

How Can LLMs Help Science Communication Writers?

Preamble

This guide was written as a set of guidelines to help produce written syntheses of research papers using large language models (LLMs) such as ChatGPT or DeepSeek, just to name these two. I have personally explored how to tailor specific prompts towards producing useful written syntheses of research papers for my Knowledge Mobilization (KMB) Practicum project, which resulted in a suite of Instagram posts on the Bishop's Research account called the *Purple Patch Papers*. In the following, I present a situation similar to the exercises I have practiced for my KMB project to help you understand the potential uses of this method.

Context

A faculty member at Bishop's University has recently published a research paper and you are tasked to share the exciting work they have conducted!¹ However, you unfortunately do not have the leisure to fully read their paper, and even if you did, you are not certain you would be able to appropriately distill it into a concise and appealing summary. To help you in that task, you inquire the services of artificial intelligence (AI). Then, the question is: how can you use AI to help you produce a synthesis of the research paper? Here's a step-by-step method:

Step #1: Define the architecture of the research synthesis

If you ask an AI to produce a synthesis of a research paper, it could be organized in many different ways. Thus, you need to think about how you want your synthesis to be organized. Here's an example of an architecture for the synthesis of a research paper:

First paragraph: introduction

- Hook followed by a presentation of the big picture science of the paper

Second paragraph: summary of the paper's methods

Third paragraph: summary of the paper's results and implications

Fourth paragraph: conclusion

- Why do those results matter?

¹ Let's say that for this exercise, we have to synthesize an astrophysics research paper written by Dr. John Ruan and his student Olivier Gilbert called "A Host Galaxy Morphology Link Between Quasi-Periodic Eruptions and Tidal Disruption Events" [<https://arxiv.org/abs/2409.10486>]

Step #2: Define the evaluation criteria of the synthesis

The research synthesis now has a base organization. That's good! However, the AI could write the synthesis in that organization in many different ways. I have identified three categories of evaluation criteria that I think are important to think about the quality of the text generated by AI: **accuracy**, **structure**, and **style**.

In the following table, the first column indicates the specific criteria to be judged in the form of a question. The first, second, and third columns describe how the text fulfills the specific criterion by answering the question asked. The three columns correspond to poor, moderate and excellence performance. In some cases, the question results more in a binary-type answer, such that there is no middle column answer.

ACCURACY

Is there any wrong information in the synthesis?	The synthesis contains erroneous information	The synthesis contains no erroneous information, but some may be too loosely imprecise	The synthesis contains no information that is erroneous or too loosely imprecise
Is the synthesis complete in terms of information?	The synthesis is missing information that is critical to the message of the paper	The synthesis contains the most important information but lacks specific details, or contains unnecessary information	The synthesis communicates all the relevant information, and all the presented information is relevant

STRUCTURE

How is the hook?	The first sentence does not draw the reader in		The first sentence draws the reader in
How is the introduction?	The introduction does not introduce the scientific theme	The introduction presents the scientific theme of the paper, but lacks conveying a sense of importance	The introduction presents the scientific theme, and conveys a sense of relevance
How is the flow between paragraphs?	The text jumps abruptly between the paragraphs	The text goes from one idea to the other, but lacks globality	The text threads a coherent narrative from the beginning to the end, almost like a story
How is the conclusion?	The conclusion does not remind the reader of the importance of the paper		The conclusion reminds the reader of the importance of the paper

STYLE

How is the vocabulary?	The vocabulary uses jargon that is inappropriate for the target audience	The vocabulary uses words that are at times a bit too technical	The vocabulary uses no words that are too technical for the target audience
How are the scientific/abstract concepts presented?	The abstract concepts are presented as they are in the paper	The abstract concepts are presented with little modification, but lack adaptation to make them more understandable to the target audience	The abstract concepts are presented by using analogies, comparisons, parallels or metaphors when needed
How is the overall message of the paper communicated?	The message boils down to the science and lacks perspective	The message communicates the science and meaning of the paper, but the reader has trouble identifying it	The message communicates and highlights the overall importance of the paper through its science with clear takeaways

To define your evaluation criteria, you can select specific criteria that you seek in your synthesis from the table above or define your own! Naturally, because we want the AI to be successful at each criterion, we would instruct it to produce the synthesis at the ‘excellent’ performance level, which corresponds to the last column. If this is unclear, it should hopefully make more sense when comes the time to write the prompt. In the meantime, let’s say that we want the AI to write a synthesis that obeys all of the criteria mentioned above:

- The synthesis contains no information that is erroneous or too loosely imprecise
- The synthesis communicates all the relevant information, and all the presented information is relevant
- The first sentence draws the reader in
- The introduction presents the scientific theme, and conveys a sense of relevance
- The text threads a coherent narrative from the beginning to the end, almost like a story
- The conclusion reminds the reader of the importance of the paper
- The vocabulary uses no words that are too technical for the target audience
- The abstract concepts are presented by using analogies, comparisons, parallels or metaphors when needed
- The message communicates and highlights the overall importance of the paper through its science with clear takeaways

Step #3: Identify the input information for your prompt

Before we go ahead and write our prompt, we need to make sure that we have all of the information that the AI needs to write its synthesis. This notably includes:

- the PDF of the research paper, naturally;
- the target audience, which is critical, and will determine the technicality and turns of phrases that the AI will use;
- the length of the synthesis, which is an important constraint to give to the AI.

For our exercise here, let’s say that the target audience are university students, and that the length of the synthesis must be about 200 words.

Step #4: Write the prompt

Now that we have all of the information we need, let's put it all together to write the prompt we will give to the AI. You can add specifications on top of everything that has been mentioned so far, based on your specific requirements. For example, in our situation here, I want the synthesis to make the point that the research I will be communicating about has been conducted by researchers at Bishop's University, so I am specifying it. Those are details that you can also modify yourself after having generated the text.

Note that in the prompt, I am telling the AI to use the architecture from step #1, the criteria from step #2. I am also informing it of the relevant information I have defined in step #3.

Here is the AI prompt, which I am giving in this particular case to DeepSeek.

Input:

I would like you to write a synthesis of a research paper, which I am attaching as a PDF. The synthesis must be about 200 words in length, and its target audience is university students (it must accommodate students with no background in physics). Notably, while the authors of this paper have multiple affiliations from different institutions, I would like you to stress that this research was conducted by members of Bishop's University (in particular, by Dr. John Ruan).

I would like your synthesis to have the following structure:

First paragraph: introduction

- Hook followed by a presentation of the big picture science of the paper

Second paragraph: summary of the paper's methods

Third paragraph: summary of the paper's results and implications

Fourth paragraph: conclusion

- Why do those results matter?

I would also like for the synthesis to obey the following criteria:

- The synthesis contains no information that is erroneous or too loosely imprecise
- The synthesis communicates all the relevant information, and all the presented information is relevant
- The first sentence draws the reader in
- The introduction presents the scientific theme, and conveys a sense of relevance

- The text threads a coherent narrative from the beginning to the end, almost like a story
- The conclusion reminds the reader of the importance of the paper
- The vocabulary uses no words that are too technical for the target audience
- The abstract concepts are presented by using analogies, comparisons, parallels or metaphors when needed
- The message communicates and highlights the overall importance of the paper through its science with clear takeaways”

Output:

Unlocking Cosmic Mysteries: How Galaxy Shapes Link Star-Destroying Events

Imagine cosmic explosions so powerful they rip stars apart, yet their origins remain shrouded in mystery. A team led by **Dr. John Ruan** of Bishop’s University tackled this puzzle by exploring whether two rare phenomena—Quasi-Periodic Eruptions (QPEs, repeating X-ray flares) and Tidal Disruption Events (TDEs, where stars are torn apart by black holes)—are connected. Both occur in galaxy centers, but their relationship is unclear. This study investigates if their host galaxies share unique traits, offering clues to their origins.

The researchers analyzed images of 9 QPE and 13 TDE host galaxies from a large sky survey. They measured galaxy shapes, densities, and black hole sizes, comparing them to similar galaxies without such events. Think of it like studying neighborhoods where fireworks occur—do they share features like crowded streets or unusual layouts? Advanced software modeled galaxy structures to identify patterns invisible to the naked eye.

The results were striking: QPE and TDE hosts both have dense, compact centers and smaller black holes compared to typical galaxies. This similarity suggests a shared origin, such as frequent stellar collisions near underweight black holes. Like finding identical fingerprints at two crime scenes, these traits hint that QPEs and TDEs might be different outcomes of the same violent processes.

This work, spearheaded by Bishop’s University, reshapes our understanding of cosmic explosions. By linking QPEs and TDEs through galaxy morphology, it guides future searches for gravitational waves and illuminates how black holes interact with stars. For students, this study exemplifies how simple observations—like galaxy shapes—can unravel the universe’s most energetic secrets.

Step #5: Rewrite

Of course, it would be too beautiful if the raw AI output was perfect! In my reflection notes below, I explain more in depth my perception of what is AI less likely to do well. Typically, I like to highlight parts that I think are working or not working and start to correct from there. Vocabulary and jargon level is an easy and important thing to correct for (even though we asked the AI to use vocabulary that should be appropriate for the target audience, it usually still contains words or phrases that we might judge inadequate).

Here is my [quick] personal correction of the DeepSeek output above:

Our Universe is home to cosmic explosions so powerful they rip stars apart, yet their origins remain shrouded in mystery! Recently, Bishop's very own **Dr. John Ruan** tackled this puzzle by exploring whether two such phenomena—Quasi-Periodic Eruptions (QPEs, repeating X-ray flares) and Tidal Disruption Events (TDEs, where stars are torn apart by black holes)—are connected. Both occur in galaxy centers, but their relationship is unclear. Dr. Ruan and his student Olivier Gilbert investigated if their host galaxies share unique traits, offering clues to their origins.

They analyzed images of 9 QPE and 13 TDE host galaxies from a large sky survey. They measured galaxy shapes, densities, and black hole sizes, comparing them to similar galaxies without such events. Think of it like studying neighborhoods where fireworks occur—do they share features like crowded streets or unusual layouts?

The results were striking: QPE and TDE hosts both have dense, compact centers and smaller black holes compared to typical galaxies. This similarity suggests a shared origin, such as frequent stellar collisions near underweight black holes. Like finding identical fingerprints at two crime scenes, these traits hint that QPEs and TDEs might be different outcomes of the same violent processes!

The work of Dr. Ruan enabled to connect two mysterious types of explosions related to black holes, QPEs and TDEs. This is an important step towards uncovering the fundamental mechanisms of their origins, which in turn will allow us to better understand how elusive objects in the Universe such as black holes look and act the way we observe them today!

Reflection Notes + Tips

(1) The AI is good at processing factual information and summarizing it, even using clever and unexpected analogies. I like to see it as a sort of “rational” imagination. However, the performance of its writing style can vary, which I globally refer to as an “artistic” imagination. For example, a lot of the time, it used the same type of hook for many different research syntheses (they all started with “Imagine...” like “Imagine a planet far away...” or “Imagine a black hole far away...”) which can quickly get boring or come off as uninspired. Therefore, the AI is very useful in accessing the information content of a research paper that would take us a lot of time to process, but its text can definitely use some rewriting to make it more readable. Otherwise, it truly does seem like a machine wrote it! I expect that in the future, these traits will be less accentuated as LLMs become more sophisticated, but in any case, it is important to review and rewrite AI-generated text to make it more lively.

(2) In the same spirit, I strongly encourage you to proof-read your synthesis by the original author(s) or knowledgeable individuals if you cannot assess if the information processed by the AI was synthesized accurately (this might be the case if you are synthesizing a research paper in a field outside of your expertise). You can submit your synthesis to the author(s) and ask for feedback within a week (for example) and tell them that past that time, you will post it online (otherwise, they might take a long time to reply, so it is a good compromise to give them a time window).

(3) Again, if you are synthesizing research papers outside of your field, you might not know which papers in particular to highlight (that is, if you have to choose them instead of them being assigned to you). While that is an issue I have not thought through, it might be a good idea in such circumstances to contact the researchers directly and ask them what paper(s) they would like for you to highlight. Upon writing the synthesis, you can then ask them for feedback.

(4) LLMs like ChatGPT or DeepSeek typically have a “reasoning” option. If chosen, the LLM will use reinforcement learning. In essence, with reinforcement learning, LLMs will process their own output to see if it makes sense, and usually they will make it better. It avoids its output to be non-sense (commonly referred to as “AI hallucinations”). I would suggest trying it out, as you will also get to see the “reasoning” of the AI while processing your prompt.

(5) I encourage you to be critical of any “principles” I establish in this method. I am by no means an expert at this, and I am sure you have a different perspective and different ideas on how to realize what you want to realize. So, if my guide can be of help in your enterprise, I will be glad, but feel free to adapt it to your needs!

(6) While AI is a tool that can be tremendously useful, we should always keep in mind that it has an energetic consumption! This is why it is important to not mindlessly make a lot of requests to ChatGPT and other LLMs. If you need to use AI to perform a task, craft a good prompt that will require less follow-up prompts. Hopefully, this guide will give you ideas on how to make a good prompt and achieve a minimum number of requests.

(7) Finally, I just want to encourage you to adapt this method to whatever knowledge you are trying to synthesize. For example, by adapting the architecture of the synthesis in step #1, you can ask the AI to produce any kind of text. In the situation I gave here, it was a social media post, but one could have specified a script for a one-minute video! I also used the same method to generate short Instagram ads for talks given during Research Week (again, by adapting the architecture and evaluation criteria of my prompt). As long as you specify the format and any other idea you have in mind for the synthesis that you need, you can ask the AI to help you. Similarly, you can adjust/add specific criteria that you want the AI to obey in its writing (step #2). As always, be careful in using its output, re-read it carefully, and adapt it as you want. Remember that AI is a useful tool, but it is not a substitute for your brain or for your sense of creation. So, use it wisely!