

MS²: A Dataset for Multi-Document Summarization of Medical Studies

Jay DeYoung^{1*} Iz Beltagy² Madeleine van Zuylen² Bailey Kuehl² Lucy Lu Wang²
 Northeastern University¹ Allen Institute for AI²
 deyoung.j@northeastern.edu
 {beltagy,madeleinev,baileyk,lucyw}@allenai.org

Abstract

To assess the effectiveness of any medical intervention, researchers must conduct a time-intensive and manual literature review. NLP systems can help to automate or assist in parts of this expensive process. In support of this goal, we release MS² (Multi-Document Summarization of Medical Studies), a dataset of over 470k documents and 20K summaries derived from the scientific literature. This dataset facilitates the development of systems that can assess and aggregate contradictory evidence across multiple studies, and is the first large-scale, publicly available multi-document summarization dataset in the biomedical domain. We experiment with a summarization system based on BART, with promising early results, though significant work remains to achieve higher summarization quality. We formulate our summarization inputs and targets in both free text and structured forms and modify a recently proposed metric to assess the quality of our system’s generated summaries. Data and models are available at <https://github.com/allenai/ms2>.

1 Introduction

Multi-document summarization (MDS) is a challenging task, with relatively limited resources and modeling techniques. Existing datasets are either in the general domain, such as WikiSum (Liu et al., 2018) and Multi-News (Fabbri et al., 2019), or very small such as DUC¹ or TAC 2011 (Owczarzak and Dang, 2011). In this work, we add to this burgeoning area by developing a dataset for summarizing biomedical findings. We derive documents and summaries from systematic literature reviews, a type of biomedical paper that synthesizes results across many other studies. Our aim in introducing MS² is to: (1) expand MDS to the biomedical domain, (2) investigate fundamentally challenging

*Work performed during internship at AI2

¹<https://duc.nist.gov>




Figure 1: Our primary formulation (*texts-to-text*) is a *seq2seq* MDS task. Given study abstracts and a BACKGROUND statement, generate the TARGET summary.

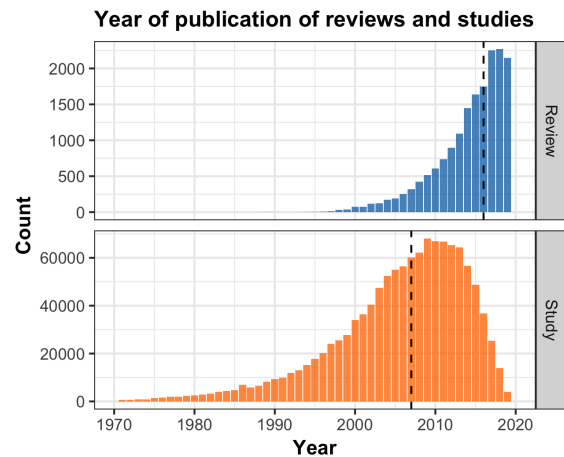


Figure 2: The distributions of review and study publication years in MS² show a clear temporal lag. Dashed lines mark the median year of publication.

issues in NLP over scientific text, such as summarization over contradictory information and assessing summary quality via a structured intermediate form, and (3) aid in distilling large amounts of biomedical literature by supporting automated generation of literature review summaries.

Systematic reviews synthesize knowledge across many studies (Khan et al., 2003), and they are so called for the systematic (and expensive) process of creating a review; each taking 1-2 years to complete (Michelson and Reuter, 2019).² As we note

²<https://community.cochrane.org/review-production/production-resources/proposing-and-registering-new-cochrane-reviews>

in Fig. 2, a delay of around 8 years is observed between reviews and the studies they cite! The time and cost of creating and updating reviews has inspired efforts at automation (Tsafnat et al., 2014; Marshall et al., 2016; Beller et al., 2018; Marshall and Wallace, 2019), and the constant deluge of studies³ has only increased this need.

To move the needle on these challenges and support further work on literature review automation, we present MS², a multi-document summarization dataset in the biomedical domain. Our contributions in this paper are as follows:

- We introduce MS², a dataset of 20K reviews and 470k studies summarized by these reviews.
- We define a *texts-to-text* MDS task (Fig. 1) based on MS², by identifying target summaries in each review and using study abstracts as input documents. We develop a BART-based model for this task, which produces fluent summaries that agree with the evidence direction stated in gold summaries around 50% of the time.
- In order to expose more granular representations to users, we define a structured form of our data to support a *table-to-table* task (§ 4.2). We leverage existing biomedical information extraction systems (Nye et al., 2018; DeYoung et al., 2020) (§3.3.1, §3.3.2) to evaluate agreement between target and generated summaries.

2 Background

Systematic reviews aim to synthesize results over all relevant studies on a topic, providing high quality evidence for biomedical and public health decisions. They are a fixture in the biomedical literature, with many established protocols around their registration, production, and publication (Chalmers et al., 2002; Starr et al., 2009; Booth et al., 2012). Each systematic review addresses one or several research questions, and results are extracted from relevant studies and summarized. For example, a review investigating the effectiveness of Vitamin B12 supplementation in older adults (Andrès et al., 2010) synthesizes results from 9 studies.

The research questions in systematic reviews can be described using the PICO framework (Zakowski et al., 2004). PICO (which stands for Population:

³Of the heterogeneous study types, randomized control trials (RCTs) offer the highest quality of evidence. Around 120 RCTs are published per day as of this writing <https://ijmarshall.github.io/sote/>, up from 75 in 2010 (Bastian et al., 2010).

who is studied? Intervention: *what intervention was studied?* Comparator: *what was the intervention compared against?* Outcome: *what was measured?*) defines the main facets of biomedical research questions, and allows the person(s) conducting a review to identify relevant studies (studies included in a review generally have the same or similar PICO elements as the review). A medical systematic review is one which reports results for applying any kind of medical or social intervention to a group of people. Interventions are wide-ranging, including yoga, vaccination, team training, education, vitamins, mobile reminders, and more. Recent work on evidence inference (DeYoung et al., 2020; Nye et al., 2020) goes beyond identifying PICO elements, and aims to group and identify overall findings in reviews. MS² is a natural extension of these paths: we create a dataset and build a system with both natural summarization targets from input studies, while also incorporating the inherent structure studied in previous work.

In this work, we use the term *review* when describing literature review papers, which provide our summary targets. We use the term *study* to describe the documents that are cited and summarized by each review. There are various study designs which offer differing levels of evidence, e.g. clinical trials, cohort studies, observational studies, case studies, and more (Concato et al., 2000). Of these study types, randomized controlled trials (RCTs) offer the highest quality of evidence (Meldrum, 2000).

3 Dataset

We construct MS² from papers in the Semantic Scholar literature corpus (Ammar et al., 2018). First, we create a corpus of reviews and studies based on the suitability criteria defined in §3.1. For each review, we classify individual sentences in the abstract to identify summarization targets (§3.2). We augment all reviews and studies with PICO span labels and evidence inference classes as described in §3.3.1 and §3.3.2. As a final step in data preparation, we cluster reviews by topic and form train, development, and test sets from these clusters (§3.4).

3.1 Identifying suitable reviews and studies

To identify suitable reviews, we apply (i) a high-recall heuristic keyword filter, (ii) PubMed filter, (iii) study-type filter, and (iv) suitability classifier, in series. The keyword filter looks for the

Label	Sentence
BACKGROUND	... AREAS COVERED IN THIS REVIEW The objective of this review is to evaluate the efficacy of oral cobalamin treatment in <i>elderly patients</i> .
OTHER	To reach this objective , PubMed data were systematically searched for <i>English and French articles published from January 1990 to July 2008</i>
TARGET	The efficacy was particularly highlighted when looking at the marked improvement in serum vitamin B12 levels and hematological parameters , for example hemoglobin level , mean erythrocyte cell volume and reticulocyte count .
OTHER	The effect of oral cobalamin treatment in <i>patients</i> presenting with severe neurological manifestations has not yet been adequately documented

Table 1: Abbreviated example from [Andrès et al. \(2010\)](#) with predicted sentence labels (full abstract in Tab. 11, App. D.3). Spans corresponding to *Population*, *Intervention*, and *Outcome* elements are tagged and surrounded with special tokens.

phrase “systematic review” in the title and abstracts of all papers in Semantic Scholar, which yields 220K matches. The PubMed filter, yielding 170K matches, limits search results to papers that have been indexed in the PubMed database, which restricts reviews to those in the biomedical, clinical, psychological, and associated domains. We then use citations and Medical Subject Headings (MeSH) to identify input studies via their document types and further refine the remaining reviews, see App. A for details on the full filtering process.

Finally, we train a suitability classifier as the final filtering step, using SciBERT ([Beltagy et al., 2019](#)), a BERT ([Devlin et al., 2019](#)) based language model trained on scientific text. Details on classifier training and performance are provided in Appendix C. Applying this classifier to the remaining reviews leaves us with 20K candidate reviews.

3.2 Background and target identification

For each review, we identify two sections: 1) the BACKGROUND statement, which describes the research question, and 2) the overall effect or findings statement as the TARGET of the MDS task (Fig. 1). We frame this as a sequential sentence classification task ([Cohan et al., 2019](#)): given the sentences in the review abstract, classify them as BACKGROUND, TARGET, or OTHER. All BACKGROUND sentences are aggregated and used as input in modeling. All TARGET sentences are aggregated and form the summary target for that review. Sentences classified as OTHER may describe the methods used to conduct the review, detailed findings such as the number of included studies or numerical results, as well as recommendations for practice. OTHER sentences are not suitable for modeling because they

Intervention	Outcome	Evidence sentence	Direction
oral cobalamin therapy	effect	The effect of oral cobalamin treatment in patients presenting with severe neurological manifestations has not yet been adequately documented .	<i>no_change</i>
oral cobalamin therapy	discomfort, inconvenience and cost	Oral cobalamin treatment avoids the discomfort , inconvenience and cost of monthly injections .	<i>decreases</i>
oral cobalamin therapy	serum vitamin b12 levels and hematological parameters	The efficacy was particularly highlighted when looking at the marked improvement in serum vitamin B12 levels and hematological parameters , for example hemoglobin level , mean erythrocyte cell volume and reticulocyte count .	<i>increases</i>

Table 2: Sample Intervention, Outcome, evidence statement, and identified effect directions from a systematic review investigating the effectiveness of vitamin B12 therapies in the elderly ([Andrès et al., 2010](#)).

either contain information specific to the review, as in methods; too much detail, in the case of results; or contain guidance on how medicine should be practiced, which is both outside the scope of our task definition and ill-advised to generate.

Five annotators with undergraduate or graduate level biomedical background labeled 3000 sentences from 220 review abstracts. During annotation, we asked annotators to label sentences into 9 classes (which we collapse into the 3 above; see App. D for detailed info on other classes). Two annotators then reviewed all annotations and corrected mistakes. The corrections yield a Cohen’s κ ([Cohen, 1960](#)) of 0.912. Though we retain only BACKGROUND and TARGET sentences for modeling, we provide labels to all 9 classes in our dataset.

Using SciBERT ([Beltagy et al., 2019](#)), we train a sequential sentence classifier. We prepend each sentence with a [SEP] token and use a linear layer followed by a softmax to classify each sentence. A detailed breakdown of the classifier scores is available in Tab. 9, App. D. While the classifier performs well (94.1 F1) at identifying BACKGROUND sentences, it only achieves 77.4 F1 for TARGET sentences. The most common error for TARGET sentences is confusing them for results from individual studies or detailed statistical analysis. Tab. 1 shows example sentences with predicted labels. Due to the size of the dataset, we cannot manually annotate sentence labels for all reviews, so we use the sentence classifier output as silver labels in the training set. To ensure the highest degree of accuracy for the summary targets in our test set, we manually review all 4519 TARGET sentences in the 2K reviews of the test set, correcting 1109 sentences. Any re-

views without TARGET sentences are considered unsuitable and are removed from the final dataset.

3.3 Structured form

As discussed in §2, the key findings of studies and reviews can be succinctly captured in a structured representation. The structure consists of PICO elements (Nye et al., 2018) that define what is being studied, in addition to the effectiveness of the intervention as inferred through Evidence Inference (§3.3.2). In addition to the textual form of our task, we construct this structured form and release it with MS² to facilitate investigation of consistency between input studies and reviews, and to provide additional information for interpreting the findings reported in each document.

3.3.1 Adding PICO tags

The Populations, Interventions, and Outcomes of interest are a common way of representing clinical knowledge (Huang et al., 2006). Recent work (Nye et al., 2020) has found that the Comparator is rarely mentioned explicitly, so we exclude it from our dataset. Previous summarization work has shown that tagging salient entities, especially PIO elements (Wallace et al., 2020), can improve summarization performance (Nallapati et al., 2016a,b), so we mark PIO elements with special tokens added to our model vocabulary: <pop>, </pop>, <int>, </int>, <out>, and </out>.

Using the EBM-NLP corpus (Nye et al., 2018), a crowd-sourced collection of PIO tags,⁴ we train a token classification model (Wolf et al., 2020) to identify these spans in our study and review documents. These span sets are denoted $\mathbf{P} = \{P_1, P_2, \dots, P_{\bar{P}}\}$, $\mathbf{I} = \{I_1, I_2, \dots, I_{\bar{I}}\}$ and $\mathbf{O} = \{O_1, O_2, \dots, O_{\bar{O}}\}$. At the level of each review, we perform a simple aggregation over these elements. Any P, I, or O span fully contained within any other span of the same type is removed from these sets (though they remain tagged in the text). Removing these contained elements reduces the number of duplicates in our structured representation. Our dataset has an average of 3.0 P, 3.5 I, and 5.4 O spans per review.

3.3.2 Adding Evidence Inference

We predict the direction of evidence associated with every Intervention-Outcome (I/O) pair found

in the review abstract. Taking the product of each I_i and O_j in the sets \mathbf{I} and \mathbf{O} yields all possible I/O pairs, and each I/O pair is associated with an evidence direction d_{ij} , which can take on one of the values in $\{\textit{increases}, \textit{no_change}, \textit{decreases}\}$. For each I/O pair, we also derive a sentence s_{ij} from the document supporting the d_{ij} classification. Each review can therefore be represented as a set of tuples \mathbf{T} of the form $(I_i, O_j, s_{ij}, d_{ij})$ and cardinality $\bar{I} \times \bar{O}$. See Tab. 2 for examples. For modeling, as in PICO tagging, we surround supporting sentences with special tokens <evidence> and </evidence>; and append the direction class with a <sep> token.

We adapt the Evidence Inference (EI) dataset and models (DeYoung et al., 2020) for labeling. The EI dataset is a collection of RCTs, tagged PICO elements, evidence sentences, and overall evidence direction labels *increases*, *no_change*, or *decreases*. The EI models are composed of 1) an evidence identification module which identifies an evidence sentence, and 2) an evidence classification module for classifying the direction of effectiveness. The former is a binary classifier on top of SciBERT, whereas the latter is a softmax distribution over effectiveness directions. Using the same parameters as DeYoung et al. (2020), we modify these two modules to function solely over I and O spans.⁵ The resulting 354k EI classifications for our reviews are 13.4% *decreases*, 57.0% *no_change*, and 29.6% *increases*. Of the 907k classifications over input studies, 15.7% are *decreases*, 60.7% *no_change*, and 23.6% *increases*. Only 53.8% of study classifications match review classifications, highlighting the prevalence and challenges of contradictory data.

3.4 Clustering and train / test split

Reviews addressing overlapping research questions or providing updates to previous reviews may share input studies and results in common, e.g., a review studying the effect of Vitamin B12 supplementation on B12 levels in older adults and a review studying the effect of B12 supplementation on heart disease risk will cite similar studies. To avoid the phenomenon of learning from test data, we cluster reviews before splitting into train, validation, and test sets. We compute SPECTER paper embeddings (Cohan et al., 2020) using the title and

⁴EBM-NLP contains high quality, crowdsourced and expert-tagged PIO spans in clinical trial abstracts. See App. I for a comparison to other PICO datasets.

⁵Nye et al. (2020) found that removing Comparator elements improved classification performance from 78.0 F1 to 81.4 F1 with no additional changes or hyper-parameter tuning.

Dataset statistics	MS ²
Total reviews	20K
Total input studies	470k
Median number of studies per review	17
Median number of reviews per study	1
Average number of reviews per study	1.9
Median BACKGROUND sentences per review	3
Median TARGET sentences per review	2

Table 3: MS² dataset statistics.

Dataset	Docs	Tokens per Summary	Tokens per Document
DUC '03/'04	320	109.6	4636.2
TAC 2011	176	99.7	4695.7
WikiSum	2,332,000	10 ¹ –10 ³	10 ² –10 ⁶
Multi-News	56,216	263.7	2103.4
MS ²	470,402	61.3	365.3

Table 4: A comparison of MDS datasets; adapted from Fabbri et al. (2019). Datasets are DUC '03/'04¹, TAC 2011 (Owczarzak and Dang, 2011), WikiSum (Liu et al., 2018), and Multi-News (Fabbri et al., 2019). Note: WikiSum only provides ranges, not exact size.

abstract of each review, and perform agglomerative hierarchical clustering using the scikit-learn library (Buitinck et al., 2013). This results in 200 clusters, which we randomly partition into 80/10/10 train/development/test sets.

3.5 Dataset statistics

The final dataset consists of 20K reviews and 470k studies. Each review in the dataset summarizes an average of 23 studies, ranging between 1–401 studies. See Tab. 3 for statistics, and Tab. 4 for a comparison to other datasets. The median review has 6.7K input tokens from its input studies, while the average has 9.4K tokens (a few reviews have lots of studies). We restrict the input size when modeling to 25 studies, which reduces the average input to 6.6K tokens without altering the median.

Fig. 2 shows the temporal distribution of reviews and input studies in MS². We observe that though reviews in our dataset have a median publication year of 2016, the studies cited by these reviews are largely from before 2010, with a median of 2007 and peak in 2009. This citation delay has been observed in prior work (Shojania et al., 2007; Beller et al., 2013), and further illustrates the need for automated or assisted reviews.

4 Experiments

We experiment with a *texts-to-text* task formulation (Fig. 1). The model input consists of the BACK-




Figure 3: Two input encoding configurations. Above: LongformerEncoderDecoder (LED), where all input studies are appended to the BACKGROUND and encoded together. Below: In the BART configuration, each input study is encoded independently with the review BACKGROUND. These are concatenated to form the input encoding.

GROUND statement and study abstracts; the output is the TARGET statement. We also investigate the use of the structured form described in §3.3.2 for a supplementary *table-to-table* task, where given inputs of I/O pairs from the review; the model tries to predict the evidence direction. We provide initial results for the *table-to-table* task, although we consider this an area in need of active research.

4.1 Texts-to-text task

Our approach leverages BART (Lewis et al., 2020b), a *seq2seq* autoencoder. Using BART, we encode the BACKGROUND and input studies as in Fig. 3, and pass these representations to a decoder. Training follows a standard auto-regressive paradigm used for building summarization models. In addition to PICO tags (§3.3.1), we augment the inputs by surrounding the background and each input study with special tokens `<background>`, `</background>`, and `<study>`, `</study>`.

For representing multiple inputs, we experiment with two configurations: one leveraging BART with independent encodings of each input, and LongformerEncoderDecoder (LED) (Beltagy et al., 2020) which can encode long inputs of up to 16K tokens. For the BART configuration, each study abstract is appended to the BACKGROUND statement and encoded independently. These representations are concatenated together to form the input to the decoder layer. In the BART configuration, interactions happen only in the decoder. For the LED configuration, the input sequence starts with the BACKGROUND statement followed by a concatenation of all input study abstracts. The BACKGROUND representation is shared among all input studies; global attention allows interactions between studies, and a sliding attention window of 512 tokens allows each token to attend to its neighbors.

We train a BART-base model, with hyperparameters described in App. F. We report experimental results in Tab. 5. In addition to ROUGE (Lin, 2004), we also report two metrics derived from evidence inference: ΔEI and F1. We describe the intuition and computation of the ΔEI metric in Section 4.3; because it is a distance metric, lower ΔEI is better. For F1, we use the EI classification module to identify evidence directions for both the generated and target summaries. Using these classifications, we report a macro-averaged F1 over the class agreement between the generated and target summaries (Buitinck et al., 2013). For example generations, see Tab. 13 in App. G.

4.2 Table-to-table task

An end user of a review summarization system may be interested in specific results from input studies (including whether they agree or contradict) rather than the high level conclusions available in TARGET statements. Therefore, we further experiment with structured input and output representations that attempt to capture results from individual studies. As described in §3.3.2, the structured representation of each review or study is a tuple of the form $(I_i, O_j, s_{ij}, d_{ij})$. It is important to note that we use the same set of *I*s and *O*s from the review to predict evidence direction from all input studies.

Borrowing from the ideas of (Raffel et al., 2020), we formulate our classification task as a text generation task, and train the models described in Section 4.1 to generate one of the classes in $\{\textit{increases}, \textit{no_change}, \textit{decreases}\}$. Using the EI classifications from 3.3.2, we compute an F-score macro-averaged over the effect classes (Tab. 6). We retain all hyperparameter settings other than reducing the maximum generation length to 10.

We stress that this is a preliminary effort to demonstrate feasibility rather than completeness — our results in Tab. 6 are promising but the underlying technologies for building the structured data: PICO tagging, co-reference resolution, and PICO relation extraction, are currently weak (Nye et al., 2020). Resorting to using the full cross-product of Interventions and Outcomes results in duplicated I/O pairs as well as potentially spurious pairs that do not correspond to actual I/O pairs in the review.

4.3 ΔEI metric

Recent work in summarization evaluation has highlighted the weaknesses of ROUGE for capturing factuality of generated summaries, and has focused

Model	R-1	R-2	R-L	$\Delta\text{EI}\downarrow$	F1
BART	27.56	9.40	20.80	.459	46.51
LED	26.89	8.91	20.32	.449	45.00

Table 5: Results for the *texts-to-text* setting. We report ROUGE, ΔEI (§ 4.3), and macro-averaged F1-scores.

Model	P	R	F1
BART	50.31	67.98	65.89

Table 6: Results for the *table-to-table* setting. We report macro-averaged precision, recall, and F-scores.

on developing automated metrics more closely correlated with human-assessed factuality and quality (Zhang* et al., 2020; Wang et al., 2020a; Falke et al., 2019). In this vein, we modify a recently proposed metric based on EI classification distributions (Wallace et al., 2020), intending to capture the agreement of *I*s, *O*s, and EI directions between input studies and the generated summary.

For each I/O tuple (I_i, O_j) , the predicted direction d_{ij} is actually a distribution of probabilities over the three direction classes $P_{ij} = (p_{\textit{increases}}, p_{\textit{decreases}}, p_{\textit{no_change}})$. If we consider this distribution for the gold summary (P_{ij}) and the generated summary (Q_{ij}), we can compute the Jensen-Shannon Distance (JSD) (Lin, 1991), a bounded score between $[0, 1]$, between these distributions. For each review, we can then compute a summary JSD metric, which we call ΔEI , as an average over the JSD of each I/O tuple in that review:

$$\sum_{i=1}^{\bar{I}} \sum_{j=1}^{\bar{J}} \text{JSD}(P_{ij}, Q_{ij}) \quad (1)$$

Different from Wallace et al. (2020), ΔEI is an average over all outputs, attempting to capture an overall picture of system performance,⁶ and our metric retains the directionality of *increases* and *decreases*, as opposed to collapsing them together.

To facilitate interpretation of the ΔEI metric, we offer a degenerate example. Given the case where all direction classifications are certain, and the probability distributions P_{ij} and Q_{ij} exist in the space of $(1, 0, 0)$, $(0, 1, 0)$, or $(0, 0, 1)$, ΔEI takes on the following values at various levels of consistency

⁶Wallace et al. (2020) only report correlation of a related metric with human judgments.

Generated	Gold			
	<i>inc.</i>	<i>no_c.</i>	<i>dec.</i>	<i>insuff.</i>
<i>increases</i>	56	3	1	19
<i>no_change</i>	13	1	1	10
<i>decreases</i>	0	0	5	2
<i>insufficient</i>	5	1	0	5
<i>skip</i>	8	0	0	3

Table 7: Confusion matrix for human evaluation results

between P_{ij} and Q_{ij} for the input studies:

100% consistent	$\Delta EI = 0.0$
50% consistent	$\Delta EI = 0.42$
0% consistent	$\Delta EI = 0.83$

In other words, in both the standard BART and LED setting, the evidence directions predicted in relation to the generated summary are slightly less than 50% consistent with the direction predictions produced relative to the gold summary.

4.4 Human evaluation & error analysis

We randomly sample 150 reviews from the test set for manual evaluation. For each generated and gold summary, we annotate the primary effectiveness direction in the summary to the following classes: (i) *increases*: intervention has a positive effect on the outcome; (ii) *no_change*: no effect, or no difference between the intervention and the comparator; (iii) *decreases*: intervention has a negative effect on the outcome; (iv) *insufficient*: insufficient evidence is available; (v) *skip*: the summary is disfluent, off topic, or does not contain information on efficacy.

Here, *increases*, *no_change*, and *decreases* correspond to the EI classes, while we introduce *insufficient* to describe cases where insufficient evidence is available on efficacy, and *skip* to describe data or generation failures. Two annotators provide labels, and agreement is computed over 50 reviews (agreement: 86%, Cohen’s κ : 0.76). Of these, 17 gold summaries lack an efficacy statement, and are excluded from analysis. Tab. 7 shows the confusion matrix for the sample. Around 50% (67/133) of generated summaries have the same evidence direction as the gold summary. Most confusions happen between *increases*, *no_change*, and *insufficient*.

Tab. 8 shows how individual studies can provide contradictory information, some supporting a positive effect for an intervention and some observing no or negative effects. EI may be able to capture some of the differences between these input studies. From observations on limited data: while studies

Effectiveness	Example statements from studies
Positive effect	Adjuvant vinorelbine plus cisplatin extends survival in patients with completely resected NSCLC... Our results suggest that patients with NSCLC at pathologic stage I who have undergone radical surgery benefit from adjuvant chemotherapy.
No effect or negative effect	No survival benefit for CAP vs no-treatment control was found in this study. Therefore, adjuvant therapy with CAP should not be recommended for patients with resected early-stage non-small cell lung cancer . On the basis of this trial , adjuvant therapy with CAP should not be recommended for patients with resected stage I lung cancer .

Table 8: Text from the input studies to [Petrelli and Barni \(2013\)](#), a review investigating the effectiveness of cisplatin-based (CAP) chemotherapy for non-small cell lung cancer (NSCLC). Input studies vary in their results, with some stating a positive effect for adjuvant chemotherapy, and some stating no survival benefit.

with positive effect tend to have more EI predictions that were *increases* or *decreases*, those with no or negative effect tended to have predictions that were mostly *no_change*. However, more work is needed to better understand how to capture these directional relations and how to aggregate them into a coherent summary.

5 Related Work

NLP for scientific text has been gaining interest recently with work spanning the whole NLP pipeline: datasets (S2ORC ([Lo et al., 2020](#)), [CORD-19 \(Wang et al., 2020b\)](#)), pretrained transformer models (SciBERT ([Beltagy et al., 2019](#)), BioBERT ([Lee et al., 2020](#)), ClinicalBERT ([Huang et al., 2019](#)), SPECTER ([Cohan et al., 2020](#))), NLP tasks like NER ([Nye et al., 2018](#); [Li et al., 2016](#)), relation extraction ([Jain et al., 2020](#); [Luan et al., 2018](#); [Kringelum et al., 2016](#)), QA ([Abacha et al., 2019](#)), NLI ([Romanov and Shivade, 2018](#); [Khot et al., 2018](#)), summarization ([Cachola et al., 2020](#); [Chandrasekaran et al., 2019](#)), claim verification ([Wadden et al., 2020](#)), and more. MS² adds a MDS dataset to the scientific document NLP literature.

A small number of MDS datasets are available for other domains, including MultiNews ([Fabbri et al., 2019](#)), WikiSum ([Liu et al., 2018](#)), and Wikipedia Current Events ([Gholipour Ghalandari et al., 2020](#)). Most similar to MS² is MultiNews,

where multiple news articles about the same event are summarized into one short paragraph. Aside from being in a different textual domain (scientific vs. newswire), one unique characteristic of MS² compared to existing datasets is that MS² input documents have contradicting evidence. Modeling in other domains has typically focused on straightforward applications of single-document summarization to the multi-document setting (Lebanoff et al., 2018; Zhang et al., 2018), although some methods explicitly model multi-document structure using semantic graph approaches (Baumel et al., 2018; Liu and Lapata, 2019; Li et al., 2020).

In the systematic review domain, work has typically focused on information retrieval (Boudin et al., 2010; Ho et al., 2016; Znaidi et al., 2015; Schoot et al., 2020), extracting findings (Lehman et al., 2019; DeYoung et al., 2020; Nye et al., 2020), and quality assessment (Marshall et al., 2015, 2016). Only recently in Wallace et al. (2020) and this work has consideration been made for approaching the entire system as a whole. We refer the reader to App. I for more context regarding the systematic review process.

6 Discussion

Though MDS has been explored in the general domain, biomedical text poses unique challenges such as the need for domain-specific vocabulary and background knowledge. To support development of biomedical MDS systems, we release the MS² dataset. MS² contains summaries and documents derived from biomedical literature, and can be used to study literature review automation, a pressing real-world application of MDS.

We define a *seq2seq* modeling task over this dataset, as well as a structured task that incorporates prior work on modeling biomedical text (Nye et al., 2018; DeYoung et al., 2020). We show that although generated summaries tend to be fluent and on-topic, they only agree with the evidence direction in gold summaries around half the time, leaving plenty of room for improvement. This observation holds both through our ΔEI metric and through human evaluation of a small sample of generated summaries. Given that only 54% of study evidence directions agree with the evidence directions of their review, modeling contradiction in source documents may be key to improving upon existing summarization methods.

Limitations Challenges in co-reference resolution and PICO extraction limit our ability to generate accurate PICO labels at the document level. Errors compound at each stage: PICO tagging, taking the product of Is and Os at the document level, and predicting EI direction. Pipeline improvements are needed to bolster overall system performance and increase our ability to automatically assess performance via automated metrics like ΔEI . Relatedly, automated metrics for summarization evaluation can be difficult to interpret, as the intuition for each metric must be built up through experience. Though we attempt to facilitate understanding of ΔEI by offering a degenerate example, more exploration is needed to understand how a practically useful system would perform on such a metric.

Future work Though we demonstrate that *seq2seq* approaches are capable of producing fluent and on-topic review summaries, there are significant opportunities for improvement. Data improvements include improving the quality of summary targets and intermediate structured representations (PICO tags and EI direction). Another opportunity lies in linking to structured data in external sources such as various clinical trial databases^{7,8,9} rather than relying solely on PICO tagging. For modeling, we are interested in pursuing joint retrieval and summarization approaches (Lewis et al., 2020a). We also hope to explicitly model the types of contradictions observed in Tab. 8, such that generated summaries can capture nuanced claims made by individual studies.

7 Conclusion

Given increasing rates of publication, multi-document summarization, or the creation of literature reviews, has emerged as an important NLP task in science. The urgency for automation technologies has been magnified by the COVID-19 pandemic, which has led to both an accelerated speed of publication (Horbach, 2020) as well as proliferation of non-peer-reviewed preprints which may be of lower quality (Lachapelle, 2020). By releasing MS², we provide a MDS dataset that can help to address these challenges. Though we demonstrate that our MDS models can produce fluent text, our results show that there are significant outstanding challenges that remain unsolved, such as PICO

⁷<https://clinicaltrials.gov/>

⁸<https://www.clinicaltrialsregister.eu/>

⁹<https://www.gsk-studyregister.com/>

tuple extraction, co-reference resolution, and evaluation of summary quality and faithfulness in the multi-document setting. We encourage others to use this dataset to better understand the challenges specific to MDS in the domain of biomedical text, and to push the boundaries on the real world task of systematic review automation.

Acknowledgements

This project is supported in part by NSF Grant OIA-2033558. We thank Ani Nenkova, Byron Wallace, Dan Weld, the reviewers, and members of the Semantic Scholar team for their valuable feedback.

Ethical Concerns and Broader Impact

We believe that automation in systematic reviews has great potential value to the medical and scientific community; our aim in releasing our dataset and models is to facilitate research in this area. Given unresolved issues in evaluating the factuality of summarization systems, as well as a lack of strong guarantees about what the summary outputs contain, we do not believe that such a system is ready to be deployed in practice. Deploying such a system now would be premature, as without these guarantees we would be likely to generate plausible-looking but factually incorrect summaries, an unacceptable outcome in such a high impact domain. We hope to foster development of useful systems with correctness guarantees and evaluations to support them.

References

- Asma Ben Abacha, Chaitanya Shivade, and Dina Demner-Fushman. 2019. Overview of the mediq 2019 shared task on textual inference, question entailment and question answering. In *BioNLP@ACL*.
- Waleed Ammar, Dirk Groeneveld, Chandra Bhagavathula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu A. Ha, Rodney Michael Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler C. Murray, Hsu-Han Ooi, Matthew E. Peters, Joanna L. Power, Sam Skjonsberg, Lucy Lu Wang, Christopher Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. 2018. Construction of the literature graph in semantic scholar. In *NAACL*.
- E. Andr es, H. Fothergill, and M. Mecili. 2010. Efficacy of oral cobalamin (vitamin b12) therapy. *Expert Opinion on Pharmacotherapy*, 11:249 – 256.
- Edoardo Aromataris and Zachary Munn, editors. 2020. *JBIManual for Evidence Synthesis*. JBI.
- Hilda Bastian, Paul Glasziou, and Iain Chalmers. 2010. [Seventy-five trials and eleven systematic reviews a day: How will we ever keep up?](#) *PLOS Medicine*, 7(9):1–6.
- Tal Baumel, Matan Eyal, and Michael Elhadad. 2018. Query focused abstractive summarization: Incorporating query relevance, multi-document coverage, and summary length constraints into seq2seq models. *ArXiv*, abs/1801.07704.
- E. Beller, J. K. Chen, U. Wang, and P. Glasziou. 2013. Are systematic reviews up-to-date at the time of publication? *Systematic Reviews*, 2:36 – 36.
- E. Beller, J. Clark, G. Tsafnat, C. Adams, H. Diehl, H. Lund, M. Ouzzani, K. Thayer, J. Thomas, T. Turner, J. Xia, K. Robinson, and P. Glasziou. 2018. Making progress with the automation of systematic reviews: principles of the international collaboration for the automation of systematic reviews (icasr). *Systematic Reviews*, 7.
- 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.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *ArXiv*, abs/2004.05150.
- Alison Booth, Michael Clarke, Gordon Dooley, Davina Ghera, David Moher, Mark Petticrew, and Lesley A Stewart. 2012. The nuts and bolts of PROSPERO: an international prospective register of systematic reviews. *Systematic Reviews*, 1:2 – 2.
- Florian Boudin, Jian-Yun Nie, and Martin Dawes. 2010. [Clinical information retrieval using document and PICO structure](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 822–830, Los Angeles, California. Association for Computational Linguistics.
- Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake VanderPlas, Arnaud Joly, Brian Holt, and Ga el Varoquaux. 2013. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, pages 108–122.
- 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.

- Iain Chalmers, Larry V Hedges, and Harris Cooper. 2002. A brief history of research synthesis. *Evaluation & the Health Professions*, 25:12 – 37.
- Muthu Kumar Chandrasekaran, Michihiro Yasunaga, Dragomir R. Radev, D. Freitag, and Min-Yen Kan. 2019. Overview and results: CI-scisumm shared task 2019. In *BIRNDL@SIGIR*.
- 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, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. [SPECTER: Document-level representation learning using citation-informed transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282, Online. Association for Computational Linguistics.
- A. Cohen, N. Smalheiser, M. McDonagh, C. Yu, C. Adams, J. M. Davis, and Philip S. Yu. 2015. Automated confidence ranked classification of randomized controlled trial articles: an aid to evidence-based medicine. *Journal of the American Medical Informatics Association : JAMIA*, 22:707 – 717.
- J. Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20:37 – 46.
- J. Concato, N. Shah, and R. Horwitz. 2000. Randomized, controlled trials, observational studies, and the hierarchy of research designs. *The New England journal of medicine*, 342(25):1887–92.
- J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Jay DeYoung, Eric Lehman, Benjamin Nye, Iain Marshall, and Byron C. Wallace. 2020. [Evidence inference 2.0: More data, better models](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 123–132, Online. Association for Computational Linguistics.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- WA Falcon. 2019. Pytorch lightning. *GitHub*. Note: <https://github.com/PyTorchLightning/pytorch-lightning>, 3.
- Tobias Falke, Leonardo F. R. Ribeiro, Prasetya Ajie Utama, Ido Dagan, and Iryna Gurevych. 2019. [Ranking generated summaries by correctness: An interesting but challenging application for natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2214–2220, Florence, Italy. Association for Computational Linguistics.
- Demian Gholipour Ghalandari, Chris Hokamp, Nghia The Pham, John Glover, and Georgiana Ifrim. 2020. [A large-scale multi-document summarization dataset from the Wikipedia current events portal](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1302–1308, Online. Association for Computational Linguistics.
- Charles R. Harris, K. Jarrod Millman, Stéfan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. 2020. [Array programming with NumPy](#). *Nature*, 585:357–362.
- Julian P.T. Higgins, James Thomas, Jacqueline Chandler, Miranda Cumpston, Tianjing Li, Matthew J. Page, and Vivian A. Welch, editors. 2019. *Cochrane Handbook for Systematic Reviews of Interventions*. Wiley.
- G. J. Ho, S. M. Liew, C. Ng, Ranita Hisham Shunmugam, and P. Glasziou. 2016. Development of a search strategy for an evidence based retrieval service. *PLoS ONE*, 11.
- S. Horbach. 2020. Pandemic publishing: Medical journals strongly speed up their publication process for covid-19. *Quantitative Science Studies*, 1:1056–1067.
- Brian Howard, Jason Phillips, Arpit Tandon, Adyasha Maharana, Rebecca Elmore, Deepak Mav, Alex Sedykh, Kristina Thayer, Alex Merrick, Vickie Walker, Andrew Rooney, and Ruchir Shah. 2020. [Swift-active screener: Accelerated document screening through active learning and integrated recall estimation](#). *Environment International*, 138:105623.
- Kexin Huang, Jaan Altosaar, and R. Ranganath. 2019. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *ArXiv*, abs/1904.05342.
- X. Huang, J. Lin, and Dina Demner-Fushman. 2006. Evaluation of pico as a knowledge representation for clinical questions. *AMIA ... Annual Symposium proceedings. AMIA Symposium*, pages 359–63.
- Sarthak Jain, Madeleine van Zuylen, Hannaneh Hajishirzi, and Iz Beltagy. 2020. Scirex: A challenge

- dataset for document-level information extraction. In *ACL*.
- Di Jin and Peter Szolovits. 2018. [PICO element detection in medical text via long short-term memory neural networks](#). In *Proceedings of the BioNLP 2018 workshop*, pages 67–75, Melbourne, Australia. Association for Computational Linguistics.
- K. Khan, R. Kunz, J. Kleijnen, and G. Antes. 2003. Five steps to conducting a systematic review. *Journal of the Royal Society of Medicine*, 96 3:118–21.
- Tushar Khot, A. Sabharwal, and Peter Clark. 2018. Scitail: A textual entailment dataset from science question answering. In *AAAI*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- J. Kringelum, Sonny Kim Kjærulff, S. Brunak, O. Lund, T. Oprea, and O. Taboureau. 2016. Chemprot-3.0: a global chemical biology diseases mapping. *Database: The Journal of Biological Databases and Curation*, 2016.
- F. Lachapelle. 2020. Covid-19 preprints and their publishing rate: An improved method. *medRxiv*.
- Logan Lebanoff, Kaiqiang Song, and Fei Liu. 2018. [Adapting the neural encoder-decoder framework from single to multi-document summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4131–4141, Brussels, Belgium. Association for Computational Linguistics.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, D. Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*.
- Eric Lehman, Jay DeYoung, Regina Barzilay, and Byron C. Wallace. 2019. [Inferring which medical treatments work from reports of clinical trials](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3705–3717, Minneapolis, Minnesota. Association for Computational Linguistics.
- M. Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020a. Pre-training via paraphrasing. *ArXiv*, abs/2006.15020.
- M. Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, A. Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020b. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *ArXiv*, abs/1910.13461.
- J. Li, Yueping Sun, Robin J. Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, A. P. Davis, C. Mattingly, Thomas C. Wieggers, and Zhiyong Lu. 2016. Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database: The Journal of Biological Databases and Curation*, 2016.
- Wei Li, Xinyan Xiao, Jiachen Liu, Hua Wu, Haifeng Wang, and Junping Du. 2020. [Leveraging graph to improve abstractive multi-document summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6232–6243, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *ACL 2004*.
- J. Lin. 1991. Divergence measures based on the shannon entropy. *IEEE Trans. Inf. Theory*, 37:145–151.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. [Generating wikipedia by summarizing long sequences](#). In *International Conference on Learning Representations*.
- Yang Liu and Mirella Lapata. 2019. [Hierarchical transformers for multi-document summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5070–5081, Florence, Italy. Association for Computational Linguistics.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. [S2ORC: The semantic scholar open research corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4969–4983, Online. Association for Computational Linguistics.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- I. Marshall, Joël Kuiper, and Byron C. Wallace. 2015. Automating risk of bias assessment for clinical trials. *IEEE journal of biomedical and health informatics*, 19 4:1406–12.
- I. Marshall, Joël Kuiper, and Byron C. Wallace. 2016. Robotreviewer: evaluation of a system for automatically assessing bias in clinical trials. *Journal of the American Medical Informatics Association : JAMIA*, 23:193 – 201.
- I. Marshall and Byron C. Wallace. 2019. Toward systematic review automation: a practical guide to using machine learning tools in research synthesis. *Systematic Reviews*, 8.

- M. Meldrum. 2000. A brief history of the randomized controlled trial. from oranges and lemons to the gold standard. *Hematology/oncology clinics of North America*, 14(4):745–60, vii.
- M. Michelson and K. Reuter. 2019. The significant cost of systematic reviews and meta-analyses: A call for greater involvement of machine learning to assess the promise of clinical trials. *Contemporary Clinical Trials Communications*, 16.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016a. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Ramesh Nallapati, Bowen Zhou, C. D. Santos, Çağlar Gülçehre, and B. Xiang. 2016b. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *CoNLL*.
- Christopher Norman. 2020. *Systematic review automation methods*. Ph.D. thesis, Université Paris-Saclay; Universiteit van Amsterdam.
- Benjamin Nye, Junyi Jessy Li, Roma Patel, Yinfei Yang, Iain Marshall, Ani Nenkova, and Byron Wallace. 2018. [A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 197–207, Melbourne, Australia. Association for Computational Linguistics.
- Benjamin E. Nye, Jay DeYoung, Eric Lehman, Ani Nenkova, Iain J. Marshall, and Byron C. Wallace. 2020. [Understanding clinical trial reports: Extracting medical entities and their relations](#).
- Karolina Owczarzak and Hoa Trang Dang. 2011. Overview of the tac 2011 summarization track: Guided task and aesop task.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- F. Petrelli and S. Barni. 2013. Non-cancer-related mortality after cisplatin-based adjuvant chemotherapy for non-small cell lung cancer: a study-level meta-analysis of 16 randomized trials. *Medical Oncology*, 30:1–10.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, W. Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Alexey Romanov and Chaitanya Shivade. 2018. Lessons from natural language inference in the clinical domain. In *EMNLP*.
- R. V. D. Schoot, J. D. Bruin, R. Schram, Parisa Zahedi, J. D. Boer, F. Weijdemans, B. Kramer, Martijn Huijts, M. Hoogerwerf, Gerbrich Ferdinands, Albert Harkema, Joukje Willemsen, Yongchao Ma, Qixiang Fang, L. Tummers, and D. Oberski. 2020. [As-review: Open source software for efficient and transparent active learning for systematic reviews](#). *ArXiv*, abs/2006.12166.
- Darsh J Shah, Lili Yu, Tao Lei, and Regina Barzilay. 2021. [Nutribullets hybrid: Multi-document health summarization](#).
- Weixiang Shao, C. Adams, A. Cohen, J. Davis, M. McDonagh, S. Thakurta, Philip S. Yu, and N. Smalheiser. 2015. [Aggregator: a machine learning approach to identifying medline articles that derive from the same underlying clinical trial](#). *Methods*, 74:65–70.
- K. Shojania, M. Sampson, M. Ansari, J. Ji, S. Doucette, and D. Moher. 2007. How quickly do systematic reviews go out of date? a survival analysis. *Annals of Internal Medicine*, 147:224–233.
- Nitish Srivastava, Geoffrey E. Hinton, A. Krizhevsky, Ilya Sutskever, and R. Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15:1929–1958.
- Mark G Starr, Iain Chalmers, Mike Clarke, and Andrew David Oxman. 2009. The origins, evolution, and future of the cochrane database of systematic reviews. *International journal of technology assessment in health care*, 25 Suppl 1:182–95.
- G. Tsafnat, P. Glasziou, Miew Keen Choong, A. Dunn, Filippo Galgani, and E. Coiera. 2014. Systematic review automation technologies. *Systematic Reviews*, 3:74 – 74.
- Özlem Uzuner, Brett R South, Shuying Shen, and Scott L DuVall. 2011. [2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text](#). *Journal of the American Medical Informatics Association*, 18(5):552–556.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas,

- Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*. *Nature Methods*, 17:261–272.
- David Wadden, Kyle Lo, Lucy Lu Wang, Shanchuan Lin, Madeleine van Zuylen, Arman Cohan, and Hananeh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. In *EMNLP*.
- Byron C. Wallace, Sayantan Saha, Frank Soboczenski, and I. Marshall. 2020. Generating (factual?) narrative summaries of rcts: Experiments with neural multi-document summarization. *ArXiv*, abs/2008.11293.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020a. Asking and answering questions to evaluate the factual consistency of summaries. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020, Online. Association for Computational Linguistics.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Doug Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Michael Kinney, Yunyao Li, Ziyang Liu, William Merrill, Paul Mooney, Dewey A. Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, Brandon Stilson, Alex D. Wade, Kuansan Wang, Nancy Xin Ru Wang, Christopher Wilhelm, Boya Xie, Douglas M. Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. 2020b. *CORD-19: The COVID-19 open research dataset*. In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*, Online. Association for Computational Linguistics.
- Chih-Hsuan Wei, Yifan Peng, Robert Leaman, Alan Peter Davis, Carolyn J Mattingly, Jiao Li, Thomas C Wieggers, and Zhiyong Lu. 2015. Overview of the biocreative v chemical disease relation (cdr) task. In *Proceedings of the fifth BioCreative challenge evaluation workshop*, volume 14.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. *Transformers: State-of-the-art natural language processing*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- L. Zakowski, C. Seibert, and S. Vaneyck. 2004. Evidence-based medicine: answering questions of diagnosis. *Clinical medicine & research*, 2(1):63–9.
- J. Zhang, Jiwei Tan, and Xiaojun Wan. 2018. Towards a neural network approach to abstractive multi-document summarization. *ArXiv*, abs/1804.09010.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Bertscore: Evaluating text generation with bert*. In *International Conference on Learning Representations*.
- Eya Znaidi, L. Tamine, and C. Latiri. 2015. Answering pico clinical questions: A semantic graph-based approach. In *AIME*.




Figure 4: Five fold cross-validation results from training a binary SciBERT classifier on the annotations. Precisions increase following a logistic curve over threshold choices; recalls decrease.

Sentence class	P	R	F
BACKGROUND	0.972	0.952	0.958
RECOMMENDATION	0.418	0.296	0.338
EVIDENCE_QUALITY	0.580	0.528	0.550
EFFECT	0.752	0.800	0.774
METHODS	0.938	0.944	0.940
GOAL	0.916	0.936	0.924
DETAILED_FINDINGS	0.856	0.860	0.858
ETC	0.406	0.322	0.338
FURTHER_STUDY	0.756	0.864	0.804

Table 9: Precision, Recall, and F1-scores for all annotation classes, averaged over five folds of cross validation.

A MeSH Filtering

For each candidate review, we extract its cited papers and identify the study type of each cited paper using MeSH publication type,¹⁰ keeping only studies that are clinical trials, cohort studies, and/or observational studies (see Appendix A.1 for full list of MeSH terms). We exclude case reports, which usually report findings on one or a small number of individuals. We observe that publication type MeSH terms tend to be under-tagged.¹¹ Therefore, we also use ArrowSmith trial labels (Cohen et al., 2015; Shao et al., 2015) and a keyword heuristic (the span “randomized” occurring in the title or abstract) to identify additional RCT-like studies.¹² Candidate reviews are culled to retain only those that cite at least one suitable study and no case

¹⁰<https://www.nlm.nih.gov/mesh/pubtypes.html>

¹¹From a cursory inspection of a random sample of studies, this problem seems to be widespread.

¹²RCTs provide the highest quality of evidence so we strive to include as many as possible as inputs in our dataset.

studies, leaving us with 30K reviews.

A.1 Suitability MeSH Terms

We use the following publication type MeSH terms to decide whether a review’s input document is a study of interest:

1. ‘Clinical Study’
2. ‘Clinical Trial’
3. ‘Controlled Clinical Trial’
4. ‘Randomized Controlled Trial’
5. ‘Pragmatic Clinical Trial’
6. ‘Clinical Trial, Phase I’
7. ‘Clinical Trial, Phase II’
8. ‘Clinical Trial, Phase III’
9. ‘Clinical Trial, Phase IV’
10. ‘Equivalence Trial’
11. ‘Comparative Study’
12. ‘Observational Study’
13. ‘Adaptive Clinical Trial’

And we exclude any reviews citing studies with the following publication type MeSH terms:

1. ‘Randomized Controlled Trial, Veterinary’
2. ‘Clinical Trial, Veterinary’
3. ‘Observational Study, Veterinary’
4. ‘Case Report’

B Suitability Annotation

The annotation guidelines for review suitability are given below. Each annotator was tasked with an initial round of annotation, followed by a round of review, then further annotation.

B.1 Suitability Guidelines

A systematic review is a document resulting from an in-depth search and analysis of all the literature relevant to a particular topic. We are interested in systematic reviews of medical literature, specifically those that assess varying treatments and the outcomes associated with them.

There are many different types of reviews, and many types of documents that look like reviews. We need to identify only the “correct” types of reviews. Sometimes this can be done from the title alone, sometimes one has to read the review itself.

The reviews we are interested in:

- Must study a human population (no animal, veterinary, or environmental studies)

	BACK-GROUND	GOAL	METHODS	DETAILED_FINDINGS	FURTHER_STUDY	RECOMMENDATION	EVIDENCE_QUALITY	EFFECT	ETC
BACKGROUND	0.952	0.017	0.008	0.003	0.0	0.003	0.003	0.008	0.006
GOAL	0.009	0.935	0.056	0.0	0.0	0.0	0.0	0.0	0.0
METHODS	0.003	0.011	0.944	0.034	0.0	0.0	0.0	0.006	0.001
DETAILED_FINDINGS	0.001	0.0	0.027	0.862	0.0	0.001	0.027	0.079	0.002
FURTHER_STUDY	0.0	0.0	0.0	0.0	0.857	0.061	0.01	0.02	0.051
RECOMMENDATION	0.0	0.0	0.021	0.021	0.277	0.298	0.021	0.277	0.085
EVIDENCE_QUALITY	0.0	0.0	0.016	0.227	0.008	0.016	0.523	0.148	0.062
EFFECT	0.004	0.0	0.002	0.114	0.013	0.015	0.032	0.8	0.021
ETC	0.051	0.051	0.064	0.128	0.077	0.038	0.09	0.218	0.282

Table 10: Full 9-class sentence classification confusion matrix, averaged over five folds of cross validation.

Label	Sentence
BACKGROUND	IMPORTANCE OF THE FIELD Cobalamin (vitamin B12) deficiency is particularly common in the <i>elderly (> 15 %)</i> .
BACKGROUND	Management of cobalamin deficiency with <i>cobalamin injections</i> is well codified at present , but new routes of <i>cobalamin</i> administration (oral and nasal) are being studied , especially oral <i>cobalamin</i> therapy for food-cobalamin malabsorption .
BACKGROUND	AREAS COVERED IN THIS REVIEW The objective of this review is to evaluate the <i>efficacy of oral cobalamin treatment</i> in <i>elderly patients</i> .
OTHER	To reach this objective , PubMed data were systematically searched for <i>English and French articles published from January 1990 to July 2008</i> .
OTHER	Data from our research group on cobalamin deficiency (Groupe d'Etude des CAREnce vitamine B12 - CARE B12) were also analyzed .
OTHER	WHAT THE READER WILL GAIN Three prospective and randomized studies , a systematic review by the Cochrane group and five prospective cohort studies were found and provide evidence that <i>oral cobalamin treatment</i> may adequately treat cobalamin deficiency .
TARGET	The <i>efficacy</i> was particularly highlighted when looking at the marked improvement in <i>serum vitamin B12 levels and hematological parameters</i> , for example <i>hemoglobin level , mean erythrocyte cell volume and reticulocyte count</i> .
OTHER	The <i>effect</i> of <i>oral cobalamin treatment</i> in <i>patients</i> presenting with <i>severe neurological manifestations</i> has not yet been adequately documented .
TARGET	<i>Oral cobalamin treatment</i> avoids the <i>discomfort , inconvenience and cost</i> of monthly injections .
TARGET	TAKE HOME MESSAGE Our experience and the present analysis support the use of <i>oral cobalamin therapy</i> in clinical practice

Table 11: Example review abstract from [Andrès et al. \(2010\)](#) with predicted sentence labels. Spans corresponding to *Population*, *Intervention*, and *Outcome* elements are tagged and surrounded with special tokens.

- Must review studies involving multiple participants. We are interested in reviews of trials or cohort studies. We are *not* interested in reviews of case studies - which describe one or a few specific people.
- Must study an explicit population or problem (P from PICO)
 - Example populations: women > 55 old with breast cancer, migrant workers, elementary school children in Spokane, WA, etc.

- Must compare one or more medical interventions

- Example interventions: drugs, vaccines, yoga, therapy, surgery, education, annoying mobile device reminders, professional naggers, personal trainers, and more! Note: placebo / no intervention is a type of intervention.
- Comparing the effectiveness of an intervention against no intervention is okay.
- Combinations of interventions count as comparisons (e.g. yoga vs. yoga + therapy).
- Two different dosages also count (e.g. 500ppm fluoride vs 1000ppm fluoride in toothpaste).
- Must have an explicit outcome measure
- Example outcome measures: survival time, frequency of headaches, relief of depression, survey results, and many other possibilities.

- The outcome measure must measure the effectiveness of the intervention.

C Suitability Classifier

Four annotators with biomedical background labeled 879 reviews sampled from the candidate pool (572 suitable, 307 not, Cohen's Kappa: 0.55) according to the suitability criteria (guidelines in Appendix B). We aim to include reviews that perform an aggregation over existing results, such as reporting on how a medical or social intervention affects a group of people, while excluding reviews that make new observations, such as identifying novel disease co-morbidities or those that synthesize case studies.

For our suitability classifier, we finetune SciBERT ([Beltagy et al., 2019](#)) using standard parameters; using five-fold cross validation we find that

a threshold of 0.75 provides a precision of greater than 80% while maintaining an adequate recall (Figure 4).

Though there are a fairly large number of false positives by this criteria, we note that these false positive documents *are* generally reviews; however, they may not investigate an intervention, or may not have suitable target statements. In the latter case, target identification described in § 3.2 helps us further refine and remove these false positives from the final dataset.

D Sentence Annotation

Sentence annotation guidelines and detailed scores are below. Each annotator was tasked with annotating 50-100 sentences, followed by a round of review, before being asked to annotate more.

D.1 Sentence Annotation Guidelines

A systematic review is a document resulting from an in-depth search and analysis of all the literature relevant to a particular topic. We are interested in systematic reviews of medical literature, specifically those that assess varying treatments and the outcomes associated with them. Ignore any existing labels; these are automatically produced and error prone. If something clearly fits into more than one category, separate the labels by commas (annoying, we know, but it can be important). For sentences that are incorrectly broken in a way that makes them difficult to label, skip them (you can fix them, but they'll be programmatically ignored). For reviews that don't meet suitability guidelines, also skip them. We want to identify sentences within these reviews as belonging to one of several categories:

- **BACKGROUND:** Any background information not including goals.
- **GOAL:** A high level goal sentence, describing the aims or purposes of the review.
- **METHODS:** Anything describing the particular strategies or techniques for conducting the review. This includes methods for finding and assessing appropriate studies to include, e.g., the databases searched or other characteristics of the searched literature. A characteristic might be a study type, it might be other details, such as criteria involving the study participants, what interventions (treatments)

were studied or compared, or what outcomes are measured in those studies. This may also include whether or not a meta-analysis is performed.

- **DETAILED_FINDINGS:** Any sections reporting study results, often includes numbers, p-values, etc. These will frequently include statements about a subset of the trials or the populations.
- **GENERAL FINDINGS:** There are four types of general findings we would like you to label. These do not include things like number of patients, or a p-value (that's DETAILED FINDINGS). Not all of these four subtypes will always be present in a paper's abstract. Some sentences will contain information about more than one subtype, and some sentences can contain information about some of these subtypes as well as DETAILED FINDINGS.
 - **EFFECT:** Effect of the intervention, may include a statement about significance. These can cover a wide range of topics, including public health or policy changes.
 - **EVIDENCE_QUALITY:** Commentary about the strength or quality of evidence pertaining to the intervention.
 - **FURTHER_STUDY:** These statements might call for more research in a particular area, and can include hedging statements, e.g.:
 - * "More rigorously designed longitudinal studies with standardized definitions of periodontal disease and vitamin D are necessary."
 - * "More research with larger sample size and high quality in different nursing educational contexts are required."
 - * "However, this finding largely relies on data from observational studies; high-quality RCTs are warranted because of the potential for subject selection bias."
 - **RECOMMENDATION:** Any kind of clinical or policy recommendation, or recommendations for use in practice. This must contain an explicit recommendation, not a passive statement saying

that a treatment is good. “Should” or “recommend” are good indicators. These may not always be present in an abstract. E.g.:

* “Public policy measures that can reduce inequity in health coverage, as well as improve economic and educational opportunities for the poor, will help in reducing the burden of malaria in SSA.”

– **ETC**: Anything that doesn’t fit into the categories above.

All sentences appear in the context of their review. Some of the selected reviews might not actually be reviews; these were identified by accident. These should be excluded from annotation - either make a comment on the side (preferred) or delete the rows belonging to the non-review.

Examples follow. Please ask questions - these guidelines are likely not perfect and we’ll have missed many edge cases

Examples:

BACKGROUND A sizeable number of individuals who participate in population-based colorectal cancer (CRC) screening programs and have a positive fecal occult blood test (FOBT) do not have an identifiable lesion found at colonoscopy to account for their positive FOBT screen.

GOAL To determine the effect of integrating informal caregivers into discharge planning on post-discharge cost and resource use in older adults.

METHODS MAIN OUTCOMES Clinical status (eg, spirometric measures); functional status (eg, days lost from school); and health services use (eg, hospital admissions). Studies were included if they had measured serum vitamin D levels or vitamin D intake and any periodontal parameter.

DETAILED_FINDINGS Overall, 27 studies were included (13 cross-sectional studies, 6 case-control studies, 5 cohort studies, 2 randomized clinical trials and 1 case series study). Sixty-five percent of the cross-sectional studies reported significant associations between low vitamin D levels and poor periodontal parameters. Analysis of group cognitive-behavioural therapy (CBT) v. usual care alone (14 studies) showed a significant effect in favour of group CBT immediately post-treatment (standardised mean difference (SMD) -0.55 (95% CI -0.78 to -0.32)).

EFFECT This review identified short-term benefits of technology-supported self-guided interven-

tions on the physical activity level and fatigue and some benefit on dietary behaviour and HRQoL in people with cancer. However, current literature demonstrates a lack of evidence for long-term benefit.

EVIDENCE_QUALITY Interpretation of findings was influenced by inadequate reporting of intervention description and compliance.

No meta-analysis was performed due to high variability across studies.

RECOMMENDATION The decision to perform EGD should be individualized and based on clinical judgement.

ETC PROSPERO CRD42017080346; https://www.crd.york.ac.uk/prospero/display_record.php?RecordID=80346.

D.2 Detailed Sentence Breakdown Scores

Sentence classification scores for 9 classes are given in Table 9. The corresponding confusion matrix can be found in Table 10.

D.3 Sentence Classification Results

Table 11 provides an example of sentence classification results over 3 classes.

E Dataset Contradiction Scores

The confusion matrix between review effect findings and input study effect findings is given in Table 12.

F Hyperparameters and Modeling Details

We implement our models using PyTorch (Paszke et al., 2019), the HuggingFace Transformers (Wolf et al., 2020) and PyTorch lightning (Falcon, 2019) libraries, starting from the BART-base checkpoint (Lewis et al., 2020b). All models were trained using FP16, using NVidia RTX 8000 GPUs (GPUs with 40G or more of memory are required for most *texts-to-text* configurations). All models are trained for eight epochs as validation scores diminished over time; early experiments ran out to approximately fifty epochs and showed little sensitivity to other hyperparameters. We use gradient accumulation to reach an effective batch size of 32. We use the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 1e-5, an epsilon of 1e-8, and a linear learning rate schedule with 1000 steps of warmup. We ran a hyperparameter sweep over decoding parameters on the validation set for 4, 6,

	<i>decreases</i>	<i>no_change</i>	<i>increases</i>	count
<i>decreases</i>	.338	.540	.122	1202991
<i>no_change</i>	.144	.659	.197	5375546
<i>increases</i>	.096	.529	.376	2490229

Table 12: Confusion matrix between review effect findings and input study effect findings. Each row corresponds to the fraction of the effect direction found in the review with the fraction of that direction accounted for in the study. The most frequent confusion is with *no_change*, as opposed to flipping the overall direction of the finding.

BACKGROUND	TARGET	GENERATED
OBJECTIVE To explore the evidence for the effectiveness of acupuncture for nonspecific low back pain (LBP). SUMMARY OF BACKGROUND DATA Since the most recent systematic review s on RCTs on acupuncture for LBP, 6 RCTs have been published, which may impact on the previous conclusions.	There is moderate evidence that acupuncture is more effective than no treatment, and strong evidence of no significant difference between acupuncture and sham acupuncture , for short-term pain relief.	The is insufficient evidence to support the use of acupuncture for LBP. CONCLUSIONS There is limited evidence for the effectiveness of acupuncture in LBP in the short term.
Objectives : To provide a quantitative analysis of all r and omized controlled trials design ed to determine the effectiveness of physical interventions for people with spinal cord injury (SCI).	There is initial evidence supporting the effectiveness of some physical interventions for people with SCI.	The Results : This systematic review provides evidence that physical interventions for people with SCI are effective in improving muscle strength and function in the short term.
BACKGROUND Neuroendocrine tumours (NET) most commonly metastasize to the liver. Hepatic resection of NET hepatic metastases (NETHM) has been shown to improve symptomology and survival. METHODS A systematic review of clinical studies before September 2010 was performed to examine the efficacy of hepatic resection for NETHM.	Poor histologic grade, extra-hepatic disease and a macroscopically incomplete resection were associated with a poor prognosis. CONCLUSION Hepatic resection for NETHM provides symptomatic benefit and is associated with favourable survival outcomes although the majority of patients invariably develop disease progression	Theatic resection of NETHM has been shown to improve survival in patients with advanced, well-differentiated NETs.
The aim of this systematic review and meta- analysis was to assess the efficacy on an intervention on breastfeeding self-efficacy and perceived insufficient milk supply outcomes.	Although significant effect of the interventions in improving maternal breastfeeding self-efficacy was revealed by this review, there is still a paucity of evidence on the mode, format, and intensity of interventions.	The findings of this systematic review and meta- analysis suggest that breastfeeding education is an effective intervention for improving breastfeeding self-efficacy and breastfeeding duration among primiparous women.

Table 13: Example summaries from the test set generated using the BART configuration.

and 8 beams; maximum lengths of 64, 128, and 256 wordpieces; and length penalties of 1, 2, and 4. We find little qualitative or quantitative variation between runs and select the setting with the highest Rouge1 scores: 6 beams, a length penalty of 2, and 128 tokens for output maximum lengths. We use an attention dropout (Srivastava et al., 2014) of 0.1. Optimizer hyperparameters, as well as any hyperparameters not mentioned, used defaults corresponding to their libraries. Training requires approximately one day on two GPUs. Due to memory constraints, we limit each review to 25 input documents, with a maximum of 1000 tokens per input document.

We make use of NumPy (Harris et al., 2020) in our models and evaluation, as well as scikit-learn (Buitinck et al., 2013), and the general SciPy framework (Virtanen et al., 2020) for evaluation.

G Example generated summaries

See Table 13 for examples of inputs, targets, and generations.

Model	R-1	R-2	R-L	ΔEI	F1
BART	26.66	9.04	19.78	.447	49.68
LED	25.82	8.44	19.29	.482	47.09

Table 14: *texts-to-text* results on the validation set. We report ROUGE, ΔEI , and macro-averaged F1-scores. These are similar to test scores.

Model	P	R	F1
BART	46.98	45.39	46.97

Table 15: *table-to-table* results on the validation set. We report precision, recall, and macro-averaged F1-scores.

H Validation Results

We provide results on the validation set in Tables 14 and 15.

I A Brief Review of Systematic Reviews

We provide a brief overview of the systematic review process for the reader. A systematic review is a thorough, evidence-based process to answer

scientific questions. In the biomedical domain, a systematic review typically consists of five steps: defining the question, finding relevant studies, determining study quality, assessing the evidence (quantitative or qualitative analysis), and drawing final conclusions. For a detailed overview of the steps, see [Khan et al. \(2003\)](#). While there are other definitions and aspects of the review process ([Aromataris and Munn, 2020](#); [Higgins et al., 2019](#)), the five-step process above is sufficient for describing reviews in the context of this work. We emphasize that this work, indeed the approaches used in this field, *cannot* replace the labor done in a systematic review, and may instead be useful for scoping or exploratory reviews.

The National Toxicology Program,¹³ part of the United States Department of Health and Human Services, conducts scoping reviews for epidemiological studies. The National Toxicology Program has actively solicited help from the natural language processing community via the Text Analysis Conference.¹⁴ Other groups conducting biomedical systematic reviews include the Cochrane Collaboration,¹⁵ the Joanna Briggs Institute,¹⁶ Guidelines International Network,¹⁷ SickKids,¹⁸ the University of York,¹⁹ and the public health agencies of various countries,²⁰ to name a few. Systematic review methodologies have also been applied in fields outside of medicine, by organizations such as the Campbell Collaboration,²¹ which conducts reviews over a wide range of areas: business, justice, education, and more.

I.1 Automation in Systematic Reviews

Automation in systematic reviews has typically focused on assisting in portions of the process: search and extraction, quality assessment, and interpreting findings. For a detailed analysis of automated approaches in aiding the systematic review process, see [Norman \(2020\)](#); [Marshall and Wallace \(2019\)](#).

Search and Extraction. Search, screening, and extracting the results of studies into a structured representation are several components of the sys-

tematic review process that have been the major focuses of natural language processing approaches. Several systems provide active-learning enhanced search ([Howard et al., 2020](#); [Schoot et al., 2020](#)), or offer screening based on study type ([Marshall et al., 2016](#)). PICO (Participants, Interventions, Controls, and Outcomes) elements can be used to assist in search and screening ([Znaidi et al., 2015](#); [Ho et al., 2016](#); [Boudin et al., 2010](#)). To this end, several datasets have been introduced. EBM-NLP ([Nye et al., 2018](#)) is a dataset of crowd-sourced PICO elements in randomized control trial abstracts. [Jin and Szolovits \(2018\)](#) provides a large-scale dataset of sentence-level PICO labels that are automatically derived using the structured abstract headers in PubMed abstracts. The Chemical-Disease Relations challenge ([Wei et al., 2015](#)) offers data for some of the PICO classes and a related relation extraction task, as does the i2b2 2010 disease-relation task ([Uzuner et al., 2011](#)). Evidence Inference ([Lehman et al., 2019](#); [DeYoung et al., 2020](#)) attempts to automate detecting the direction of conclusions given PICO elements of interest; e.g., [Nye et al. \(2020\)](#) starts from RCTs, finds PICO elements, and then finds conclusions associated with those PICO elements. Many review tools^{22,23,24,25} incorporate workflow management tools for manual extraction of these elements and associated conclusions.

Quality Assessment. Relatively few tools focus on quality assessment. The primary tool seems to be RobotReviewer ([Marshall et al., 2016](#)), which assesses Risk of Bias in trial results, which is one aspect of quality. There are opportunities for quality assessment that focus on automatically assessing statistical power or study design.

Interpretation. The interpretation step of the systematic review process involves drawing overall conclusions about the interventions studied: how effective is the intervention, when should it be used, what is the overall strength of the evidence supporting the effectiveness and recommendations, and what else needs to be studied. It too has received relatively little attention from those developing assistive systems. Similar to this work, [Wallace et al. \(2020\)](#) takes advantage of structured Cochrane reviews to identify summary targets, and uses portions of the input documents as model inputs. [Shah](#)

¹³<https://ntp.niehs.nih.gov/>

¹⁴<https://tac.nist.gov/2018/SRIE/>

¹⁵<https://www.cochrane.org/>

¹⁶<https://jbi.global/>

¹⁷<https://www.g-i-n.net>

¹⁸<https://www.sickkids.ca/>

¹⁹<https://www.york.ac.uk/crd/>

²⁰<https://www.canada.ca/en/public-health/services/reports-publications.html>

²¹<https://www.campbellcollaboration.org>

²²<https://www.evidencepartners.com/>

²³<https://www.covidence.org/reviewers/>

²⁴<https://sysrev.com/>

²⁵<https://www.jbisumari.org/>

et al. (2021) extracts relations from nutritional literature, and uses content planning methods to generate summaries highlighting contradictions in the relevant literature.