

# Good Idea or Not, Representation of LLM Could Tell

Anonymous ACL submission

## Abstract

In the ever-expanding landscape of academic research, the proliferation of ideas presents a significant challenge for researchers: discerning valuable ideas from the less impactful ones. The ability to efficiently evaluate the potential of these ideas is crucial for the advancement of science and paper review. In this work, we focus on idea assessment, which aims to leverage the knowledge of large language models to assess the merit of scientific ideas. First, we investigate existing text evaluation research and define the problem of quantitative evaluation of ideas. Second, we curate and release a benchmark dataset from nearly four thousand manuscript papers with full texts, meticulously designed to train and evaluate the performance of different approaches to this task. Third, we establish a framework for quantifying the value of ideas by employing representations in a specific layer of large language models. Experimental results show that the scores predicted by our method are relatively consistent with those of humans. Our findings suggest that the representations of large language models hold more potential in quantifying the value of ideas than their generative outputs, demonstrating a promising avenue for automating the idea assessment process.

## 1 Introduction

The rapid pace of scientific research in various disciplines has given rise to an overwhelming number of academic papers (Tabah, 1999; Bornmann and Mutz, 2015; Xu et al., 2023). Typically, ideas are commonly conveyed through these papers, which reviewers must carefully scrutinize to grasp the ideas authors present. However, amidst the vast volume of paper submissions, the review process becomes slow, labor-intensive, and less precise (Checco et al., 2021), making it a challenge to identify valuable ideas. Fortunately, with the advent of large language models (LLMs), we are

presented with an unprecedented opportunity to revolutionize how we assess the merit of scientific ideas, and this work explores their use as a knowledge tool for such evaluation. In order to make idea more concrete, **we take unpublished manuscripts or the latest papers as the research object.**

While the generative capabilities of LLMs have been widely recognized, their potential as idea (or paper) evaluative instruments has remained relatively underexplored. Recent studies (Yuan et al., 2022; Liang et al., 2023; Liu and Shah, 2023; Agarwal et al., 2024) have begun to harness LLMs for the automatic generation of paper reviews, aiming for outputs that are both informative and emulate human feedback. These efforts primarily utilize the text generation capability of LLM, which are highly dependent on the scale of model parameters. For instance, Liu and Shah (2023) proves that GPT-4 surpasses other open-source LLMs, such as LLaMA (Touvron et al., 2023), in generating reviews. Meanwhile, crafting intricate prompts containing specific commands and inquiries is essential for LLMs to produce meaningful reviews. Nonetheless, LLM-generated reviews can still reflect models’ subjective biases and occasionally produce hallucinated contents (Zhang et al., 2023; Manakul et al., 2023). There is currently no research that quantitatively evaluate ideas in an objective way.

According to Geva et al. (2021); Zou et al. (2023), the representations of different layers in LLM contain different semantic information, and in some tasks, the performance of the last layer is not the best. Based on this point, we suppose that the representations of LLMs encapsulate a detailed and nuanced comprehension of text, which can be leveraged to construct a systematic and objective framework for the assessment of ideas. Our research thus focuses on the **quantitative evaluation of ideas** through LLMs, an approach we argue is more objective than generative techniques. It is worth not-

ing that LLMs’ generative method is not inherently adept at processing numerical data (digits), which is also why we turn to their representations as a means to quantify the value of ideas. We delve into the latent knowledge embedded within the representations of LLMs for this purpose. Specifically, we first define the problem of quantitative idea evaluation, and construct a benchmark dataset comprised of nearly 4k scientific manuscripts with full texts in the discipline of computer science. This dataset serves as the bedrock for training and testing various LLM-based approaches to idea assessment. In the subsequent phase of our research, we develop a framework for quantifying the value of ideas. This framework leverages the representations produced by LLMs, which encode semantic features of the text in a high-dimensional space. Finally, by training a downstream evaluator using these representations, our framework can identify patterns and signals within these representations that align with the value of scientific ideas as determined by the distribution of expert consensus.

We perform extensive experiments on our benchmark dataset. The results reveal that the representations generated by LLMs are inherently indicative of the potential value of scientific ideas, especially using the representations in the middle and rear layers, which demonstrate a high degree of consistency with human judgements. Meanwhile, we also found that high consistency can be achieved with only a small amount of data, thanks to the pretrained knowledge in LLMs. The contributions of our paper are summarized as follows:

- We conduct a thorough investigation of current methods for paper (idea) assessment and pioneer the study of quantitative evaluation of scientific ideas.
- We meticulously curate a new benchmark dataset, complete with human-assigned scores on overall quality, novelty, and correctness. We are making it publicly available. This dataset is comprised of the full texts of 3,795 papers in PDF format.
- We propose a new framework that quantifies the value of ideas by utilizing the deep representations produced by LLMs.
- We conduct extensive experiments to verify the performance of our method. The results indicate that utilizing the representations from

LLMs aligns more closely with human judgments than using generative textual outputs of LLMs when assessing the quality of ideas.

## 2 Related Work

Yuan et al. (2022) have explored the use of various NLP techniques to produce decisive and comprehensive reviews of academic papers from multiple perspectives. The work of ReviewerGPT (Liu and Shah, 2023) studies the application of LLMs as review assistants, focusing on tasks such as error identification, checklist verification, and comparative paper analysis. Another innovative approach by researchers (Liang et al., 2023) involves an automated system that employs GPT-4 to create review comments and suggestions for revisions, which are then benchmarked against feedback from human reviewers. What’s more, LitLLM (Agarwal et al., 2024) equips Retrieval Augmented Generation (RAG) module to address the hallucination problem in review generation.

In addition to the methods mentioned above, some researches rely on external data to assess the quality of a paper. For instance, Thelwall et al. (2023) have developed a framework that predicts article quality scores using a range of bibliometric and metadata indicators, including citation counts, journal impact factors, and institutional rankings. Similarly, KQI (Wang et al., 2023) leverages the structure of citation networks to quantify the knowledge contribution of a paper. This kind of approaches, however, pertains to post-publication evaluation. Unlike our approach, which is based solely on the text of the paper itself, it does not require information about the paper’s acceptance or publication status.

## 3 Idea Assessment

In this work, we focus on quantitative evaluation of ideas. Let  $D = \{d_1, d_2, \dots, d_n\}$  be a dataset consisting of  $n$  scientific manuscripts (papers), each representing a distinct scientific idea. The direct scoring of an idea involves mapping each manuscript  $d_i$  to a quantitative score  $s_i$  based on a predefined criterion  $c \in C$ . The criterion set  $C$  serves as the basis for the assessment and is essential for guiding the evaluation process. It can encompass various aspects of potential impacts of an idea, such as novelty, correctness, excitement, soundness, or alignment with current research trends.

We first define a function  $A_c : D \rightarrow \mathbb{R}$  such that for any manuscript  $d_i \in D$  and criterion  $c \in C$ ,  $A_c(d_i)$  produces a scalar value  $s_i \in \mathbb{R}$  which quantifies the value of the idea  $d_i$  with respect to  $c$ . The quantitative evaluation of idea  $d_i$  with respect to criterion  $c$  can be expressed as:

$$s_i = A_c(d_i), \quad (1)$$

where the function  $A_c$  is the evaluator that assesses the idea based on the text of the manuscript  $d_i$  and the specified criterion  $c$ .

In our work, we utilize the representations of LLM to quantify the value of an idea. Let  $M$  be an LLM that transforms textual data into a high-dimensional representation space. We define an encoding function  $Rep : D \rightarrow \mathbb{R}^{l \times m}$  such that for any manuscript  $d_i \in D$ , there is a hidden representation  $Rep(d_i) \in \mathbb{R}^{l \times m}$ , where  $l$  is the number of tokens in  $d_i$  and  $m$  is the dimension of the representation space in LLM  $M$ . Now, we revise the evaluator function  $A_c : \mathbb{R}^{l \times m} \rightarrow \mathbb{R}$  that maps the representation to a scalar value  $s_i$ :

$$s_i = A_c(Rep(d_i)). \quad (2)$$

The evaluator,  $A_c$ , is designed to be flexible and adaptable to different criteria and can be trained using annotated data that provide ground truth measures of the ideas’ impact with respect to the chosen criterion. To the best of our knowledge, we are the first to quantify the value of an idea.

## 4 Dataset

To ensure that the idea evaluator is well-calibrated, the benchmark idea dataset  $D$  should be representative of the scientific community and contain cutting-edge knowledge in academia. To this end, we have compiled a collection of 3,795 manuscripts that are available in PDF format from the International Conference on Learning Representations (ICLR) 2023. For the extraction of full texts from these PDFs, we employed GROBID (Lopez, 2009), a sophisticated tool for parsing academic PDF documents.

Additionally, the metadata of these papers includes comprehensive evaluation criteria from official reviewers, encompassing scores for overall quality, correctness, technical and empirical novelty, providing a rich ground truth for training and validation. It is possible that an idea is interesting but the paper score of a criterion such as correctness is low because there are flaws in the experiments.

Therefore, in our work, we mainly investigate the criterion *overall quality* and take it as the overall score of an idea.

	ICLR23-low-std	ICLR23-all
# paper	1901	3795
overall quality	5.52 ± 0.61	5.41 ± 1.06
correctness	3.09 ± 0.44	3.09 ± 0.49
technical novelty	2.59 ± 0.43	2.59 ± 0.48
empirical novelty	2.56 ± 0.41	2.56 ± 0.47

Table 1: Statistics of benchmarks.

Considering the different consistencies of human-rated data, which can have impacts on different evaluation models, we choose papers with highly consistent human-rated scores from the original dataset *ICLR23-all* and construct dataset *ICLR23-low-std*, where the standard deviation (std) of overall quality scores for each paper is relatively lower. The statistics are listed in Table 1.

## 5 Methodology

The purpose of our method is to train an evaluator  $A_c$  to score ideas, which consists of four steps: consistency sorting, layer selection, token selection, and evaluator training. Figure 1 shows the pipeline of quantitative evaluation of ideas using the representation of LLMs. It should be highlighted that the steps of layer and token selection only exist in training process, which are determined during the inference process.

### 5.1 Consistency Sorting

In our scenario, we anticipate that models can learn the rating standard from human-rated data. Specifically, the human-assigned scores for each paper in the training set should exhibit a high level of consistency; that is, the more uniform the scores for each paper are (reflected by lower variance), the more suitable the data is for model training. Therefore, our method employs a consistency-based sorting mechanism to construct the training and testing sets. We commence by ordering the papers according to the variance in human scores for a given criterion  $c$ . Subsequently, based on a predetermined threshold for training set partitioning, papers that demonstrate high consistency (low variance) are allocated to the training set, while the remainder are designated for the testing set. This mechanism facilitates a more straightforward learning process for models to grasp the human rating standards.

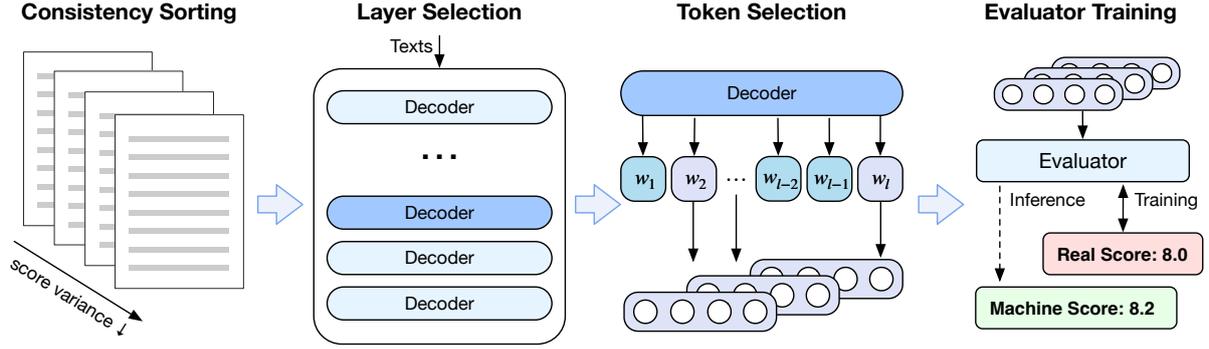


Figure 1: Pipeline of quantitative evaluation of ideas using the representations of LLMs.

Conversely, a high degree of variance in human-assigned scores suggests that the paper (or idea) is controversial, rendering the learning of standards from such data as potentially futile.

## 5.2 Layer Selection

As claimed by Geva et al. (2021), lower layers of LLMs tend to capture shallow data patterns, while upper layers contain more semantic knowledge. This hierarchical processing of information within LLMs suggests that the utility of representations may vary across layers. Further to this, RepE (Zou et al., 2023) explores the relationship between layer depth and performance in utility estimation tasks, finding that middle-layer representations often yield the highest accuracy.

Inspired by these findings, our approach involves identifying the optimal layer within an LLM that provides the most effective representations for constructing an accurate evaluator. We hypothesize that a specific intermediate layer may offer the ideal balance between capturing fundamental linguistic features and the nuanced semantic understanding necessary for assessing the quality of scientific ideas. Our experiments are thus tailored to pinpoint this layer by analyzing the performance of given data across all layers. Then, we leverage its representations to enhance the performance of our idea evaluation framework.

## 5.3 Token Selection

Considering that a manuscript  $d_i$  is composed of  $l$  sequential tokens, the semantic information of these token representations varies significantly. Due to the fact that most LLMs are auto-regressive models, the last token aggregates the attention information of all previous tokens (Zou et al., 2023). With a slight abuse of notation, by default, we use

the last token representation  $Rep(d_i, -1) \in \mathbb{R}^m$  to symbolize the entirety of the manuscript  $d_i$ .

Nevertheless, when dealing with lengthy input texts, such as full-text manuscript  $d_i$ , there are two issues with the default approach. For one thing, memory optimization mechanism such as vLLM (Kwon et al., 2023) should be adopted to prevent GPU from running out of memory. For another thing, the representation of the last token may become diluted or overly abstracted owing to the extensive accumulation of attention, potentially leading to a loss of specific semantic details pertinent to the overall idea assessment. To address these issues, we explore alternative strategies for token selection that aim to maintain the richness of semantic information while ensuring computational feasibility.

We consider a manuscript  $d_i$  to be a composition of distinct sections. We select the last token representations from each section, and concatenate these to form a composite representation. The approach allows us to capture the essence of each section. Formally, if a manuscript  $d_i$  is divided into  $r$  sections, and  $Rep(d_{i,j}, -1)$  represents the last token of the  $j^{th}$  section, then the combined representation  $Rep(d_i)$  is given by:

$$Rep(d_i) = \bigoplus_{j=1}^r Rep(d_{i,j}, -1), \quad (3)$$

where  $\bigoplus$  denotes the concatenation operation, and  $Rep(d_i)$  is in  $\mathbb{R}^{r \times m}$ . Similarly, we can take into account the full length of the manuscript and divide it into equidistant segments based on a predefined length to obtain  $Rep(d_i)$ . By experimenting with these strategies, we aim to refine our approach to token selection and optimize the representation of manuscript for more accurate idea assessment.

## 5.4 Evaluator Training

In this part, we use the pre-processed  $Rep(d_i)$  to train an idea evaluator  $A_c$ . Let  $s_i$  be the average score given by humans for manuscript  $d_i$ , reflecting its overall quality according to the criterion  $c$ . The average score serves as the ground truth in our training process. The evaluator  $A_c$  is instantiated as an Multilayer Perceptron (MLP) with one hidden layer. The MLP is tasked with learning the mapping from the representation space to the scalar scores, which takes as input the representation  $Rep(d_i)$  for each manuscript  $d_i$  and outputs a predicted score  $\hat{s}_i$ . To optimize all parameters of the MLP, we employ the Mean Squared Error (MSE) loss function:

$$\mathcal{L} = \frac{1}{n} \sum_{i=1}^n (\hat{s}_i - s_i)^2. \quad (4)$$

By minimizing  $\mathcal{L}$ , the MLP learns to approximate the human-assigned scores as closely as possible. Through this training process, we aim to calibrate the evaluator  $A_c$  such that it can reliably predict the value of new or unseen ideas based on their textual representations.

## 6 Experiments

This section presents a series of experiments to verify the performance of LLMs in the task of quantitative evaluation of ideas. Our main focus is on the mean value of the criteria **overall quality**, which is used as the training objective for the idea evaluator. Through our released benchmark and the experimental methodology, we answer the following four research questions (RQs):

- **RQ1:** To what extent do the representations from LLMs correlate with human judgements in the evaluation of scientific ideas? Additionally, is the LLM generation method suitable for this task?
- **RQ2:** What is the impact of choosing different layers and tokens for LLM representations on the performance of idea evaluation?
- **RQ3:** How significantly does the consistency of human judgements influence the performance of LLM representations in this context?
- **RQ4:** How does the size of training set impact the correlation between  $A_c$  evaluations and human judgments in idea assessment?

## 6.1 Baselines

We categorize the baselines into three distinct groups: LLM Generation, LLM Representation, and Human Evaluation. The first category involves LLMs generating numerical scores in response to textual descriptions of ideas. This category includes models such as GPT-3.5-turbo, LLaMa-2-7b-base (Touvron et al., 2023), and Baichuan-2-7b-base (Yang et al., 2023), which are fine-tuned using techniques like LoRA (Hu et al., 2021) or with full parameter updates. The prompts we choose are presented in Appendix A. In the LLM Representation category, we evaluate models like BERT (Kenton and Toutanova, 2019), SciBERT (Beltagy et al., 2019), and RePE (Zou et al., 2023). For BERT and SciBERT, we also apply our proposed framework to quantify the value of ideas, with the primary distinction being in the token selection strategy. Specifically, we used the  $[CLS]$  token as the representation of an idea, and if the length of a section exceeds 512 tokens, we will divide it into equidistant subsections to apply the token selection strategy for BERT-like models. Moreover, we also analyze the performance of human evaluators through randomly selecting one score from the human-rated list against other scores.

## 6.2 Training Settings and Evaluation Details

In our implementation, our method employs LLaMA-2-7b-base as the foundational model. In order to make our experiments more solid and validate our framework is model-agnostic, we also use Baichuan-2-7b-base as the base model, the results of which are provided in Appendix C. We use the grid search to find appropriate sets of hyperparameters for baselines and our proposed method. For the configuration of the MLP evaluator, we choose a batch size of 32, a hidden layer dimension of 1024, a learning rate of 0.001, a dropout rate of 0.2, and employ the Adam optimizer. We limit the training to 20 epochs. More detailed settings are documented in Appendix B. Each experiment is executed three times with random initializations, and the mean results are reported. We use the results of the training set for model selection. All experiments are first conducted to evaluate the efficacy of our framework using the **abstracts of papers** for all research questions. We also explore the effects of using the **full texts of papers** as the training inputs for the token selection in RQ2.

To gauge the alignment of scores generated by

Method	ICLR23-low-std				ICLR23-all			
	train: 5%		train: 30%		train: 5%		train: 30%	
	corr	layer	corr	layer	corr	layer	corr	layer
<b>LLM Generation</b>								
LLaMA-2-Full-SFT	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
LLaMA-2-LoRA-SFT	-0.0513	N/A	0.0820	N/A	0.0634	N/A	0.0692	N/A
Baichuan-2-LoRA-SFT	0.1391	N/A	0.2054	N/A	0.1413	N/A	0.1867	N/A
GPT-3.5-turbo	0.1290	N/A	0.1375	N/A	0.0874	N/A	0.0719	N/A
<b>LLM Representation</b>								
BERT	0.1986	-3	0.2515	-4	0.1907	-1	0.2326	-1
SciBERT	0.2677	-3	0.3314	-3	0.2447	-2	0.2584	-3
RePE-with-prompt	0.0820	-31	0.0993	-31	0.0738	-21	0.0880	-31
RePE-no-prompt	0.1020	-1	0.0605	-31	0.0635	-2	0.0508	-2
<b>Ours</b>	<b>0.3441</b>	-20	<b>0.3880</b>	-9	<b>0.2783</b>	-1	<b>0.3366</b>	-4
<b>Human Evaluation</b>								
Human	0.8175	N/A	0.7648	N/A	0.4174	N/A	0.3290	N/A

Table 2: Spearman correlations with humans of different methods on ICLR23 datasets. N/A in *corr* column means its corresponding *pvalue* > 0.05. There is no need for LLM Generation baselines to select layers. The human performance is evaluated by randomly selecting one score from the human-rated list against other reviews.

various methods with human-assigned scores, we report a widely-used metric called *Spearman Correlation* (Spearman, 1961). The correlation *corr* with human is defined as:

$$corr = \rho([s_1, s_2, \dots, s_n], [\hat{s}_1, \hat{s}_2, \dots, \hat{s}_n]), \quad (5)$$

where  $\rho$  is the Spearman Correlation function,  $s_i$  is the average score given by humans, and  $\hat{s}_i$  is the  $A_c$  predicted score for  $d_i$  in the testing set. Since Spearman correlation is invariant under affine transformations, we also provide score distribution and the absolute error between human-rated scores and our LLM representation scores in Section 6.4 and Appendix E.

### 6.3 Comparative Experiments (RQ1)

According to the principle of consistency sorting in Section 5.1, we construct training sets using the top 5% and top 30% ratios from ICLR23-low-std and ICLR23-all datasets respectively to preliminarily exclude the influence of dataset proportion on the conclusion, and take the rest of each dataset as the testing set. Table 2 shows the Spearman correlations with humans of different methods on these two datasets. We also provide indexes for the layers with the highest correlation.

It can be observed that our proposed method achieves the best performance in all settings, where the performance on ICLR23-all is **at most 30%** higher than the second best method SciBERT. As expected, the correlation of ICLR23-low-std

among human scores is close to 1, which is attributed to our data partitioning strategy. It should be noted that the correlation of our method on ICLR23-all dataset exceed the result of humans, when the training ratio is 30%, proving the feasibility of our method and its potential ability to be applied to real-world review scenarios. Moreover, in terms of different layers' performance, the middle and back layers of most models may achieve better results.

For the LLM Generation baselines, the fine-tuned LLaMA-2 is worse than Baichuan-2, especially for the LLaMA-2-Full-SFT, fine-tuned with full parameters, lacking the capability of effective evaluation since its *pvalue* > 0.05. Due to the inability of GPT-3.5 being fine-tuned, we adopt the zero-shot setting, which is only for sketchy reference. We also try k-shot setting for GPT-3.5, but it only generates the most frequent scores from the given examples. Overall, the LLM Generation methods are not competent for the quantitative evaluation of ideas. By analyzing the generated results, we believe there are two possible reasons. One is that the amount of training data is relatively small, and models are prone to overfitting. Apart from that, LLM generation is not sensitive to digital numbers, and the semantic knowledge is hidden in its representations, which should be guided through appropriate means.

Furthermore, our experiment studies the degree of consistency between the predicted score and that

assigned by the human reviewer whose score most closely match the predicted one. As depicted in Table 3, the ICLR23-all dataset exhibits a higher consistency with the closest human-rated scores compared to the ICLR23-low-std dataset. This suggests that, despite the higher variance in human scores of the ICLR23-all dataset, our proposed method is adept at mirroring the evaluation of the most similar human reviewer.

Dataset	Method	Training Ratio	
		5%	30%
ICLR23-low-std	SciBERT	0.4299	0.5093
	<b>Ours</b>	<b>0.5469</b>	<b>0.5605</b>
ICLR23-all	SciBERT	0.6381	0.5112
	<b>Ours</b>	<b>0.6462</b>	<b>0.5617</b>

Table 3: Spearman correlations with the closest human-rated score.

#### 6.4 Score Distribution (RQ1)

We also examine the difference (absolute error) between human-rated scores and the predicted scores on ICLR23-low-std dataset. The results are shown in the pie chart of Figure 2. We can see that **86.8%** paper scores generated by our method are close to the human-rated scores, where the differences between them are lower than 2. Additionally, the distributions are shown in the right part of Figure 2. The distribution of scores predicted by our idea evaluator is normal distribution as expected while the human reviewers tend to give more higher or lower scores. More analysis can be found in Appendix E.

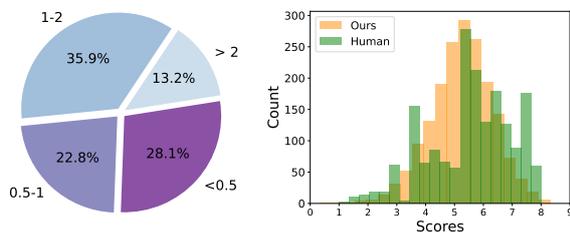


Figure 2: The difference (absolute error) between human-rated scores and the predicted scores (left subfigure). The distributions of human-rated scores and the predicted scores (right subfigure).

#### 6.5 Influence of Layer Selection (RQ2)

We analyzed the representational efficacy across various layers of LLM and SciBERT. As illustrated in Figure 3, it is evident that for both LLM and

SciBERT, the representations from the middle to later layers outperform those from other layers. Obviously, the very last layers do not typically yield the best performance. This may be attributed to the specific semantic information encapsulated in different layers. The last layer is inherently born to facilitate generation tasks, rather than tasks that require more discriminative capabilities, like classification or regression.

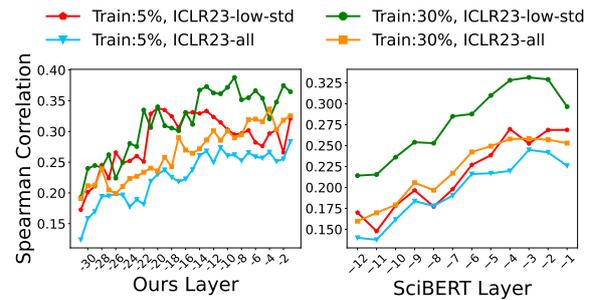


Figure 3: Spearman correlations of different layers in LLM and SciBERT.

Given the nuanced role of layer-specific representations in the context of assessing the merit of scientific ideas, we propose a layer selection of representations for the task at hand. Specifically, we advocate for the utilization of representations from the layers situated in the last one-third of the model’s depth. Such choice is informed by the empirical evidence suggesting that these layers strike a balance between retaining rich semantic content and providing the necessary abstraction for discriminative tasks.

#### 6.6 Influence of Token Selection (RQ2)

In this part, we use the ICLR23-all dataset to investigate the influence of token selection. In terms of the paper abstract inputs, we first test the correlation results of using the last token. Subsequently, we expand our scope to include both the middle and last tokens (middle + last token) of the input text. The findings, as presented in Table 4, indicate that solely relying on the last token yields superior results compared to combining it with the middle token. The latter approach appears to introduce a surplus of redundant information that may hinder the downstream performance of evaluator.

When dealing with the full text of papers, we implement two token selection strategies outlined in Section 5.3: the amalgamation of last tokens from equidistant segments and the aggregation of last tokens from distinct paper sections. The ex-

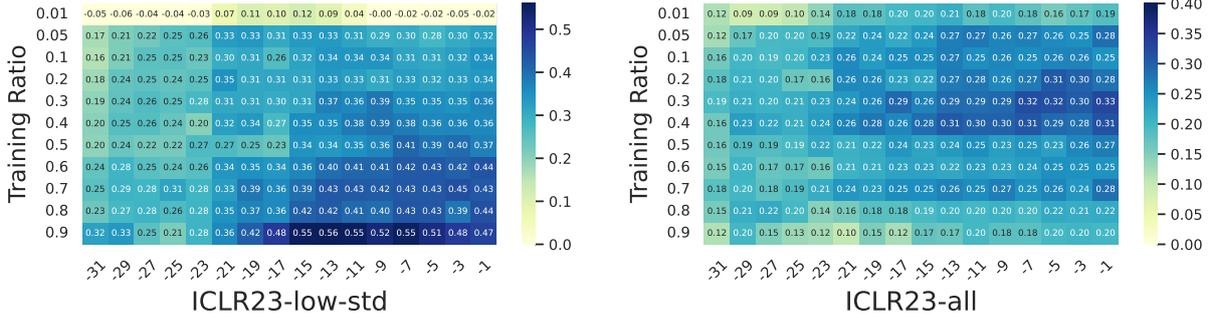


Figure 4: Spearman correlations varying with layers under different training ratios on ICLR23 datasets.

Token Selection Strategy	Training Ratio	
	5%	30%
<b>Abstract</b>		
last token	0.2783	0.3366
middle + last token	0.2162	0.2772
<b>Full-text</b>		
segment tokens	0.1306	0.2597
section tokens	0.3258	0.3821

Table 4: Spearman correlations of different token selection strategies.

560 experimental outcomes suggest that merging last tokens  
561 from all segments does not effectively capture  
562 the necessary semantic content. This is largely  
563 due to the fact that many segments are incomplete  
564 sentences, arbitrarily truncated from the original  
565 texts. On the other hand, the section-based strategy  
566 successfully compiles comprehensive information  
567 from each complete section. Overall, we recom-  
568 mend using the representations of the last tokens or  
569 section-based strategy to train the idea evaluator.

### 570 6.7 Analysis of Training Set (RQ3 & RQ4)

571 Figure 4 shows the Spearman correlations vary-  
572 ing with layers under different training ratios on  
573 ICLR23 datasets. Due to space limitations, we  
574 only present the results of odd numbered layers.  
575 See Appendix C for more results. The manuscript  
576 papers are sorted in ascending order of variance  
577 in their corresponding human-rated scores to parti-  
578 tion datasets according to the consistency sorting  
579 in Section 5.1.

580 For ICLR23-low-std dataset, the human-rated  
581 scores in the dataset are highly consistent, and it  
582 is observed that the Spearman correlation tends to  
583 improve in tandem with increases in the training  
584 set size. Notably, when the proportion of data used  
585 for training surpasses the 50% threshold, the cor-  
586 relation between the scores predicted by our idea

587 evaluator and those assigned by human experts be-  
588 comes moderate, exceeding 0.4. Furthermore, our  
589 analysis reveals that even a relatively small subset  
590 of the data (with a training ratio of 5%) is capable  
591 of yielding positive performance.

592 As to ICLR23-all dataset, the outcomes indi-  
593 cate that an increase in the volume of training data  
594 does not necessarily correspond to a higher align-  
595 ment with human evaluations in the testing set.  
596 The phenomenon can be attributed to the dimi-  
597 nishing consistency of human scores as the dataset  
598 expands; that is, the variance in human-assigned  
599 scores grows with the size of the dataset. It be-  
600 comes evident that while a larger training set gener-  
601 ally provides more information, it also introduces  
602 a greater diversity of human judgment, which may  
603 not always be conducive to improving the ability  
604 of evaluator to mimic human scoring behavior.

## 605 7 Conclusion and Future Work

606 The study focuses on the quantitative evaluation of  
607 scientific ideas. We have reviewed existing method-  
608 ologies for paper and idea evaluation and have bro-  
609 ken new ground by focusing on the quantitative  
610 aspect of idea evaluation. Specifically, we first in-  
611 troduce a comprehensive benchmark dataset, acces-  
612 sible to the research community. Then, we develop  
613 a new framework that leverages the token repre-  
614 sentations of specific layers in LLM to quantify  
615 the value of ideas. Through rigorous experiments,  
616 we demonstrate that LLM representations correlate  
617 more strongly with human judgments compared to  
618 generative text outputs. Additionally, in our bench-  
619 mark, the predicted scores of more than 80% papers  
620 are close to human-rated scores. In the future, we  
621 will broaden the scope of our research to encom-  
622 pass diverse disciplines with balanced data ratios,  
623 including the exact and social sciences, to further  
624 validate and refine our evaluative framework.

## 625 **Limitations**

### 626 **Discipline**

627 The scope of our research is confined to the field of  
628 computer science, which may restrict the broader  
629 applicability of our framework. The generalization  
630 performance of our model across different scientific  
631 disciplines remains an open question. Future  
632 research endeavors should aim to adapt and vali-  
633 date the framework in diverse fields, ranging from  
634 the exact sciences to the humanities.

### 635 **Criterion**

636 Our experiments have primarily focused on the  
637 overall quality score of manuscript papers, which  
638 is a composite yet somewhat abstract. Important as-  
639 pects such as the correctness of the presented work  
640 and its novelty are equally critical in determining  
641 a paper’s impact and significance. In forthcoming  
642 studies, we plan to dissect these individual crite-  
643 ria, developing a more granular approach to idea  
644 evaluation.

### 645 **Model Scale**

646 The impact of model scale on performance is an  
647 aspect that has not been extensively explored in  
648 our research. The performance of LLMs is often  
649 closely tied to the number of parameters they con-  
650 tain; thus, models with different sizes may yield dif-  
651 ferent results in the task of idea evaluation. Larger  
652 models may have the capacity to encode more nu-  
653 anced representations of text, potentially leading to  
654 more accurate assessments of scientific ideas. Con-  
655 versely, they may also introduce complexities that  
656 do not necessarily translate to better performance,  
657 such as overfitting or increased computational costs.  
658 The trade-offs between model size, accuracy, and  
659 efficiency are still an area ripe for exploration.

### 660 **Ethics Statement**

661 The dataset used in our study consists of publicly  
662 available academic papers. We have ensured that  
663 all data was collected and handled in a manner that  
664 respects the privacy and intellectual property rights  
665 of the authors. No personal data was used, and all  
666 information is attributed to its original source.

667 We are committed to transparency in our re-  
668 search process. To this end, we have made our  
669 benchmark dataset publicly available and have pro-  
670 vided detailed descriptions of our methodologies  
671 and experimental setups to facilitate reproducibility  
672 by other researchers.

We recognize the importance of the human el-  
ment in the evaluation of scientific ideas. Our  
framework is designed to assist, rather than replace,  
human judgment. We believe that the most effec-  
tive use of our model is as a tool to support and  
enhance the work of human reviewers, not to sup-  
plant them.

## 680 **References**

- 681 Shubham Agarwal, Issam H Laradji, Laurent Char-  
682 lin, and Christopher Pal. 2024. Litllm: A toolkit  
683 for scientific literature review. *arXiv preprint*  
684 *arXiv:2402.01788*.
- 685 Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. Scibert:  
686 A pretrained language model for scientific text. In  
687 *Proceedings of the 2019 Conference on Empirical*  
688 *Methods in Natural Language Processing and the 9th*  
689 *International Joint Conference on Natural Language*  
690 *Processing (EMNLP-IJCNLP)*, pages 3615–3620.
- 691 Lutz Bornmann and Rüdiger Mutz. 2015. Growth rates  
692 of modern science: A bibliometric analysis based  
693 on the number of publications and cited references.  
694 *Journal of the association for information science*  
695 *and technology*, 66(11):2215–2222.
- 696 Alessandro Checco, Lorenzo Bracciale, Pierpaolo  
697 Loreti, Stephen Pinfield, and Giuseppe Bianchi. 2021.  
698 Ai-assisted peer review. *Humanities and Social Sci-*  
699 *ences Communications*, 8(1):1–11.
- 700 Mor Geva, Roei Schuster, Jonathan Berant, and Omer  
701 Levy. 2021. Transformer feed-forward layers are  
702 key-value memories. In *Proceedings of the 2021*  
703 *Conference on Empirical Methods in Natural Lan-*  
704 *guage Processing*, pages 5484–5495.
- 705 Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu,  
706 Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen,  
707 et al. 2021. Lora: Low-rank adaptation of large lan-  
708 guage models. In *International Conference on Learn-*  
709 *ing Representations*.
- 710 Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina  
711 Toutanova. 2019. Bert: Pre-training of deep bidirec-  
712 tional transformers for language understanding. In  
713 *Proceedings of NAACL-HLT*, pages 4171–4186.
- 714 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying  
715 Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gon-  
716 zalez, Hao Zhang, and Ion Stoica. 2023. Efficient  
717 memory management for large language model serv-  
718 ing with pagedattention. In *Proceedings of the 29th*  
719 *Symposium on Operating Systems Principles*, pages  
720 611–626.
- 721 Weixin Liang, Yuhui Zhang, Hancheng Cao, Binglu  
722 Wang, Daisy Ding, Xinyu Yang, Kailas Vodrahalli,  
723 Siyu He, Daniel Smith, Yian Yin, et al. 2023. Can  
724 large language models provide useful feedback on  
725 research papers? a large-scale empirical analysis.  
726 *arXiv preprint arXiv:2310.01783*.

727	Ryan Liu and Nihar B Shah. 2023. Reviewergpt? an exploratory study on using large language models for paper reviewing. <i>arXiv preprint arXiv:2306.00622</i> .	781	
728		782	
729		783	
730	Patrice Lopez. 2009. Grobid: Combining automatic bibliographic data recognition and term extraction for scholarship publications. In <i>Research and Advanced Technology for Digital Libraries: 13th European Conference, ECDL 2009, Corfu, Greece, September 27-October 2, 2009. Proceedings 13</i> , pages 473–474. Springer.	784	
731		785	
732			
733		Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, et al. 2023. Representation engineering: A top-down approach to ai transparency. <i>arXiv preprint arXiv:2310.01405</i> .	786
734		787	
735		788	
736		789	
737	Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models. In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing</i> , pages 9004–9017, Singapore. Association for Computational Linguistics.	790	
738		791	
739			
740			
741			
742			
743			
744	Charles Spearman. 1961. The proof and measurement of association between two things.		
745			
746	Albert N Tabah. 1999. Literature dynamics: Studies on growth, diffusion, and epidemics. <i>Annual review of information science and technology (ARIST)</i> , 34:249–86.		
747			
748			
749			
750	Mike Thelwall, Kayvan Kousha, Paul Wilson, Meiko Makita, Mahshid Abdoli, Emma Stuart, Jonathan Levitt, Petr Knuth, and Matteo Cancellieri. 2023. Predicting article quality scores with machine learning: The uk research excellence framework. <i>Quantitative Science Studies</i> , 4(2):547–573.		
751			
752			
753			
754			
755			
756	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .		
757			
758			
759			
760			
761			
762	Xinbing Wang, Huquan Kang, Luoyi Fu, Ling Yao, Jixin Ding, Jianghao Wang, Xiaoying Gan, Chenghu Zhou, and John E Hopcroft. 2023. Quantifying knowledge from the perspective of information structuring. <i>Plos one</i> , 18(1):e0279314.		
763			
764			
765			
766			
767	Yi Xu, Shuqian Sheng, Bo Xue, Luoyi Fu, Xinbing Wang, and Chenghu Zhou. 2023. Exploring and verbalizing academic ideas by concept co-occurrence. In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 13001–13027, Toronto, Canada. Association for Computational Linguistics.		
768			
769			
770			
771			
772			
773			
774	Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, et al. 2023. Baichuan 2: Open large-scale language models. <i>arXiv preprint arXiv:2309.10305</i> .		
775			
776			
777			
778	Weizhe Yuan, Pengfei Liu, and Graham Neubig. 2022. Can we automate scientific reviewing? <i>Journal of Artificial Intelligence Research</i> , 75:171–212.		
779			
780			

## A Prompts for LLM Generation

For the baselines involving LLM generation, we design a prompt to elicit the evaluation of a manuscript paper’s **overall quality** based on its abstract. The prompt, which demonstrated optimal performance when applied to GPT-3.5-turbo, is structured as follows:

**Evaluate** the quality (overall quality score) of the following manuscript paper based on its abstract.

**title:** {title}

**abstract:** {abstract}

The score should be between 1 and 10, with 1 being the lowest and 10 being the highest. Just output your score, no more other words.

It is important to note that while this prompt is most effective for GPT-3.5-turbo, its influence on the performance of other fine-tuned models, such as LLaMA-2 and Baichuan-2, is less pronounced. These models have been trained to adapt to the distribution of scores in the training set, which mitigates the impact of the prompt’s phrasing on their generative capabilities.

## B Hyper Parameters

We first declare that the reason for using **LLaMA-2 instead of LLaMA-3** or other updated models is because we are concerned that new models may be pretrained using papers from ICLR23, resulting in a data leakage problem.

Parameter	Value
learning rate	2e-5
epoch	3
weight decay	0
warmup ratio	0.03
bf16	True

Table 5: Hyper-parameters of LLaMA-2-Full-SFT.

We detail the hyper-parameters for the LLM generation baseline, LLaMA-2-Full-SFT, in Table 5. Additionally, the hyper-parameters for models trained with the LoRA technique, specifically

LLaMA-2-LoRA-SFT and Baichuan-2-LoRA-SFT, are outlined in Table 6.

Parameter	LLaMA-2	Baichuan-2
learning rate	2e-5	2e-5
epoch	10	10
weight decay	0	0
warmup ratio	0	0
bf16	True	True
LoRA modules	q_proj, v_proj	W_pack
LoRA r	8	16
LoRA alpha	16	32
LoRA dropout	0.05	0.1

Table 6: Hyper-parameters of LLaMA-2-LoRA-SFT and Baichuan-2-LoRA-SFT.

For the representation-based evaluation method RePE (Zou et al., 2023), we employ Principal Component Analysis (PCA) as the embedding evaluation mechanism, in line with the recommendations provided in the original paper.

## C Expanded Results

Building upon the experimental setup detailed in Section 6.2, we examine the performance across **all layers** of the foundational model, LLaMA-2-7b-base. The detailed Spearman correlation results, which consider the full spectrum of layers under various training ratios, are illustrated in Figure 5 for the ICLR23-low-std dataset and in Figure 6 for the ICLR23-all dataset.

In our pursuit to validate the robustness of our findings, we conducted parallel experiments using Baichuan-2-7b-base as an alternative base model. The corresponding Spearman correlation results are depicted in Figure 7 and Figure 8. The patterns observed with Baichuan-2-7b-base are found to be in harmony with those from the LLaMA-2-7b-base model, lending credence to the consistency and reliability of our conclusions.

The experiments across different models not only reinforces the validity of our initial observations but also suggests that the underlying phenomena we have identified are model-agnostic to a certain extent. Such findings are indicative of the potential generalizability of our framework, hinting at its applicability across a variety of LLMs. Future work may delve deeper into the comparative analysis of additional models, further expanding our understanding of the relationship between the base model and the efficacy of idea evaluation.

## D Case Study

We list four cases to show the performance and drawbacks of our method in Table 7. All these cases are selected from the domain of reinforcement learning. The first two cases are correctly predicted, and the human-rated scores and LLM representation scores are very close. As to the third case, although our method gives an overestimated score, the final score is not enough to make it acceptable. The fourth case is underestimated. One possible reason is that in such cases, our method may lack more contextual information to make a decision, such as tables and figures in papers, which is also something we need to consider in the future.

## E Domain Analysis

To see how our method rates ideas on popular topics or less trendy domains, we analyze the score distributions and differences between human-rated scores and LLM representation scores in 14 domains divided by ICLR-2023 program committee. The results are shown in Table 8. On the whole, the differences in mean scores of most domains are less than 10%. However, there are three domains (*Theory, Neuroscience and Cognitive Science, Infrastructure*) where the mean values of human-rated scores are relatively higher than average, and the differences exceed 10%. We believe these three domains are distinguished from other domains since others frequently focus on *Learning and Optimization*, which makes our evaluator overfitting on these data. Therefore, it is necessary to train a domain-specific evaluator, while also maintaining a balance in the content of the dataset. We will address these issues in our future work.

## F Frequently Asked Questions

### Numerical Processing Limitations in LLMs

A critical issue faced by LLMs is their inherent difficulty in processing numerical data, such as digits. This limitation stems from the finite-sized vocabulary and tokenization strategies used by these models, affecting both encoder and decoder architectures. This impacts the ability to perform tasks that require precise numerical understanding. Therefore, we leverage the deep, contextual representations within LLMs to quantify the value of scientific ideas. These representations encapsulate rich semantic and contextual information that extends beyond the superficial token sequences.

Our approach diverges significantly from the strategy of merely adding task-specific heads to the model. Instead, our approach involves a strategic selection of layers and tokens. By leveraging the hierarchical processing capabilities of LLMs, we can harness the most relevant and informative features for idea evaluation. This approach contrasts with using the entire weight set of the LLM (adding head to LLM), which might not be as efficient or effective for capturing the specific attributes necessary for assessing the value of scientific ideas.

### Framework

The current design of our framework shows considerable promise in the automated assessment of scientific ideas, yet there are avenues for further enhancing the evaluator’s performance. The quantitative assessment of scientific ideas is inherently complex, involving a blend of objective metrics and subjective judgments. Our method, leveraging the representations of large language models, demonstrates the potential to approximate human judgment to a significant degree, which will provide human reviewers with an objective score, rather than replacing them to give subjective comments from multiple perspectives.

<b>Title</b>	<b>SMART: Self-supervised Multi-task pretrAining with contRol Transformers</b>
<b>Abstract</b>	Self-supervised pretraining has been extensively studied in language and vision domains, where a unified model can be easily adapted to various downstream tasks by pretraining representations without explicit labels. When it comes to sequential decision-making tasks, however, it is difficult to properly design such a pretraining approach that can cope with both high-dimensional perceptual information and the complexity of sequential control over long interaction horizons ...
<b>Scores</b>	<b>Human-rated Score: 7.50</b> <b>LLM Representation Score: 7.22</b>
<b>Title</b>	<b>Ensemble Homomorphic Encrypted Data Classification</b>
<b>Abstract</b>	Homomorphic encryption (HE) is encryption that permits users to perform computations on encrypted data without first decrypting it. HE can be used for privacy-preserving outsourced computation and analysis, allowing data to be encrypted and outsourced to commercial cloud environments for processing while encrypted or sensitive data. HE enables new services by removing privacy barriers inhibiting data sharing or increasing the security of existing services ...
<b>Scores</b>	<b>Human-rated Score: 1.50</b> <b>LLM Representation Score: 1.61</b>
<b>Title</b>	<b>Comparative Analysis between Vision Transformers and CNNs from Neuroscience</b>
<b>Abstract</b>	Neuroscience has provide many inspirations for the development of artificial intelligence, especially for neural networks for computer vision tasks. Recent research on animals' visual systems builds the connection between neural sparsity and animals' levels of evolution, based on which comparisons between two most influential vision architecture, Transformer and CNN, are carried out. In particular, the sparsity of attentions in Transformers is comprehensively studied, and previous knowledge on sparsity of ...
<b>Scores</b>	<b>Human-rated Score: 2.50</b> <b>LLM Representation Score: 4.80 (Over Estimated)</b>
<b>Title</b>	<b>Neural Causal Models for Counterfactual Identification and Estimation</b>
<b>Abstract</b>	Evaluating hypothetical statements about how the world would be had a different course of action been taken is arguably one key capability expected from modern AI systems. Counterfactual reasoning underpins discussions in fairness, the determination of blame and responsibility, credit assignment, and regret. In this paper, we study the evaluation of counterfactual statements through neural models ...
<b>Scores</b>	<b>Human-rated Score: 7.33</b> <b>LLM Representation Score: 5.21 (Under Estimated)</b>

Table 7: Case study of the comparison between human-rated scores and LLM representation scores



Figure 5: Spearman correlations varying with layers under different training ratios of **LLaMA-2-7b-base** on **ICLR23-low-std** dataset.

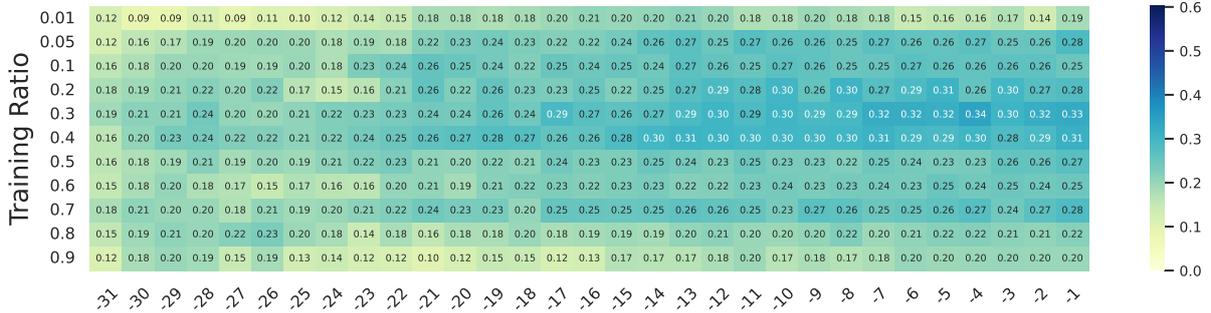


Figure 6: Spearman correlations varying with layers under different training ratios of **LLaMA-2-7b-base** on **ICLR23-all** dataset.

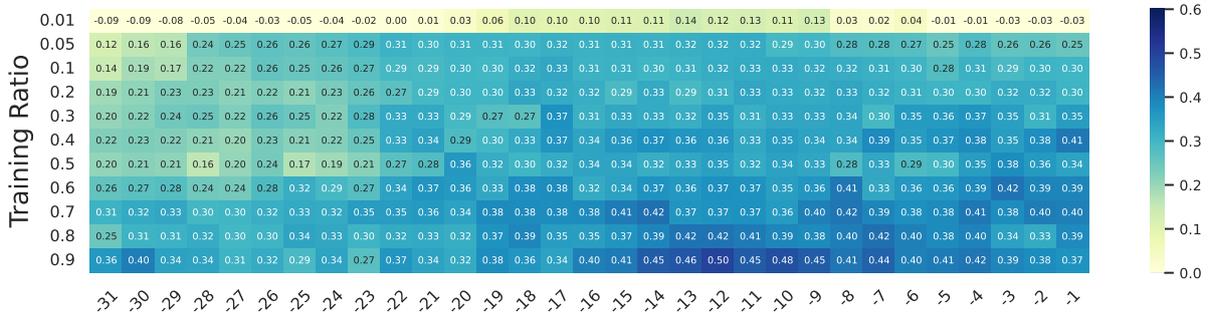
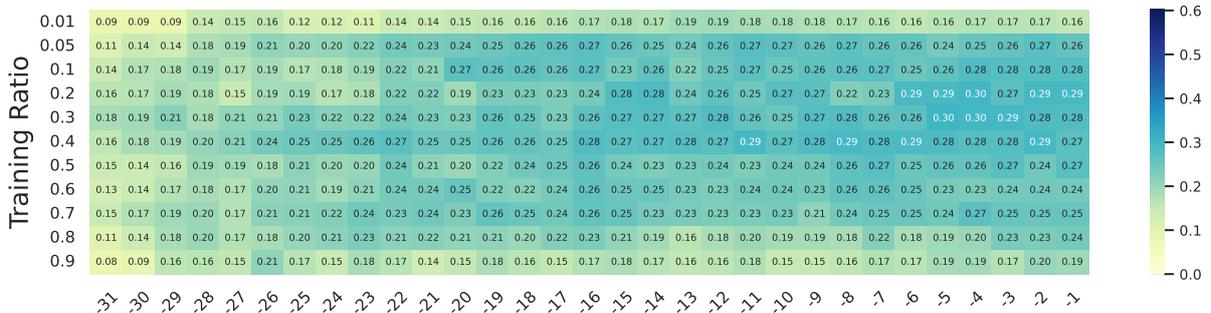


Figure 7: Spearman correlations varying with layers under different training ratios of **Baichuan-2-7b-base** on **ICLR23-low-std** dataset.



Domain	Num	Mean		Diff		Min		Max	
		Human Mean	Ours Mean	Human Mean	Ours Mean	Human Min	Ours Min	Human Max	Ours Max
Deep Learning and representational learning	527	5.5131 ± 1.3967	5.3385 ± 1.0125	3.17%	1.6667	1.7936	8.0000	8.0428	
Social Aspects of Machine Learning	168	5.4062 ± 1.4132	5.1867 ± 1.0791	4.06%	1.5000	1.0039	8.0000	7.2320	
Reinforcement Learning	232	5.4297 ± 1.5289	5.2762 ± 1.1101	2.83%	1.5000	2.2600	8.0000	8.2055	
General Machine Learning	111	5.4955 ± 1.2828	5.1640 ± 1.0825	6.03%	2.3333	1.7422	8.0000	7.4929	
Applications	249	5.6285 ± 1.4025	5.1934 ± 1.0932	7.73%	1.0000	2.3488	8.0000	7.6288	
Theory	86	6.1155 ± 1.2607	5.3296 ± 1.0076	12.85%	2.0000	3.5655	8.0000	8.3156	
Generative models	72	5.6245 ± 1.3514	5.4084 ± 1.0636	3.84%	1.5000	1.5492	8.0000	7.7018	
Unsupervised and Self-supervised learning	95	5.5672 ± 1.2967	5.2882 ± 0.9917	5.01%	2.0000	2.9334	8.0000	7.2551	
Optimization	76	5.3518 ± 1.3639	5.0908 ± 0.9983	4.88%	1.5000	1.5517	8.0000	7.2960	
Machine Learning for Sciences	96	5.2392 ± 1.6701	5.0426 ± 1.1133	3.75%	1.6667	2.3116	8.0000	7.3444	
Neuroscience and Cognitive Science	23	6.0667 ± 1.3181	5.1115 ± 0.9040	15.75%	3.0000	2.7473	8.0000	6.6993	
Probabilistic Methods	49	5.6442 ± 1.4135	5.2455 ± 0.9943	7.06%	2.5000	3.3245	8.0000	7.9918	
Infrastructure	18	5.9648 ± 1.2578	5.2065 ± 1.0139	12.71%	3.5000	3.0230	8.0000	7.1837	
None	5	3.5500 ± 1.6763	3.6157 ± 1.8071	1.85%	1.0000	0.34500	5.7500	5.3299	

Table 8: The detailed distribution of LLM representation scores and human-rated scores in 14 domains.