

Highlighting Case Studies in LLM Literature Review of Interdisciplinary System Science

Lachlan McGinness^{1,2}[0000–0002–3231–4827], Peter
Baumgartner^{2,1}[0000–0002–6559–9654]
Esther Onyango³[0000–0003–0116–4923] and Zelalem Lema⁴[0000–0001–9275–8828]

¹ Australian National University, first.last@anu.edu.au

² Data61|CSIRO, first.last@data61.csiro.au

³ CSIRO, Agriculture and Food, Esther.Onyango@csiro.au

⁴ CSIRO, Agriculture and Food, Zelalem.Moti@csiro.au

Abstract. Large Language Models (LLMs) were used to assist four Commonwealth Scientific and Industrial Research Organisation (CSIRO) researchers to perform systematic literature reviews (SLR). We evaluate the performance of LLMs for SLR tasks in these case studies. In each, we explore the impact of changing parameters on the accuracy of LLM responses. The LLM was tasked with extracting evidence from chosen academic papers to answer specific research questions. We evaluate the models’ performance in faithfully reproducing quotes from the literature and subject experts were asked to assess the model performance in answering the research questions. We developed a semantic text highlighting tool to facilitate expert review of LLM responses.

We found that state of the art LLMs were able to reproduce quotes from texts with greater than 95% accuracy and answer research questions with an accuracy of approximately 83%. We use two methods to determine the correctness of LLM responses; expert review and the cosine similarity of transformer embeddings of LLM and expert answers. The correlation between these methods ranged from 0.48 to 0.77, providing evidence that the latter is a valid metric for measuring semantic similarity.

Keywords: Systematic Literature Review · Large Language Models · Highlighting.

1 Introduction

The scientific community is currently full of hype and hope for the use of Artificial Intelligence (AI) to accelerate research [28]. Messeri and Crockett claim that scientists are too trusting and lack awareness of the biases and errors of Large Language Models (LLMs) [26]. They present ‘AI as Oracle’ as a vision of the future where LLMs overcome the problem of too much literature to digest by efficiently searching and summarising information [26]. Many research groups are optimising and improving LLM tools for literature review [7,18,2] including the CSIRO’s (Commonwealth Scientific and Industrial Research Organisation’s) Science Digital program [10].

Many tools exist to enhance or automate literature review, including LitLLM [1], Scite [7], Elicit [18] and Scopus AI [2]. The techniques of these tools remain undisclosed as commercial secrets. However they all appear to use a combination of the same strategies: calling LLMs through APIs, prompt engineering (in-context learning), Retrieval Augmented Generation (RAG) and fine tuning [6]. The lack of transparency about these tools make it difficult to determine and accelerate best practice in AI systematic review methods.

Systematic Literature Review (SLR) was developed in the field of Evidence-Based Medicine as a method of reducing bias by sticking to strict protocols [6]. It has since been adopted in many disciplines including social science, education and environmental science [6]. SLR tasks include planning, searching, screening, extraction and synthesis[6]. In the extraction phase, desired information is extracted from a set of selected studies. Few studies have attempted to objectively measure the capability of LLMs for SLR [36,41,11,32,33] and even fewer for the extraction and screening phases. Previous studies found that LLMs are unreliable SLR tools as they ‘hallucinate’ references that do not exist [35].

In this paper we evaluate the performance of GPT-3.5 Turbo, and GPT-4 Turbo, one of the best models currently available, on SLR tasks. We will refer to these models as GPT3 and GPT4 respectively. We present four case studies where LLMs were used to assist CSIRO interdisciplinary systems science researchers, including the last two authors of this paper, in different stages of SLR. In each we systematically explore the impact of changing a parameter on the accuracy of LLM responses.

Although automatically checking LLM responses is highly desirable, currently there are no tools that can perfectly check the correctness or semantic similarity of texts. An important and tedious step of using LLMs in SLR is verifying their responses. To make this task easier we contribute a highlighting algorithm, in analogy to highlighting with a text marker on paper. It aims to provide a human reader with visual clues to quickly scan generated text. The algorithm is driven by a small set of user-supplied keywords provided by a domain expert. We describe the algorithm and report on experiments and experiences with application to our case studies.

The rest of this paper is structured as follows. Section 2 summarizes our methodology for application to the case studies chosen. This includes statistical evaluation methods and LLM techniques. Section 3 introduces an algorithm for explainable text relevance in terms of semantic similarity, and communicating it through text highlighting. Section 4 reports on experimental results, and Section 5 discusses these results and Section 6 summarises the conclusions.

2 Methodology

We explore four case studies with interdisciplinary system scientists. All four studies used Microsoft Azure endpoints to call either GPT3 and GPT4. The first case study focuses on the health impacts of agri-food transitions. In the second case study, the researcher had performed an SLR extracting the enablers

and constraints from twelve papers on coordinated responses to crises. They were interested to know if they had missed any key points when completing the SLR. The third case study verifies an SLR involving sixty papers on a sustainable transitions SLR task [27]. The final case study focuses on screening papers for an SLR on the use of generative AI for marking student responses to exam questions. The studies investigate the impact on overall performance when using different models, asking LLMs to provide evidence for their answers, and splitting tasks into several calls.

An example research question from the first case study is “What are the health outcomes of an agri-food transition”? The LLM was tasked with answering the question and finding evidence from an academic paper to support its answer. Evidence would be a quote from the paper such as “An overabundance of food supply alone has been identified as a key cause of the obesity epidemic” and the LLM answer could be “Obesity”.

2.1 Statistical and Evaluation Techniques

In this section we outline the statistical methods used to analyse the results. In cases where multiple similar data are available these results are summarised using mean (μ_x) and population standard deviation (σ_x) defined as usual.

In our analysis we use two methods to determine the correctness of LLM responses; expert review and automated similarity metrics of LLM and expert answers. The first automated similarity metric is SpaCy Semantic Similarity [14]. This method compares two strings the average embedding vector of the tokens in each of the strings is calculated and the cosine similarity is taken. The second method is the cosine similarity between transformer-based embeddings [30,38,19], which we will refer to as ‘transformer similarity’. Transformer based embeddings have the ability to take order of words into account.

To compare these two metrics we use Pearson correlation coefficient shown in Equation 1.

$$\text{Correlation} = r = \frac{\sum_{i=1}^n (\mu_x - x_i)(\mu_y - y_i)}{\sigma_x \sigma_y} \quad (1)$$

Uncertainty in correlation values was calculated using the Fisher transformation of correlation [13] with a 95% confidence interval ($Z = 1.96$) using Equation 2. In order to make these values more comparable to standard deviation they were divided by Z .

$$\text{Correlation Uncertainty} = \Delta r = \frac{1}{Z} \tanh \left(\operatorname{arctanh}(r) \pm \frac{Z}{\sqrt{n-3}} \right) \quad (2)$$

One of the SLR tasks involves the using an LLM to screen papers and determine their relevance. If a paper is relevant we refer to this as positive. A true positive (TP) result is when a model classifies a relevant paper (as determined by experts) as relevant. A false positive (FP) occurs when an irrelevant paper is classified as relevant. True negatives (TN) and false negatives (FN) are defined similarly. False positive and False Negative rates in the normal way.

2.2 LLM Techniques

We converted the research papers from PDF to text with the `pdftotext` utility⁵. Instead of using standard RAG techniques, we provided the entire paper to the LLM as the context windows were large enough. The research scientists were concerned by the biases of LLMs [5,31,4]. Here we outline these and our approach to minimising them.

Selection bias: The LLM might favour, for example, well known authors or recent papers. To avoid this type of bias, the research scientist selected the papers for the SLRs.

Inability to understand nuance: There are many studies which note that Large Language Models are effective for general tasks but can struggle with domain specific knowledge and the nuance of specialised tasks [9,24,21]. We were concerned that LLMs may lack the nuance to determine whether information relates to the study in the paper or referenced studies. To overcome this issue, researchers removed specific pages based on the task that the LLM needed to perform. For example, when asking about the geographic location of the paper, only the abstract and introduction were provided. We also addressed this issue by using in-context learning to provide the model with specific definitions provided by domain experts.

Lack of domain specific knowledge: The domain or even subject specific vocabulary in scientific literature poses challenges for LLM-based analysis. LLMs are trained on vast corpuses of which the subject specific matter is only a small portion. In agri-food transitions and co-ordinated responses to crises research literature terms have very specific meanings and LLMs would miss the nuances of these words. We mitigated this problem by providing the LLMs with definitions of subject specific vocabulary as part of the prompts.

Hallucinations: It is well known that LLMs can ‘make up’ information [17]. To manage this, tasks were broken into two steps. First, the Large Language Model was asked to find quotes that provide evidence of the desired information. Binary (yes/no) verification of the quotes failed because of trivial errors related to unicode characters. Therefore quotes were instead verified using fuzzy text matching using `thefuzz` implementation⁶ of the Levenshtein distance metric [20]. The LLM was then given the quotes and asked for a final answer.

3 Semantic Text Highlighting

In this section we introduce a method for semantic text highlighting, or **highlighting** for short. The idea of semantic highlighting is not new and was originally formulated for information retrieval [16]. In our context we apply highlighting to LLM retrieved evidence. Because the amount of retrieved evidence can still be

⁵ <https://www.xpdfreader.com/pdftotext-man.html>

⁶ <https://github.com/seatgeek/thefuzz>

overwhelming, further support is needed to aid a human reviewer. This is where highlighting comes in.

Like with a text marker on paper, highlighting does not need to be perfect to be useful. However, the highlighted words should be semantically related to a given specific SLR research question. Our method requires a set of keywords representing the research question. Highlighting then boils down to determining which words in a given sentence are semantically related to the keywords and conveying the findings in a useful way. For that, we rely on readily available information-theoretic similarity measures and a carefully curated corpus of English, which requires no training and is explainable.

Our method works as follows. The SLR researcher provides sets of keywords comprising of entities E , relations R and properties P as nouns, (possibly transitive) verbs and adjectives or adverbs, respectively. For example, in our first SLR case study on the topic influences of agri-food transitions on health outcomes the following were used:

Entities E : health, disease, outcome, food, lifestyle

Relations R : explain, affect, improve, stimulate

Properties P : environmental

A word w in a given text is highlighted if it is deemed related to a keyword according to the procedure $\text{Similarity}(w, C, t)$ in Algorithm 1, where $C = E \cup R \cup P$. The given text is first parsed with the part-of-speech (POS) parser SpaCy [15]. It assigns a grammatic role t to every word w for determining the best suited subset of C for similarity. The algorithm computes two similarity scores between w and that subset in the range $[0, 1]$. One of them is word vector similarity, as readily provided by spaCy, and the other is Wu-Palmer similarity [40,34] based on hypernym-reachability in WordNet [12].

Using WordNet for word similarity is an established and well-researched topic [23]. In our algorithm, the search for similar words is broad and includes (one-step) synonyms, pertainyms, and related derived forms, weighted for each category. This was a design choice motivated by the case study in Section 4.2, where the researcher wants to ensure they did not miss any key points in their manual review. Similarity takes scores higher than given thresholds to determine whether w should be highlighted or not. We found that Wu-Palmer similarity often produces results more similar to what a human user expects from a highlighter. Hence, vector similarity acts only as a fall-back. Here are some examples for highlighted evidence text with the keywords above.

1. It is also likely that **climate change** will **contribute** to **novel occurrences** of **disease emergence** and transmission.
2. **Foodborne illnesses** significantly **influence** individuals **nutritional status**.
3. **Changing lifestyles**, mainly due to **work commitment**, **have fuelled** the **increase** in numbers eating out and the need for **convenience foods**.
4. **Significant changes** **have occurred** in **food systems** in the last decades that **have contributed** to widen **such 'holes'** in the barriers from phase to phase: agricultural intensification and industrialization **causing** major environmental deterioration, the

Algorithm 1 Similarity

Similarity(w, C, t)**Input:** w word, C keywords in canonical form (lemmas), t type of w (noun, verb, adjective ...)**Output:** Similarity score for w P-Weight \leftarrow 0.95 RF-Weight \leftarrow 0.95 WUP-Threshold \leftarrow 0.8 VEC-Threshold \leftarrow 0.95best_{wup} \leftarrow max _{$c \in C$} WUP-X(w, c, t) **or else** 0.0best_{vec} \leftarrow max _{$c \in C$} VEC(w, c) **or else** 0.0**if** best_{wup} \geq WUP-Threshold **and** best_{wup} \geq best_{vec} **then return** best_{wup}**elif** best_{vec} \geq VEC-Threshold **then return** best_{vec}**else return** 0.0WUP-X(w, c, t) // Extended Wu-Palmer similarity, considers reachable words from w and c $S_w = \text{Extend}(\text{WN-Synsets}(w, t))$ // WN-Synsets returns synonyms of w $S_c = \text{Extend}(\text{WN-Synsets}(c, t))$ **return** max _{$(s_w \xrightarrow{\omega_w} r_w, s_c \xrightarrow{\omega_c} r_c) \in S_w \times S_c$} $\omega_w \cdot \omega_c \cdot \text{WN-WUP}(r_w, r_c, t)$ // Wu-Palmer from WordNetVEC(w, c) // Vector similarity**if** $w = c$ **then return** 1.0**elif** both w and c have vector embeddings **then** **return** cosine-similarity of the embeddings of w and of c **else return** 0.0Extend(S, t)**Input:** S a set of WordNet synsets**Output:** Weighted extension of S by pertainyms and derivationally related forms $R \leftarrow \emptyset$ // Result relation**for** $s \in S$ **do** $R \leftarrow R \cup \{s \xrightarrow{1.0} s\}$ // R is reflexive **for** $l \in \text{WN-Lemmas}(s)$ **do** $R \leftarrow R \cup \{s \xrightarrow{\text{P-Weight}} \text{WN-SynSet}(v) \mid v \in \text{WN-Pertainyms}(l)\}$ $R \leftarrow R \cup \{s \xrightarrow{\text{RF-Weight}} \text{WN-SynSet}(v) \mid v \in \text{WN-RelatedForms}(l)\}$ **return** R

increasing distance traveled by food in global markets, and the nutrition transition towards diets rich in ultra - processed food and animal protein are the three cornerstones of such changes.

Entities (nouns, noun chunks) are colored red and relations (verbs) are colored blue. Additional colors are used for supporting words according to their grammatical roles. Properties (adjectives) of colored entities are purple.

For each word, the algorithm can provide an explanation of why it is highlighted. These explanations are helpful for customising parameter settings; for example, we get:

2. Foodborne illnesses (NCP(Foodborne illnesses, [SimilarTo('disease', 0.95, 'wup')])) significantly influence (SimilarTo('affect', 0.84, 'wup')) individuals nutritional status (NCP(nutritional status, [SimilarTo('food', 0.91, 'wup')])).

In these annotations, NCP means 'NounChunkPart', and the similarity of the highlighted word(s) to keyword(s) is indicated as in SimilarTo(keyword, similarity, kind), where 'wup' is Wu-Palmer similarity.

Highlighting can be useful beyond marking up text excerpts. In one of our case studies below we take the ‘highlighting rate’ as a statistical measure to assess the similarity between the LLM’s response and the researcher’s benchmark evaluation.

4 Results

4.1 Case Study 1: Similarity Metrics in Agri-Food Transition SLR

In this first case study we determine the accuracy of the chosen metrics and explore if there is a significant difference in performance between GPT3 and GPT4. There were eight tasks (as shown in the first column of Table 1) and ten papers. GPT3 and GPT4 were given identical instructions to answer each task in separate calls. We compared the researcher’s and models’ answers.

Table 1. Case Study 1: Average and standard deviation for similarity scores, fuzzy matching scores, average word count and expert judged model accuracy are presented for GPT3 and GPT4.

Information	Complexity	Model	Quote Fuzzy Matching Score	Model Average Word Count	Expert Average Word Count	Transformer Similarity	SpaCy Similarity	Model Accuracy
Global Context	Low	GPT-3	97 ± 5	3.5	4.8	0.74 ± 0.25	0.58 ± 0.16	0.9
		GPT-4	98 ± 6	5.2	4.8	0.84 ± 0.06	0.57 ± 0.14	1.0
Associated Health Focus	Low	GPT-3	97 ± 5	5.6	1.6	0.82 ± 0.13	0.57 ± 0.16	0.75
		GPT-4	98 ± 4	7	1.6	0.85 ± 0.05	0.60 ± 0.14	0.94
Transition Pathway	Moderate	GPT-3	95 ± 10	13.4	5.4	0.81 ± 0.04	0.66 ± 0.08	0.65
		GPT-4	97 ± 8	23	5.4	0.85 ± 0.06	0.65 ± 0.21	1.0
Agri-food Boundary	Moderate	GPT-3	97 ± 4	32	17	0.83 ± 0.04	0.77 ± 0.14	0.5
		GPT-4	98 ± 6	50.6	17	0.87 ± 0.03	0.79 ± 0.12	0.85
Public Health Risk	Moderate	GPT-3	99 ± 7	8.8	6.5	0.85 ± 0.05	0.59 ± 0.16	0.7
		GPT-4	97 ± 6	20.5	6.5	0.87 ± 0.06	0.74 ± 0.17	0.95
Synergies	High	GPT-3	97 ± 5	31.3	26.4	0.83 ± 0.03	0.84 ± 0.07	0.25
		GPT-4	98 ± 5	58	26.4	0.81 ± 0.05	0.83 ± 0.07	0.1
Constraints	High	GPT-3	97 ± 5	35	18	0.82 ± 0.02	0.81 ± 0.08	0.44
		GPT-4	98 ± 5	59	18	0.84 ± 0.02	0.83 ± 0.07	1.0
Integrated Solutions	High	GPT-3	97 ± 4	28	30	0.89 ± 0.04	0.90 ± 0.07	0.88
		GPT-4	99 ± 3	50	30	0.89 ± 0.05	0.89 ± 0.07	1.0

The models were asked to record three quotes (evidences), then give a final answer in a second call. Cases where the average fuzzy string similarity was less than 90% were manually investigated. There were 25 and 38 cases where this occurred for GPT4 and GPT3 respectively, resulting in overall error rates of 2% and 5%.

The semantic similarity between LLM and expert answers was calculated using transformer similarity and SpaCy similarity. An example of an LLM/expert answer could be “The Global context is Africa” which has a word count of 5.

The correlation between the Transformer similarity score and the expert’s judgment of the model answers was 0.48 ± 0.09 and there was almost no correlation (-0.07 ± 0.08) between the SpaCy similarity and expert judgment.

4.2 Case Study 2: Impact of Evidence on Coordinated Response to Crisis SLR

In this case study we compare LLM output based on two methods. In the first method (‘evidence’) the LLM first obtains quotes to support its answer to the question and then writes it’s answer. In the second method (‘direct’) the LLM writes an answer without searching for or providing evidence.

Providing evidence is a good way to increase the trustworthiness of LLM responses. However it will increase the number of completion tokens and therefore cost. We test our highlighting algorithm as an automated similarity metric by calculating the correlation between the highlighted fraction of each expert answer and model answers.

Table 2 shows that the evidence method results in slightly lower SpaCy Semantic Score, vector embedding cosine similarity, human judged accuracy and highlighting correlation. The highlighting correlation score changes by a more significant margin, but also has a greater uncertainty than the other measures.

Table 2. Case Study 2: Prompt tokens, completion tokens, SpaCy similarity, transformer similarity, human judged accuracy and Highlighting Correlation are presented for comparison of direct and evidence-based conditions.

Experimental Condition	Prompt Tokens ($\times 10^3$)	Completion Tokens	SpaCy Similarity	Transformer Similarity	Human Judged Accuracy	Highlighting Correlation
Evidence	20.2 ± 5.7	842 ± 397	0.85 ± 0.06	0.87 ± 0.06	69%	-0.18 ± 0.24
Direct	20.0 ± 5.7	213 ± 64	0.88 ± 0.06	0.90 ± 0.04	72%	0.13 ± 0.25

The number of prompt tokens is not significantly changed by asking the model for evidence, but the number of completion tokens increases fourfold. As the majority of tokens are prompt tokens, it might be expected that the number of completion tokens would have a small impact on the overall cost. However the computational cost of running a transformer in this case is proportional to the number of completion tokens. A technique to more effectively reduce the cost using an LLM for literature review is grouping multiple tasks into a single call as explored in the next case study.

4.3 Case Study 3: Impact of Combining Tasks on Sustainable Transitions SLR

In this case study, we compare conditions that we name ‘separate’ and ‘together’. For the separate condition, GPT4 is called ten times. In each call the paper is provided and the LLM is asked to find an answer and evidence for a research

question. In the together condition, the model is given all tasks in one call. The results for both experimental conditions are in Table 3. Note that papers were classified into three groups according to their geographical scale as requested by the researcher. There were twenty papers in each category.

Table 3. Case Study 3: Quote and final answer metrics for both conditions. The expert found significant errors in the first paper from the global scale responses for the together condition. They decided that the together method was not worth pursuing and did not evaluate the remaining together responses.

Experimental Condition	Scale of paper	Prompt Tokens $\times 10^5$	Completion Tokens $\times 10^3$	Fuzzy Text Matching	SpaCy Similarity	Transformer Similarity	Instances of failing to find quotes	Expert Judged Accuracy
Separate	Global	1.1 ± 0.2	1.5 ± 0.2	98.4 ± 2.5	0.84 ± 0.4	0.85 ± 0.01	0	82%
	International/national	4.7 ± 1.7	1.3 ± 0.1	95.3 ± 6.8	0.79 ± 0.12	0.86 ± 0.01	5	84%
	Subnational/Local	2.2 ± 0.7	1.4 ± 0.2	99.4 ± 0.6	0.83 ± 0.03	0.84 ± 0.06	2	75%
Together	Global	0.12 ± 0.02	0.97 ± 0.37	99.5 ± 0.4	0.79 ± 0.5	0.83 ± 0.03	0	N/A
	International/national	0.48 ± 0.16	1.3 ± 0.2	93.1 ± 7.2	0.77 ± 0.16	0.84 ± 0.01	9	N/A
	Subnational/Local	0.22 ± 0.07	0.8 ± 0.2	98.0 ± 1.0	0.81 ± 0.03	0.84 ± 0.01	1	N/A

The researcher did not like the results of the together condition because GPT4 jumbled its responses resulting in a frame-shift. Despite the near tenfold increase in prompt tokens and cost, they preferred the separate condition. In the separate condition the model had no awareness of its answers to the other questions resulting in overlap of the ideas presented for each sub-question.

4.4 Case Study 4: Prompt Variations for Automated Screening

In the final case study the researcher wanted to automatically extract information from papers for the purpose of screening. The researcher had manually reviewed 14 papers from the arXiv and 20 from Pubmed of which 12 and 3 were relevant respectively. These were used as datasets to measure the performance of LLM information extraction. These datasets provide opposite extremes, one where the LLM needs to accept nearly all of the papers as relevant and another where it needs to reject nearly all the papers.

This allowed for a study in simultaneously avoiding false positives and false negatives, see Table 4. Three prompts were used: ‘prompt relevant’ which indicated that the paper was likely to be relevant, ‘prompt irrelevant’, and a ‘neutral prompt’ which makes no indication of the paper’s relevance.

The false positive and false negative rates in Table 4 show the LLM was likely to state that a paper was relevant when in fact it was not. The correlations show

Table 4. Case Study 4: SpaCy and transformer similarity, expert verified accuracy, false positive and false negative rates for three different prompts on both datasets. The correlation between expert accuracy and the similarities were calculated to measure the quality of these metrics.

Condition	Dataset	SpaCy Similarity	Transformer Similarity	Expert Accuracy	Spacy-Expert Correlation	Transformer-Expert Correlation	False Positive Rate	False Negative Rate
Prompt Relevant	arXiv	0.87 ± 0.08	0.94 ± 0.05	0.83	0.64 ± 0.08	0.77 ± 0.06	1.00	0.00
	PubMed	0.70 ± 0.17	0.90 ± 0.07	0.86	0.23 ± 0.28	0.69 ± 0.21	0.94	0.00
Prompt Neutral	arXiv	0.88 ± 0.7	0.94 ± 0.04	0.82	0.66 ± 0.09	0.66 ± 0.09	1.00	0.08
	PubMed	0.68 ± 0.17	0.90 ± 0.08	0.83	-0.10 ± 0.25	0.44 ± 0.27	0.88	0.00
Prompt Irrelevant	arXiv	0.87 ± 0.08	0.94 ± 0.04	0.85	0.55 ± 0.1	0.67 ± 0.08	0.00	0.00
	PubMed	0.70 ± 0.15	0.89 ± 0.08	0.83	-0.05 ± 0.26	0.44 ± 0.27	0.12	0.00

transformer similarity correlated better with expert review than SpaCy similarity. It was found that SpaCy similarity is heavily impacted by changes in capitalisation. In some cases changes in capitalisation alone reduced the similarity to 0.3.

5 Discussion

In the first case study, GPT4 made less mistakes than GPT3 when extracting exact quotes from documents. This may be because of the quality of the model or GPT3’s shorter maximum context length. The longest context length available for GPT3 models was 16,000 tokens which was not always sufficient to read an entire paper. When the paper was broken into multiple sections the model was less likely to find relevant quotes. We noticed that GPT3 was more likely to select quotes from the beginning of the context window than GPT4, whose quotes were more evenly distributed from the entire paper.

Across all studies it was found that transformer similarity correlates more strongly with expert opinion than than SpaCy similarity. This indicates that transformer embeddings are a better metric. This is expected given that it is able to take the positions of words into account. The correlation between amounts of text highlighted increased with increased human judged accuracy, showing the correct trend as a measure of semantic similarity. However the uncertainty values were very high and more work would need to be performed to determine if this is a suitable metric.

The highlighting tool’s primary purpose is to aid a researcher in sifting through evidence and other LLM response text. The researchers anecdotally confirmed its value.

We originally anticipated that SpaCy’s similarity would be a better measure of semantic similarity than transformer similarity because the transformer similarity scores only ranged between 0.7 and 1, while SpaCy similarities scores had much larger ranges. Contrary to our original expectations, transformer similarity ubiquitously correlated more strongly with expert opinion. To make transformer similarity more interpretable for humans we recommend scaling these values be-

fore interpretation. This is because a transformer similarity of 0.8 is actually low, despite normal human expectations.

The second case study found that when a model was asked to provide evidence for its claims, it was slightly less accurate on all metrics including expert judgement. This is an unexpected result as normally ‘chain of thought’ or asking a model to explain its reasoning improves performance [39,25]. As trustworthiness is increased when the model provides verifiable evidence for its answers, this result indicates that there is an unfortunate trade off between accuracy and trustworthiness. The reason could be that the model is ‘overloaded’ when it needs to focus on multiple tasks at once. This was confirmed by the third case study (Section 4.3) where model performance also decreased when the number of tasks it was asked to complete in a single call increased.

In third third case study, as expected the number of prompt tokens is approximately ten times higher for the separate condition compared to the together condition. As far as the researcher was concerned the reduced number of errors in the results was worth the the extra cost. The major issue with the together condition was the possibility for tasks to be jumbled. One area for future work is to apply more advanced parsing techniques to avoid frameshift errors and therefore make the computationally cheaper technique more desirable for researchers. Another area for future work is to provide the model with specific keywords for each call to guide the LLM responses.

The fourth case study demonstrated that SpaCy semantic similarity can be heavily impacted by the capitalisation of words for medium sized models. This is not a desirable property for a system measuring semantic similarity; we argue that writing a sentence in all capital letters makes little change to the semantic meaning.

Readers may be tempted to think that scores such as SpaCy or transformer similarity could be used as better alternatives to subject expert review as they do not contain human biases. However, one needs to be careful in assuming that there is no bias when using metrics like these. There can be an illusion of objectivity, when in fact these models have been trained and validated on data which contains significant unknown biases. In this study, we highly value the opinion of experts who are part of the active research community and have observed that they demonstrate a strong awareness of their own biases.

False positive rates were consistently higher than false negative rates; GPT4 was more likely to think that an irrelevant paper was in fact relevant and most accurately screened papers when prompted to expect that papers may be irrelevant.

Overall GPT3 and GPT4 were able to find and correctly reproduce quotes from a text with 95% and 98% accuracy respectively. For low complexity tasks like finding the title or location of a paper, GPT4 performed with close to 100% accuracy, but accuracy was lower for more nuanced tasks such as identifying enablers in agri-food transitions. The overall approximate average accuracy of GPT4 in answering research questions was 83%.

Our highlighting workflow requires keyword calibration to determine a suitable set of keywords for each research question in an SLR. Currently, keyword calibration is semi-automated process with an expert in the loop. The expert proposes keywords, runs highlighting experiments on sample papers and adjusts the set of keywords according to their observations.

From our experiments we were able to derive some guidelines for calibration. A keywords set yielding a highlighting rate of around 0.4 ± 0.1 on evidence texts often seems to be a good compromise. If it is much higher, often too many irrelevant words are highlighted. If much lower, the domain has not been covered sufficiently and relevant keywords are missing. It is better to use evidence text than expert answers for keyword calibration as they are almost always proper English sentences. Expert assessments however can vary widely and sometimes are just lists of keywords. Hence the proportion of highlighted words is less reliable in this case, resulting in a less meaningful hit rate.

If the hit rate is unusually high this could be because of denser writing (less filler words) or because of denser information. The latter includes the possibility that surprising, additional insights have been unveiled. This helps to get a more complete picture of the problem. Conversely, we found that a much lower highlighting rate typically applies to irrelevant texts.

6 Conclusion

Large Language Models (LLMs) were used to assist interdisciplinary system scientists to conduct four Systematic Literature Reviews (SLR). The topics of the reviews were agri-food system transitions, coordinated responses to crises, sustainable transitions and automated marking. GPT-3.5 Turbo and GPT-4 Turbo had error rates of 5% and 2% when extracting exact quotes from research papers. Levenshtein distance accurately determined the faithfulness of quotes produced by LLMs and was robust to unexpected unicode characters and hyphenations.

When GPT-4 Turbo completed multiple tasks in a single prompt the number of prompt tokens decreased tenfold but with significant losses in accuracy in some cases due to frameshift errors. One area for future work would be to use more advanced parsing techniques to avoid frameshift or other ‘jumbling’ errors.

The accuracy of the models’ answers was found to decrease with complexity of the task. For very simple tasks, expert rating of LLM answer correctness was close to 100% while for highly nuanced tasks it could be as low as 10%. On average it was found that GPT-4 Turbo was able to extract information from papers with approximately 83% accuracy. When screening papers it was found that GPT-4 Turbo was more likely to include irrelevant papers than exclude relevant papers. One area for future work is to provide the model with specific keywords to focus its answers, this may help a model focus on the desired ideas while trying to complete a nuanced task.

It was found that taking the cosine similarity of transformer embeddings of expert and LLM answers was a measure of accuracy that correlated more strongly with expert opinion than SpaCy’s semantic similarity score. Nearly

all of these transformer embedding cosine scores were in the range of 0.7 to 0.95. In order to make these cosine similarities more human interpretable, we recommend scaling them to take up the full range between 0 and 1. An area of future work would be to use cosine similarity of transformer embeddings for sentence-wise comparison of researcher and LLM answers in order to determine if any important pieces of information are missing.

Although highlighting is designed to assist researchers with manual checking of answers, correlation between amounts of highlighted text is showing some promise as an automated method for measuring the quality of LLM responses. Additional research would need to be conducted to see if this can be used as a valid similarity metric. Another idea for future work is to re-formulate the semantic similarity algorithm (Algorithm 1) with probabilistic logic programming. This would allow for a more flexible and expressive framework. In addition, weight parameters could be rephrased as probabilities and be learned from examples by maximum likelihood estimation.

Acknowledgements. This research was supported by funding from CSIRO Data61 and its Valuing Sustainability Future Science Platform initiative. We thank Enayat Moallemi for providing knowledge, expertise and data for the third case study. We thank Stephen Wan and Shima Khanehzar for helpful discussions.

References

1. Agarwal, S., Laradji, I.H., Charlin, L., Pal, C.: LitLLM: A Toolkit for Scientific Literature Review. 2024. <https://arxiv.org/abs/2403.08399>
2. Aguilera-Cora, E., Lopezosa, C., Fernández-Cavia, J., Codina, L.: Accelerating research processes with Scopus AI: A place branding case study. *Revista Panamericana de Comunicación* 6(1) (2024). <https://doi.org/10.21555/rpc.v6i1.3088>
3. Antu, S.A., Chen, H., Richards, C.K.: Using LLM (Large Language Model) to Improve Efficiency in Literature Review for Undergraduate Research. *WS on Empowering Education with LLMs* (2023)
4. Bender, E.M., Gebru, T., McMillan-Major, A., Shmitchell, S.: On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In: *Proc. FAccT '21, ACM* (2021). <https://doi.org/10.1145/3442188.3445922>
5. Blodgett, S.L., Barocas, S., Daumé III, H., Wallach, H.: Language (Technology) is Power: A Critical Survey of "Bias" in NLP. In: *Proc. 58th Annual Meeting of the ACL. ACL* (2020). <https://doi.org/10.18653/v1/2020.acl-main.485>
6. Bolanos, F., Salatino, A., Osborne, F., Motta, E.: Artificial Intelligence for Literature Reviews: Opportunities and Challenges. *arXiv:2402.08565 [cs]* (2024)
7. Brody, S.: Scite. *Journal of the Medical Library Association* 109(4), 707–710 (2021). <https://doi.org/10.5195/jmla.2021.1331>
8. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language Models are Few-Shot Learners. In: *Proc. NeurIPS* 2020.

9. Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y., Ye, W., Zhang, Y., Chang, Y., Yu, P.S., Yang, Q., Xie, X.: A Survey on Evaluation of Large Language Models. *ACM Trans. Intell. Syst. Technol.* 15(3) (2024). <https://doi.org/10.1145/3641289>
10. Trust in CINTEL. Collaborative Intelligence Future Science Platform (2024)
11. De Silva, A., Wijekoon, J.L., Liyanarachchi, R., Panchendrarajan, R., Rajapaksha, W.: AI Insights: A Case Study on Utilizing ChatGPT Intelligence for Research Paper Analysis. arXiv:2403.03293 [cs] (2024)
12. Fellbaum, C.: *WordNet: An electronic lexical database*. Bradford Books (1998)
13. Fisher, R.A.: Frequency Distribution of the Values of the Correlation Coefficient in Samples from an Indefinitely Large Population. *Biometrika* 10(4), 507–521 (1915). <https://doi.org/10.2307/2331838>
14. Honnibal, M., Montani, I.: *Linguistic Features · spaCy Usage Documentation*. <https://spacy.io/usage/linguistic-features>
15. Honnibal, M., Montani, I.: *spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing* (2017)
16. Hussam, A., Ford, B., Hyde, J., Merayyan, A., Plummer, B., Anderson, T.: Semantic highlighting. In: *CHI 98 Conference Summary on Human Factors in Computing Systems. CHI '98, ACM* (1998). <https://doi.org/10.1145/286498.286667>
17. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A., Fung, P.: Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* 55(12) (2023). <https://doi.org/10.1145/3571730>
18. Kung, J.: Elicit (product review). *Journal of the Canadian Health Libraries Association / Journal de l'Association des bibliothèques de la santé du Canada* 44(1) (2023). <https://doi.org/10.29173/jchla29657>
19. Laskar, M.T.R., Huang, J.X., Hoque, E.: Contextualized Embeddings based Transformer Encoder for Sentence Similarity Modeling in Answer Selection Task. In: *Proc. Twelfth LREC. ELRA* (2020)
20. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady* 10, 707 (1966)
21. Li, T.O., Zong, W., Wang, Y., Tian, H., Wang, Y., Cheung, S.C., Kramer, J.: Nuances are the Key: Unlocking ChatGPT to Find Failure-Inducing Tests with Differential Prompting. In: *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pp. 14–26. (2023). <https://doi.org/10.1109/ASE56229.2023.00089>
22. Li, Y., Chen, L., Liu, A., Yu, K., Wen, L.: ChatCite: LLM Agent with Human Workflow Guidance for Comparative Literature Summary. arXiv:2403.02574 [cs] (2024)
23. Manna, S., Mendis, B.S.U.: Fuzzy word similarity: A semantic approach using WordNet. In: *International Conference on Fuzzy Systems*. pp. 1–8. IEEE, Barcelona, Spain (2010). <https://doi.org/10.1109/FUZZY.2010.5584785>
24. Meng, X., Yan, X., Zhang, K., Liu, D., Cui, X., Yang, Y., Zhang, M., Cao, C., Wang, J., Wang, X., Gao, J., Wang, Y.G.S., Ji, J., Qiu, Z., Li, M., Qian, C., Guo, T., Ma, S., Wang, Z., Guo, Z., Lei, Y., Shao, C., Wang, W., Fan, H., Tang, Y.D.: The application of large language models in medicine: A scoping review. *iScience* 27(5) (2024). <https://doi.org/10.1016/j.isci.2024.109713>
25. McGinness, L., Baumgartner, P.: Automated Theorem Provers Help Improve Large Language Model Reasoning. In: Bjørner, N., Heule, M., Voronkov, A. (eds.) *Proceedings of 25th Conference on Logic for Programming, Artificial Intelligence and Reasoning. EPiC Series in Computing*, vol. 100, pp. 51–69. EasyChair (2024). <https://doi.org/10.29007/2n9m>

26. Messeri, L., Crockett, M.J.: Artificial intelligence and illusions of understanding in scientific research. *Nature* 627(8002), 49–58 (2024). <https://doi.org/10.1038/s41586-024-07146-0>
27. Moallemi, E., Miller, M., Szetey, K., Chakori, S., Palmer, J., Battaglia, M., Bryan, B.A., Gao, L., Hall, A., Leith, P., Raven, R., Reed, P.M.: Entry points for accelerating transitions towards a more sustainable future (2024). *EarthArXiv pre-print*, <https://doi.org/10.31223/X5C68X>
28. Nature Editorial: Why scientists trust AI too much — and what to do about it. *Nature* 627(8003), 243–243 (2024). <https://doi.org/10.1038/d41586-024-00639-y>
29. OpenAI: GPT-4 Technical Report (2023)
30. Ormerod, M., Martínez del Rincón, J., Devereux, B.: Predicting Semantic Similarity Between Clinical Sentence Pairs Using Transformer Models: Evaluation and Representational Analysis. *JMIR Medical Informatics* 9(5), e23099 (2021).
31. Rogers, A.: Changing the World by Changing the Data. In: *Proc. 59th Annual Meeting of the ACL and the 11th IJCNLP*. ACL (2021). <https://doi.org/10.18653/v1/2021.acl-long.170>
32. Sami, A.M., Rasheed, Z., Kemell, K.-K., Waseem, M., Kilamo, T., Saari, M., Duc, A.N., Systä, K., Abrahamsson, P.: System for systematic literature review using multiple AI agents: Concept and an empirical evaluation. 2024. <https://arxiv.org/abs/2403.08399>
33. Shaib, C., Li, M., Joseph, S., Marshall, I., Li, J.J., Wallace, B.: Summarizing, Simplifying, and Synthesizing Medical Evidence using GPT-3 (with Varying Success). In: *Proc. 61st Annual Meeting of the ACL (Short Papers)*. ACL (2023).
34. Shenoy, M.: A New Similarity measure for taxonomy based on edge counting. *International journal of Web & Semantic Technology* 3(4), 23–30 (2012). <https://doi.org/10.5121/ijwest.2012.3403>
35. Smith, L.C.: Reviews and Reviewing: Approaches to Research Synthesis. An Annual Review of Information Science and Technology (ARIST) paper. *Journal of the ASIS&T* 75(3), 245–267 (2024). <https://doi.org/10.1002/asi.24851>
36. Spillias, S., Tuohy, P., Andreotta, M., Annand-Jones, R., Boschetti, F., Cvitanovic, C., Duggan, J., Fulton, E., Karcher, D., Paris, C., Shellock, R., Trebilco, R.: Human-AI Collaboration to Identify Literature for Evidence Synthesis. *Research Square* (2023). <https://doi.org/10.21203/rs.3.rs-3099291/v1>
37. de la Torre-López, J., Ramírez, A., Romero, J.R.: Artificial intelligence to automate the systematic review of scientific literature. *Computing* 105(10), 2171–2194 (2023). <https://doi.org/10.1007/s00607-023-01181-x>
38. Turton, J., Smith, R.E., Vinson, D.: Deriving Contextualised Semantic Features from BERT (and Other Transformer Model) Embeddings. In: *Proc. RepL4NLP-2021*. ACL (2021). <https://doi.org/10.18653/v1/2021.repl4nlp-1.26>
39. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, b., Xia, F., Chi, E.H., Le, Q.V., Zhou, D.: Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In: *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)* (2022)
40. Wu, Z., Palmer, M.: Verb Semantics and Lexical Selection. In: *32nd Annual Meeting of the ACL*. ACL (1994). <https://doi.org/10.3115/981732.981751>
41. Ye, A., Maiti, A., Schmidt, M., Pedersen, S.J.: A Hybrid Semi-Automated Workflow for Systematic and Literature Review Processes with Large Language Model Analysis. *Future Internet* 16(5), 167 (2024). <https://doi.org/10.3390/fi16050167>
42. Zou, Y., Liu, H., Gui, T., Wang, J., Zhang, Q., Tang, M., Li, H., Wang, D.: Divide and Conquer: Text Semantic Matching with Disentangled Keywords and Intents. In: *Findings of the ACL*. ACL (2022).