

# On the Effectiveness of Large Language Models in Automating Categorization of Scientific Texts

Gautam Kishore Shahi<sup>a</sup> and Oliver Hummel<sup>b</sup>

*University of Applied Sciences, Mannheim, Germany*  
{g.shahi, o.hummel}@hs-mannheim.de

**Keywords:** Large Language Models, Field of Research Classification, Prompt Engineering, Scholarly Publications.

**Abstract:** The rapid advancement of Large Language Models (LLMs) has led to a multitude of application opportunities. One traditional task for Information Retrieval systems is the summarization and classification of texts, both of which are important for supporting humans in navigating large literature bodies as they e.g. exist with scientific publications. Due to this rapidly growing body of scientific knowledge, recent research has been aiming at building research information systems that not only offer traditional keyword search capabilities, but also novel features such as the automatic detection of research areas that are present at knowledge-intensive organizations in academia and industry. To facilitate this idea, we present the results obtained from evaluating a variety of LLMs in their ability to sort scientific publications into hierarchical classifications systems. Using the FORC dataset as ground truth data, we have found that recent LLMs (such as Meta's Llama 3.1) are able to reach an accuracy of up to 0.82, which is up to 0.08 better than traditional BERT models.


## 1 INTRODUCTION


The amount of scholarly texts is consistently increasing; around 2.5 million research articles are published yearly (Rabby et al., 2024). Due to this enormous increase, the classification of (scientific) texts has been attracting even more attention in recent years (Bornmann et al., 2021). Classifying the research area of scientific texts requires significant domain knowledge in various complex research fields. Hence, manual classification is challenging and time-consuming for librarians and limits the number of texts that can be classified manually (Zhang et al., 2023). Moreover, due to complex hierarchical classification schemes and their existing variety, classification of publications is also an unbeloved activity for researchers. Prominent examples of classification schemes include the Open Research Knowledge Graph (ORKG) (Auer and Mann, 2019), Microsoft Academic Graph (Wang et al., 2020), the Semantic Scholar Academic Graph (Kinney et al., 2023), ACM computing classification system (Rous, 2012), Dewey Decimal Classification (DDC) (Scott, 1998), and the ACL Anthology (Bird et al., 2008). Moreover, the coverage of these schemes is often subject-specific, for example, the

well-known ACM classification is merely limited to computer science topics. As another example, ORKG currently has no in-depth classification for the top-level domain Arts and Humanities in its taxonomy.<sup>1</sup>

Additional challenges with the existing systems in terms of scalability are highlighted by the following examples. First, consider ORKG, which was only recently created by volunteers who assigned tags to scientific texts and were merely able to classify a few thousand publications so far. Thus, an automated classification engine could significantly help to increase its coverage more quickly. Similarly, Microsoft Academic Graph (MAG) has only been relying on existing fields of study for scientific texts from the Microsoft Academic website (Herrmannova and Knoth, 2016), but did not explicitly analyze or apply them. DDC, eventually, still has trouble dealing with new research topics and interdisciplinary fields (Wang, 2009).

Thus, within organizations such as universities, research institutes, or even large companies where numerous researchers and other knowledge creators are working in multiple diverse domains, categorizing texts still requires considerable manual effort, making it challenging to deal with the huge volume of

<sup>a</sup>  <https://orcid.org/0000-0001-6168-0132>

<sup>b</sup>  <https://orcid.org/0009-0007-3826-9477>

<sup>1</sup><https://huggingface.co/spaces/rabuahmad/forcI-taxonomy/blob/main/taxonomy.json>

created texts and the contained knowledge. Consequently, there is a need for automated subject tagging systems to efficiently manage the steadily increasing volume of scientific texts and general knowledge contained in institutional repositories and comprehensive digital archives.

With the growth in generative artificial intelligence (GAI), especially, Large Language Models (LLMs) (Zhao et al., 2023), a new opportunity to automate this tedious task has become tangible. LLMs are Artificial Intelligence (AI) systems that are specialized in generating human-like text for tasks such as summarization, translation, content creation, and even coding. LLMs have already been applied for several use cases, such as analyzing scientific documents (Giglou et al., 2024), writing scientific reviews (Mahapatra et al., 2024), or information extraction (Pertsas et al., 2024). LLMs can be configured by setting parameters such as the so-called quantization for reduced resource usage or their temperature, which controls the degree of creativity in an LLM’s answer. LLMs are applied to their respective tasks by using so-called prompts, which are essentially textual commands describing the desired task at hand. The proper engineering of these prompts plays an important role in achieving the desired results with a model invocation (Gao, 2023).

## 1.1 Research Goals

In this study, we aim to better understand the benefits and quality currently achievable when using “off-the-shelf” LLMs for the classification of scientific texts and hence propose the following research question for our study.

**RQ.** How can LLMs be effectively used to perform accurate tagging of research areas based on existing taxonomies?

To answer this question, we decided to utilize an existing classification as well as an existing dataset that has been recently published – our experiments are based upon the ORKG taxonomy.<sup>2</sup> and the Field of Research Classification (FoRC) Shared Task dataset (Abu Ahmad et al., 2024). The FORC dataset has been compiled by collecting manuscripts from ORKG and arXiv and was categorized into five top-level domains taken from ORKG, since the ORKG taxonomy provides a proven – although not yet fully complete – hierarchical structure for the classification of scientific texts from various domains.

For our evaluation, we used a number of publicly available LLMs to evaluate their performance in terms of finding the (presumably) correct classification that

the human volunteers have attributed to each publication from the dataset. The candidate LLMs each classified 59,344 scientific texts based on their titles and abstracts with different temperatures by applying two types of prompts – zero-shot and few-shot prompts. In a zero-shot prompt, an LLM is primed with limited information, namely merely by providing the task and the requirement to identify the research area, while in a few-shot prompt, we explained the task and also provided an example of a scientific paper together with an appropriate research area. A detailed description of our prompts is provided in section 3. We applied these prompts with different contemporary LLMs, such as Gemma or Llama 3.1, and gauged the results with precision and recall, to finally calculate the accuracy for each model. Hence, the key contributions of this paper are as follows; it presents:

- finding the research areas of scientific texts scientific documents as a novel application for LLMs
- an investigation of the influence of prompt engineering and parameter tuning in optimizing the results
- initial results on the performance of recent open-source LLMs for the classification for scientific texts.

In the remainder of this paper, we discuss the state of the art in section 2 and the proposed approach itself in section 4. After that, we discuss the implementation of the proposed approach in section 5 and present results in section 6. Finally, we conclude our work and discuss future work in section 8.

## 2 RELATED WORK

Document classification is one of the primary tasks for classifying scholarly research that is usually either performed by librarians or by subject experts, where both groups are faced with individual challenges: while the former are usually no subject experts, the latter are normally not trained for using document classification schemes. Multiple such schemes have been developed in recent decades to structure and classify the growing amount of scientific and subject-specific documents, for instance, the ACM computing classification system (Rous, 2012), the Dewey Decimal Classification (DDC) (Scott, 1998), or the ACL Anthology (Bird et al., 2008). However, despite (or because of) this variety of existing taxonomies, manual subject tagging still remains challenging and especially time-consuming. For example, previous research reported that applying the Dewey Decimal Classification to a diverse dataset taken from

<sup>2</sup><https://orkg.org/fields>

the Library of Congress (Frank and Paynter, 2004) took librarians roughly five minutes per publication, as they were only able to assign DDC categories to 10.92 publications per hour (Wiggins, 2009).

Until now, the automated classification of scientific articles into their respective research fields is – despite decades of research – still rather an emerging discipline (Desale and Kumbhar, 2014) than a proven practice that can be applied in libraries, universities, or the knowledge management in large corporations. In previous works, multiple approaches have been applied for this challenge, for instance, the work of (Wang, 2009) used a supervised machine-learning approach for assigning DDC identifiers to documents collected from the Library of Congress (Wiggins, 2009). (Golub et al., 2020) used six different machine learning algorithms to classify the documents from the Swedish library, where a Support Vector Machine gave the best results in terms of accuracy of classifiers. (Jiang et al., 2020) used BERT, an early transformer model for the identification of the research area on previously annotated data. However, up until today, the automation of deriving document classifications has mainly been a supervised learning task that requires specific training data and a thorough validation of results.

In addition, additional challenges, such as the deep nesting of many classification taxonomies and data sparseness in certain classes need to be taken into account when implementing classification with “traditional” supervised learning. In general, this is especially challenging due to the need for a significant amount of labelled training data (Kalyan, 2023) that is still hard to find today. However, the recent generation of Large Language Models is pre-trained on extensive, unlabelled text data and hence is supposed to be more proficient in generating high-quality results in text classification without additional training or finetuning.

Thus, with the recent advancement of Large Language Models, LLMs have already been tested for several generic tasks in scholarly writing and provided promising results. In one recent study, ChatGPT has been used for automated classification of undergraduate courses in an e-learning system and improved overall performance in terms of accuracy significantly (Young and Lammert, 2024). In another study, (Abburri et al., 2023) used an LLM for the automatic evaluation of scientific texts (in German language) written by students to assign a grade. (Pal et al., 2024) proposed an approach for using ChatGPT to develop an algorithm for plagiarism-free scientific writing. (Mosca et al., 2023) built a data set to detect machine-generated scientific papers and com-

pared results with other benchmark models. However, to our knowledge, so far, LLMs have not been tested for the identification of research areas of scientific texts, and hence our work provides a novel insight for the current performance of off-the-shelf LLMs in this area. In a recent publication, we have already proposed a search engine for indexing scientific documents enhanced with research areas, and demonstrated the practical usability of such subject tagging (Shahi and Hummel, 2024), e.g. in a search for domain experts.

### 3 FOUNDATIONS

The key aspect of GAI that separates it from other forms of artificial intelligence (AI), is that it is not primarily dedicated to analysing (numerical) data or acting based on such data like “traditional” AI (i.e. machine learning approaches) that has been used for this purpose in the past. Instead, GAI focuses on creating new content by using the data it was trained on (Hacker et al., 2023; Murphy, 2022). *The term GAI thus refers to computational approaches which are capable of producing apparently new, meaningful content such as text, images, or audio* (Feuerriegel et al., 2024).

Modern GAI for texts utilizes so-called Large Language Models that are trained on massive datasets to acquire different capabilities such as content generation or text summarization by learning statistical relationships of words (Wang et al., 2024). Modern LLMs are developed based on the so-called transformer architecture and trained on extensive corpora collected from public sources such as Web Crawls or GitHub using self- and human-supervised learning, enabling them to capture complex language patterns and contextual relationships (Perełkiewicz and Poświata, 2024). Hence, LLMs can also be used for other quite diverse applications in natural language processing such as text summarization or data annotation. The well-known ChatGPT, launched by OpenAI, based on the Generative Pre-trained Transformer (GPT) architecture (Nah et al., 2023) is one such GAI that has been trained on a huge body (i.e., a significant part of the public WWW) of text.

Two important factors that can influence the perceived performance of LLMs for a certain task are prompt engineering and the so-called temperature used by the model. Prompting strategies in LLMs include writing instructions for the models that are intended to guide responses effectively. Common techniques include providing context, step-by-step instructions, and examples to improve accuracy and rel-

evance. A more detailed description of prompting strategies can be found in (Al Nazi et al., 2025). In our experiments, we used the following two prompting strategies for evaluating the LLM and provide some examples fitting our context in Table 1.

- **Zero-Shot.** In this approach, we ask LLMs to annotate the research area without providing any description or examples, which employs a simple and straightforward approach for extracting the research area. Zero-shot is also known as Vanilla Prompt, which does not take any prior knowledge or specific training on that task. It uses the pre-trained general “intelligence” of an LLM to obtain the research area for a scientific text.
- **Few-Shot.** In this case, learning is done based on context, where the model takes some description and an example for the research area as defined by ORKG taxonomy (Auer and Mann, 2019) to better understand its task. The model takes this input and provides answers based on the given information in conjunction with its general knowledge.

Moreover, we have employed the LLMs with varying temperatures, which adjusts the randomness of the responses given by an LLM. Lower temperatures give more focused and deterministic results, while higher temperatures generate more diverse and “creative” results. The value of the parameter starts from 0; however, we limited it to the range of [0,1] as temperatures above 1 result in a very high degree of randomness and neither coherence nor good reproducibility.

Different LLMs are trained in different objectives and with different training datasets, which is likely affecting their strengths in producing helpful results for our context. Hence, we used a set of temperatures and prompts for different LLMs, such as Llama and Gemma (more details follow in section 5.2), to let them identify the research area of scientific texts from the test data set.

## 4 APPROACH

To address the research question in our present study, we propose a methodology comprising data collection, data cleaning, and preprocessing, followed by the application of prompt engineering to classify the research areas of scientific texts. This process is summarized in Figure 1 and explained in more detail in the following subsections.

Our approach is divided into three main parts, i.e., data collection, data cleaning, preprocessing, and classification and analysis. The first step involves

the collection of the required dataset suitable for the study; in this case, we analysed scientific texts extracted from the FORC dataset; a detailed description of the data is provided in section 3. The second step involves data cleaning and preprocessing, which includes removing unwanted information such as formatting information before the texts were fed to the LLMs. Finally, the third and most crucial step involves the application of LLMs to identify the research area and analyze its results. At present, we aim at the prediction of the top-level domain from the ORKG classification as explained in Section 3.

We employed four contemporary LLMs (cf. Table 2) with a small and medium amount of parameters ranging from 3.82b to 70.4b to classify the research areas of the selected texts. Each LLM was obtained from and executed with Ollama<sup>4</sup>. To assess the performance of LLMs, we compared their results with those of traditional Bidirectional Encoder Representations from Transformers (BERT) models (Devlin et al., 2018). A detailed explanation of the experimental setup is provided in Section 5, while the results are discussed in Section 6.

## 5 EXPERIMENTS

For our experiments, we have been using Ollama frameworks (Morgan and Chiang, 2024), an open source application that allows the easy running of LLMs on local hardware. Ollama provides an easy opportunity to run the model locally with a simple command-line interface that directly interacts with the LLMs and allows easy installation and implementation. Ollama allows downloading models with a given number of parameters. Currently, there are more than 3,100 models registered on Ollama (accessed on 6th August 2024)<sup>9</sup> by numerous different users.

For this experiment, we utilized LLaMa (70 billion parameters), Mistral Nemo (12 billion parameters), Gemma (27 billion parameters), and Phi (3.8 billion parameters) as our foundational models. The respective model sizes are 2.4 GB for Phi, 15 GB for Gemma, 4.1 GB for Mistral Nemo, and 39 GB for LLaMa. The computational infrastructure consists of an in-house server equipped with four NVIDIA RTX A6000 GPUs, each with 48 GB of VRAM, 512 GB of system memory, and 8 TB of storage, ensuring sufficient resources for efficient model execution and experimentation.

<sup>4</sup><https://ollama.com/library>

<sup>9</sup><https://ollama.com/search>



Table 1: Prompting strategies for determining research area from scientific texts.

Zero Shot (Prompt 1)	Few Shot (Prompt 2)
<p>Suppose you are a data annotator who finds the research area of scientific texts.</p> <p>Scientific text to annotate is <i>factors influencing the behavioral intention to adopt a technological innovation from a developing country context: the case of mobile augmented reality games</i></p> <p>Assign a research area to the given scientific texts and provide it as output</p> <p><b>Expected Output:</b> [Social and Behavioral Sciences] <b>Output:</b> [Social Science]</p>	<p>Suppose you are a data annotator who finds the research area of scientific texts.</p> <p>You are provided with scientific texts. Your task is to read texts and determine which research area from the list best represents the content of the scientific texts. Here is the hierarchy for each research: taxonomy of research field extracted from ORKG<sup>3</sup>)</p> <p>scientific texts to annotate is <i>comparative analysis of algorithms for identifying amplifications and deletions in array cgh data</i></p> <p>Assign a research area from the given taxonomy above and provide it as output</p> <p><b>Expected Output:</b> [Physical Sciences &amp; Mathematics] <b>Output:</b> [Physical Sciences &amp; Mathematics]</p>

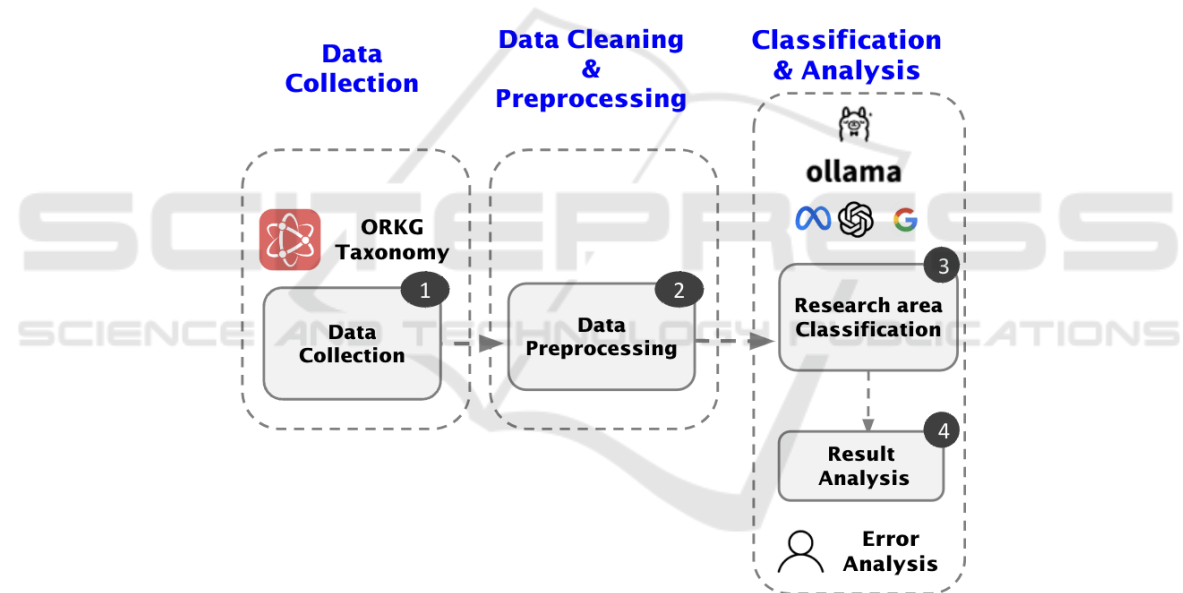


Figure 1: Methodology used in the identification of research area.

We developed a Python program utilizing Ollama and LangChain, an open-source framework for building large language model (LLM) applications, to execute the LLM models for research area annotation. The generated results were systematically stored for evaluation. For LLM optimization, key parameters such as temperature and prompt strategies play a crucial role. To evaluate different configurations, we implemented the models with various combinations of temperature settings and prompts. In the zero-shot setting, only the temperature and task description were provided, whereas the few-shot setting involved different prompt combinations to refine the outputs.

An example of zero-shot and few-shot prompt strategies is presented in Table 1.

5.1 Dataset

We used the scientific texts collected by the FORC shared task (Abu Ahmad et al., 2024), which is based upon the ORKG taxonomy. FORC consists of scientific texts, mainly research papers with DOI, research area, abstract, title, and author information. The FORC initiative compiled scientific texts from open-source resources such as ORKG (CC0 1.0 Universal) and arXiv (CC0 1.0), whereas scientific text

Table 2: A short description of LLM models used in the study.

Model)	Description	No. of Parameters	Release Date
Gemma 2 <sup>5</sup>	Gemma2 is a lightweight, state-of-the-art open models	parameters-27.2B & quantization-Q4.0	June 2024
Llama 3.1 <sup>6</sup>	Llama 3.1 70B is a multilingual model that has a significantly longer context length of 128K, state-of-the-art tool use, and overall stronger reasoning capabilities	parameters-70.4B and quantization-Q4.0	July 2024
Mistral Nemo <sup>7</sup>	Mistral NeMo offers a large context window of up to 128k tokens. Its reasoning, world knowledge, and coding accuracy are state-of-the-art in its size category	parameters-12.2B & quantization-Q4.0	July 2024
Phi <sup>8</sup>	Phi 3.5 is a lightweight, state-of-the-art open model built upon synthetic datasets	parameters-3.82B & quantization-Q4.0	August 2024

with non-English titles or abstracts were excluded. Each scientific text has been assigned a field of research based on ORKG taxonomy.<sup>10</sup>

The ORKG taxonomy<sup>11</sup> provides a structured framework for the systematic classification and exploration of research domains. This taxonomy is organized into five primary domains: *Arts and Humanities*, *Engineering*, *Life Sciences*, *Physical Sciences and Mathematics*, and *Social and Behavioral Sciences*. Each of these top-level domains is hierarchically structured into two additional levels: subdomains and subjects. At the first sub-level, each primary domain is subdivided into specific research areas, which are further refined into specialized subjects. For instance, within the *Physical Sciences and Mathematics* domain, the *Computer Science* subdomain includes *Artificial Intelligence* as a subject. In this study, for the time being, we consider the top-level domains to maintain a high-level perspective on the classification of research fields.

Overall, FORC provides a collection of 59,344 scientific texts, each categorized using a taxonomy of 123 Fields of Research (FoR). These are organized across three hierarchical levels and grouped into five top-level categories: *Physical Sciences and Mathematics*, *Engineering*, *Life Sciences*, *Social and Behavioral Sciences*, *Arts and Humanities*. For each scientific text, we filtered only meaningful information useful for us, as described in table 5.1. We used DOI as a unique identifier and title with abstract for our classification model to predict the research area.

<sup>10</sup><https://orkg.org/fields>

<sup>11</sup><https://orkg.org/fields>

### 5.1.1 Data Cleaning and Preprocessing

After collecting the dataset, we cleaned and preprocessed the data to remove unwanted information from the scientific text. This section describes the steps involved in the data preprocessing and analysis. LLMs work as a black-box algorithm (Liu et al., 2024), and we do not have an internal functional model to provide the research area of scientific texts, so we provided preprocessed data for all models to maintain fairness. After collecting the dataset, we removed unwanted information, such as URLs mentioned in the text, special characters in abstracts, and authors of the publication. We used title and abstract from all 59,344 scientific texts tagged by FORC to identify the research area and derived the accuracy to evaluate our LLM-based prediction model.

## 5.2 Evaluated LLMs

We employed four open-source LLMs for our classification experiments, namely, Llama (Meta), Gemma (Google), Nemo (Mistral), and Phi (Microsoft). A detailed description of each selected LLM is given below, together with important metadata, which is provided in Table 2. After that, the experimental setup is explained in section 5.

- **Gemma.** We used the recent version Gemma 2 (Team et al., 2024), which is a family of lightweight, state-of-the-art open-source models that are advertised as high-performing and efficient models by Google. They are currently available in two sizes; we have used Gemma 2 with 27 million parameters. Gemma was trained on web

Table 3: Description of collected dataset.

Field	Description	Example
DOI	A DOI (Digital Object Identifier) is a standardized unique number given to paper, and we used it as a unique identifier of paper	10.1145/2736277.2741136
Title	Title describes the title of the paper	Crowd Fraud Detection in Internet Advertising
Abstract	Abstract of the paper describing the summary of paper	"the rise of crowdsourcing brings new types of malpractices in internet advertising. one can easily hire web workers through malicious crowdsourcing platforms to attack other advertisers....
Research area	research area defined based on the ORKG taxonomy, and it is a dependent variable for our prediction model	Engineering

documents and using mathematics, outperforming other models in 11 of 18 text-based tasks in terms of efficiency (Team et al., 2024).

- **Llama.** We used the latest version, which was Llama 3.1 at the time of writing. Llama is developed and released by Meta (Touvron et al., 2023); there are currently three versions of Llama with different sizes of 8b, 70b, and 405b parameters; we have used Llama 3.1 with 70b parameters. Llama is trained on publicly available data without resorting to proprietary datasets. For the training, different data sources, such as Common-Crawl and GitHub, were used.
- **Mistral Nemo.** is the latest LLM developed jointly by Mistral AI and NVIDIA AI with 12B parameters and a context window of up to 128k tokens. Mistral Nemo outperformed the prior Mistral model Llama 3 and Gemma 2 in terms of efficiency and effectiveness despite having fewer parameters.<sup>12</sup>
- **Phi.** is a family of lightweight, open large language models developed by Microsoft that are designed to be efficient and accessible. The "Phi-3" family includes models with 3 billion (3B) and 14 billion (14B) parameters, classified as "Mini" and "Medium" respectively. Phi outperforms<sup>13</sup> Gemini 1.0 Pro, and the model is trained on high-quality educational data, newly created synthetic, "textbook-like" data, which should make it espe-

cially suitable for use for classification tasks of the scientific domain.

### 5.3 Baseline Models

Given our collected corpus of scientific texts, we chose the following two classification models as state state-of-the-art methods for result comparison: **BERT** (Devlin et al., 2018) is a widely used pre-trained model for text classification. The model has been applied to various classification tasks and evaluated across multiple domains, including the classification of text related to COVID-19 (Shahi and Nandini, 2020). BERT utilizes a bidirectional transformer mechanism, allowing it to capture contextual relationships in text more effectively than traditional models. It has demonstrated state-of-the-art performance in numerous natural language processing (NLP) benchmarks, making it a strong candidate for research area classification.

**BiLSTM** (Bidirectional Long Short-Term Memory) (Huang et al., 2015) is a recurrent neural network (RNN) designed for text classification, capturing input flows in both forward and backward directions. It has been successfully applied to various NLP tasks, including the classification of scientific texts (Enamoto et al., 2022). BiLSTM enhances sequential data processing by preserving long-range dependencies, reducing the vanishing gradient problem, and improving contextual understanding. Its ability to capture bidirectional dependencies makes it effective in tasks requiring nuanced text comprehension.

For the implementation of the baseline models, we developed a Python program and retrieved the pre-

<sup>12</sup><https://mistral.ai/news/mistral-nemo/>

<sup>13</sup><https://azure.microsoft.com/de-de/blog/new-models-added-to-the-phi-3-family-available-on-microsoft-azure/>

trained models from Hugging Face.<sup>14</sup> Both BERT and BiLSTM were implemented using the models available on this platform.

## 6 RESULTS

We evaluated the baseline models and the selected LLMs according to the two explained prompts and temperature values, from 0.2 up to 1.0 with all nearly 60,000 titles and abstracts from the FORC dataset. Depending on the size of the model, this took about 3 hours (for Phi) and about 22 hours (for Llama) with the others in between. The accuracy obtained by the state of the art baseline models is shown in Table 4, the accuracy for the LLMs in Table 5 and 6 respectively.

Overall, the Llama model achieved the best results, and few-shot prompt strategies outperformed the models with zero-shot prompts. Also, increasing the temperature seems to help in obtaining better results; however, after reaching the best performance at 0.8, quality starts decreasing with the one larger temperature value we have tested.

Table 4: Accuracy of State of Arts models.

Model	Accuracy
BERT	0.74
BiLSTM	0.66

Table 5: Accuracy of LLMs according to prompt 1 (Zero-shot) with different temperatures.

Parameter Model	0.2	0.4	0.6	0.8	1
Gemma	0.18	0.24	0.30	0.42	0.50
Llama	0.24	0.28	0.48	<b>0.62</b>	0.60
Mistral Nemo	0.10	0.18	0.34	0.52	0.44
Phi	0.30	0.18	0.24	0.48	0.42

Table 6: Accuracy of different LLMs according to prompt 2 (Few-shot) with different temperatures.

Parameter Model	0.2	0.4	0.6	0.8	1
Gemma	0.18	0.24	0.40	0.66	0.62
Llama	0.34	0.38	0.64	<b>0.82</b>	0.72
Mistral Nemo	0.14	0.38	0.44	0.76	0.62
Phi	0.30	0.18	0.44	0.58	0.62

As is visible in the tables, the relatively high temperature of 0.8 seems to work best with most models, sometimes 1.0 provides even better results. Not surprisingly, the model largest in terms of parameters,

i.e. Llama with 70 billion parameters, delivered the best results in our experiments and achieved an Accuracy of 0.82 for the few-shot prompt.

### 6.1 Error Analysis

In addition to the automated evaluation presented earlier, we conducted a manual error analysis to assess the performance of the best-performing model. Specifically, we focused on analyzing the misclassified instances produced by LLaMa 3.1 to gain deeper insights into its errors. To achieve this, we randomly sampled 100 scientific texts from the incorrectly classified results and manually analyzed them. This manual review provided qualitative insights into common causes of misclassification (such as short abstracts, missing abstracts, or missing titles), helping to identify potential reasons for incorrect predictions. The findings support the automated evaluation, indicating that LLaMa 3.1 produces reliable and usable results.

## 7 SUMMARY

In the present study, we have tested a number of state-of-the-art open-source LLMs for the identification of research areas of scientific documents based on 59,344 abstracts taken from the FORC dataset (Abu Ahmad et al., 2024). Overall, the LLMs are capable of identifying the research areas pretty well, as underlined by accuracy of up to 0.82. However, for the time being, our study merely tackled the highest level of the FORC classification, leaving a lot of room for future work aiming on the lower levels as well.

In general, automatic tagging of scientific (and other) texts is still an ongoing challenge that will require future work, as there still exists a lack of cross-domain datasets that cover common subject areas. Hence, another limitation of our study is that the dataset we have used was taken from a previous work with a different goal, so that it was not a perfect fit for the task at hand in terms of overall coverage. E.g., we found that some of the classes are currently not well covered in the dataset, as there are currently no sub-classes for Arts and Humanities, which would make it hard to generalize our topic classification results for general libraries that have to deal with texts from virtually all knowledge areas.

Several other taxonomy systems exist for tagging scientific texts, such as the ACM classification system or the Dewey Decimal Classification (DDC). However, they still suffer from similar limitations, such as non-existent gold standards. Consequently, evaluating LLMs for these taxonomies must be covered in

<sup>14</sup><https://huggingface.co/>



future work that is also likely to feed back interesting improvement potential for the taxonomies.

Although large language models have exhibited remarkable efficacy in addressing a wide range of challenges, their deployment for classification tasks remains fraught with significant challenges. Specifically, LLMs require substantial computational resources, including high-performance GPUs and extensive memory capacities, resulting in considerable economic implications, which is another indicator for the substantial efforts that are required in the near future to achieve a solid understanding of how well LLMs can be used for various tasks.

## 8 CONCLUSION & FUTURE WORK

This study systematically investigates the application of large language models (LLMs) for automated research area classification in scientific literature. The proposed methodology was implemented on the FORC dataset, taken from the ORKG initiative, employing two distinct prompt engineering strategies while optimizing the temperature parameter to enhance classification performance. The study used the five top-level domains of ORKG taxonomies to classify research domains, which can obviously be extended to predict subdomains and subjects of scientific text in the near future. Four LLMs, namely Gemma, Llama, Nemo and Phi were rigorously evaluated against two baseline BERT models using a large dataset of almost 60,000 publications. Results are indicating that modern LLMs are superior to previous models and that few-shot prompting significantly improves classification accuracy. Among the models tested, Llama achieved the highest accuracy, making it the most effective for research area identification.

Future research directions include leveraging additional state-of-the-art LLMs, addressing more fine-granular taxonomy levels, and integrating alternative classification schemes, such as the ACM Computing Classification System and the Dewey Decimal Classification (DDC), to refine scientific text classification even further. The proposed method can also be deployed in institutional research centers and academic libraries to systematically identify and categorize forthcoming scholarly publications, enhancing knowledge organization and retrieval in academic and industrial research environments.

## DATA SHARING

We have conducted all experiments on a macro level following strict data access, storage, and auditing procedures for the sake of accountability. We release the processed data used in the study along with minimal code to replicate the model for the community. The code and the dataset are available at GitHub here.<sup>15</sup>

## ACKNOWLEDGEMENTS

The work has been carried out under the TransforMA project. Authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article. This project has received funding from the federal-state initiative "Innovative Hochschule" of the Federal Ministry of Education and Research (BMBF) in Germany.

## REFERENCES

- Abhuri, H., Suesserman, M., Pudota, N., Veeramani, B., Bowen, E., and Bhattacharya, S. (2023). Generative ai text classification using ensemble llm approaches. *arXiv preprint arXiv:2309.07755*.
- Abu Ahmad, R., Borisova, E., and Rehm, G. (2024). Forc@nslp2024: Overview and insights from the field of research classification shared task. In *International Workshop on Natural Scientific Language Processing and Research Knowledge Graphs*, pages 189–204. Springer.
- Al Nazi, Z., Hossain, M. R., and Al Mamun, F. (2025). Evaluation of open and closed-source llms for low-resource language with zero-shot, few-shot, and chain-of-thought prompting. *Natural Language Processing Journal*, page 100124.
- Auer, S. and Mann, S. (2019). Towards an open research knowledge graph. *The Serials Librarian*, 76(1-4):35–41.
- Bird, S., Dale, R., Dorr, B. J., Gibson, B. R., Joseph, M. T., Kan, M.-Y., Lee, D., Powley, B., Radev, D. R., Tan, Y. F., et al. (2008). The acl anthology reference corpus: A reference dataset for bibliographic research in computational linguistics. In *LREC*.
- Bornmann, L., Haunschild, R., and Mutz, R. (2021). Growth rates of modern science: a latent piecewise growth curve approach to model publication numbers from established and new literature databases. *Humanities and Social Sciences Communications*, 8(1):1–15.
- Desale, S. K. and Kumbhar, R. M. (2014). Research on automatic classification of documents in library environment: a literature review. *KO KNOWLEDGE ORGANIZATION*, 40(5):295–304.

<sup>15</sup><https://github.com/Gautamshahi/LLM4ResearchArea>

- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Enamoto, L., Santos, A. R., Maia, R., Weigang, L., and Filho, G. P. R. (2022). Multi-label legal text classification with bilstm and attention. *International Journal of Computer Applications in Technology*, 68(4):369–378.
- Feuerriegel, S., Hartmann, J., Janiesch, C., and Zschech, P. (2024). Generative ai. *Business & Information Systems Engineering*, 66(1):111–126.
- Frank, E. and Paynter, G. W. (2004). Predicting library of congress classifications from library of congress subject headings. *Journal of the American Society for Information Science and Technology*, 55(3):214–227.
- Gao, A. (2023). Prompt engineering for large language models. Available at SSRN 4504303.
- Giglou, H. B., D’Souza, J., and Auer, S. (2024). Llms4synthesis: Leveraging large language models for scientific synthesis. *arXiv preprint arXiv:2409.18812*.
- Golub, K., Hagelbäck, J., and Ardö, A. (2020). Automatic classification of swedish metadata using dewey decimal classification: a comparison of approaches. *Journal of Data and Information Science*, 5(1):18–38.
- Hacker, P., Engel, A., and Mauer, M. (2023). Regulating chatgpt and other large generative ai models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 1112–1123.
- Herrmannova, D. and Knoth, P. (2016). An analysis of the microsoft academic graph. *D-lib Magazine*, 22(9/10):37.
- Huang, Z., Xu, W., and Yu, K. (2015). Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Jiang, M., D’Souza, J., Auer, S., and Downie, J. S. (2020). Improving scholarly knowledge representation: Evaluating bert-based models for scientific relation classification. In *Digital Libraries at Times of Massive Societal Transition: 22nd International Conference on Asia-Pacific Digital Libraries, ICADL 2020, Kyoto, Japan, November 30–December 1, 2020, Proceedings 22*, pages 3–19. Springer.
- Kalyan, K. S. (2023). A survey of gpt-3 family large language models including chatgpt and gpt-4. *Natural Language Processing Journal*, page 100048.
- Kinney, R., Anastasiades, C., Authur, R., Beltagy, I., Bragg, J., Buraczynski, A., Cachola, I., Candra, S., Chandrasekhar, Y., Cohan, A., et al. (2023). The semantic scholar open data platform. *arXiv preprint arXiv:2301.10140*.
- Liu, S., Yu, S., Lin, Z., Pathak, D., and Ramanan, D. (2024). Language models as black-box optimizers for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12687–12697.
- Mahapatra, R., Gayan, M., Jamatia, B., et al. (2024). Artificial intelligence tools to enhance scholarly communication: An exploration based on a systematic review.
- Morgan, J. and Chiang, M. (2024). Ollama. <https://ollama.com>. Online; accessed 6 August 2024.
- Mosca, E., Abdalla, M. H. I., Basso, P., Musumeci, M., and Groh, G. (2023). Distinguishing fact from fiction: A benchmark dataset for identifying machine-generated scientific papers in the llm era. In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 190–207.
- Murphy, K. P. (2022). *Probabilistic machine learning: an introduction*. MIT press.
- Nah, F., Cai, J., Zheng, R., and Pang, N. (2023). An activity system-based perspective of generative ai: Challenges and research directions. *AIS Transactions on Human-Computer Interaction*, 15(3):247–267.
- Pal, S., Bhattacharya, M., Islam, M. A., and Chakraborty, C. (2024). Ai-enabled chatgpt or llm: a new algorithm is required for plagiarism-free scientific writing. *International Journal of Surgery*, 110(2):1329–1330.
- Perełkiewicz, M. and Poświata, R. (2024). A review of the challenges with massive web-mined corpora used in large language models pre-training. *arXiv preprint arXiv:2407.07630*.
- Pertsas, V., Kasapaki, M., and Constantopoulos, P. (2024). An annotated dataset for transformer-based scholarly information extraction and linguistic linked data generation. In *Proceedings of the 9th Workshop on Linked Data in Linguistics@ LREC-COLING 2024*, pages 84–93.
- Rabby, G., Auer, S., D’Souza, J., and Oelen, A. (2024). Fine-tuning and prompt engineering with cognitive knowledge graphs for scholarly knowledge organization. *arXiv preprint arXiv:2409.06433*.
- Rous, B. (2012). Major update to acm’s computing classification system. *Communications of the ACM*, 55(11):12–12.
- Scott, M. L. (1998). Dewey decimal classification. *Libraries Unlimited*.
- Shahi, G. K. and Hummel, O. (2024). Enhancing research information systems with identification of domain experts. In *Proceedings of the Bibliometric-enhanced Information Retrieval Workshop (BIR) at the European Conference on Information Retrieval (ECIR 2024)*, CEUR Workshop Proceedings. CEUR-WS.org.
- Shahi, G. K. and Nandini, D. (2020). FakeCovid – a multilingual cross-domain fact check news dataset for covid-19. In *Proceedings of the 14th International AAAI Conference on Web and Social Media*.
- Team, G., Mesnard, T., Hardin, C., Dadashi, R., Bhupatiraju, S., Pathak, S., Sifre, L., Rivière, M., Kale, M. S., Love, J., et al. (2024). Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. (2023). Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Wang, J. (2009). An extensive study on automated dewey decimal classification. *Journal of the Ameri-*

- can Society for Information Science and Technology*, 60(11):2269–2286.
- Wang, K., Shen, Z., Huang, C., Wu, C.-H., Dong, Y., and Kanakia, A. (2020). Microsoft academic graph: When experts are not enough. *Quantitative Science Studies*, 1(1):396–413.
- Wang, S., Hu, T., Xiao, H., Li, Y., Zhang, C., Ning, H., Zhu, R., Li, Z., and Ye, X. (2024). Gpt, large language models (llms) and generative artificial intelligence (gai) models in geospatial science: a systematic review. *International Journal of Digital Earth*, 17(1):2353122.
- Wiggins, B. J. (2009). Acquisitions and bibliographic access directorate library of congress report of fiscal year 2009 (fiscal year ended september 30, 2009).
- Young, J. S. and Lammert, M. (2024). Chatgpt for classification: Evaluation of an automated course mapping method in academic libraries.
- Zhang, C., Tian, L., and Chu, H. (2023). Usage frequency and application variety of research methods in library and information science: Continuous investigation from 1991 to 2021. *Information Processing & Management*, 60(6):103507.
- Zhao, W. X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., et al. (2023). A survey of large language models. *arXiv preprint arXiv:2303.18223*.

