

# Scientific Opinion Summarization: Paper Meta-review Generation Dataset, Methods, and Evaluation

Author Name

Affiliation

email@example.com

## Abstract

Opinions in scientific research papers can be divergent, leading to controversy or consensus among reviewers. However, most existing datasets for opinion summarization are centered around product reviews and assume that the analyzed opinions are non-controversial, failing to account for the variability seen in other contexts such as academic papers, political debates, or social media discussions. To address this gap, we propose the task of scientific opinion summarization, where research paper reviews are synthesized into meta-reviews. To facilitate this task, we introduce the ORSUM dataset covering 15,062 paper meta-reviews and 57,536 paper reviews from 47 conferences. Furthermore, we propose the Checklist-guided Iterative Introspection (CGI<sup>2</sup>) approach, which breaks down scientific opinion summarization into several stages, iteratively refining the summary under the guidance of questions from a checklist. Our experiments show that (1) human-written summaries do not always satisfy all necessary criteria such as depth of discussion, and identifying consensus and controversy for the specific domain, and (2) the combination of task decomposition and iterative self-refinement shows strong potential for enhancing the opinions and can be applied to other complex text generation using black-box LLMs.

## 1 Introduction

Opinion Summarization traditionally targets product reviews, aiming to distill representative opinions on key product aspects such as product quality and price. This assumes a dominant, singular opinion within the texts being summarized [Hu and Liu, 2006; Amplayo *et al.*, 2021b; Angelidis and Lapata, 2018; Suhara *et al.*, 2020]. However, this approach often overlooks the nuanced and multi-faceted nature of discussions in scientific documents, where multiple viewpoints may coexist and no single opinion dominates.

Furthermore, most opinion summarization datasets in the product domain for abstractive summarization are synthetic, containing redundant cut-and-paste extracts built by combining extracted snippets, or by sampling a review from the

Domain	Reviews	Meta-reviews
Product	I love these protein bars in the vanilla flavor. They taste like Rice Krispies treats with vanilla frosting ...    Nugo bars are great for breakfast, lunch or a snack ... Eat them with a tall glass of water and they will keep you satisfied for hours.    ...	These bars are fantastic and taste great like a Rice Krispy treat. Good for morning, lunch or afternoon snack and a good way to get your protein in-take. They keep you full for a long time especially if you are out and about ...
Paper	It is unclear why this work is needed. Why not use ...    The paper is well written and the math seems to be sound ... The empirical evaluation of the method is not overwhelming ...    The work appears to be sound ...	Two of the reviews suggest that the technical aspects of the paper are sound, while one reviewer questions the need for the proposed approach ... While some reviewers raised concerns about ... the majority of reviewers acknowledge the ... In light of these findings, I recommend rejection ...

Figure 1: Product meta-reviews and paper meta-reviews have different compositions: A product meta-review presents the most prominent opinion instead of summarizing opinions, while a paper meta-review summarizes different opinions and makes recommendations.

collection and pretending that it is a gold-standard meta-review [Amplayo *et al.*, 2021b].

To address this gap, we introduce the new task of **Scientific Opinion Summarization**, where a set of opinions must be synthesized into a meta-opinion that justifies a decision. Scientific Opinion Summarization aims to provide a succinct synopsis for scientific documents, helping readers to recap salient information and understand the professional discussion. Scientific meta-reviews, in particular, summarize the *controversies* and *consensuses* in the reviews, guiding decision making such as the acceptance or rejection of a paper. Taking research paper meta-review generation as a typical scenario, we build the **ORSUM** dataset by collecting open-sourced paper and meta-reviews from OpenReview<sup>1</sup>, covering 15,062 meta-reviews and 57,536 reviews from 47 conference venues. Compared to synthetic datasets from product review domains, ORSUM is built upon large-scale real-world data, enabling applications of supervised abstractive summarization methods and more fine-grained tex-

<sup>1</sup><https://openreview.net/>

tual analysis. In addition to meta-review generation, ORSUM’s structured content, including ratings on different aspects such as if agreements/disagreements are present alongside strengths/weaknesses and multi-turn discussions, will benefit a wide range of related tasks, such as review generation [Wang *et al.*, 2020], recommendation prediction [Deng *et al.*, 2020; Friedl *et al.*, 2021], review rating prediction [Li *et al.*, 2017; Chan *et al.*, 2020], and argument pair extraction [Cheng *et al.*, 2020].

The task of Scientific Opinion Summarization presents a distinct set of challenges, including (1) *Decision Consistency*: Whether the Meta-review aligns with the decision, which guides opinion selection and discussion in the meta-review. Generated scientific meta-reviews should reflect these decisions. (2) *Discussion involvement*: Unlike product meta-reviews that rely on majority voting, scientific meta-reviews assess both the pros and cons, as well as opinion agreement and disagreement, to evaluate the paper from the perspective of a more senior reviewer.

To tackle these challenges, we propose Checklist-guided Iterative Introspection (CGI<sup>2</sup>). CGI<sup>2</sup> first breaks the task of scientific opinion summarization into multiple steps, constantly requesting evidence to mitigate both LLMs’ inability to follow complicated instructions and their tendency to produce hallucinations. To enhance discussion involvement, CGI<sup>2</sup> iteratively revises the generated meta-review based on a predefined checklist. Finally, we identify key aspects a meta review should satisfy to be of high quality, and propose ways to evaluate these aspects using reference-free LLM-based metrics.

Our contributions include the following:

- We introduce the task of scientific opinion summarization and construct the ORSUM dataset, which contains 15,062 meta-reviews and 57,536 reviews from 47 conferences on OpenReview. It is currently the largest paper meta-review dataset.
- We propose Checklist-guided Iterative Introspection (CGI<sup>2</sup>), which breaks down the task of scientific opinion summarization into several stages and iteratively refines the summary under the guidance of questions from a checklist.
- We construct a comprehensive evaluation framework for meta-review generation and assess the different summarization paradigms on ORSUM.

## 2 Related Work

### 2.1 Opinion Summarization

The task of opinion summarization is typically decomposed into three stages: aspect extraction, which identifies the specific features discussed in reviews; polarity identification, which assesses whether the sentiment towards each aspect is positive, negative, or neutral; and summary generation, which compiles these aspects and sentiments into a cohesive summary of the opinions [Hu and Liu, 2006]. The lack of parallel data in review summaries limits most methodologies into the few-shot abstractive setting [Brazinskas *et al.*, 2020a;

Brazinskas *et al.*, 2022], or unsupervised extractive setting [Angelidis and Lapata, 2018; Angelidis *et al.*, 2020; Chowdhury *et al.*, 2022] where the aspects and sentiments from the input reviews are collected, selected, and rearranged into the output meta-reviews.

Only a few previous opinion summarization datasets [Wang and Ling, 2016] contain gold-standard summaries and support supervised training of abstractive models [Amplayo and Lapata, 2019]. Pretrained aspect-based sentiment analysis [Suhara *et al.*, 2020], variational autoencoders [Brazinskas *et al.*, 2020b; Chu and Liu, 2019; Iso *et al.*, 2021; Isonuma *et al.*, 2021] and large language models [Bhaskar *et al.*, 2022] enable unsupervised abstractive approaches, where the generated summaries are validated to be more fluent, informative, coherent, and concise compared to traditional extractive summaries.

To support the training and evaluation of supervised methods, recent work constructs synthetic datasets by random sampling [Shen *et al.*, 2023], adding noise to the sampled summary to generate documents [Amplayo and Lapata, 2020], searching for relevant reviews to act as the input document set [Elsahar *et al.*, 2021], or sampling with trained models [Amplayo *et al.*, 2021a; Amplayo *et al.*, 2021b]. However, synthetic pseudo-summaries in the product review domain are known to be detached from real-world distributions, be possibly irrelevant or inconsistent with input documents, and are known to ignore important underlying details.

### 2.2 Meta-review Generation

The first attempt to generate paper meta-reviews is MetaGen [Bhatia *et al.*, 2020], which generates an extractive summary draft then uses a fine-tuned model for decision prediction and abstractive review generation. [Kumar *et al.*, 2021] emphasizes decision awareness, proposing a model for decision prediction and subsequent meta-review generation. The most similar work to ours is MR<sub>ED</sub> [Shen *et al.*, 2022], where 7,089 paper meta-reviews from ICLR 2018 - 2021 are manually annotated with sentence-level structure labels. These structure labels categorize sentences based on their function in the document, such as summary, evaluation, or recommendation. The difference between their work and ours is that they focus on structure-controlled text generation, while our work 1) enables scientific opinion summarization with a larger corpus, 2) provides a prompting-based solution, and 3) performs broader evaluations. Note that while there are other concurrent efforts to collect paper meta-reviews or reviews [Dycke *et al.*, 2023], we are the first to model meta-review generation as scientific opinion summarization and to offer a unified dataset covering a broad range of conference venues.

## 3 Task Formulation

Given a research paper’s title, abstract, and set of reviews, the goal of **Scientific Opinion Summarization** is to generate a meta-review summarizing the reviews’ opinions in order to make a decision recommendation for acceptance or rejection.

As noted by ACL’s area chair guidance<sup>2</sup>, meta-reviews

<sup>2</sup><https://aclrollingreview.org/aetutorial>

Dataset	Collection	Count(SRC)	Count(TRG)	Len(SRC)	Len(TRG)	Novel 4-gram	NID
RT	Human	246,164	3,731	20.57	21.4	97.10	0.1615
Copycat	AMT	480	180	42.63	54.33	89.62	0.2506
OPOSUM	AMT	600	60	43.51	67.77	85.92	0.1260
Yelp	AMT	3,200	200	65.25	61.15	93.26	0.1661
DENOISESUM	Synthetic	73282	837	24.32	26.45	94.12	0.2270
PLANSUM	Synthetic	249,844	869	42.81	97.2	91.40	0.2395
SPACE	Human	5000	1050	34.27	54.38	90.38	0.1671
<b>ORSUM</b>	Human	57,536	15,062	376.36	141.76	99.89	0.1572

Table 1: We compare ORSUM with existing opinion summarization datasets that contain gold-standard summaries. SRC refers to the source or input reviews. TRG refers to the target or output meta-reviews. A higher novel 4-gram score suggests better abstractiveness, while a lower NID score implies less redundancy.

171 summarize reviews by aggregating opinions to support the  
 172 decision. The task entails summarizing the paper’s key  
 173 strengths and weaknesses and explicitly evaluating whether  
 174 those strengths surpass the weaknesses.

## 4 ORSUM Dataset

### 4.1 Dataset Collection and Preprocessing

175 To facilitate the task of scientific opinion summarization, we  
 176 collect the **ORSUM** dataset which consists of human-written  
 177 meta-reviews from OpenReview. The dataset contains each  
 178 paper’s URL, title, abstract, decision, meta-review from the  
 179 area chair, and reviews from individual reviewers. We crawl  
 180 15,062 paper meta-reviews and 57,536 individual reviews  
 181 from 47 conference venues. Papers with meta-reviews shorter  
 182 than 20 tokens and comments made by non-official review-  
 183 ers are excluded. The data format is unified across venues,  
 184 and we provide train/validation/test splits with 9,890/549/550  
 185 samples for convenient usage by future works.

### 4.2 Dataset Comparison

186 We compare ORSUM with existing opinion summarization  
 187 datasets (or their subsets) with gold-standard summaries, in-  
 188 cluding The Rotten Tomatoes (RT) [Wang and Ling, 2016],  
 189 Copycat [Brazinskas *et al.*, 2020b], OPOSUM [Angelidis  
 190 and Lapata, 2018], Yelp [Chu and Liu, 2019], DENOIS-  
 191 ESUM [Amplayo and Lapata, 2020], PLANSUM [Amplayo  
 192 *et al.*, 2021b], and SPACE [Angelidis *et al.*, 2021] datasets.  
 193 To perform a quantitative comparison, we utilize two key  
 194 metrics:

195 **Abstractiveness.** The percentage of novel n-grams in a  
 196 meta-review is defined by the ratio of n-grams which do not  
 197 appear in the source reviews, to the total number of n-grams  
 198 in the meta review. This metric intuitively measures the ab-  
 199 stractiveness of the summaries [Chen *et al.*, 2021]. Table 1  
 200 indicates a greater degree of abstractiveness in ORSUM.

201 **Redundancy.** To examine the presence of insightful in-  
 202 formation in the input reviews, we assess redundancy using  
 203 the Normalized Inverse of Diversity (NID) score [Xiao and  
 204 Carenini, 2020] This score is calculated as the inverse of  
 205 the diversity metric, which measures the variability of infor-  
 206 mation in the reviews with length normalization:  $NID =$   
 207  $1 - \frac{\text{entropy}(D)}{\log(|D|)}$ . A higher NID signifies greater redundancy.  
 208 Table 1 shows lower redundancy in ORSUM, which can be




Figure 2: Meta-review composition. The scores range from 0 to 2: 0 indicates that the meta-review does not address the discussion at all. 1 signifies that the meta-review incorporates the discussion but lacks concrete evidence. 2 denotes that the meta-review involves a detailed discussion. Only 47.7% and 35.0% of meta-reviews meet the fundamental criteria for discussions of advantages and disadvantages, and consensus and controversy, respectively.

212 attributed to the fact that many reviews address distinct as-  
 213 pects of their papers.

### 4.3 Composition Analysis

214 To investigate whether ORSUM’s human-authored meta-  
 215 reviews discuss both a paper’s pros/cons and the reviews’  
 216 level of agreement/disagreement, we conduct a human evalua-  
 217 tion focused on meta-review composition. Three annotators  
 218 are asked to assess the meta-reviews in terms of **discussion**  
 219 **involvement**: how effectively a summary engages with the  
 220 content by discussing the paper’s advantages/disadvantages,  
 221 and by discussing the agreements/disagreements of the re-  
 222 views. Annotation scores range from 0 (no involvement) to 2  
 223 (detailed involvement).

224 Our evaluation results depicted in Figure 2 reveal that only  
 225 20.7% of meta-reviews include an assessment of both advan-  
 226 tages/disadvantages and review agreements/disagreements,  
 227 regardless of their length. For each category, 47.7%, and  
 228 35.0% of meta-reviews meet the criteria of containing dis-  
 229 cussions of advantages and disadvantages and discussions of  
 230 agreements/disagreements, respectively. Based on these re-  
 231 sults, we conclude that *human-written meta-reviews do not*  
 232 *always meet the necessary criteria for an effective meta re-*  
 233 *view, and they may be unsuitable for developing summariza-*  
 234 *tion models as supervised training signals. The low percent-*  
 235 *age of comprehensive reviews highlights a gap in coverage*  
 236 *and thoroughness that can affect the performance and reli-*  
 237 *ability of models trained on these summaries.*  
 238

## 5 Checklist-guided Iterative Introspection Method for Meta-review Generation

Motivated by the unreliability of human-written meta-reviews, we turn to Large Language Models (LLMs) like ChatGPT [OpenAI, 2021] for meta-review generation. We choose LLMs for their world knowledge, and their potential to generate reviews efficiently and scalably. However, LLMs struggle to follow complicated instructions, and have a tendency to produce hallucinations. To mitigate these deficiencies, we propose to break the task of scientific review generation into multiple steps, consistently requesting evidence for each step. To enhance discussion involvement and evidence-based coherence in the generation process, we further introduce a checklist-guided self-feedback mechanism. Our method is similar to the process of self-refinement [Madaan *et al.*, 2023], which involves the LLM iteratively revising the generated meta-review based on its own feedback. Unlike prior work, however, our checklist-guided self-feedback uses self-feedback derived from questions in a predefined checklist, ensuring that the revision process progresses towards our desired criteria. Figure 3 illustrates our proposed Checklist-guided Iterative Introspection (CGI<sup>2</sup>) method.

**Initial Run.** Given a paper’s title, abstract, and set of reviews, CGI<sup>2</sup> generates a draft of the meta-review in four steps: (1) For each review, we prompt the LLM to extract and rank opinions, while including sentiment, aspect, and evidence. Due to the input length constraint, each review is truncated to 300 tokens. (2) Based on the extracted opinions, we prompt the LLM to list the paper’s most important advantages and disadvantages, the evidence for those statements, and those statements’ corresponding reviewers. (3) We prompt the LLM to list the consensuses and controversies in the reviews, the evidence for those statements, and their corresponding reviewers. (4) Given the paper’s acceptance or rejection decision, we prompt the LLM to write a meta-review based on the information extracted in steps (1)–(3).

**Iterative Runs.** With the meta-review draft from the initial four-step run, CGI<sup>2</sup> iteratively poses questions, obtains self-feedback, and requests further refinement. In each run, we first select an assessment question from a pre-constructed list of questions, as shown in Table 2. This checklist, customized for meta-review generation, covers the four most crucial aspects of meta-reviews. The checklist can also easily be expanded and adapted to other complex text generation tasks. After prompting the LLM with the assessment questions, we collect the refinement suggestions from the LLM’s. These refinement suggestions are used as prompts to generate a revised version of the meta-review. The checklist questions are posed sequentially in one iterative run, with the number of iterations set as a hyper-parameter in CGI<sup>2</sup>.

Our proposed approach offers two key benefits. First, it eliminates the need for external scoring functions that demand training data or human annotations. Second, it provides a general solution for employing LLMs as black boxes in complex text generation tasks.

## 6 Evaluation

Meta-review generation requires a system to accurately summarize opinions, highlight reviewer consensuses and controversies, offer judgments, and make recommendations. The task’s complexity thus requires an evaluation that is multi-faceted and goes beyond n-gram similarity. However, current evaluation metrics for long text generation are inadequate to measure the particular requirements of meta-review generation. To address this gap, we propose a comprehensive evaluation framework that combines standard evaluation metrics with LLM-based evaluation metrics.

### 6.1 Standard Metrics

We apply standard metrics in natural language generation to assess *relevance*, *factual consistency*, and *semantic coherence*. For relevance, ROUGE-L [Lin, 2004] quantifies the similarity between the generated and reference texts by calculating the longest common subsequence, while BERTScore [Zhang *et al.*, 2020] offers a more nuanced relevance evaluation by leveraging contextualized embeddings without relying on n-gram overlaps. For factual consistency, FACTCC [Kryscinski *et al.*, 2019] checks whether a given claim in the generated text is consistent with the facts presented in the source document, while SummaC [Laban *et al.*, 2021] utilizes sentence-level natural language inference models for inconsistency detection. For semantic coherence, DiscoScore [Zhao *et al.*, 2022] presents six BERT-based model variants to measure discourse coherence. We average the scores from these six models as the coherence indicator. The references used in our reference-free evaluation metrics are sourced from a test subset of our dataset, where the instances are chosen for their relevance and quality. These references provide a practical benchmark that mirrors current standards in meta-review generation at top conferences.

### 6.2 LLM-based Metrics

The aforementioned methods do not evaluate discussion involvement or evidence-decision consistency. Some reference summaries may not include discussions or utilize evidence to substantiate decisions. To address this, we propose supplementary measures for this task that can be assessed and quantified using reference-free LLM-based metrics. We aim to assess the following key aspects:

- **Discussion involvement:** whether the meta-review discusses the paper’s strengths and weaknesses, and the paper’s agreements and disagreements amongst reviewers.
- **Opinion Faithfulness:** whether the meta-review contradicts reviewers’ opinions.
- **Decision Consistency:** whether the meta-review accurately reflects the final decision.

Despite its prevalence, the GPTScore [Fu *et al.*, 2023] metric requires its criteria to be described as a single word, a requirement incompatible with our detailed criteria. On the other hand, G-EVAL [Liu *et al.*, 2023] assesses the quality of NLG outputs by utilizing chain-of-thought (CoT) and a form-filling paradigm. It has been shown to have a very high correlation with human-based judgments. G-EVAL uses carefully

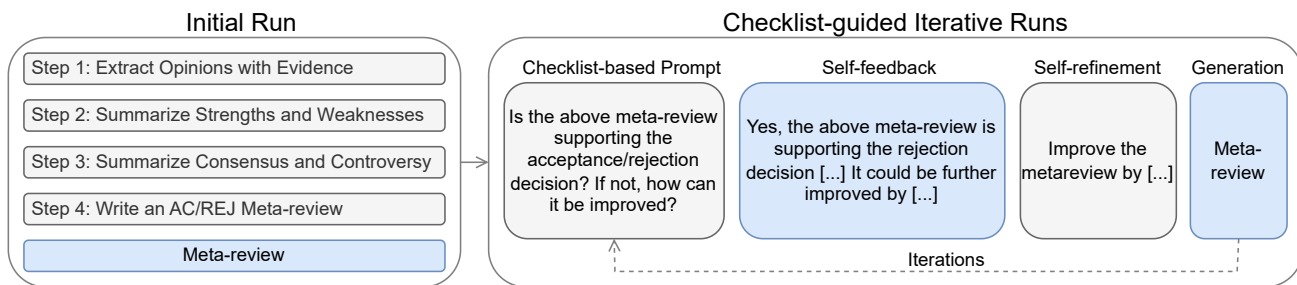


Figure 3: Our proposed CGI<sup>2</sup> framework operates through multiple iterations. In the initial iteration, the task is divided into four steps: (1) Review Opinion Extraction, (2) Strength and Weakness Synthesis, (3) Consensus and Controversy Analysis, and (4) Meta-review Drafting. For subsequent iterations, we present the black-box LLM with a query from a predefined list, acquire self-feedback, and request additional refinements.

- 
1. Are the most important advantages and disadvantages discussed in the above meta-review? If not, how can it be improved?
  2. Are the most important consensus and controversy discussed in the above meta-review? If not, how can it be improved?
  3. Is the above meta-review contradicting reviewers' comments? If so, how can it be improved?
  4. Is the above meta-review supporting the acceptance/rejection decision? If not, how can it be improved?
- 

Table 2: The extensible and easily adaptable checklist for Meta-review Generation accesses the essential aspects of self-consistency, faithfulness, and active engagement in discussions.

#### G-EVAL

You will be given one metareview written for reviews by the committee on a paper. Your task is to rate the metareview on one metric. Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.

Evaluation Criteria: Quality of Metareview (1-5) - the collective quality of all sentences. We align this dimension with the DUC quality question of structure and coherence whereby the metareview should be well-structured and well-organized. The metareview should always discuss the disadvantages and advantages of a paper and have a clear scope of the accept/reject decision. The metareview should have concrete evidence from the papers reviews and concrete comments as well.

Evaluation Steps:

1. Read the reviews carefully and identify the main topic and key points.
2. Read the metareview and compare it to the reviews. Check if the metareview covers the main topic, discusses advantages and disadvantages, if the most important advantages and disadvantages discussed in the above meta-review, if the most important advantages and disadvantages discussed in the above meta-review, if the most important consensus and controversy discussed in the above meta-review, if the above meta-review contradicting reviewers' comments, if the above meta-review supporting the acceptance/rejection decision, and if it presents them in a clear and logical order.
3. Assign a score for the quality of the meta-review on a scale of 1 to 5, where 1 is the lowest and 5 is the highest based on the Evaluation Criteria.

Source Text: {Reviews} Metareview: {Meta-review} Evaluation Form (scores ONLY): - Quality of metareview :

#### Likert scale scoring with ChatGPT

Imagine you are a human annotator now. You will evaluate the quality of metareviews written for a conference by giving a mean value from 1 to 5 and no other explanation. Please follow these steps:

1. Carefully read the reviews, and be aware of the information it contains.
2. Read the proposed metareview.
3. Rate the summary on three dimensions: 'Discussion Involvement', 'Opinion Faithfulness' and 'Decision Consistency'. You should rate on a scale from 1 (worst) to 5 (best) and give me an average of these scores over all aspects from 1 to 5 calculated by the mean of all aspects.

Definitions are as follows:

- (1) Discussion Involvement: Whether the meta-review discusses the paper's strengths and weaknesses, as well as agreements and disagreements among reviewers,
- (2) Opinion Faithfulness: Whether the meta-review contradicts reviewers' comments,
- (3) Decision Consistency: Whether the meta-review accurately reflects the final decisions.

Only generate the mean rating as a number on the likert scale, nothing else.

Figure 4: We customize the prompts in G-EVAL and GPTLikert for evaluating meta-review generation to assess discussion involvement, opinion faithfulness, and decision consistency.

349 constructed instructions for GPT models to follow, yielding a  
 350 rating on the Likert scale ranging from 1 to 5. Likert scoring  
 351 with ChatGPT (GPTLikert), a human-like evaluation method  
 352 introduced by [Gao *et al.*, 2023], follows a similar evalua-  
 353 tion protocol, outperforming many standard metrics on text  
 354 summarization as measured by human correlation. We are

the first to adapt these methods to meta-review generation by 355  
 356 modifying the prompts as shown in Figure 4. The combi-  
 357 nation of standard metrics like ROUGE-L and BERTScore  
 358 with LLM-based metrics such as G-EVAL and GPTLikert  
 359 ensures a comprehensive evaluation, capturing nuances that  
 360 traditional metrics may overlook. This multifaceted approach

361 not only adheres to current evaluation methodologies, but also  
362 enhances them by introducing metrics that demonstrate a high  
363 correlation with human annotations.

## 364 7 Experiments

### Human

The paper unanimously receives positive rates thanks to strong motivations and interesting results. As the reviews show satisfaction on the authors' feedback, the final draft needs to respect it accordingly, for example, about the limitations of this research.

### Vanilla

This paper could have benefited from a clearer and more concise presentation of the main findings. It is sometimes difficult to follow the paper's chain of thought, and the authors could have presented their results in a more accessible way. The experiments are limited to a small set of environments, and it would be interesting to see how the results generalize to other tasks. Additionally, it would be useful to have a deeper analysis of why iGPT does not perform well, as this could help inform future research on how to improve image-based pretraining.

### CGI<sup>2</sup>

After carefully considering the reviews provided, there is consensus that the paper presents an important and thought-provoking analysis of representation dynamics between pretraining and finetuning in the context of reinforcement learning. Both reviewers agree that the paper addresses an important question and presents insightful experiments. One reviewer highlights the experiments as a strong point in the paper, while the other reviewer raises a concern about the paper's tendency to present hypotheses without exploring them further.

Overall, the strengths of the paper outweigh its weaknesses, and it has the potential to make a significant contribution to the field of reinforcement learning. Therefore, I recommend accepting the paper with a minor revision to address the concern raised by the reviewer. Specifically, the paper should explore the hypotheses presented in more depth to provide a more comprehensive analysis. If the authors can address this concern, the paper can make a valuable contribution to the field of reinforcement learning. In summary, I recommend accepting the paper with a minor revision.

### Initial run

Based on the reviews, I recommend accepting the paper. The paper addresses an important and under-explored research question and presents a detailed and thought-provoking analysis of the pretraining-finetuning paradigm for Transformer models in the context of reinforcement learning. The paper also identifies relevant techniques to perform analysis of representation similarity and training dynamics. While there are some weaknesses pointed out, such as the need for further exploration of some hypotheses and occasional difficulty in following the writing and presentation, the strengths of the paper outweigh these concerns. Therefore, I recommend accepting the paper with minor revisions to address the weaknesses pointed out by the reviewers.

Figure 5: We show the meta-reviews from human, vanilla, CGI<sup>2</sup>, and CGI<sup>2</sup> without iterative runs for the same paper. The yellow background indicates hallucinated content. The green background indicates redundant content.

### 365 7.1 Baselines

366 We compare our proposed CGI<sup>2</sup> method with methods of dif-  
367 ferent paradigms. Results in Table 3 are averaged across three  
368 random runs.

369 **Abstractive Methods.** PlanSum [Amplayo *et al.*, 2021b]  
370 uses a Condense-Abstract Framework, where reviews are  
371 condensed and used as input to an abstractive summarization  
372 model. OpinionDigest [Suhara *et al.*, 2020] extracts opinions  
373 from input reviews and trains a seq2seq model that gener-  
374 ates a summary from this set of opinions. MeanSum [Chu

and Liu, 2019] is an unsupervised multi-document abstrac- 375  
376 tive summarizer that minimizes a combination of reconstruc-  
377 tion and vector similarity losses. LED [Beltagy *et al.*, 2020] is  
378 a Longformer [Beltagy *et al.*, 2020] variant supporting long  
379 document generative sequence-to-sequence tasks.

**Extractive Methods.** LexRank [Erkan and Radev, 2004] is 380  
381 an unsupervised extractive summarization method that selects  
382 sentences based on centrality scores calculated with graph-  
383 based sentence similarity. MemSum [Gu *et al.*, 2022] mod-  
384 els extractive summarization as a multi-step episodic Markov  
385 Decision Process of scoring and selecting sentences.

**Prompting Methods.** All prompting methods are initi- 386  
387 ated with the GPT-3.5-turbo model with a temperature of 0.7.  
388 3Sent [Goyal *et al.*, 2022] applies a simple prompt "Summary  
389 of document in 3 sentences". TCG [Bhaskar *et al.*, 2022] ex-  
390 plores a four-step generation pipeline involving topic classi-  
391 fication, sentence grouping by topic, generating chunk-wise  
392 summary, and generating the final summary. We also ex-  
393 plore In Context Learning (ICL) [Brown *et al.*, 2020], where  
394 a highly rated meta-review alongside the reviews is given as  
395 part of the model's prompt. This meta-review is manually  
396 picked based on adherence to the previously defined check-  
397 list, and is chosen for its fulfillment of the criteria that define  
398 a high-quality meta-review. Vanilla uses "Generate a metare-  
399 view" as the prompt. InstructPrompt provides more detailed  
400 step by step instructions and specifies the criteria for writing  
401 a metareview.

### 7.2 Automatic Evaluation 402

Higher standard metric scores indicate better summarization, 403  
404 but not necessarily better opinion summarization. ROUGE-  
405 L, BERTScore, SummaC, and DiscoScore do not consider  
406 the multifaceted nature of meta-review, which goes be-  
407 yond summarization. Our method performs near average  
408 in BERTScore and SummaC, and the highest in ROUGE-L  
409 and DiscoScore amongst the prompting methods. Compared  
410 to extractive and abstractive methods, our method achieves  
411 lower scores as some metrics measure semantic similarity  
412 which a high-quality measure review with its variability may  
413 not score well in. Additionally due to the multifaceted na-  
414 ture of opinion summarization, reference-based metrics such  
415 as Rouge-L can be biased towards the reference, thus the ele-  
416 vated scores of the summarization methods.

Evaluators like G-Eval and GPTLikert favor specific di- 417  
418 mensions given in their prompts. Our method shows promis-  
419 ing results in both G-Eval and GPTLikert due to the carefully  
420 constructed and revised prompts. Most prompting methods  
421 also outperform extractive and abstractive methods.

Human meta-reviews in the dataset scored among the low- 422  
423 est in all categories, signifying the unreliability of some  
424 human-written meta-reviews and the need for an automatic,  
425 or auxiliary, writing process. When compared by seman-  
426 tic similarity, extractive methods outperform both abstractive  
427 and prompting methods with the exception of Plansum. This  
428 is due to the nature of content planning in Plansum which is  
429 central to the task of meta-review generation.

Models	ROUGE-L	BERTScore	FactCC	SummaC	DiscoScore	G-EVAL	GPTLikert
Human	-	-	0.538	0.368	0.740	0.731	0.607
<i>Abstractive Methods</i>							
PlanSum	<b>0.465</b>	0.785	0.608	0.533	0.911	0.731	0.608
OpinionDigest	0.124	0.838	0.612	0.575	0.862	0.762	0.618
MeanSum	0.132	0.827	0.559	0.464	0.900	0.767	0.622
LED	0.161	0.846	0.618	0.785	0.958	0.731	0.624
LED-finetuned	0.221	0.853	0.634	0.795	0.961	0.751	0.649
<i>Extractive Methods</i>							
LexRank	0.433	<b>0.881</b>	<b>0.729</b>	<b>0.937</b>	<b>1.256</b>	0.726	0.656
MemSum	0.337	0.827	0.683	0.825	0.989	0.711	0.628
<i>Prompting Methods</i>							
Vanilla	0.174	0.817	0.498	0.423	0.808	0.752	0.626
3Sent	0.109	0.783	0.562	0.503	0.667	0.758	0.661
InstructPrompt	0.208	0.823	0.543	0.449	0.862	0.751	0.646
TCG	0.189	0.847	0.544	0.466	0.895	0.761	0.632
ICL	0.192	0.847	0.578	0.470	0.871	0.756	0.612
CGI <sup>2</sup> (ours)	0.199	0.836	0.559	0.320	0.906	<b>0.770</b>	<b>0.687</b>
CGI <sup>2</sup> w/o Iterative Runs	0.118	0.830	0.536	0.332	0.849	0.732	0.629

Table 3: ROUGE-L and BERTScore assess semantic similarity with reference text. FactCC and SummaC detect factual consistency. DiscoScore measures coherence. G-EVAL and GPTLikert are GPT-based comprehensive evaluation measures for discussion involvement, opinion faithfulness, and decision consistency.

Model	Informativeness	Soundness	Self-consistency	Faithfulness
Human	0.71	0.68	0.67	-
LED-finetuned	0.56	0.46	0.21	0.73
LexRank	0.87	0.94	0.16	-
CGI <sup>2</sup> (ours)	<b>0.98</b>	<b>0.92</b>	<b>0.84</b>	<b>0.79</b>
CGI <sup>2</sup> w/o Iterative Runs	0.97	0.76	0.48	0.74

Table 4: Human annotation results on meta-reviews for 50 challenging papers from the test set.

### 7.3 Human Evaluation

We conduct a human annotation on 50 challenging papers from the test set which have average review scores on the borderline of acceptance. Five anonymized outputs from Human, LED-finetuned, LexRank, CGI<sup>2</sup>, and CGI<sup>2</sup> without iterative runs, are shown to three annotators. Annotators are asked to provide binary labels for informativeness, soundness, self-consistency, and faithfulness for each meta-review. Informativeness measures whether the meta-review involves a discussion of both strengths and weaknesses. Soundness examines whether the meta-review provides evidence to support the discussed strengths and weaknesses. Decision consistency indicates whether the recommendation decision is clearly written and consistent with the comments in the meta-review. Faithfulness evaluates whether the meta-review contains hallucinations. We assume Human and the extractive LexRank framework have perfectly faithful summaries.

Results shown in Table 4 validate the effectiveness of our proposed method. The extractive method (LexRank) is easily biased toward one reviewer and involves no discussion or decision, but generates no hallucinations by construction. The abstractive method (LED-finetuned) learns to copy the sentences in the input and form a short meta-review with little discussion, sometimes hallucinating or generating repetitive outputs. Our prompting-based method exhibits less hallucination due to the evidence requirements in our

prompts. Compared to human-written meta-reviews, all automatic methods are less capable of generating in-depth analyses, a deficiency which calls for knowledge enhancement that happens a LLM enhanced with reviews.

We also observe that hallucinations in LLMs are more likely to happen when summarizing consensuses and controversies, which require information from the paper itself. By contrast, the abstractive methods’ hallucinations were more likely to be general comments, whereas extractive methods tend to misrepresent the context by selecting irrelevant or less important sections. Despite our method’s improvements in this area, hallucination detection for scientific opinion summarization remains an open problem.

### 7.4 Case Study

Figure 5 presents the meta-reviews from human, vanilla, CGI<sup>2</sup>, and CGI<sup>2</sup> without iterative runs for a random paper<sup>3</sup>.

We make the following general observations: (1) The hallucination problem is alleviated in CGI<sup>2</sup> as the model is constantly asked for evidence. (2) CGI<sup>2</sup>’s summary sentences are redundant. (3) The vanilla prompting baseline does not make recommendations and involve discussion, as the model fails to fully understand the complex task requirement. (4) Iterative refinement sometimes improves the concreteness of

<sup>3</sup>[https://openreview.net/forum?id=9GXoMs\\_ckJ](https://openreview.net/forum?id=9GXoMs_ckJ)

479 opinion discussion. However, there are two problems with it-  
480 erative refinements. First, suggestions provided by the large  
481 language model are usually generic and less useful for further  
482 refinement. Second, more self-refinement iterations cause the  
483 model to forget the initial instructions for opinion extraction  
484 and discussion.

## 485 8 Conclusions and Future Work

486 In this paper, we introduced the task of scientific opinion  
487 summarization, in which research paper reviews are synthe-  
488 sized into meta-reviews. To facilitate this task, we introduce  
489 the ORSUM dataset, an evaluation framework, and an ap-  
490 proach that we call Checklist-Guided Iterative Introspection.  
491 We conduct an empirical analysis of methods from differ-  
492 ent paradigms, concluding that human-written summaries do  
493 not always satisfy the criteria of an ideal meta-review, and  
494 that the combination of task decomposition and iterative self-  
495 refinement shows promise in on this task.

496 Direct extensions of this work include the incorporation of  
497 author rebuttals into the input data to enhance the model's  
498 ability to generate more balanced meta-reviews, and intro-  
499 ducing an effective and efficient hallucination detection tool  
500 for scientific opinion summarization.

## 501 Limitations

502 This work on scientific opinion summarization has limitations  
503 in terms of data scope and task configuration. As the dataset is  
504 collected from OpenReview, the majority of meta-reviews are  
505 in Machine Learning, and many papers have been accepted.  
506 Conclusions drawn from this data distribution might not be  
507 applicable to datasets in other domains. Furthermore, to sim-  
508 plify the task setting, author rebuttals have not been included  
509 as input, which may also constrain the extent of discussion in-  
510 volvement in generating meta-reviews. section\*Ethics State-  
511 ment

512 We acknowledge the following potential ethical concerns  
513 that may arise. First, the meta-reviews generated by LLMs  
514 may contain hallucinations, which may lead to misunder-  
515 standings of the original research paper or reviewers' opin-  
516 ions. Therefore, users should be cautious when using system-  
517 generated meta-reviews for recommendation decisions. Sec-  
518 ond, the use of black-box LLMs for meta-review genera-  
519 tion may raise concerns about the transparency of the de-  
520 cision process. Though our method improves explainabil-  
521 ity by prompting an LLM to provide supporting evidence  
522 for the recommendation decision, the evidence may not per-  
523 fectly reflect the decision-making process. Third, the dataset  
524 used in this study mainly focuses on machine learning pa-  
525 pers, which could introduce biases to the recommendation  
526 decisions. Hence, it is critical to consider these biases when  
527 applying our method to generate meta-reviews for research  
528 papers in other domains.

## 529 References

530 [Amplayo and Lapata, 2019] Reinald Kim Amplayo and  
531 Mirella Lapata. Informative and controllable opinion sum-  
532 marization. *CoRR*, abs/1909.02322, 2019.

[Amplayo and Lapata, 2020] Reinald Kim Amplayo and 533  
Mirella Lapata. Unsupervised opinion summarization with 534  
noising and denoising. In Dan Jurafsky, Joyce Chai, Na- 535  
talie Schluter, and Joel R. Tetreault, editors, *Proceedings 536*  
*of the 58th Annual Meeting of the Association for Computa- 537*  
*tional Linguistics, ACL 2020, Online, July 5-10, 2020,* 538  
pages 1934–1945. Association for Computational Linguis- 539  
tics, 2020. 540

[Amplayo *et al.*, 2021a] Reinald Kim Amplayo, Stefanos 541  
Angelidis, and Mirella Lapata. Aspect-controllable opin- 542  
ion summarization. In *Proceedings of the 2021 Confer- 543*  
*ence on Empirical Methods in Natural Language Process- 544*  
*ing*, pages 6578–6593, Online and Punta Cana, Dominican 545  
Republic, November 2021. Association for Computational 546  
Linguistics. 547

[Amplayo *et al.*, 2021b] Reinald Kim Amplayo, Stefanos 548  
Angelidis, and Mirella Lapata. Unsupervised opinion 549  
summarization with content planning. In *Thirty-Fifth 550*  
*AAAI Conference on Artificial Intelligence, AAAI 2021,* 551  
*Thirty-Third Conference on Innovative Applications of Ar- 552*  
*tificial Intelligence, IAAI 2021, The Eleventh Symposium 553*  
*on Educational Advances in Artificial Intelligence, EAAI 554*  
*2021, Virtual Event, February 2-9, 2021*, pages 12489– 555  
12497. AAAI Press, 2021. 556

[Angelidis and Lapata, 2018] Stefanos Angelidis and 557  
Mirella Lapata. Summarizing opinions: Aspect extraction 558  
meets sentiment prediction and they are both weakly 559  
supervised. In Ellen Riloff, David Chiang, Julia Hock- 560  
enmaier, and Jun'ichi Tsujii, editors, *Proceedings of 561*  
*the 2018 Conference on Empirical Methods in Natural 562*  
*Language Processing, Brussels, Belgium, October 31 - 563*  
*November 4, 2018*, pages 3675–3686. Association for 564  
Computational Linguistics, 2018. 565

[Angelidis *et al.*, 2020] Stefanos Angelidis, Reinald Kim 566  
Amplayo, Yoshihiko Suhara, Xiaolan Wang, and Mirella 567  
Lapata. Extractive opinion summarization in quantized 568  
transformer spaces. *CoRR*, abs/2012.04443, 2020. 569

[Angelidis *et al.*, 2021] Stefanos Angelidis, Reinald Kim 570  
Amplayo, Yoshihiko Suhara, Xiaolan Wang, and Mirella 571  
Lapata. Extractive opinion summarization in quantized 572  
transformer spaces. *Transactions of the Association for 573*  
*Computational Linguistics*, 9:277–293, 2021. 574

[Beltagy *et al.*, 2020] Iz Beltagy, Matthew E. Peters, and Ar- 575  
man Cohan. Longformer: The long-document trans- 576  
former. *CoRR*, abs/2004.05150, 2020. 577

[Bhaskar *et al.*, 2022] Adithya Bhaskar, Alexander R. Fab- 578  
bri, and Greg Durrett. Zero-shot opinion summarization 579  
with GPT-3. *CoRR*, abs/2211.15914, 2022. 580

[Bhatia *et al.*, 2020] Chaitanya Bhatia, Tribikram Pradhan, 581  
and Sukomal Pal. Metagen: An academic meta-review 582  
generation system. In Jimmy X. Huang, Yi Chang, Xueqi 583  
Cheng, Jaap Kamps, Vanessa Murdock, Ji-Rong Wen, and 584  
Yiqun Liu, editors, *Proceedings of the 43rd International 585*  
*ACM SIGIR conference on research and development in 586*  
*Information Retrieval, SIGIR 2020, Virtual Event, China,* 587  
*July 25-30, 2020*, pages 1653–1656. ACM, 2020. 588



- 589 [Brazinskas *et al.*, 2020a] Arthur Brazinskas, Mirella Lap- 646  
590 ata, and Ivan Titov. Few-shot learning for opinion sum- 647  
591 marization. In Bonnie Webber, Trevor Cohn, Yulan He, 648  
592 and Yang Liu, editors, *Proceedings of the 2020 Confer- 649*  
593 *ence on Empirical Methods in Natural Language Pro- 650*  
594 *cessing, EMNLP 2020, Online, November 16-20, 2020,* 651  
595 *pages 4119–4135.* Association for Computational Linguis- 652  
596 *tics, 2020.* 653
- 597 [Brazinskas *et al.*, 2020b] Arthur Brazinskas, Mirella Lap- 654  
598 ata, and Ivan Titov. Unsupervised opinion summariza- 655  
599 tion as copycat-review generation. In Dan Jurafsky, Joyce 656  
600 Chai, Natalie Schluter, and Joel R. Tetreault, editors, *Pro- 657*  
601 *ceedings of the 58th Annual Meeting of the Association for 658*  
602 *Computational Linguistics, ACL 2020, Online, July 5-10,* 659  
603 *2020, pages 5151–5169.* Association for Computational 660  
604 *Linguistics, 2020.* 661
- 605 [Brazinskas *et al.*, 2022] Arthur Brazinskas, Ramesh Nalla- 662  
606 pati, Mohit Bansal, and Markus Dreyer. Efficient few-shot 663  
607 fine-tuning for opinion summarization. In Marine Carpuat, 664  
608 Marie-Catherine de Marneffe, and Iván Vladimir Meza 665  
609 Ruíz, editors, *Findings of the Association for Computa- 666*  
610 *tional Linguistics: NAACL 2022, Seattle, WA, United 667*  
611 *States, July 10-15, 2022, pages 1509–1523.* Association 668  
612 *for Computational Linguistics, 2022.* 669
- 613 [Brown *et al.*, 2020] Tom B. Brown, Benjamin Mann, Nick 670  
614 Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhari- 671  
615 wal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, 672  
616 Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, 673  
617 Gretchen Krueger, Tom Henighan, Rewon Child, Aditya 674  
618 Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Win- 675  
619 ter, Christopher Hesse, Mark Chen, Eric Sigler, Ma- 676  
620 teusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, 677  
621 Christopher Berner, Sam McCandlish, Alec Radford, Ilya 678  
622 Sutskever, and Dario Amodei. Language models are few- 679  
623 shot learners. In Hugo Larochelle, Marc’Aurelio Ran- 680  
624 zato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien 681  
625 Lin, editors, *Advances in Neural Information Processing 682*  
626 *Systems 33: Annual Conference on Neural Information 683*  
627 *Processing Systems 2020, NeurIPS 2020, December 6-12,* 684  
628 *2020, virtual, 2020.* 685
- 629 [Chan *et al.*, 2020] Hou Pong Chan, Wang Chen, and Irwin 686  
630 King. A unified dual-view model for review summariza- 687  
631 tion and sentiment classification with inconsistency loss. 688  
632 In Jimmy X. Huang, Yi Chang, Xueqi Cheng, Jaap Kamps, 689  
633 Vanessa Murdock, Ji-Rong Wen, and Yiqun Liu, editors, 690  
634 *Proceedings of the 43rd International ACM SIGIR confer- 691*  
635 *ence on research and development in Information Retri- 692*  
636 *eval, SIGIR 2020, Virtual Event, China, July 25-30,* 693  
637 *2020, pages 1191–1200.* ACM, 2020. 694
- 638 [Chen *et al.*, 2021] Yulong Chen, Yang Liu, Liang Chen, and 695  
639 Yue Zhang. Dialogsum: A real-life scenario dialogue sum- 696  
640 marization dataset. In Chengqing Zong, Fei Xia, Wenjie 697  
641 Li, and Roberto Navigli, editors, *Findings of the Associa- 698*  
642 *tion for Computational Linguistics: ACL/IJCNLP 2021,* 699  
643 *Online Event, August 1-6, 2021, volume ACL/IJCNLP 700*  
644 *2021 of Findings of ACL, pages 5062–5074.* Association 701  
645 *for Computational Linguistics, 2021.* 702
- [Cheng *et al.*, 2020] Liying Cheng, Lidong Bing, Qian Yu, 646  
Wei Lu, and Luo Si. APE: argument pair extraction from 647  
peer review and rebuttal via multi-task learning. In Bon- 648  
nie Webber, Trevor Cohn, Yulan He, and Yang Liu, ed- 649  
itors, *Proceedings of the 2020 Conference on Empirical 650*  
*Methods in Natural Language Processing, EMNLP 2020,* 651  
*Online, November 16-20, 2020, pages 7000–7011.* Asso- 652  
ciation for Computational Linguistics, 2020. 653
- [Chowdhury *et al.*, 2022] Somnath Basu Roy Chowdhury, 654  
Chao Zhao, and Snigdha Chaturvedi. Unsupervised ex- 655  
tractive opinion summarization using sparse coding. In 656  
Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, 657  
editors, *Proceedings of the 60th Annual Meeting of 658*  
*the Association for Computational Linguistics (Volume 1: 659*  
*Long Papers), ACL 2022, Dublin, Ireland, May 22-27,* 660  
*2022, pages 1209–1225.* Association for Computational 661  
*Linguistics, 2022.* 662
- [Chu and Liu, 2019] Eric Chu and Peter J. Liu. Meansum: 663  
A neural model for unsupervised multi-document abstrac- 664  
tive summarization. In Kamalika Chaudhuri and Ruslan 665  
Salakhutdinov, editors, *Proceedings of the 36th Interna- 666*  
*tional Conference on Machine Learning, ICML 2019, 9- 667*  
*15 June 2019, Long Beach, California, USA, volume 97 of 668*  
*Proceedings of Machine Learning Research, pages 1223– 669*  
*1232.* PMLR, 2019. 670
- [Deng *et al.*, 2020] Zhongfen Deng, Hao Peng, Congying 671  
Xia, Jianxin Li, Lifang He, and Philip S. Yu. Hierarchi- 672  
cal bi-directional self-attention networks for paper review 673  
rating recommendation. In Donia Scott, Núria Bel, and 674  
Chengqing Zong, editors, *Proceedings of the 28th Interna- 675*  
*tional Conference on Computational Linguistics, COLING 676*  
*2020, Barcelona, Spain (Online), December 8-13, 2020,* 677  
*pages 6302–6314.* International Committee on Computa- 678  
*tional Linguistics, 2020.* 679
- [Dycke *et al.*, 2023] Nils Dycke, Ilia Kuznetsov, and Iryna 680  
Gurevych. NLPeer: A unified resource for the compu- 681  
tational study of peer review. In Anna Rogers, Jordan 682  
Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings 683*  
*of the 61st Annual Meeting of the Association for Computa- 684*  
*tional Linguistics (Volume 1: Long Papers), pages 5049– 685*  
*5073, Toronto, Canada, July 2023.* Association for Com- 686  
*putational Linguistics.* 687
- [Elsahar *et al.*, 2021] Hady Elsahar, Maximin Coavoux, Jos 688  
Rozen, and Matthias Gallé. Self-supervised and controlled 689  
multi-document opinion summarization. In *Proceedings of 690*  
*the 16th Conference of the European Chapter of the As- 691*  
*sociation for Computational Linguistics: Main Volume,* 692  
*pages 1646–1662, Online, April 2021.* Association for 693  
*Computational Linguistics.* 694
- [Erkan and Radev, 2004] Günes Erkan and Dragomir R. 695  
Radev. Lexrank: Graph-based lexical centrality as salience 696  
in text summarization. *J. Artif. Intell. Res.*, 22:457–479, 697  
2004. 698
- [Friedl *et al.*, 2021] Korbinian Friedl, Georgios Rizos, Lukas 699  
Stappen, Madina Hasan, Lucia Specia, Thomas Hain, and 700  
Björn W. Schuller. Uncertainty aware review hallucina- 701  
tion for science article classification. In Chengqing Zong, 702

- 703 Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 5004–5009. Association for Computational Linguistics, 2021.
- 708 [Fu *et al.*, 2023] Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. Gptscore: Evaluate as you desire. *CoRR*, abs/2302.04166, 2023.
- 711 [Gao *et al.*, 2023] Mingqi Gao, Jie Ruan, Renliang Sun, Xunjian Yin, Shiping Yang, and Xiaojun Wan. Human-like summarization evaluation with chatgpt. *CoRR*, abs/2304.02554, 2023.
- 715 [Goyal *et al.*, 2022] Tanya Goyal, Junyi Jessy Li, and Greg Durrett. News summarization and evaluation in the era of GPT-3. *CoRR*, abs/2209.12356, 2022.
- 718 [Gu *et al.*, 2022] Nianlong Gu, Elliott Ash, and Richard Hahnloser. MemSum: Extractive summarization of long documents using multi-step episodic Markov decision processes. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6507–6522, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- 725 [Hu and Liu, 2006] Minqing Hu and Bing Liu. Opinion extraction and summarization on the web. In *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*, pages 1621–1624. AAAI Press, 2006.
- 732 [Iso *et al.*, 2021] Hayate Iso, Xiaolan Wang, Yoshihiko Suhara, Stefanos Angelidis, and Wang-Chiew Tan. Convex aggregation for opinion summarization. *CoRR*, abs/2104.01371, 2021.
- 736 [Isonuma *et al.*, 2021] Masaru Isonuma, Junichiro Mori, Danushka Bollegala, and Ichiro Sakata. Unsupervised abstractive opinion summarization by generating sentences with tree-structured topic guidance. *Transactions of the Association for Computational Linguistics*, 9:945–961, 2021.
- 742 [Kryscinski *et al.*, 2019] Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. Evaluating the factual consistency of abstractive text summarization. *CoRR*, abs/1910.12840, 2019.
- 746 [Kumar *et al.*, 2021] Asheesh Kumar, Tirthankar Ghosal, and Asif Ekbal. A deep neural architecture for decision-aware meta-review generation. In J. Stephen Downie, Dana McKay, Hussein Suleman, David M. Nichols, and Faryaneh Poursardar, editors, *ACM/IEEE Joint Conference on Digital Libraries, JCDL 2021, Champaign, IL, USA, September 27-30, 2021*, pages 222–225. IEEE, 2021.
- 753 [Laban *et al.*, 2021] Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. Summac: Revisiting nli-based models for inconsistency detection in summarization. *CoRR*, abs/2111.09525, 2021.
- [Li *et al.*, 2017] Piji Li, Zihao Wang, Zhaochun Ren, Lidong Bing, and Wai Lam. Neural rating regression with abstractive tips generation for recommendation. In Noriko Kando, Tetsuya Sakai, Hideo Joho, Hang Li, Arjen P. de Vries, and Ryen W. White, editors, *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Shinjuku, Tokyo, Japan, August 7-11, 2017*, pages 345–354. ACM, 2017.
- [Lin, 2004] Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics.
- [Liu *et al.*, 2023] Yang Liu, Dan Iter, Yichong Xu, Shuo-hang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: NLG evaluation using GPT-4 with better human alignment. *CoRR*, abs/2303.16634, 2023.
- [Madaan *et al.*, 2023] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Sean Welleck, Bodhisattwa Prasad Majumder, Shashank Gupta, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback, 2023.
- [OpenAI, 2021] OpenAI. Chatgpt-3.5-turbo. <https://openai.com/research/>, 2021. Accessed: 2023-05-20.
- [Shen *et al.*, 2022] Chenhui Shen, Liying Cheng, Ran Zhou, Lidong Bing, Yang You, and Luo Si. Mred: A meta-review dataset for structure-controllable text generation. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2521–2535. Association for Computational Linguistics, 2022.
- [Shen *et al.*, 2023] Ming Shen, Jie Ma, Shuai Wang, Yogarshi Vyas, Kalpit Dixit, Miguel Ballesteros, and Yasmine Benajiba. Simple yet effective synthetic dataset construction for unsupervised opinion summarization. *CoRR*, abs/2303.11660, 2023.
- [Suhara *et al.*, 2020] Yoshihiko Suhara, Xiaolan Wang, Stefanos Angelidis, and Wang-Chiew Tan. Opiniondigest: A simple framework for opinion summarization. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5789–5798. Association for Computational Linguistics, 2020.
- [Wang and Ling, 2016] Lu Wang and Wang Ling. Neural network-based abstract generation for opinions and arguments. In Kevin Knight, Ani Nenkova, and Owen Rambow, editors, *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 47–57. The Association for Computational Linguistics, 2016.
- [Wang *et al.*, 2020] Qingyun Wang, Qi Zeng, Lifu Huang, Kevin Knight, Heng Ji, and Nazneen Fatema Rajani. Reviewrobot: Explainable paper review generation based on

813 knowledge synthesis. In Brian Davis, Yvette Graham,  
814 John D. Kelleher, and Yaji Sripada, editors, *Proceedings*  
815 *of the 13th International Conference on Natural Language*  
816 *Generation, INLG 2020, Dublin, Ireland, December 15-*  
817 *18, 2020*, pages 384–397. Association for Computational  
818 Linguistics, 2020.

819 [Xiao and Carenini, 2020] Wen Xiao and Giuseppe  
820 Carenini. Systematically exploring redundancy re-  
821 duction in summarizing long documents. In Kam-Fai  
822 Wong, Kevin Knight, and Hua Wu, editors, *Proceedings*  
823 *of the 1st Conference of the Asia-Pacific Chapter of the*  
824 *Association for Computational Linguistics and the 10th*  
825 *International Joint Conference on Natural Language Pro-*  
826 *cessing, ACL/IJCNLP 2020, Suzhou, China, December*  
827 *4-7, 2020*, pages 516–528. Association for Computational  
828 Linguistics, 2020.

829 [Zhang *et al.*, 2020] Tianyi Zhang, Varsha Kishore, Felix  
830 Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore:  
831 Evaluating text generation with BERT. In *8th Inter-*  
832 *national Conference on Learning Representations, ICLR*  
833 *2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenRe-  
834 view.net, 2020.

835 [Zhao *et al.*, 2022] Wei Zhao, Michael Strube, and Steffen  
836 Eger. Discoscore: Evaluating text generation with BERT  
837 and discourse coherence. *CoRR*, abs/2201.11176, 2022.