionally, PLOS is significantly larger than those currently available (over 4× larger than CDSR) and eLife contains sum-

Dataset	# Docs	Doc		Summary	
		# words	# words	# sents	# sents
LaySumm	572	4,426.1	82.15	3.8	
Eureka-Alert	5,204	5,027.0	635.6	24.3	
CDSR	6,695	576.0	338.2	16.1	
<b>PLOS</b>	27,525	5,366.7	175.6	7.8	
<b>eLife</b>	4,828	7,806.1	347.6	15.7	

Table 1: Statistics of lay summarisation datasets, with ours given in **bold**. Words and sentences (sents) are average values.

maries written by expert editors. Furthermore, our work is the first to provide two datasets with different levels of readability, thus supporting the needs of different audiences and applications. Through each of these factors, we hope to enable the creation of more usable lay summarisation models.

### 3 Our Datasets

We introduce two datasets from different biomedical journals (PLOS and eLife), each containing full scientific articles paired with manually-created lay summaries. For each data source, articles were retrieved in XML format and parsed using Python to retrieve the lay summary, abstract, and article text.<sup>1</sup> In line with previous datasets for scientific summarisation (Cohan et al., 2018), the article text is separated into sections, and the heading of each section is also retrieved. Sentences are segmented using the PySBD rule-based parser (Sadvilkar and Neumann, 2020), which we empirically found to outperform neural alternatives. We separate our datasets into training, validation, and testing splits at a ratio of 90%/5%/5%. Statistics describing the contents of our datasets and that of past lay summarisation datasets are given in Table 1.

**PLOS** The Public Library of Science (PLOS) is an open-access publisher that hosts influential peer-reviewed journals across all areas of science and medicine. Several of these journals require authors to submit an *author summary* alongside their work, defined as a 150-200 word non-technical summary aimed at making the findings of a paper accessible to a wider audience, including non-scientists.<sup>2</sup> The journals in question focus specifically on Biology,

<sup>1</sup>For each article, we also retrieve a number of *keywords* from the meta-data, providing an indication of the high-level topics covered within the article.

<sup>2</sup>Source of PLOS author summary definition: <https://journals.plos.org/plosgenetics/s/submission-guidelines>

Metric	Abstract		Lay Summary	
	PLOS	eLife	PLOS	eLife
FKGL↓	15.04	15.57	14.76	10.92
CLI↓	16.39	17.68	15.90	12.51
DCRS↓	11.06	11.78	10.91	8.83
WordRank↓	9.08	9.21	8.98	8.68

Table 2: Mean readability scores for abstracts and lay summaries from our datasets. For all metrics, a lower score indicates greater readability.

Computational Biology, Genetics, Pathogens, and Neglected Tropical Diseases.

**eLife** eLife is an open-access peer-reviewed journal with a specific focus on biomedical and life sciences. Of the articles published in eLife, some are selected to be the subject of a *digest*, a simplified summary of the work written by expert editors based on both the article itself and questions answered by its author. Similarly to PLOS, these digests aim to explain the background and significance of a scientific article in language that is accessible to non-experts (King et al., 2017).

## 4 Dataset Analysis

We carry out several analyses comparing the lay summaries of our datasets to the respective technical abstracts. Through these analyses, we seek to highlight and quantify the key differences between these two different types of summary, as well as those present between the lay summaries of our two datasets. Specifically, we focus on readability (§4.1), rhetorical structure (§4.2), vocabulary sharing (§4.3), and abstractiveness (§4.4).

### 4.1 Readability

We assess the readability of our lay summaries and abstracts using several established metrics. Specifically, we employ Flesch-Kincaid Grade Level (FKGL), Coleman-Liau Index (CLI), Dale-Chall Readability Score (DCRS), and WordRank score.<sup>3</sup> FKGL, CLI, and DCRS provide an approximation of the (US) grade level of education required to read a given text. The formula for FKGL surrounds the total number of sentences, words, and syllables present within the text, whereas CLI is based on the number of sentences, words, and characters. Alternatively, DCRS measures readability using the average sentence length and the number of *familiar*

<sup>3</sup>Computed using the `textstat` and `EASSE` (Alva-Manchego et al., 2019) packages.

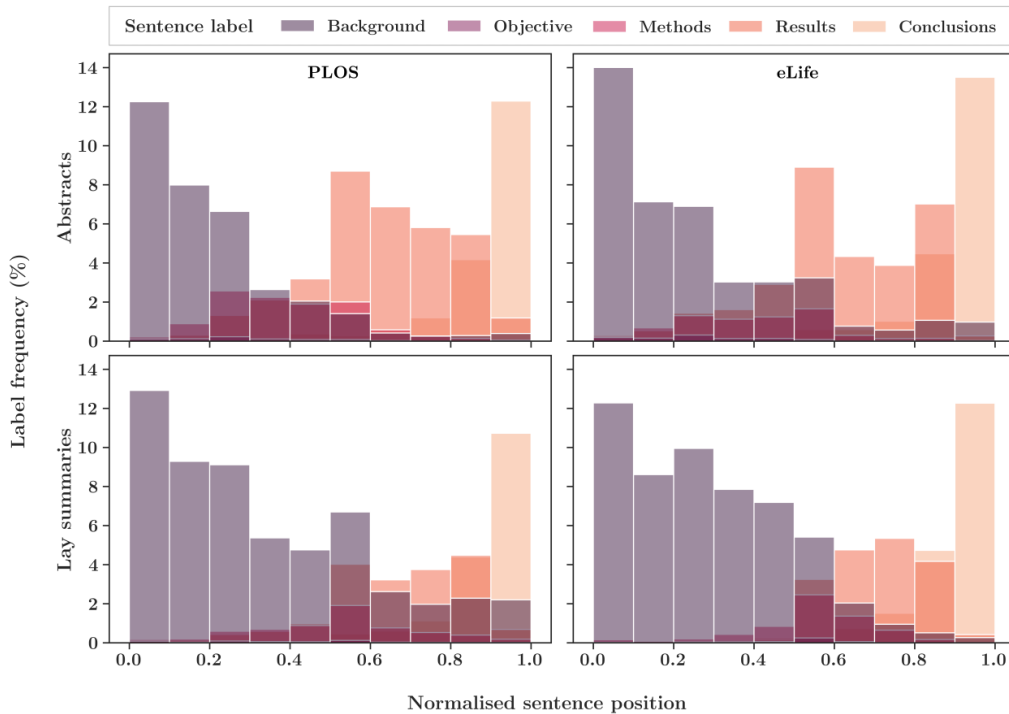


Figure 2: Barplot visualising the rhetorical class distributions in our abstracts and lay summaries.

words present, using a lookup table of the 3,000 most commonly used English words. Similarly, WordRank estimates the lexical complexity of a text based on how common the language is, using a frequency table derived from English Wikipedia.

The scores given in Table 2 show that the lay summaries of both datasets are consistently more readable than their respective abstracts across all metrics. Although these differences are small in some cases, in line with the findings of previous works (Devaraj et al., 2021), we find them all to be **statistically significant** by way of Mann–Whitney U tests ( $p < 0.05$ ). These results indicate that lay summaries are more readable than technical abstracts in terms of both syntactic structure and lexical intelligibility. Additionally, the lay summaries from eLife obtain lower readability scores than those of PLOS across all metrics, confirming our expectation that they are suitable for less technical audiences.<sup>4</sup>

## 4.2 Rhetorical Structure

Rhetoric is another important factor when assessing the comprehensibility of a text. Specifically, a lay person will require a much larger focus on the background of a scientific article than an expert

<sup>4</sup>Manual inspection of the summaries from each dataset also support this.

Label	Abstract		Lay Summary	
	PLOS	eLife	PLOS	eLife
Background	35.40	41.05	58.11	55.03
Objective	0.76	1.06	0.54	0.47
Methods	10.26	6.73	6.24	6.23
Results	34.75	30.60	17.86	18.23
Conclusions	18.83	20.55	17.26	18.83

Table 3: Mean percentage of each rhetorical label within our abstracts and lay summaries.

in order to understand the significance of its findings (King et al., 2017), thus we would expect lay summaries to focus more on such aspects.

To provide further insight into the structural differences between abstracts and lay summaries, we classify all sentences within each based on their rhetorical status. To do this, we make use of PubMed RTC (Dernoncourt and Lee, 2017), a dataset containing the 20,000 biomedical abstracts retrieved from PubMed, with each sentence labelled according to its rhetorical role (roles: Background, Objective, Methods, Results, Conclusions). We use PubMed RTC to train the BERT-based sequential classifier introduced by Cohan et al. (2019) due to its strong reported performance (92.9 micro F1-score), before applying this model to lay summary and abstract sentences from our datasets.

Figure 2 provides a visualisation of how the fre-

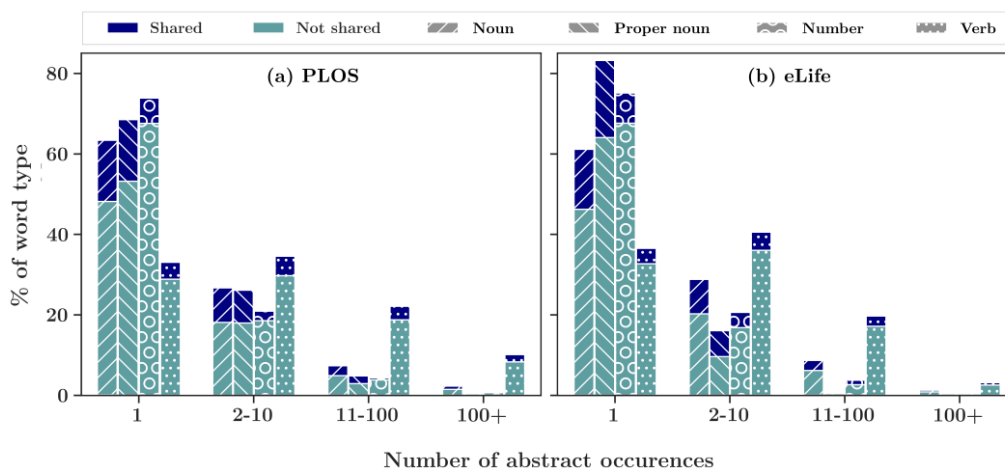


Figure 3: Stacked barplot showing how regularly (on average) abstract content words are shared with the respective lay summaries (as a % of all words of that type), separated by number of abstract occurrences.

quency of each rhetorical class changes according to the sentence position within our summaries. For each sub-graph, observing the pattern of most frequent labels (tallest bars) across all positions allows us to get an idea of the dominant rhetorical structure. In Table 3, we further quantify the difference in structure by giving the average percentage of each label present in the different summaries.

For both datasets, we see a similar pattern when comparing abstract and lay summary distributions. Specifically, a much greater portion of lay summary sentences is dedicated to explaining the relevant background information (“Background”). This is unsurprising, as such information is essential to understanding the motivation and significance of any work and, thus, would be of great value to a non-expert. This additional focus on “Background” comes at the expense of sentences focusing on “Results” and (to a lesser extent) “Methods”, which are less frequent within lay summaries. Again, this is to be expected, as these details are less meaningful to an audience without domain expertise.

### 4.3 Content Words

Aiming to determine what terminology is shared between summary types, we analyse the frequency at which *content words* occur simultaneously within abstracts and lay summaries. We treat nouns, proper nouns, verbs, and numbers as content words, and we extract these from the summaries using ScispaCy (Neumann et al., 2019), a library that specialises in the processing of biomedical texts.<sup>5</sup> Figure 3 shows the results of our analysis, visual-

<sup>5</sup>en\_core\_sci\_scibert model used for POS tagging, derived from SciBERT (Beltagy et al., 2019).

ising the average rate at which different types of content word from the abstracts are shared with lay summaries. We divide our analysis based on the number of abstracts content words occur in, allowing us to observe how the ubiquity of a word within the corpus affects the rate at which it is shared.<sup>6</sup>

Generally, we observe similar patterns for content words between datasets. Firstly, regardless of word type or number of abstract occurrences, we find that abstract content words are rarely shared with lay summaries (i.e., ‘shared’ % < ‘not shared’ % for all bars). This is indicative of a clear shift in content and/or vocabulary when it comes to the creation of a lay summary. For each dataset, we can also see that the vast majority of content words of all types, except verbs, occur in 10 or fewer abstracts (> 90% on average for both datasets), with most of these occurring within a single abstract. We found that the majority of these low-frequency content words are highly specific to the topic of a single article or small group of articles. In the case of nouns and proper nouns, we empirically observed these instances to often be highly technical terms (e.g., specific chemicals, lesser-known diseases, etc.), whereas numbers are typically exact numerical figures. It is understandable, therefore, that single-use content words of these types (noun, proper noun, and number) are rarely included in the lay summary, as they will likely be meaningless to a lay reader. The pattern exhibited by verbs differs significantly from that of other word types, as they

<sup>6</sup>E.g., For nouns within PLOS, 15.3% occur in a single abstract and are shared with lay summaries, and 48.2% occur in a single abstract and are not shared (i.e., total percentage of nouns that occur in a single abstract = 15.3 + 48.2 = 63.5%).

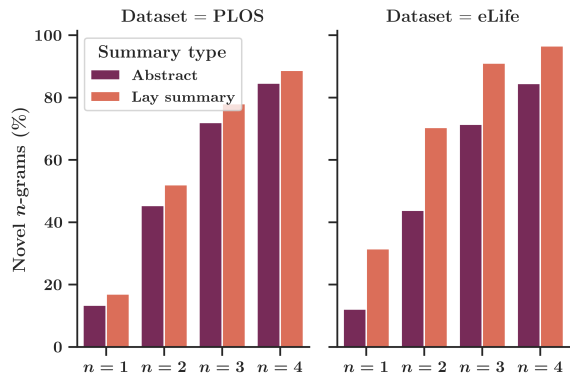


Figure 4: Barplot showing the percentage of novel  $n$ -grams for each summary type.

typically occur in a greater number of abstracts (most commonly being present within 2-10). For content words of all types, we observe that the ratio of ‘shared’ to ‘not shared’ generally increases in line with the number of abstract occurrences.<sup>7</sup>

#### 4.4 Abtractiveness

We follow the example of prior works (Sharma et al., 2019; See et al., 2017) by calculating the *abtractiveness* of our summaries using  $n$ -gram novelty, thus providing a measurement of the degree to which the summary uses different language to describe the content of the article. Specifically, for both abstracts and lay summaries, we compute the percentage of summary  $n$ -grams which are absent from their respective article. The results of this analysis are presented in Figure 4, where we can observe that lay summaries consistently contain more novel  $n$ -grams than abstracts across both datasets. However, the lay summaries of eLife, in addition to being approximately twice as long as those of PLOS (Table 1), appear to be significantly more abtractive. Alongside differences in readability (highlighted in §4.1), we believe these to be important distinctions that should be considered in determining a suitable dataset for a particular use case or application.

## 5 Experiments and Results

To help facilitate future work, we benchmark our datasets using popular heuristics-based, unsupervised, and supervised summarisation approaches (§5.1). Additionally, we provide further insight into these results via a detailed discussion (§5.2) and an expert-based manual evaluation (§5.3).

<sup>7</sup>Better illustrated by Table 7 in Appendix A, which gives

### 5.1 Baseline Approaches

For our heuristics-based approaches, we include the widely-used LEAD-3 baseline which simply uses the first three sentences of the main body of the text. As our lay summaries typically consist of more than three sentences, we also include LEAD- $k$ , with  $k$  being equal to the average lay summary length for each dataset (Table 1). Additionally, we include the scores obtained by the technical abstracts (ABSTRACT) and ORACLE\_EXT, a greedy extractive oracle (Nallapati et al., 2017) that provides an upper bound for the expected performance of extractive models.<sup>8</sup>

We benchmark four unsupervised extractive approaches: LSA (Steinberger and Jezek, 2004), LEXRANK (Erkan and Radev, 2004), TEXTRANK (Mihalcea and Tarau, 2004), and HIPORANK (Dong et al., 2021). For supervised models, we use the transformer-based BART base model (Lewis et al., 2020), which we fine-tune on our own datasets. To assess how the use of additional data in various forms can benefit performance, we include several other variants of this model which are described in the remainder of this subsection.<sup>9</sup>

**Additional training** As our datasets remain smaller than those used in other forms of summarisation, we experiment with BART<sub>PubMed</sub>, which is previously trained on the PubMed abstract generation dataset (Cohan et al., 2018) and fine-tuned on our own datasets. Aiming to assess how well models trained on PLOS can generalise to eLife and vice versa, we also include BART<sub>Cross</sub>, which is initially trained on the opposite dataset to that which it is eventually fine-tuned and evaluated on.

**Scaffolding** We also experiment with artificially enlarging our training data by way of a scaffold task. Inspired by CATTS (Cachola et al., 2020), we remove the article’s abstract from the input text and train the model to generate the abstract as a scaffold task to lay summarisation. In addition to showing whether training for abstract generation can benefit lay summarisation, results for this model will provide an indication of the baseline BART model’s reliance on the abstract content. Specifically, we include two copies of every article within our training data - one using the abstract as the reference summary and the other using the lay summary. We

the exact percentages shown in Figure 3.

<sup>8</sup>The extractive oracle is further explained in Appendix A.

<sup>9</sup>For all variants, we truncate the input to 1,024 tokens.

Approach	PLOS					eLife					
	R-1	R-2	R-L	FKGL	DCSR	R-1	R-2	R-L	FKGL	DCSR	
Heuristic	LEAD-3	25.46	6.35	22.89	15.08	12.66	17.93	3.66	16.45	13.30	12.65
	LEAD-K	35.50	8.69	32.33	14.94	11.88	34.12	6.73	32.06	11.89	10.58
	ABSTRACT	47.07	18.60	43.51	14.98	11.10	28.92	6.19	27.04	15.35	11.87
	ORACLE_EXT	56.43	31.24	52.88	15.28	11.20	46.38	11.48	43.82	13.18	10.51
Unsupervised	LSA	35.14	6.92	31.46	17.39	12.41	34.57	5.69	32.27	16.50	11.52
	LEXRANK	36.53	9.43	33.09	14.90	<b>10.51</b>	30.32	5.29	28.40	13.45	9.75
	TEXTRANK	36.03	8.71	32.04	20.42	11.83	31.45	5.28	29.22	19.13	11.28
	HIPORANK	42.16	11.41	38.01	14.91	11.82	29.95	5.38	27.44	13.58	12.32
Supervised	BART	<b>42.35</b>	12.96	38.57	14.62	12.07	<b>46.57</b>	<b>11.65</b>	<b>43.70</b>	10.94	9.36
	BART <sub>PubMed</sub>	42.12	12.70	38.34	14.75	12.09	46.20	11.36	43.21	11.24	9.64
	BART <sub>Cross</sub>	42.24	<b>13.52</b>	<b>38.63</b>	<b>14.24</b>	12.18	46.22	11.53	43.33	10.81	9.31
	BART <sub>Scaffold</sub>	39.47	9.73	35.79	14.45	11.74	45.28	10.99	42.51	<b>10.65</b>	<b>9.19</b>

Table 4: Performance of summarisation models on the test splits of each dataset (**R** = average ROUGE F1-score). The best non-heuristic scores for each metric are given in **bold**.

distinguish between the two by prepending the input document with the control tokens (`<|ABSTRACT|>`) or (`<|SUMMARY|>`). Documents within the validation and test splits are prepended the (`<|SUMMARY|>`) code to induce lay summary generation. This model is denoted by BART<sub>Scaffold</sub>.

## 5.2 Discussion

Table 4 presents the performance of the aforementioned approaches on the PLOS and eLife test splits using automatic metrics. In line with common practice for summarisation, we report the F1-scores of ROUGE-1, 2, and L (Lin, 2004). Additionally, we include FKGL and DCRS scores of the generated output (see §4.1), providing an assessment of the syntactic and lexical complexity, respectively.

**The importance of the abstract** Based on the ROUGE scores obtained by the ABSTRACT baseline, we can safely assume that the lay summaries of PLOS are much closer in resemblance to their respective abstracts than those of eLife. The importance of the abstract for lay summary generation is further highlighted by the ROUGE scores of the BART<sub>Scaffold</sub> model, which performs notably worse than the standard BART model on PLOS and slightly worse on eLife. These results suggest that having the abstract included within the model input provides significantly more benefit than using abstract generation as an auxiliary training signal.

**Extractive vs abstractive** In general, we would expect abstractive methods to have greater application for the task of lay summarisation due to their ability to transform (and thus, simplify) an input text. However, abstractive approaches have a ten-

dency to generate hallucinations, resulting in factual inconsistencies between the source and output that damages their usability (Maynez et al., 2020). Therefore, extractive approaches may still have utility for the task, especially if the comprehensibility of selected sentences is directly considered.

For ROUGE scores, we find that extractive baselines (i.e., all unsupervised and heuristic approaches) perform significantly better on PLOS than on eLife, aligning with our previous analysis (§4.4) which identified PLOS as the less abstractive dataset. Interestingly, readability scores achieved by extractive models on PLOS match and sometimes exceed those of abstractive BART models, although they are inconsistent. For eLife, abstractive methods (i.e., BART models) generally obtain superior scores for both ROUGE and readability metrics. In fact, the ROUGE scores achieved by BART exceed those obtained by ORACLE\_EXT, further indicating that abstractive methods have greater potential for this dataset.

**Use of additional data** As previously mentioned, artificially creating more data via an abstract-generation scaffold task results in a decrease in ROUGE scores for both datasets, indicating a reliance on the abstract content for lay summarisation. We also find that pretraining BART on Pubmed (BART<sub>PubMed</sub>) does very little to affect the performance, suggesting that habits learned for abstract generation do not transfer well to lay summarisation. Similarly, BART<sub>Cross</sub> achieves a performance close to that of the standard BART model. Overall, these results indicate that additional out-of-domain training does provide much benefit for

Dataset	Comp.	Layness	Factuality
PLOS	3.7	3.0	3.0
eLife	3.1	3.0	3.0

Table 5: Mean evaluator ratings (1-5) obtained by BART outputs for each metric.

lay summarisation, and alternative modelling approaches that make better use of available data may be a more promising route for future work.

### 5.3 Human evaluation

To further assess the usability of the generated abstractive summaries, we perform an additional human evaluation of our standard BART baseline model using two domain experts.<sup>10</sup>

Our evaluation uses a random sample of 10 articles from the test split of each dataset. Alongside each model-generated summary, judges are presented with both the abstract and reference lay summary of the given article. Using a 1-5 *Likert* scale, the annotators are asked to rate the model output based on three criteria: (1) *Comprehensiveness* – to what extent does the model output contain all information that might be necessary for a non-expert to understand the high-level topic of the article and the significance of the research; (2) *Layness* – to what extent is the content of the model output comprehensible (or readable) to a non-expert, in terms of both structure and language; (3) *Factuality* – to what extent is the model output factually consistent with the two other provided summaries. We choose not to provide judges with the full article text in an effort to minimise the complexity of the evaluation and the cognitive burden placed upon them.

Table 5 presents the average ratings from our manual evaluation. We calculate Krippendorff’s  $\alpha$  to measure inter-rater reliability, where we obtain values of 0.78 and 0.54 for PLOS and eLife, respectively. In addition to providing ratings, evaluators also provided comments on the general performance on each criterion for both datasets, providing further insights into model performance.

**Comprehensiveness** We can see from Table 5 that model outputs on PLOS are judged to be more comprehensive than on eLife. From evaluators’ comments, we understand that this largely results from extensive use of abstract content for PLOS, which is sometimes copied directly (or with minor

<sup>10</sup>Both judges have experience in scientific research and hold at least a bachelor’s degree in Biomedical Science.

edits) to the lay summary. For eLife, it was observed that new information (i.e., not contained in the reference abstract or lay summary) was often introduced which was irrelevant or confusing, potentially affecting the understanding of a lay reader.

**Layness** Interestingly, given the previously highlighted differences in readability, the average layness of the model output is judged to be equal for both datasets (3.0), suggesting a reasonable degree of content simplification. However, evaluators’ comments indicate that model outputs for each dataset were penalised for different reasons. For PLOS, the aforementioned use of abstract content often resulted in the inclusion of jargon terms that a lay reader would struggle to interpret. Alternatively, the language of eLife outputs was observed to be better suited to a lay audience but was sometimes simplified to a point that it could be misconstrued and mislead a reader, occasionally containing grammatical errors, typos, and repeated content.

**Factuality** Again, we find an equal average score of 3.0 given for factuality, suggesting the model struggles to produce factually correct outputs for both datasets. In fact, we found no output from either dataset was given a perfect score by both annotators, indicating that simplifying technical content accurately is a consistent problem. Evaluators’ comments highlight contradictions, unclear phrasing, and misrepresentation of entities as key contributing factors to factual inconsistencies. We regard this as an integral obstacle to overcome in the development of usable lay summarisation models and an essential focus for future research.

## 6 Conclusion

In this work, we have introduced PLOS and eLife, two new datasets for the lay summarisation of biomedical research articles. Compared to currently available resources, these datasets possess source article and summary formats that are more broadly applicable to wider literature, with PLOS also being larger by a significant margin. A thorough analysis of our lay summaries highlights key differences between datasets, enabling them to cater to the needs of different audiences and applications. Specifically, in addition to being approximately twice as long as those of PLOS, we find eLife summaries to be both more readable and abstractive, thus better suited to a less technical au-



dience. To facilitate future research, we benchmark our datasets with popular summarisation models using automatic metrics and conduct an expert-based human evaluation, providing further insight into the intricacies of model performance on our datasets and highlighting key challenges for the task of lay summarisation.

## 7 Limitations

Although we introduce the largest dataset available to the task of lay summarisation, our datasets remain smaller than those available for other forms of summarisation (e.g., abstract generation), where there exists datasets containing 100,000+ articles. This is largely due to the fact that lay summaries are less ubiquitous than other forms of summary (e.g., the abstract), only being used in a relatively small number of journals, of which only some are open-access and available to be utilised for such purposes as dataset creation.

On a related note, another potential limitation of our datasets is the fact they only cover a single broad domain - biomedicine. Again, this comes down to the availability of data, and the fact that the use of lay summaries is much less common in other scientific domains (e.g. Computer Science). There is, however, a reason for this disparity in the adoption of lay summaries between domains, as it is generally considered more important that the public have an awareness and understanding of research breakthroughs in health-related areas such as biomedicine. Therefore, we believe it is in these domains that automatic lay summarisation can provide the greatest benefit, although we also hope to address the lay summarisation of other domains in future work.

## 8 Acknowledgements

This work was supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation [grant number EP/S023062/1].

## References

Mathieu Albert, Elise Paradis, and Ayelet Kuper. 2015. [Interdisciplinary promises versus practices in medicine: The decoupled experiences of social sciences and humanities scholars](#). *Social Science and Medicine*, 126:17–25.

Mathieu Albert, Paula Rowland, Farah Friesen, and Suzanne Laberge. 2022. [Barriers to cross-disciplinary knowledge flow: The case of medical education research](#). *Perspect. Med. Educ.*, 11(3):149–155.

Fernando Alva-Manchego, Louis Martin, Carolina Scarton, and Lucia Specia. 2019. [EASSE: Easier automatic sentence simplification evaluation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 49–54, Hong Kong, China. Association for Computational Linguistics.

Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China. Association for Computational Linguistics.

Isabel Cachola, Kyle Lo, Arman Cohan, and Daniel Weld. 2020. [TLDR: Extreme summarization of scientific documents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4766–4777, Online. Association for Computational Linguistics.

Muthu Kumar Chandrasekaran, Guy Feigenblat, Eduard Hovy, and Anita de Waard. 2020. [Overview and insights from the shared tasks at scholarly document processing 2020: CL-SciSumm, LaySumm and](#). In *Proceedings of the First Workshop on Scholarly Document Processing*, pages 214–224. unknown.

Arman Cohan, Iz Beltagy, Daniel King, Bhavana Dalvi, and Dan Weld. 2019. [Pretrained language models for sequential sentence classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3693–3699, Hong Kong, China. Association for Computational Linguistics.

Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. [A discourse-aware attention model for abstractive summarization of long documents](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 615–621, New Orleans, Louisiana. Association for Computational Linguistics.

Franck Dernoncourt and Ji Young Lee. 2017. [PubMed 200k RCT: a dataset for sequential sentence classification in medical abstracts](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 308–313, Taipei, Taiwan. Asian Federation of Natural Language Processing.

- Ashwin Devaraj, Iain Marshall, Byron Wallace, and Junyi Jessy Li. 2021. [Paragraph-level simplification of medical texts](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4972–4984, Online. Association for Computational Linguistics.
- Yue Dong, Andrei Mircea, and Jackie Chi Kit Cheung. 2021. Discourse-Aware unsupervised summarization for long scientific documents. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1089–1102, Online. Association for Computational Linguistics.
- G. Erkan and D. R. Radev. 2004. [LexRank: Graph-based lexical centrality as salience in text summarization](#). *Journal of Artificial Intelligence Research*, 22:457–479.
- Yue Guo, Wei Qiu, Yizhong Wang, and Trevor Cohen. 2021. [Automated Lay Language Summarization of Biomedical Scientific Reviews](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(1):160–168.
- Md Saiful Islam, Tonmoy Sarkar, Sazzad Hossain Khan, Abu-Hena Mostofa Kamal, S M Murshid Hasan, Alamgir Kabir, Dalia Yeasmin, Mohammad Ariful Islam, Kamal Ibne Amin Chowdhury, Kazi Selim Anwar, Abrar Ahmad Chughtai, and Holly Seale. 2020. [COVID-19–Related Infodemic and Its Impact on Public Health: A Global Social Media Analysis](#). *The American Journal of Tropical Medicine and Hygiene*, 103(4):1621–1629.
- Stuart R F King, Emma Pewsey, and Sarah Shailes. 2017. [Plain-language Summaries of Research: An inside guide to eLife digests](#). *eLife*, 6:e25410.
- Lauren M Kuehne and Julian D Olden. 2015. [Lay summaries needed to enhance science communication](#). *Proceedings of the National Academy of Sciences*, 112(12):3585–3586.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). 7th International Conference on Learning Representations.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan Thomas McDonald. 2020. On faithfulness and factuality in abstractive summarization. In *Proceedings of The 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI’17*, page 3075–3081. AAAI Press.
- Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. [ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing](#). In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 319–327, Florence, Italy. Association for Computational Linguistics.
- Keqin Peng, Chuantao Yin, Wenge Rong, Chenghua Lin, Deyu Zhou, and Zhang Xiong. 2021. Named entity aware transfer learning for biomedical factoid question answering. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*.
- Nipun Sadvilkar and Mark Neumann. 2020. [PySBD: Pragmatic sentence boundary disambiguation](#). In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*, pages 110–114, Online. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Eva Sharma, Chen Li, and Lu Wang. 2019. [BIG-PATENT: A Large-Scale dataset for abstractive and coherent summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2204–2213, Florence, Italy. Association for Computational Linguistics.
- Barry Smith. 2006. [From concepts to clinical reality: An essay on the benchmarking of biomedical terminologies](#). *Journal of Biomedical Informatics*, 39(3):288–298.
- Josef Steinberger and Karel Jezek. 2004. Using latent semantic analysis in text summarization and summary evaluation. pages 93–100. 7th International Conference ISIM.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen,

Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. [Huggingface's transformers: State-of-the-art natural language processing](#).

Farooq Zaman, Matthew Shardlow, Saeed-Ul Hassan, Naif Radi Aljohani, and Raheel Nawaz. 2020. HTSS: A novel hybrid text summarisation and simplification architecture. *Inf. Process. Manag.*, 57(6):102351.

## A Appendix

**Availability of data sources** Both PLOS and eLife are open-access journals, with an emphasis on making scientific research accessible to a wide audience. PLOS articles are available to be mined, reused, and shared by anyone, as per their data mining policy.<sup>11</sup> eLife articles are available under the permissive CC-BY 4.0 license, and thus also available to be retrieved and shared for these purposes.<sup>12</sup> The data for PLOS and eLife was retrieved on 7/03/22 and 11/03/22, respectively. All articles and summaries are in English only. Our datasets are made available to the community to facilitate future research.

**Additional data processing details** Here we provide some additional details regarding the dataset creation process, building on the description given in §3. For eLife, the XML files retrieved were found to include multiple versions of the same article, identifiable by the article id which includes a version number. For these, we removed duplicates and kept only the most recent versions.

For both datasets, prior to extraction, we remove all Tables, Figures, and sections marked with the tag "supplementary-material". We also do not extract sections with the heading "acknowledgments". During sentence segmentation, we use a regular expression to identify and temporarily replace all "et al." occurrences with unique placeholder tokens, which are then replaced following segmentation. Following segmentation, we again use a regular expression to identify and remove all sentences which began with "DOI:" followed by a URL from both abstracts and lay summaries, as these were found to commonly occur at the end of both.

**Comparison to previous datasets** We provide a comparison of the readability and abstractiveness of lay summaries for all lay summarisation datasets in Table 6 and Figure 5, respectively.

<sup>11</sup><https://plos.org/text-and-data-mining>

<sup>12</sup><https://elifesciences.org/inside-elifesciences/6933fe8e/resources-for-developers>

**Supplementary figure details** Table 7 gives the exact percentages that are visualised in Figure 3, allowing for a more detailed analysis of content word sharing (e.g., calculating the ratio 'shared' to 'not shared' for different content word types).

**Baseline model details** Here we provide additional experimental details for our baselines approaches (Table 4). ORACLE\_EXT is a greedy oracle, which means it repeatedly extracts the next article sentence that will maximise the mean ROUGE scores (1, 2, and L) of the extracted summary, up to the maximum length (equal to the average lay summary length for a given dataset - Table 1).

For all BART models, we make use of the huggingface library (Wolf et al., 2019). Specifically, we use the "facebook/bart-base" model for baselines BART, BART<sub>Cross</sub>, and BART<sub>Scaffold</sub>, and we use the "mse30/bart-base-finetuned-pubmed" model for BART<sub>PubMed</sub>. Training was run (using 4x NVIDIA Tesla V100 SXM2 GPUs) for all models with AdamW optimisation (Loshchilov and Hutter, 2019) and an early stopping patience of 25 epochs, with the best model being selected by performance on the validation set (ROUGE-2).

All unsupervised baselines were run with default configurations.

**Automatic evaluation** For the calculation of ROUGE scores, we use the rouge-score Python package.<sup>13</sup> For FKGL and DCSR, we use the textstat Python package.<sup>14</sup>

**Human evaluation comments** Comments on the general model performance for each criterion provided by each annotator for our human evaluation are given in Figures 6 and 7 for PLOS and eLife, respectively.

**Lay summary examples** Full examples of lay summaries and their respective technical abstracts are given in Figures 8 and 9 for PLOS and eLife, respectively.

<sup>13</sup><https://pypi.org/project/rouge-score/>

<sup>14</sup><https://github.com/shivam5992/textstat>

Dataset	FKGL	CLI	DCRS	WordRank
LAYSUMM	14.81 $\pm$ 2.91	16.28 $\pm$ 2.81	11.63 $\pm$ 1.16	9.04 $\pm$ 0.60
EUREKA-ALERT	13.16 $\pm$ 2.07	14.09 $\pm$ 1.56	9.77 $\pm$ 0.85	8.79 $\pm$ 0.80
CDSR	12.72 $\pm$ 2.16	14.17 $\pm$ 0.88	9.073 $\pm$ 0.88	8.54 $\pm$ 0.32
<b>PLOS</b>	14.76 $\pm$ 2.33	15.90 $\pm$ 2.01	10.91 $\pm$ 0.85	8.98 $\pm$ 0.32
<b>ELIFE</b>	10.91 $\pm$ 1.44	12.52 $\pm$ 1.35	8.94 $\pm$ 0.53	8.68 $\pm$ 0.31

Table 6: Comparison of the lay summary readability scores for all lay summarisation datasets.

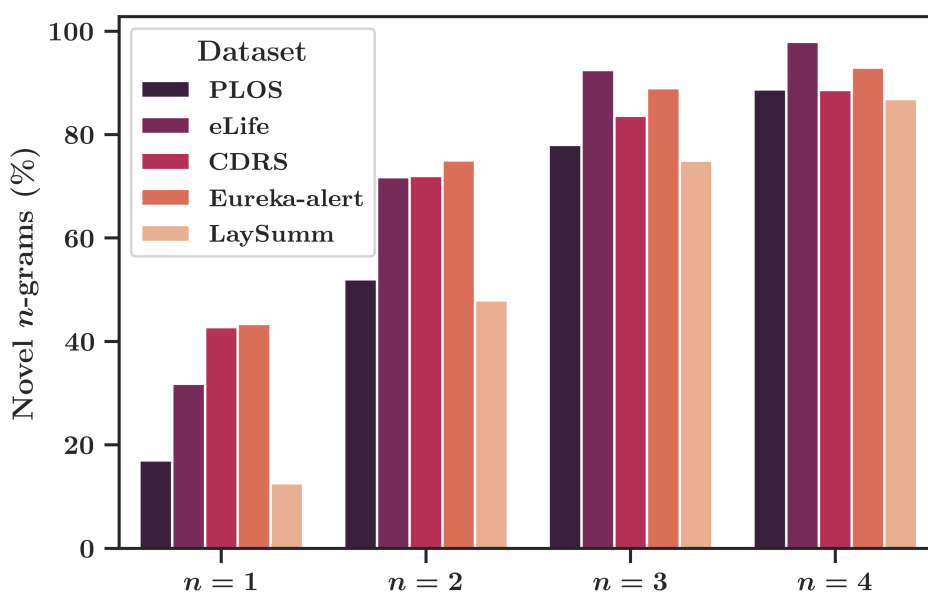


Figure 5: Comparison of lay summary  $n$ -gram novelty for all lay summarisation datasets.

Word type	# of abstract occurrences			
	1	2-10	11-100	100+
Noun	15.3 / 48.2	8.7 / 18.1	2.5 / 4.9	0.8 / 1.5
<b>PLOS</b> Proper noun	15.3 / 54.2	8.2 / 18.0	1.9 / 2.9	0.2 / 0.3
<b>PLOS</b> Number	6.3 / 67.6	2.4 / 18.6	0.5 / 3.8	0.2 / 0.6
<b>PLOS</b> Verb	4.2 / 28.9	4.9 / 29.7	3.4 / 18.9	1.8 / 8.3
Noun	14.9 / 46.2	8.6 / 20.3	2.5 / 6.2	0.5 / 0.9
<b>eLife</b> Proper noun	19.2 / 64.1	6.4 / 9.6	0.3 / 0.4	0.0 / 0.0
<b>eLife</b> Number	7.6 / 67.6	3.8 / 16.8	1.2 / 2.6	0.2 / 0.2
<b>eLife</b> Verb	3.8 / 32.7	4.5 / 36.0	2.6 / 17.2	0.6 / 2.6

Table 7: Statistics for bars in Figure 3. For each table cell, the overall percentage is split into ‘% that are shared with lay summaries’ / ‘% that are not shared with lay summaries’.

**Comprehensiveness**

*Annotator 1:* The model outputs summarised the important information, however it also cut out quite a lot of background info which is key for understanding the science. Overall, the comprehensiveness was enough for a non-expert to grasp the overall gist of the studies.

*Annotator 2:* The model seems to use parts of the abstract, and therefore seems quite comprehensive. It also does a decent job of a final "summary" sentence to the lay summary to summarize/put into context.

**Layness**

*Annotator 1:* Some abstracts contained a lot of jargon which would be confusing/off putting to a non-expert. Although I know some scientific words cannot be substituted, it would be good to have an explanation of the more complex words in brackets, for example.

*Annotator 2:* Due to the overlap with the reference abstract, the output is comprehensive but probably confusing to a lay audience, in some cases there is no introduction/background on scientific jargon (e.g. we cannot expect a lay audience to understand complex scientific techniques, cellular or molecular machinery). Use of Genus species nomenclature is also likely to confuse lay audiences, where a common name could be used instead, not as well as (e.g. 'Egyptian mosquito' instead of 'Aedes aegypti', also known as the Egyptian mosquito').

**Factuality**

*Annotator 1:* The majority of the statements were factually correct, although sometimes the meaning of the simplified language could be misinterpreted, which would result in a similar outcome to factually wrong statements.

*Annotator 2:* Some minor factual errors throughout, and mixups between gene symbols (e.g. where one letter will be changed, PMN to PMA), there are also come cases where it will pull out a % but mix up what gene / condition it is related to, ultimately leading to the formation of a sentence which is factually incorrect.

Figure 6: Human evaluation comments for PLOS.

<p><b>Comprehensiveness</b></p> <p><i>Annotator 1:</i> Overall the information contained within the model-generated summaries effectively conveyed the information in the references. However, there was a few occasions where new elements/concepts were introduced that could confuse the reader and affect their understanding (these were sometimes factual statements, sometimes seemingly made up). The model abstract did provide enough information for a general understanding of the topic and would be sufficient as a brief overview.</p> <p><i>Annotator 2:</i> The model usually picks up on the core points of the abstract but can often introduce extra information which is either off-topic or factually incorrect. The model seems to start off by introducing the topic well but struggles to hit the “what is the significance of this research?” question.</p>
<p><b>Layness</b></p> <p><i>Annotator 1:</i> The language was, in the majority, well suited to a lay person and terminology was adapted accordingly. However, at times the information was simplified to a point where it could be misconstrued which, with scientific information, is a potential risk. At times, jargon still remained and I could imagine some people being confused by this. There were a few grammatical errors and poor sentence structure, typos, repetition etc.</p> <p><i>Annotator 2:</i> Sometimes the model introduces extra information which is not suitable for a lay audience, for example: references to genome sequencing, progenitor cells, endoplasmid reticulum. There are also instances where misinterpretation by the model may mislead the lay audience, for example there was an output where Norepinephrine was said to be “a.k.a. dopamine”, which is not factually correct.</p>
<p><b>Factuality</b></p> <p><i>Annotator 1:</i> There were a few summaries which contained incorrect information, things that are well-known in the scientific community were poorly conveyed. At times, new information was introduced which contradicted earlier statements, and those of the reference abstracts/lay summaries. Of course, some information was correct. I would be concerned about the level of misinformation which could arise from these summaries, if used to educate a lay audience.</p> <p><i>Annotator 2:</i> This seemed to vary based on the abstract and how well the output started, for example if the model introduced the topic well, it would lead to more factual points. However, there were some generated summaries which were factually incorrect from the start and this lead to more errors.</p>

Figure 7: Human evaluation comments for eLife.

**Technical Abstract**

Rabies is a uniformly fatal disease, but preventable by timely and correct use of post exposure prophylaxis (PEP). Unfortunately, many health care facilities in Pakistan do not carry modern life-saving vaccines and rabies immunoglobulin (RIG), assuming them to be prohibitively expensive and unsafe. Consequently, Emergency Department (ED) health care professionals remain untrained in its application and refer patients out to other hospitals. The conventional Essen regimen requires five vials of cell culture vaccine (CCV) per patient, whereas Thai Red Cross intradermal (TRC-id) regimen requires only one vial per patient, and gives equal seroconversion as compared with Essen regimen. This study documents the cost savings in using the Thai Red Cross intradermal regimen with cell culture vaccine instead of the customary 5-dose Essen intramuscular regimen for eligible bite victims. All patients presenting to the Indus Hospital ED between July 2013 to June 2014 with animal bites received WHO recommended PEP. WHO Category 2 bites received intradermal vaccine alone, while Category 3 victims received vaccine plus wound infiltration with Equine RIG. Patients were counseled, and subsequent doses of the vaccine administered on days 3, 7 and 28. Throughput of cases, consumption utilization of vaccine and ERIG and the cost per patient were recorded. Government hospitals in Pakistan are generally underfinanced and cannot afford treatment of the enormous burden of dog bite victims. Hence, patients are either not treated at all, or asked to purchase their own vaccine, which most cannot afford, resulting in neglect and high incidence of rabies deaths. TRC-id regimen reduced the cost of vaccine to 1/5th of Essen regimen and is strongly recommended for institutions with large throughput. Training ED staff would save lives through a safe, effective and affordable technique.

**Lay Summary**

Rabies is a killer disease caused by the rabies virus that is present in the saliva of rabid animals, mainly the dog. Once symptoms become apparent, death is inevitable. However, rabies can be prevented if correct post exposure prophylaxis (PEP) is instituted as soon as possible after a dog bite and before symptoms of rabies begin. In Pakistan, government hospitals treat 50-70 new bite victims each day. Many still dispense the free but poor quality nerve tissue vaccine that is often ineffective and fraught with serious adverse reactions. Hospital administrators consider PEP too expensive to be administered free of cost. The Indus Hospital (TIH), Karachi is a private teaching hospital which provides free medical care to all. From July 2013-June 2014, 2,983 new bites were seen in the ED, and rather than use the Essen regimen of five full vial intramuscular doses per patient over 28 days, we administered the WHO-approved Thai Red Cross-intradermal (TRC-id) 4-dose regimen. The use of the TRC-id regimen resulted in 80% cost savings over the Essen regimen. In resource-poor settings, we advocate training of ED personnel in TRC-id regimen, which, ultimately, will result in less vaccine consumption, more patient compliance and complete treatment, resulting in more lives being saved.

Figure 8: PLOS lay summary example.

**Technical Abstract**

Adult stem cells are responsible for life-long tissue maintenance. They reside in and interact with specialized tissue microenvironments (niches). Using murine hair follicle as a model, we show that when junctional perturbations in the niche disrupt barrier function, adjacent stem cells dramatically change their transcriptome independent of bacterial invasion and become capable of directly signaling to and recruiting immune cells. Additionally, these stem cells elevate cell cycle transcripts which reduce their quiescence threshold, enabling them to selectively proliferate within this microenvironment of immune distress cues. However, rather than mobilizing to fuel new tissue regeneration, these ectopically proliferative stem cells remain within their niche to contain the breach. Together, our findings expose a potential communication relay system that operates from the niche to the stem cells to the immune system and back. The repurposing of proliferation by these stem cells patch the breached barrier, stoke the immune response and restore niche integrity.

**Lay Summary**

Most, if not all, tissues of an adult animal contain stem cells. These stem cells regenerate and repair damaged tissues and organs for the entire lifetime of an animal, contributing to a healthy life. They divide to make daughter cells that become either new stem cells or specialized cells of that organ. Adult stem cells exist in specific areas within tissues known as niches, where they interact with surrounding cells and molecules that inform their behavior. For example, cells and molecules within these niches can signal stem cells to remain in a ‘dormant’ state, but upon injury, they can mobilize stem cells to form new tissue and repair the wound. So far, it has been unclear how stem cells sense damage and stress and direct their efforts away from their normal duties towards repair. Here, Lay et al. studied the stem cells in the mouse skin that are responsible to regenerate hair. Every hair follicle contains a niche (the ‘bulge’), where these stem cells live and share their environment with cells that anchor the hair. The niche tethers to the stem cells through specific adhesion molecules that also help the niche to form a tight seal to prevent bacteria from entering. Lay et al. removed one of the adhesion molecules called E-cadherin, which caused a breach in the niche’s barrier. The stem cells sensed their damaged niche, prepared to multiply, and sent out stress signals to the immune system. The immune cells then arrived at the niche and sent signals back to the stem cells, prodding them to multiply and patch the barrier, while at the same time, keeping the inflammation in check. This remarkable ability of the stem cells to recruit immune cells and initiate a dialogue with them enabled the stem cells to divert their attention from regenerating hair and instead directing it towards the site of the tissue damage. Other stem cells, such as those in the lung or gut, may have similar mechanisms to detect and respond to physical damage. It will be interesting to investigate the underlying mechanism of how immune cells are involved in balancing stem cell regenerative capacity and response to physical damage. A better knowledge of these processes could help to regenerate tissues or even entire organs.

Figure 9: eLife lay summary example.