Can LLMs Generate Tabular Summaries of Science Papers? Rethinking the Evaluation Protocol

Daniel Khashabi

JHU, Assistant Professor

Task: Generating Tabular Summary for Scientific Content

- **Input:** A user prompt seeking scientific information.
- **Output:** A table that summarizes the information/insights extracted from the relevant science papers.

Prompt: Generate a table comparing video datasets

Repositories of science: arXiv, Google Scholar, etc.

Potentially relevant papers

	Dataset size	Annotation method	Intended Application	Evaluation Metric
Paper 1	1,200 video sequences	Subjectively annotated	Objective VQA method development	Subjective Mean Opinion Score
Paper 2	585 videos	Subjective video quality scores via crowdsourcing	NR video quality prediction advancement	Subjective video quality scores
Paper 3	153,841 videos	Coarsely annotated set with five quality ratings each	Deep-learning VQA model training	Spearman rank-order correlation coefficient
Paper 4	1 million YouTube videos	N/A	Large-scale video classification and action recognition	Performance improvements over baselines

Tabular summary

Why focus on this task?

- Tables are widely adopted format for scientific content.
  - Improves clarity, enables comparisons
- Integrating it into our workflow will improve quality and interpretability.

Prompt: Generate a table comparing video datasets

Repositories of science: arXiv, Google Scholar, etc.

Potentially relevant papers

	Dataset size	Annotation method	Intended Application	Evaluation Metric
Paper 1	1,200 video sequences	Subjectively annotated	Objective VQA method development	Subjective Mean Opinion Score
Paper 2	585 videos	Subjective video quality scores via crowdsourcing	NR video quality prediction advancement	Subjective video quality scores
Paper 3	153,841 videos	Coarsely annotated set with five quality ratings each	Deep-learning VQA model training	Spearman rank-order correlation coefficient
Paper 4	1 million YouTube videos	N/A	Large-scale video classification and action recognition	Performance improvements over baselines

Tabular summary

what we do, briefly

- Develop **arXiv2Table**, a framework for evaluating this task.
- Develop strong system for tackling our benchmark.

Prompt: Generate a table comparing video datasets

Repositories of science: arXiv, Google Scholar, etc.

Potentially relevant papers

	Dataset size	Annotation method	Intended Application	Evaluation Metric
Paper 1	1,200 video sequences	Subjectively annotated	Objective VQA method development	Subjective Mean Opinion Score
Paper 2	585 videos	Subjective video quality scores via crowdsourcing	NR video quality prediction advancement	Subjective video quality scores
Paper 3	153,841 videos	Coarsely annotated set with five quality ratings each	Deep-learning VQA model training	Spearman rank-order correlation coefficient
Paper 4	1 million YouTube videos	N/A	Large-scale video classification and action recognition	Performance improvements over baselines

Tabular summary

How good is the state-of-the-art at this task?

An easy prompt since it's extensively discussed in various forms (papers, Twitter, Reddit, etc.)

Create a table where the rows are recent language model and the columns are different details about their alignment process.

What I manually created for my course

Language Model	Relea se	Base	Alignment Algorithm(s) Used	Alignment Data Sources for alignment
GPT-3-instruct	2020	GPT-3	SFT --> RLHF/PPO	Curated datasets with human-labeled prompts and responses
GPT-4	2023	GPT-4 pre-trained?	SFT --> RLHF/PPO	Curated datasets with human-labeled prompts and responses
Gemini	2023	Gemini pre-trained?	SFT --> RLHF/PPO	Curated datasets with human-labeled prompts and responses
LLaMA2	2023	LLaMA2 pre-trained	SFT --> RLHF/PPO	Curated datasets with human-labeled prompts and responses
LLaMA3	2024	LLaMA3 pre-trained	Iterate: Rejection sampling -> SFT -> DPO	10 million human-annotated examples. The alignment process was conducted over multiple rounds, with each round involving the collection of new preference annotations and SFT data.
Alpacha	2023	LLAMA 1	SFT	Self-Instruct, 52,000 input-output pairs
Qwen2.5	2024	Qwen2.5 pre-trained	SFT -> DPO -> GRPO	1 million samples
Tulu 3	2024	Llama 3.1	SFT -> DPO -> RLVR	near 1 million samples
DeepSeek (V3)	2024	DeepSeek pre-trained	SFT -> GRPO	1.5 million samples (reasoning + non-reasoning tasks). Reasoning data was generated by specialized models. Non-reasoning data was produced by DeepSeek-V2.5 and validated by human reviewers.

Not able to identify a meaningful set of column values (table schema)

ChatGPT

Create a table where the rows are recent language model and the columns are different details about their alignment process.

Here's a table summarizing recent language models and their alignment processes:

nguage Model	Alignment Process	Description
InstructGPT	Reinforcement Learning from Human Feedback (RLHF)	Trained to follow instructions by learning from human-provided feedback, enhancing its ability to align with user intentions.
ChatGPT	Reinforcement Learning from Human Feedback (RLHF)	Utilizes human feedback to fine-tune responses, aiming to improve relevance and appropriateness in conversational contexts.
Claude	Reinforcement Learning from AI Feedback (RLAIF)	Employs AI-generated feedback based on constitutional principles to guide its behavior, reducing reliance on human input.
Gemini	Reinforcement Learning from Human Feedback (RLHF)	Incorporates human feedback to align its outputs with user preferences, aiming for safer and more accurate responses.
Sparrow	Reinforcement Learning from Human Feedback (RLHF)	Focuses on aligning dialogue agents through targeted human judgments to improve safety and usefulness.
DeepSeek R1	Reinforcement Learning with Novel Training Methods	Introduces innovative training techniques that prioritize correct answers over human-legible reasoning, leading to unexpected behaviors.

A lot of details are missing here! (many open-source models)

Example retrieved on 02/27/2025

Better schema
(already given in the prompt) 👍

ChatGPT

Create a table where the rows are language models: GPT-3, BLOOM, GPT-4, Gemini, DeepSeek (V3), DeepSeek (R1), LLaMA3, Qwen2, OLMo 2

The columns are different details about their alignment (after pre-training) stage. For example, what alignment algorithm they used, in what order, where did they get their alignment data, etc.

Compiling detailed information about the alignment processes for various language models can be
ing due to the proprietary nature of some models and the evolving practices in the field.
a table summarizing the available information on the alignment stages of the specified

Language Model	Alignment Algorithm(s) Used	Alignment Data Sources	Alignment Process Details
GPT-3	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
BLOOM	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
GPT-4	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
Gemini	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
DeepSeek (V3)	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
DeepSeek (R1)	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
LLaMA3	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
Qwen2	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.
OLMo 2	Not publicly disclosed	Not publicly disclosed	Limited information available on post-training alignment methods.

There is so much known about these models that are missing!

Example retrieved on 03/13/2025

“Wait but did you tr

Create a table where the rows are recent language model and the columns are different details about their alignment process.

ChatGPT **Deep Search**
(>15 mins)

Model	Alignment Methodology (Key Steps)
OpenAI InstructGPT (GPT-3.5, 2022)	- **Supervised fine-tuning** on a curated set of human-written demonstrations of desired behavior (labelers provided prompts and ideal answers) ARXIV.ORG . - **Reward model training** on human preference data: labelers ranked multiple model outputs for the same prompt, producing a dataset of comparisons ARXIV.ORG . - **RLHF**: The model was further fine-tuned with **reinforcement learning from human feedback** (using PPO) to maximize the reward model's score, aligning outputs with human-preferred responses ARXIV.ORG . - **Safety & bias mitigation**: Human labelers followed guidelines to prefer helpful, truthful, and inoffensive outputs, leading to a model that generated *more truthful and far less toxic* responses than the original GPT-3 ARXIV.ORG .
OpenAI ChatGPT / GPT-4 (GPT-3.5 & GPT-4, 2022–2023)	- **Initial fine-tuning**: Both ChatGPT (based on GPT-3.5) and GPT-4 use a similar alignment pipeline as InstructGPT, starting with SFT on demonstration data and then RLHF fine-tuning CDN.OPENAI.COM . GPT-4's post-training *“alignment process”* notably improved factuality and adherence to desired behavior CDN.OPENAI.COM . - **Safety-focused data**: OpenAI augmented GPT-4's training with an *additional set of safety-relevant prompts* during RLHF CDN.OPENAI.COM . These are adversarial or sensitive queries designed to teach the model to refuse or handle harmful requests appropriately. - **Rule-based reward modeling**: GPT-4 introduced *rule-based reward models (RBRMs)* – essentially GPT-4 classifier models that evaluate outputs against a set of written safety rules CDN.OPENAI.COM CDN.OPENAI.COM . During RLHF, the model received an extra reward signal from these classifiers, encouraging it to **refuse disallowed content** in the correct style and **not refuse** innocuous queries CDN.OPENAI.COM . - **Bias & safety mitigations**: The above methods, along with extensive red-team testing by domain experts, significantly improved GPT-4's safety. GPT-4 is far less

GPT 3.5 has two rows?

Incorrect information (**precision** issues)

Lots of text, not much structure/insight.

Example retrieved on 03/13/2025

No reference to Llama2 paper itself?

Create a table where the rows are recent language model and the columns are different details about their alignment process.

ChatGPT **Deep Search (>15 mins)**

The table has these models: GPT-3.5, GPT-4, Claude, Sparrow, LaMDA, LLaMA-2

What about the rest? **Recall issue.**

Meta LLaMA-2 Chat (2023)	- **Supervised instruction tuning**: LLaMA-2 Chat was initialized by **supervised fine-tuning** on a large collection of instruction-following data. This included prompt-response pairs from public sources and safety-centric data, so the model learned to produce helpful answers and follow user instructions while adhering to content guidelines VISO.AI VISO.AI . The SFT data covered both general helpful responses and safer completions (e.g. how to politely refuse inappropriate requests).
(continued)	- **Two-step RLHF with separate rewards**: Meta then applied **RLHF** in two dimensions HEIDLOFF.NET . They trained **two reward models** – a *helpfulness* reward model (judging how well the assistant's answer satisfies the user's request) and a *safety* reward model (judging harmful or policy-violating content) HEIDLOFF.NET . The chat model was optimized via Proximal Policy Optimization, balancing these rewards to produce responses that are both useful *and* harmless. They also employed **rejection sampling**: the model generates multiple responses and selects the one with the best combined helpfulness/safety score, which further improves alignment HEIDLOFF.NET . - **Safety context distillation**: An additional safety technique was used wherein, if the model started to produce unsafe outputs, the behavior from a higher-precision safety model or human-written safe responses were distilled back into the chat model (as extra fine-tuning) ARXIV.ORG VISO.AI . This way, the model learns to internalize safer responses for problematic prompts. - **Bias and toxicity mitigation**: The **safety reward model** was explicitly trained on detecting toxic, biased, or harmful content VISO.AI . By optimizing against this model's feedback, LLaMA-2 Chat greatly reduces toxic or biased generations. The model card reports strong performance on safety evaluations compared to previous open models HEIDLOFF.NET HEIDLOFF.NET . However, like other LLMs, it can still be adversarially prompted to reveal unsafe behavior in edge cases VISO.AI , so ongoing evaluation is necessary.

???

Example retrieved on 03/13/2025

Recent related work: ArxivDigesTables

Newman et al. ArxivDIGESTables: Synthesizing Scientific Literature into Tables using Language Models, EMNLP 2024

- Extracted 2.2K tables from existing papers.
- The table captions serve as the task prompts.
- Rows of the table correspond to individual papers (7K) papers.

	Dataset	Size	Task	Annotations
Paper 1	KoNViD-1k	1200	VQA	114
Paper 2	LIVE-VQC	585	VQA	240
Paper 3	KoNViD-150k	153,841	VQA	5
Paper 4	Sports-1M	1,133,158	Classification	- (auto)

We build upon this work by addressing their weaknesses!

Limitations of prior work

1. **The assumption that papers a carefully curated relevant papers are available.**
2. Table captions are not appropriate task prompts.
3. Rely on static embedding and human annotation to evaluate generated tables.

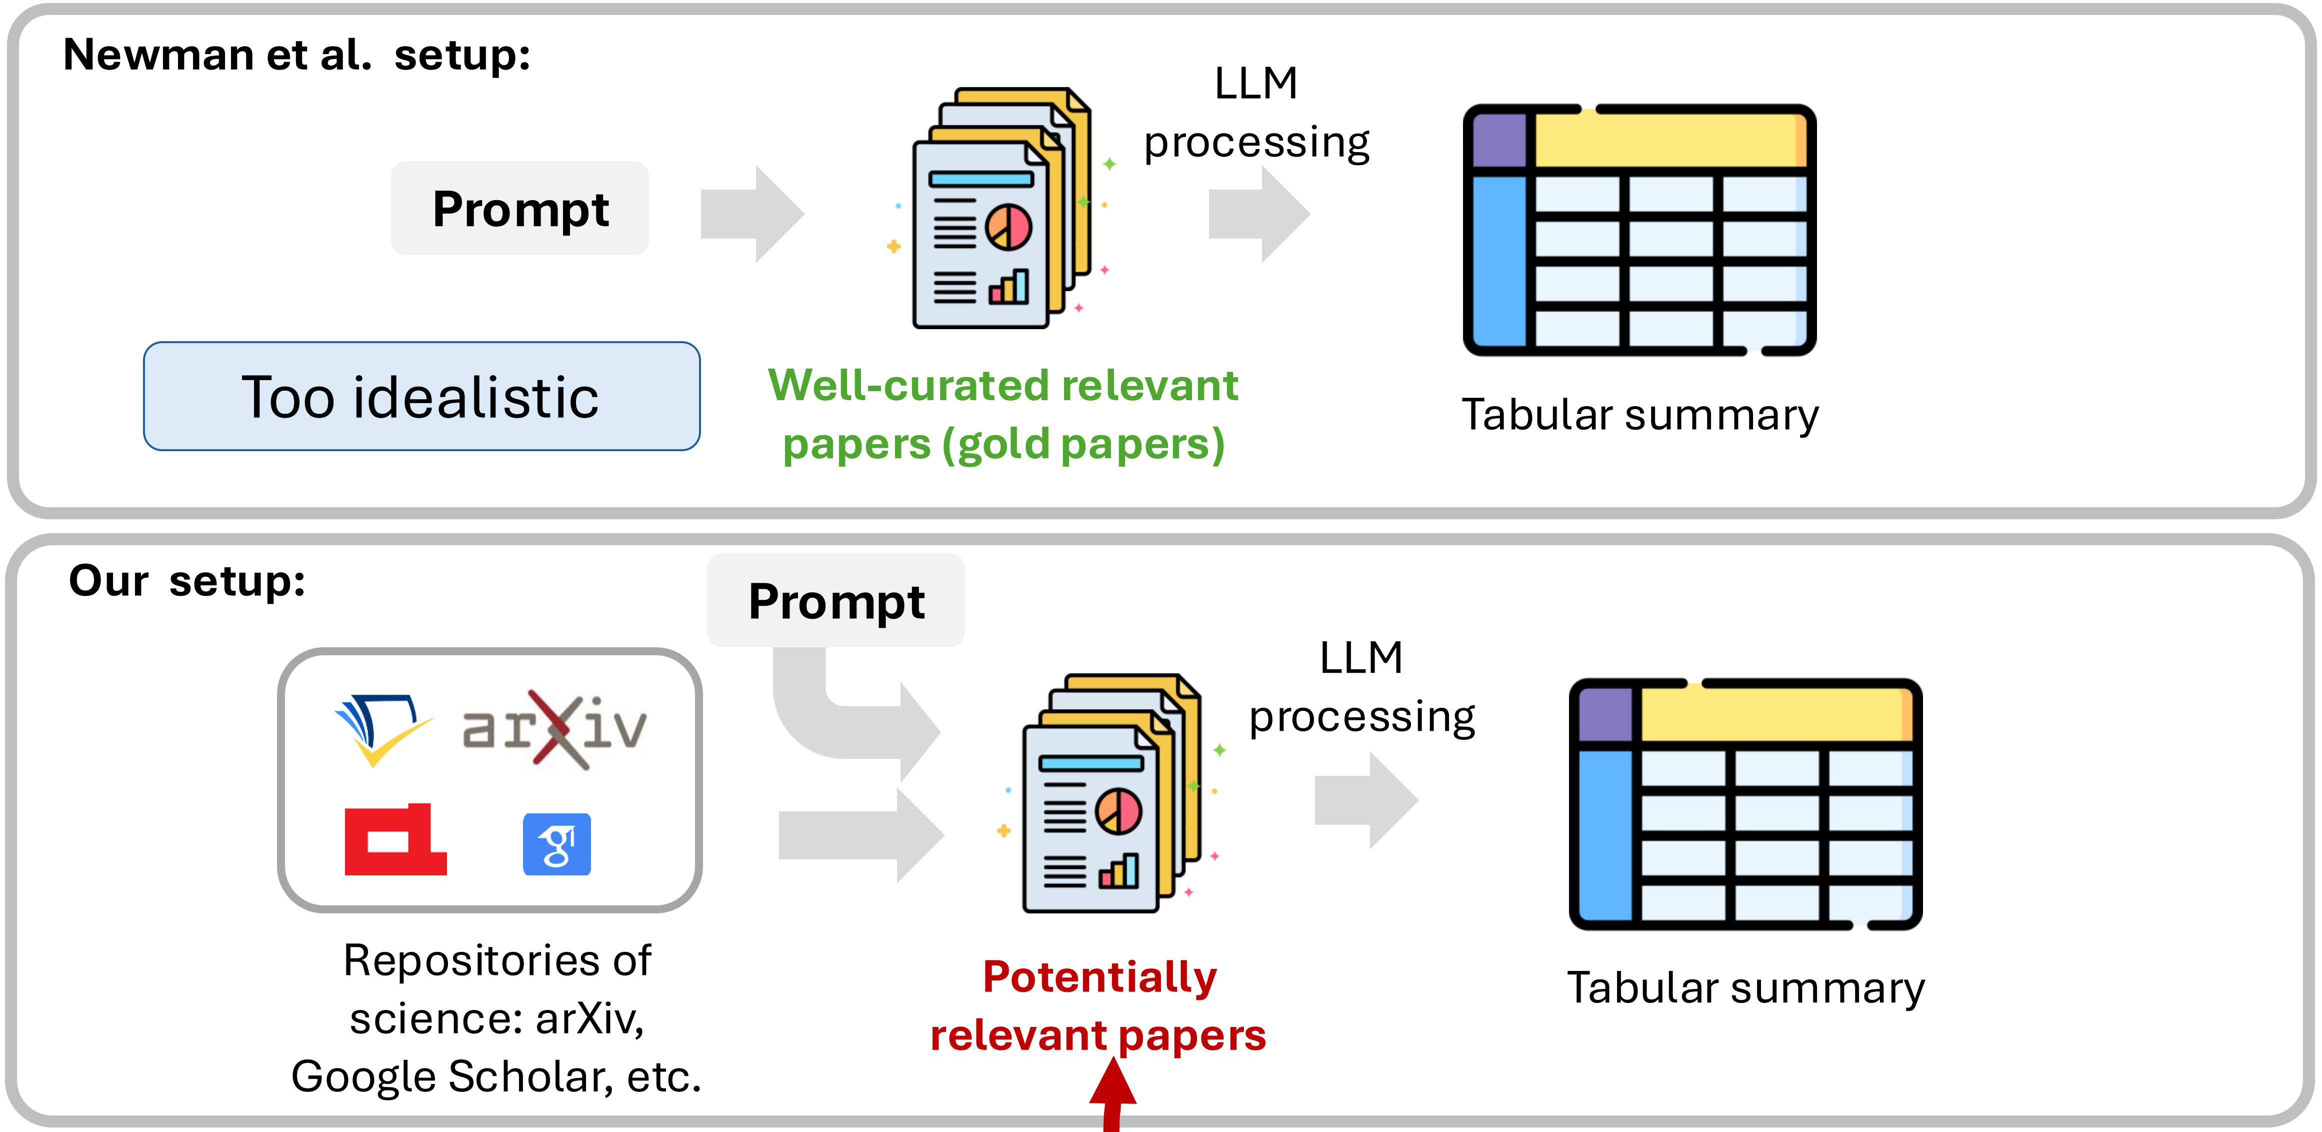

We build a retrieval engine over papers and identify hard negative candidate papers to make evaluation realistic.

Limitations of prior work

1. The assumption that papers a carefully curated relevant papers are available, is idealistic in realistic scenarios.
2. **Table captions are not appropriate as task prompts.**
3. Rely on static embedding and human annotation to evaluate generated tables.

User Demand vs. Captions

- Prompts in prior work [Newman et al.] are table captions.

Comparison of Trajectory and Path Planning Approach

Brief and ambiguous

- We replace them with **user demand** prompts:

Generate a table that compares different trajectory and path planning approaches. You can focus on their collision avoidance techniques, benefits, limitations, and applicable scenarios.

Longer and more precise

  - We collect these by careful prompting of LLMs to:
    - (1) obtain a more complete prompt while
    - (2) avoiding leakage of table schema/values.

User Demand vs. Captions

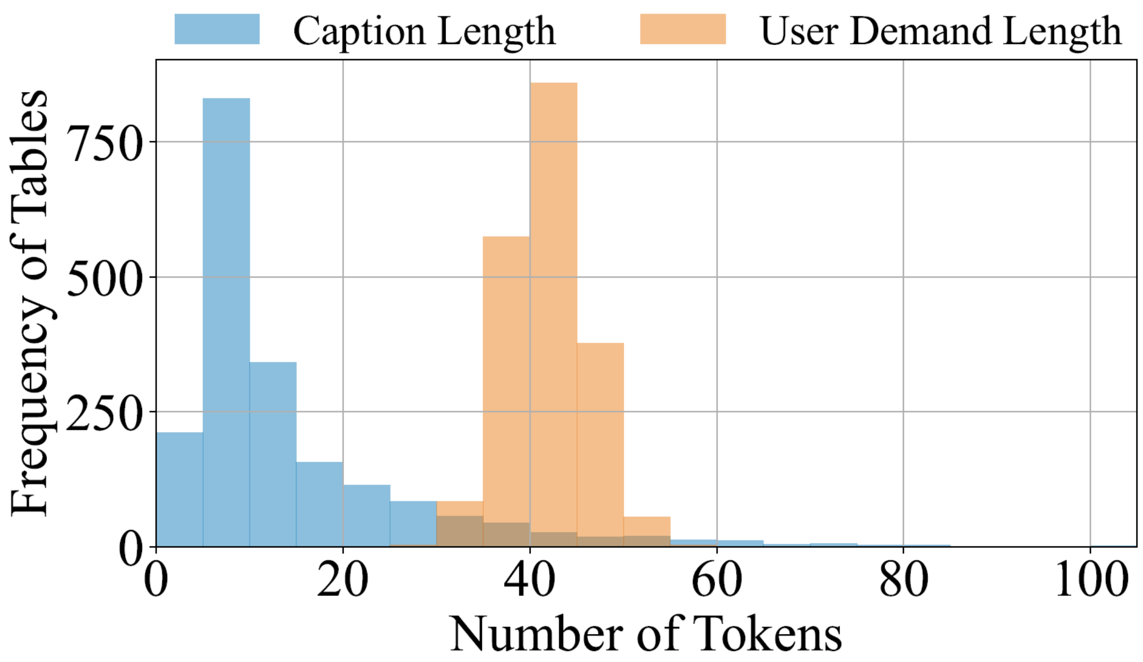

Our collected user demands feature longer context, thus including better hints to curate the table.

Limitations of prior work

1. The assumption that papers a carefully curated relevant papers are available, is idealistic in realistic scenarios.
2. Table captions are not appropriate as task prompts.
3. **Rely on static embedding and human annotation to evaluate generated tables.**

Originally, we have a ground-truth table extracted from a paper

Ground-truth Table

CBFIR Networks	*Datasets*	*Evaluation Metrics*	*Loss Function*
GAN	*DARN*	*Recall@1*	*TL, AL*
CN-LexNet	*Shopping100K*	*Recall@20*	*CL, TL*
ResNet-v2	*DeepFashion*	*Recall@1,10*	*BCE Loss*

We synthesize QA pairs from the ground-truth table about schema and values

table **schema**:
Is ***Dataset*** *included in the table schema?*

unary *(cell) values*:
Is ***CL, TL*** *the loss function for paper CN-LexNet?*

pairwise *comparisons*:
Is ResNet using ***more evaluation metrics*** *than GAN?*

Then, we ask an LLM to answer these QAs based on the generated table

The ratio of "Correct" indicates the **recall**.

Correct!

Correct!

Incorrect!

Then, we have a table that is generated by an LLM

Generated Table

Backbone Model	*Losses*	*Attributes*	*Datasets*
GAN	*TL+AL*	*Shape*	*DARN Color*
CNLexNet	*CL+TL*	*Various*	*Consumer-to-Shop*
ResNet	*Landmark*	*Various*	*DeepFashion*

Similarly, we can reverse the process by starting with the generated table.

Generated Table

Backbone Model	***Losses***	***Attributes***	***Datasets***
GAN	*TL+AL*	*Shape*	*DARN Color*
CNLexNet	*CL+TL*	*Various*	*Consumer-to-Shop*
ResNet	*Landmark*	*Various*	*DeepFashion*

Again, we synthesize QAs based on the generated table.

table **schema**:
*Is **Attributes** included in the table schema?*

***unary** (cell) values*:
*Is **DARN Color** used in GAN?*

***pairwise** comparisons*:
*Is ResNet using **fewer losses** than GAN?*

The ratio of "Correct" indicates the **precision**.

But answer them using the ground-truth table.

Incorrect! ❌ Incorrect! ❌ Correct! ✅

And answer QAs using our ground-truth table.

Ground-truth Table

CBFIR Networks	***Datasets***	***Evaluation Metrics***	***Loss Function***
GAN	*DARN*	*Recall@1*	*TL, AL*
CN-LexNet	*Shopping100K*	*Recall@20*	*CL, TL*
ResNet-v2	*DeepFashion*	*Recall@1,10*	*BCE Loss*

Our released data: **arXiv2Table**

- Expanded version of Newman et al. 2024.
- Contains
  - 2.1K user demand prompts
  - 2.1K tables (inherited from arXivDigestable).
    - Dropped few low-quality tables.
  - Each prompt comes with it a set of candidate (distractor + gold) papers.
  - Evaluation framework based on utilization.

Dataset will be on arXiv on coming weeks!

We also proposed a new approach

- An inference-time algorithm that iteratively digests and organizes papers in tabular form.

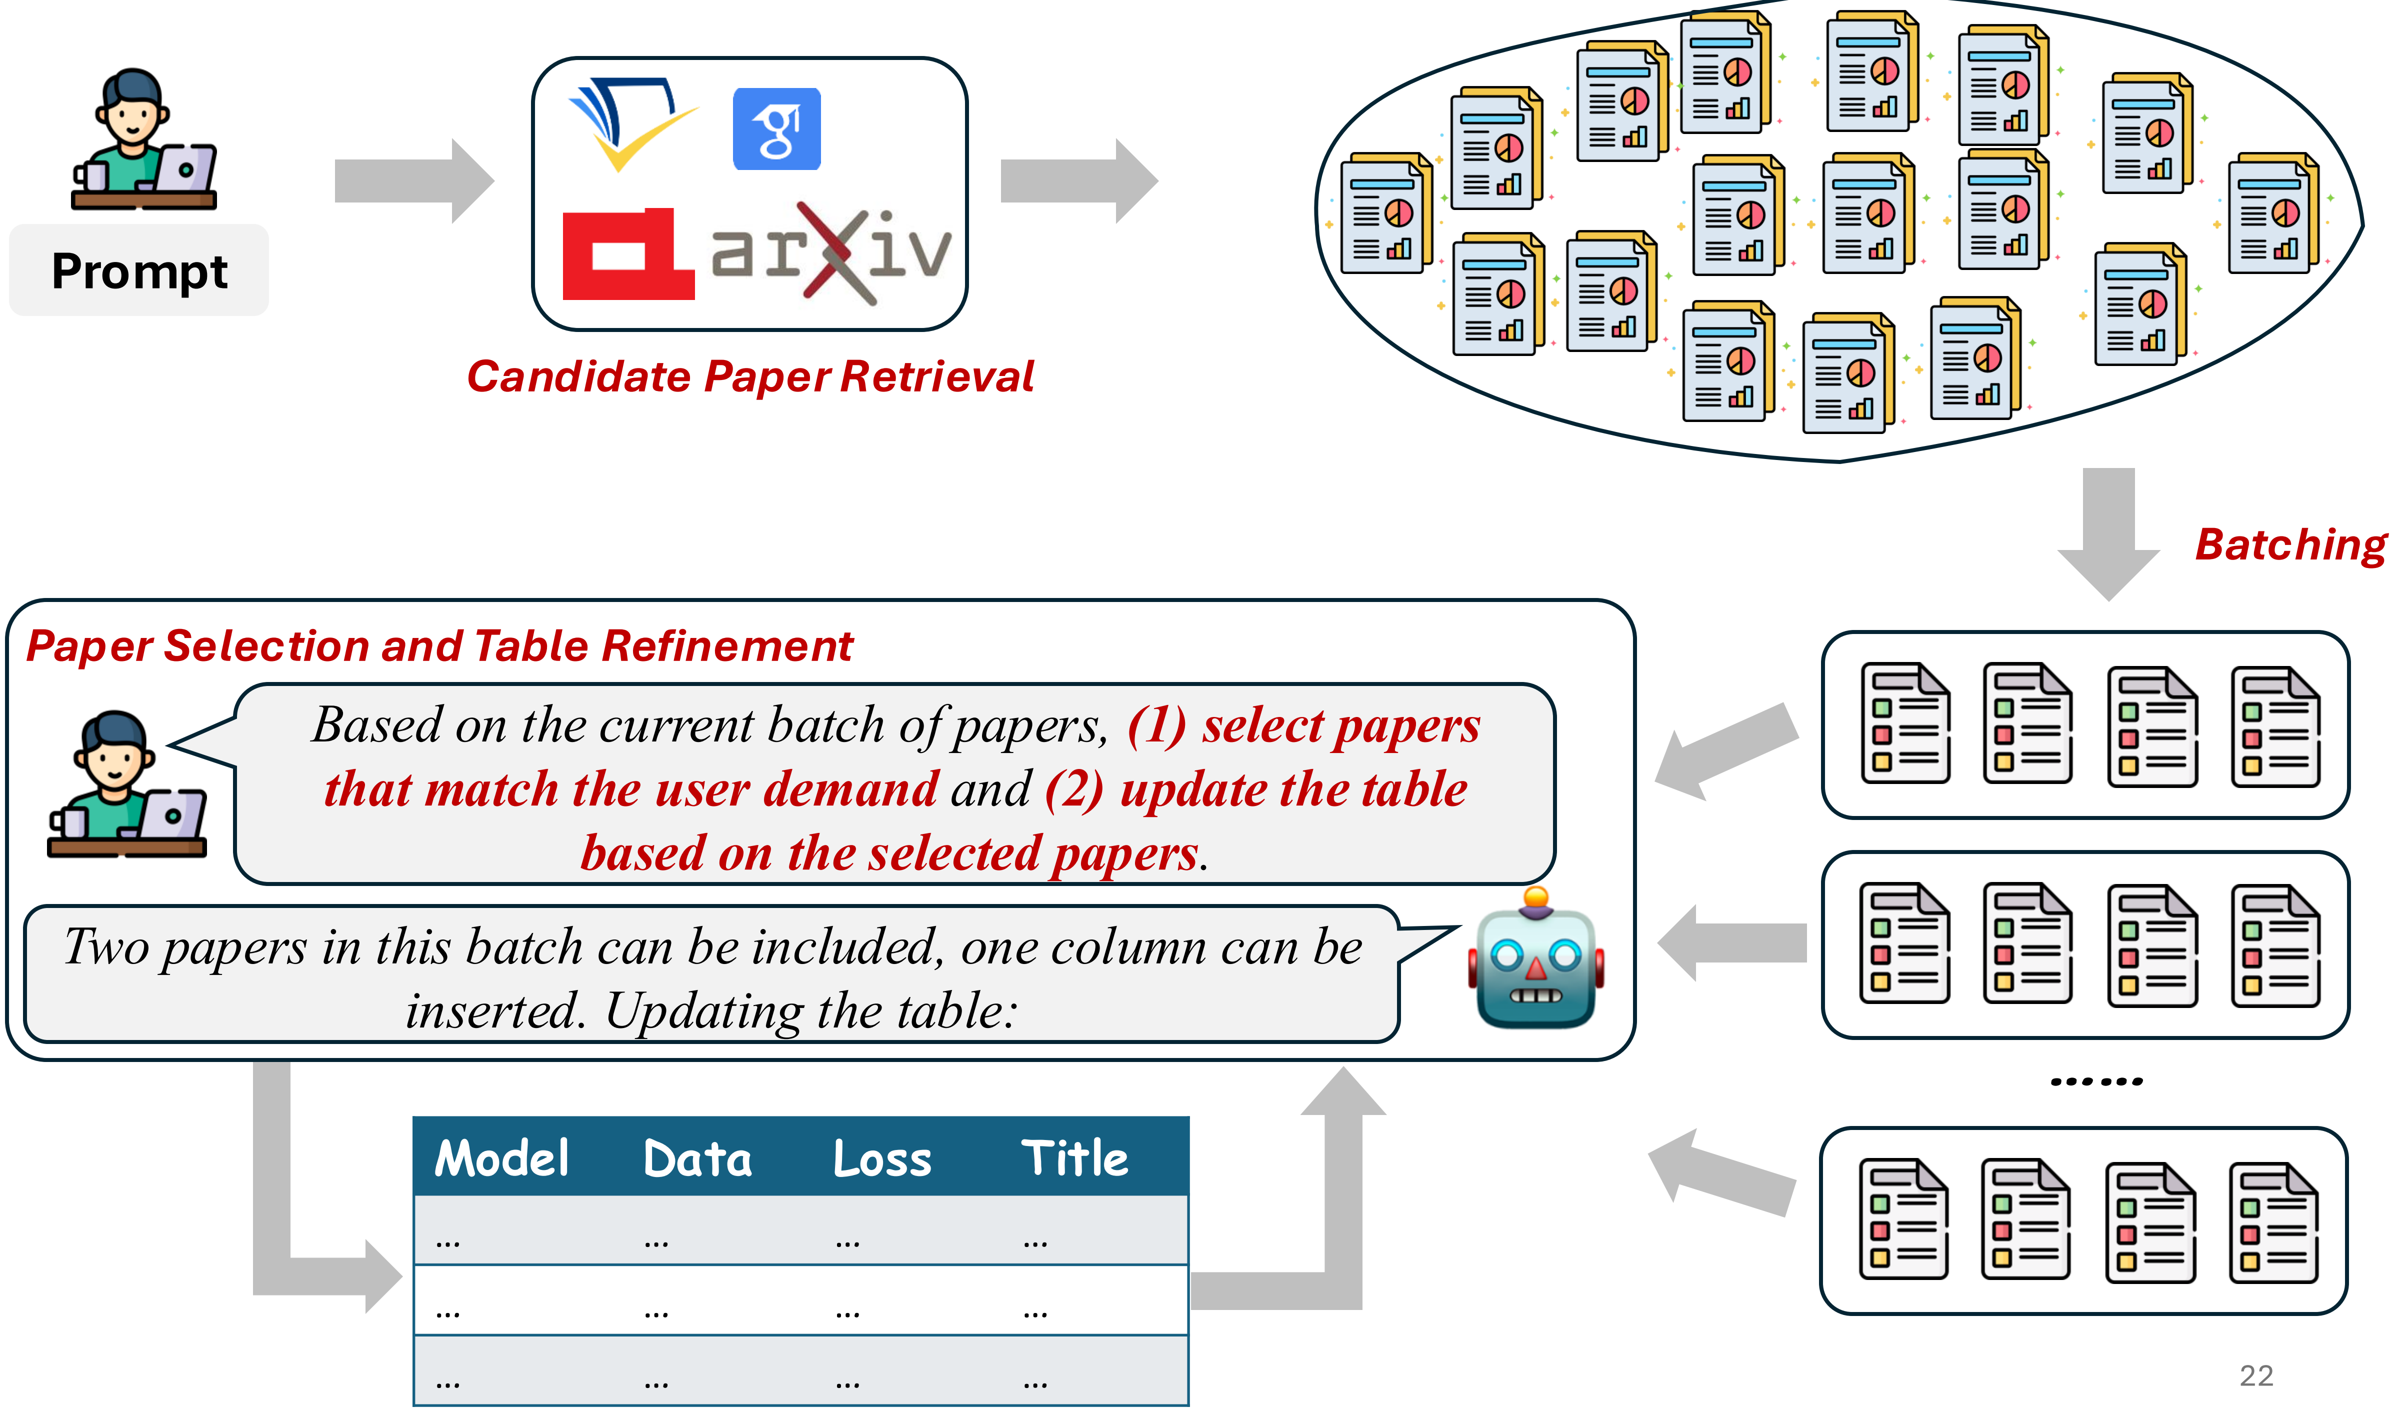
Prompt
Candidate Paper Retrieval
arXiv
Batching
Paper Selection and Table Refinement
Based on the current batch of papers, (1) select papers that match the user demand and (2) update the table based on the selected papers.
Two papers in this batch can be included, one column can be inserted. Updating the table:
Model
Data
Loss
Title
……

Evaluation vs Number of Iterations

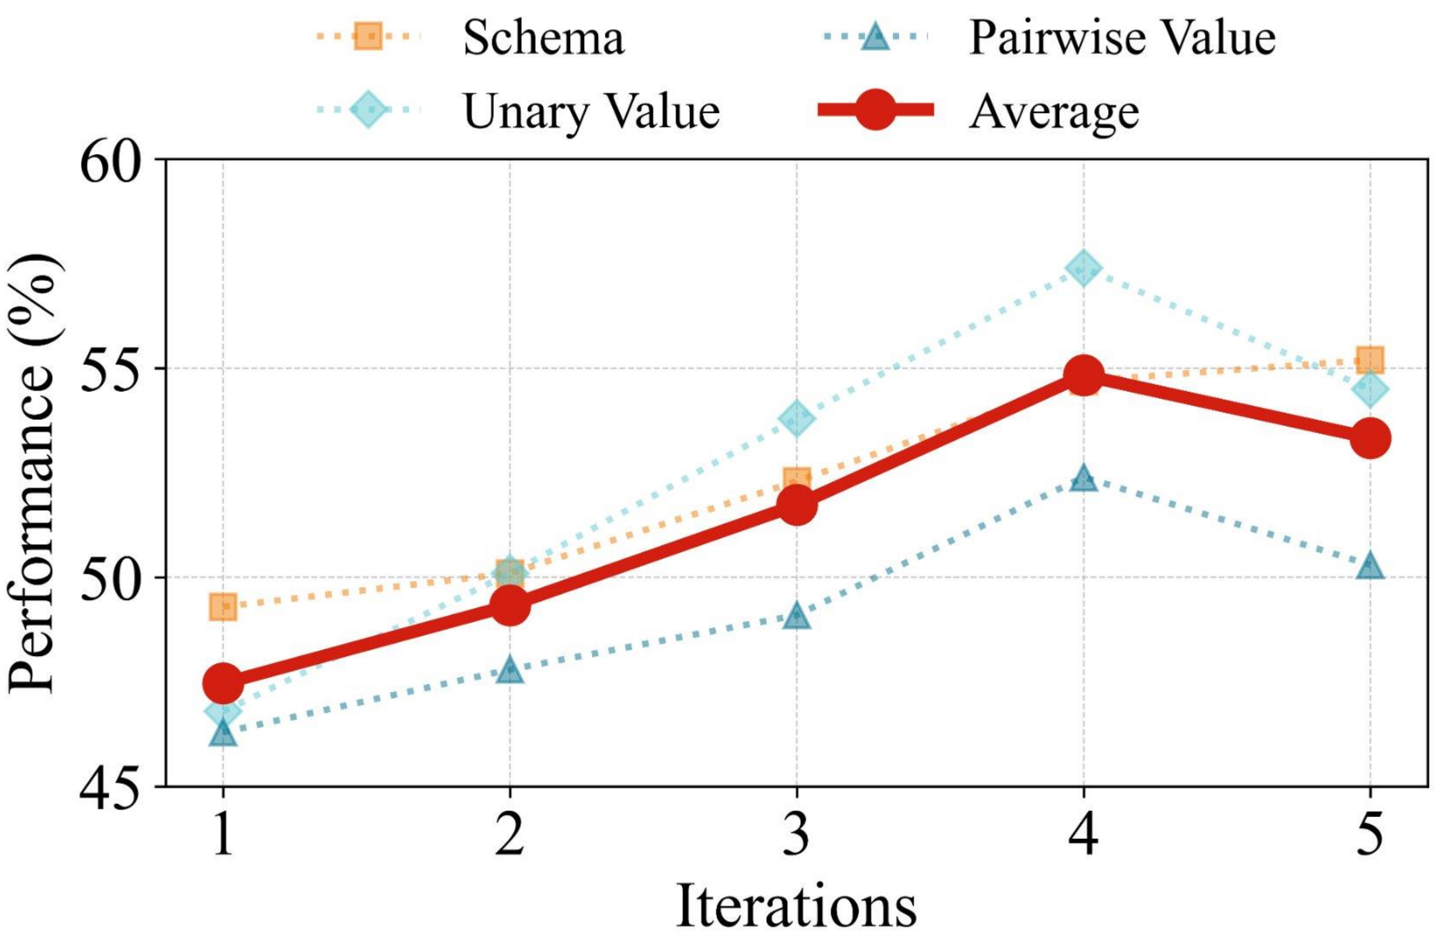

With more iterations, all aspects of the generated tables improves (up to iter ~4).

Evaluation of the end-to-end pipeline

- **Model:** Llama 3.3 (70B)
- **Baseline 1:** Read all papers in one go and write a table.
- **Baseline 2:** Read one paper at a time and incrementally form a table.
- **Newman et al.:** Two stages; define schema in the first round, then fill in the values.

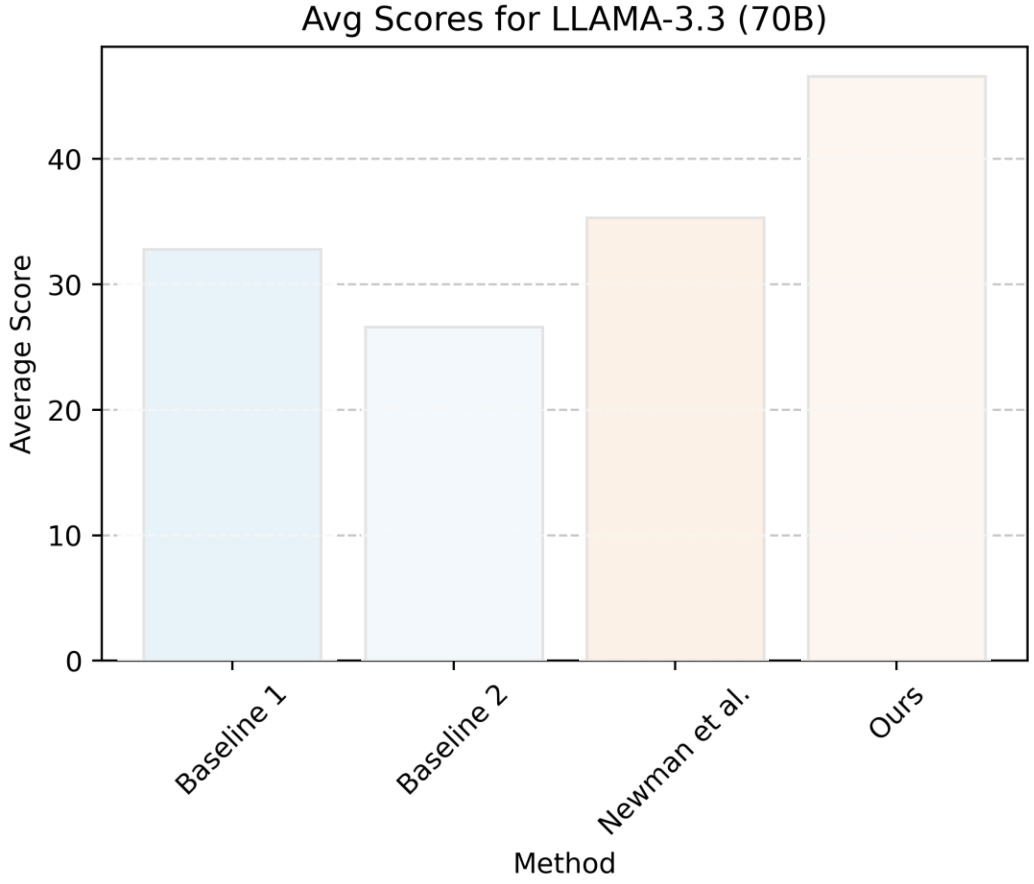

Our proposed approach outperforms existing results.

Evaluation of the end-to-end pipeline (other models)

The gains of our approach is consistent across different models.

The task remains challenging for all these approaches.

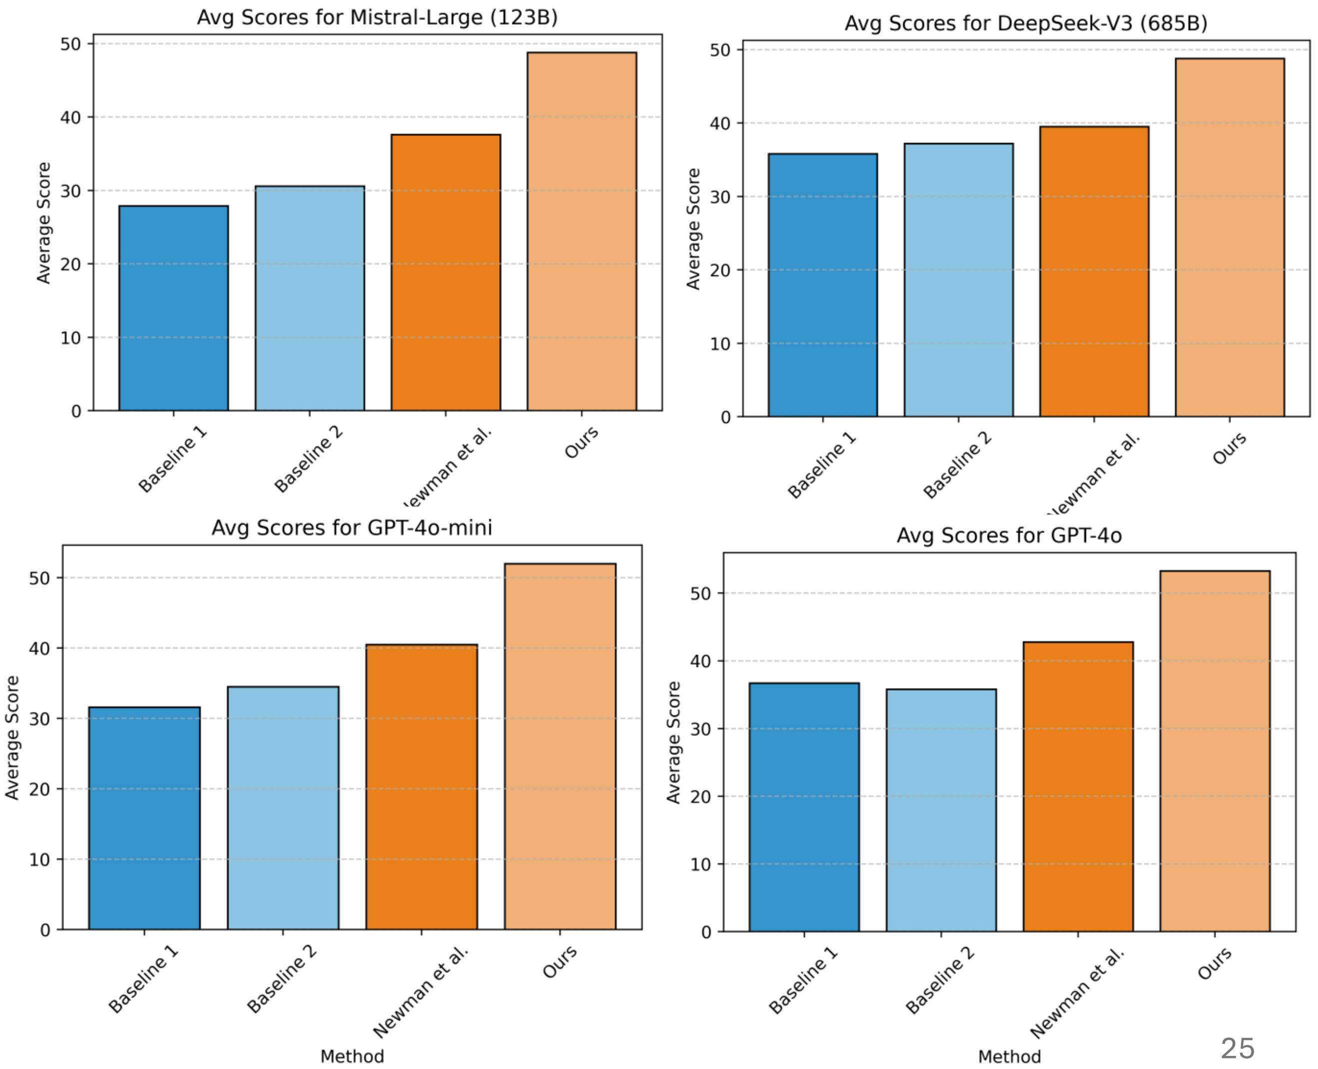

Summary and Conclusion

- **Motivation:** A more realistic pipeline for evaluating tabular summarization of science literature.
  - **Why?** Tabular summaries are crucial framework for quickly aggregating and understanding the progress in science.
- We introduce arXiv2Table, a framework for evaluating systems for tabular summarization.
- We also develop a system to address the challenge posed.
- Our benchmark is challenging! Give it a try!! 🚀