

Exploring Medical Literature Retrieval: Course Project

Assignment created by: Georgios Peikos & David La Barbera

1 Goal

Develop a search engine that retrieves and ranks medical literature in response to precision medicine information needs for oncology cases. Each search query corresponds to a synthetic patient case describing cancer type, genetic variants, demographic details, and other relevant clinical factors. The system should be capable of identifying and ranking clinical trials that discuss effective treatments for the given molecular and clinical profile. It must address challenges such as interpreting structured and semi-structured medical information, understanding genomic and clinical terminology, matching relevant treatments and variants.

For the purpose of this project, we will focus exclusively on retrieving clinical trials addressing relevant treatments for the given patient. This corresponds to the TREC 2019 Precision Medicine clinical trials subtask.

2 Project Description and Material

This project enables students to explore and apply approaches that combine Information Retrieval and Natural Language Processing techniques in practice. Students can leverage knowledge acquired from previous courses, material presented during this course, and external resources to develop their own ideas.

The project aims to help students gain a solid understanding of medical Information Retrieval systems and the specific challenges of representing domain-specific (i.e. medical language). To aid the students, dedicated resources will be provided. Students are expected to analyze the provided document collection, conduct baseline experiments using the lab materials, and then design and implement their own approach to address the task of medical literature retrieval through clinical trials. Finally, students will effectively communicate their investigation by clearly presenting their methodology, findings, and insights in both a written report and an oral presentation.

In detail, students should first read the provided material to understand the main challenges that medical IR systems face. They are encouraged to explore the dataset by analyzing term distributions, entity (e.g. disease, treatments) occurrences, and the proportions of the diseases/treatments in the provided abstracts. Possible project directions include developing and evaluating retrieval methods for medical literature retrieval, experimenting with neural, lexicon-based, or hybrid retrieval models, and assessing their performance using standard ranking metrics.

Students may also focus on query expansion and reformulation techniques, query understanding, designing approaches to identify and link medical entities. Other directions may involve exploring techniques for integrating information from external sources, such as medical knowledge graphs, to improve the retrieval effectiveness. Students are encouraged to review recent research to understand how modern NLP advancements have been applied in medical IR and consider incorporating such methods into their work.

As more people rely on conversational AI systems such as ChatGPT to explore symptoms, diseases, and possible treatments, it has become increasingly important to ensure that these systems use Retrieval-Augmented Generation approaches to improve their factually accuracy. Such approaches first retrieve documents relevant to the user's information need and then generate factually accurate answers supported by explicit evidence through citations. This project represents the first step toward that goal, as students will focus on building effective retrieval systems for medical literature retrieval.

As a potential direction for future work, students may explore integrating large language models with their retrieval systems. They could begin by using a model to generate an initial answer to a query from the collection. Next, they would use their retrieval system to find relevant supporting documents for the same query. These documents could then be added to the model's prompt to provide additional evidence. Finally, they would regenerate the answer and evaluate whether the inclusion of retrieved information improves its quality.

This project aims to strengthen both technical and analytical skills in NLP and Information Retrieval while promoting the ability to apply them to real-world challenges. Students are strongly encouraged to think creatively and innovatively when defining their approach to this search task. The project is also designed to foster collaboration, helping students develop key soft skills such as teamwork, effective communication, and an understanding of team dynamics.

Students are encouraged to build upon the discussions and materials presented in the labs when conceptualizing and developing their projects. They should also make use of the provided resources, including carefully selected academic

papers and online tutorials, to support their exploration of NLP, Information Retrieval, and financial data processing.

Resources related to the task:

- ◇ **Lab Material:** The material covered during the lab sessions serves as the primary resource for students.
- ◇ **Research Paper on Medical IR:** [An academic paper discussing advanced topics in medical information retrieval.](#)
- ◇ **Official Precision Medicine 2017 Track Website:** [TREC Precision Medicine 2017 Track](#) Here you can find the 2017 benchmark collection you will use in this project and further information about the project's task.
- ◇ **Official Precision Medicine 2018 Track Website:** [TREC Precision Medicine 2017 Track](#) Here you can find the 2018 benchmark collection you will use in this project and further information about the project's task.
- ◇ **Previous Participants in Precision Medicine 2017 Track:** [Previous Participants in 2017](#) Here you can read the papers of researchers who have developed systems for medical information retrieval and participated in TREC Precision Medicine 2017 Track.
- ◇ **Previous Participants in Precision Medicine 2018 Track:** [Previous Participants in 2018](#) Here you can read the papers of researchers who have developed systems for medical information retrieval and participated in TREC Precision Medicine 2018 Track.
- ◇ **Previous Participants in Precision Medicine 2019 Track:** [Previous Participants in 2019](#) **The benchmark you work with is the 2019, the clinical trials part.** Here you can read the papers of researchers who have developed systems for medical information retrieval and participated in TREC Precision Medicine 2018 Track.
- ◇ **Precision Medicine 2017 Track Overview Paper:** [Overview of the TREC 2017 Precision Medicine Track](#)
- ◇ **Precision Medicine 2018 Track Overview Paper:** [Overview of the TREC 2018 Precision Medicine Track](#)
- ◇ **Precision Medicine 2019 Track Overview Paper:** [Overview of the TREC 2019 Precision Medicine Track](#) As this is the collection you are working on, this is the most relevant information to your task.
- ◇ **Research on Medical Literature Retrieval in 2025:** [Research Papers published in the last 2 years and use the same collections](#) For inspiration!
- ◇ **GitHub Repository for Information Extraction in Medical IR:** [A repository containing code and resources related to information extraction in medical information retrieval.](#)
- ◇ **Tutorial in Health-Related Information Retrieval:** [A detailed tutorial covering various aspects of information retrieval in the health domain.](#)
- ◇ **Tutorial in Precision Health:** [Offers insights into the evolving field of precision health and its applications.](#)
- ◇ **PyTerrier Documentation and Experiments:** [Comprehensive documentation and guides for conducting experiments using PyTerrier, an IR research platform.](#)
- ◇ **Additional material will be provided as teams progress with their chosen project ideas (e.g. keyword extraction, usage of LLMs, etc.)..**

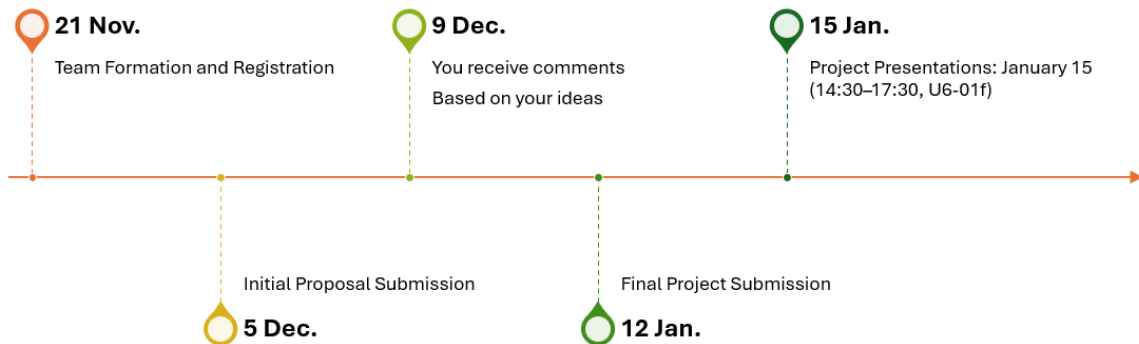
3 Participation Guidelines

- ◇ **Eligibility:** Open to all students enrolled in the course.
- ◇ **Team Formation:** Teams are required to be formed with a composition of **3 members**.
- ◇ **Communication:** Each team must appoint one member as the **communication leader**. This person is responsible for registering the team and its members with the instructor and handling all communication between the team and the course instructor.
- ◇ **Participation in Exams:** You are required to complete and submit one project to be eligible to participate in the exam! Please remember that, to be eligible to take part in an exam session, you must **first** submit the project.
 - For example, if you submit the project by 15 January, you can participate in all subsequent exam sessions.
 - If you decide to complete a project in May, you will be eligible for all sessions following that month, such as those in June, July, and September.
 - Your project grade is kept and is valid for all subsequent exam sessions.

4 Steps for Completion

4.1 Timeline

Deadlines



- Team Formation and Registration:** Deadline - **November 21**, via a single email to both: georgios.peikos@unimib.it and david.labarbera@unimib.it.
 - By the deadline, the communication leader must send an email to the instructor including:
 - ◇ A team name [Optional].
 - ◇ The selected project (for example, Financial Question Answering).
 - ◇ The name and student ID of each team member.
 - After registration, teams should start working on the Initial Proposal Submission (see Section 4.2).
 - Teams can also start working on Development Phase I (see Section 4.3).
- Initial Proposal Submission:** Deadline - **December 5**. Teams have to send via email the PDF file described in Section 4.2.
- Ideas Evaluation:** Deadline - **December 9**, each team will receive an email with further materials and comments on their proposed ideas.
 - After the email communication, teams can continue working on Development Phase I and start working on Phase II (see Sections 4.3 and 4.4).
- Final Project Submission:** Deadline - **January 12 by 23:59**; this will include a single *.zip* that contains your project, your (final/draft) presentation, and your source code. The *.zip* file must be sent via a single email at georgios.peikos@unimib.it and david.labarbera@unimib.it.
- Project Presentations:** Deadline - **January 15, from 14:30 - 17:30 in U6-01f**.

4.2 Initial Proposal Submission

This initial proposal is the first stage of the project. Your submission should be a concise yet comprehensive document, spanning 2-3 pages in the provided format (i.e. 12pt Font size). The PDF should be sent by the team's communication leader via email to georgios.peikos@unimib.it and david.labarbera@unimib.it and encompass the following components: Describe the task by outlining the situation in which it takes place and the specific information need it addresses. Explain who the intended users are, what goals they pursue, and what challenges they face when searching. Summarize the main characteristics of the document collection, such as its source, structure, and content type. This exercise will help you to understand the context in which your search engine will operate and guide you to the development process.

- ◇ **Task Description:** Provide a clear and concise explanation of the task. Describe who the intended users of your search engine are and explain why this task is valuable or beneficial for them.
- ◇ **Why the task is challenging?** Based on your opinion and your research, what makes your retrieval task challenging?

- ◇ **Why Solving This Task is Important:** Explain, based on your understanding and readings, why addressing this task has value. Describe the potential impact of your solution and identify who would benefit from it and how.

After understanding the task and its context, the next step is to start thinking of a solution, or even better, several potential ones. Let your creativity flow and explore different ways the problem could be tackled. The more ideas you come up with now, the better your chances of finding something original and effective later.

- ◇ **Report a List of Ideas:** Each team member should propose three ideas for a retrieval system that could address the chosen task. Include them in the PDF and discuss them together.
- ◇ **Select at Least One Idea to Develop:** As a team, agree on at least one idea to explore further and transform into a concrete project plan. Write in the PDF a short explanation why you selected this solution.
- ◇ **Define a Research Question:** Formulate one clear research question based on your selected idea, focusing on what you aim to investigate through your system [Can be the title of your project!]. **Never written one? Have a look <https://www.youtube.com/watch?v=42-d2HdbyS8>**

As the saying goes, a picture is worth a thousand words, often expressing ideas more clearly than lengthy explanations.

- ◇ **Method Outline (i.e. your idea/s):** Present an overview of your proposed methodology. You are encouraged to use visual aids such as diagrams or flowcharts to enhance clarity. For creating these, <https://www.draw.io/> or <https://excalidraw.com/> can be helpful tools. The latter allows online collaborative drawing.

Make sure you give proper credit to all sources, ideas, and materials that have contributed to your work.

- ◇ **Resource References:** Include citations or references to the resources that have helped you develop your ideas, your understanding of the problem, and inspired your ideas. This may include academic papers, online tutorials, or any other relevant material.

Ideas are exciting, but you should also assess their feasibility and ensure your goals can realistically be achieved within the available time!

- ◇ **Technical Aspects:** List and describe the technical elements and tools you intend to utilize in your project. This could include, GitHub libraries, frameworks, or algorithms, or tools you would like to learn.
- ◇ **Feedback:** You will receive feedback on the scope and feasibility of your ideas, such as whether your approach seems too ambitious, too simple, or well balanced for the project goals. We will also share additional resources to help you refine and implement your ideas more effectively.

4.3 Development and Implementation: Phase I

The initial phase of development focuses on implementing **three retrieval runs** that will serve as **baselines** for further progress. A baseline retrieval experiment represents a relatively simple retrieval pipeline designed to establish reference performance. For your projects, the baselines can also be the three simplest ideas you had in the initial proposal phase.

The first step in developing a search engine, after understanding your users and their context, is to become familiar with your data. In offline evaluation of Information Retrieval Systems, this includes the document collection you will index and the queries that represent how users express their information needs. Finally, you should examine the relevance assessments available in your collection to understand how documents are judged in relation to queries.

- ◇ **Collection, Query, and Relevance Analysis:** Teams can present several statistics that can inform the development of their IR system.
- ◇ **Purpose:** This analysis not only guides the design of your retrieval system but also helps you understand why the system may not perform as expected and identify areas for improvement.
- ◇ **Usage of ASPIRE:** To help you in your analysis, you can use ASPIRE or you can write your own code.

In any case, here are some simple ideas your teams can try as baselines.

- ◇ **Indexing Strategies:** Experiment with various indexing approaches. For instance, one approach could involve indexing the entire content of documents, while another could focus on indexing only titles or specific sections (if applicable in your task).

- ◇ **Query Formulation and Expansion:** Experiment with different ways of expressing user queries, ranging from simple keyword-based formulations to more structured or natural language queries. Explore query expansion techniques such as adding synonyms, related terms for query expansion, or automatically retrieved relevant words (e.g. using the RM3 model).
- ◇ **Retrieval Models:** Start with classic lexical-based retrieval using TF-IDF and BM25.
- ◇ In total, you should implement **three** simple retrieval experiments.

Here are specific guidelines to help you structure and carry out this phase effectively.

- ◇ **Experiments:** You must index **the whole collection**, i.e. all documents.
- ◇ **Evaluation Measures:** To ensure consistent comparison of retrieval effectiveness across your retrieval systems and across teams, you must report the following metrics in your experiments: **P@1, P@5, P@10, R@5, R@10, nDCG@5, nDCG@10, and MAP.**
- ◇ **Starting Resource:** Each team can start their experiments based on the colab notebook we have prepared for you here: [Download the collection here.](#)

This phase is important for setting a foundational understanding of the retrieval process, upon which more advanced techniques will be built in subsequent phases of the project. Please ensure that your final PDF report goes beyond merely stating the results by providing clear explanations and interpretations of the obtained outcomes. Use this opportunity to reflect on what your results reveal about the behavior of your retrieval system and the characteristics of your data. Discuss possible reasons behind performance differences across runs and suggest what could be improved or explored in the next phase.

4.4 Development and Implementation: Phase II

During this second phase, teams will move forward from the foundational work established in Phase I to the development and evaluation of their unique ideas.

Here are specific guidelines to help you structure and carry out this phase effectively.

- ◇ Conduct **at least two** additional experiments (E1, E2) that extend beyond the baselines. One of these (E1) can involve a neural approach, such as a bi-encoder retriever or a neural re-ranker.
- ◇ Alternatively, E1 can be an experiment using LLMs for tasks like query expansion or reformulation to enhance retrieval effectiveness.
- ◇ For the E2, you must develop and evaluate your proposed idea in depth, clearly describing the methods used, the rationale behind your choices, and the impact on retrieval performance.
- ◇ **Experiments:** You must index **the whole collection**, i.e. all available documents in the provided collection.
- ◇ **Evaluation Measures:** To ensure consistent comparison of retrieval effectiveness across your retrieval systems and across teams, you must report the following metrics in your experiments: **P@1, P@5, P@10, R@5, R@10, nDCG@5, nDCG@10, and MAP.**

5 Minimum Requirements Expected

Each project should demonstrate both understanding and experimentation. In summary, we expect you to:

- ◇ **Show understanding of the task and its context:** Explain how your search engine will operate, describe the benchmark collection, and identify the main challenges of your task.
- ◇ **Develop three baseline experiments:** Implement and evaluate **three** simple retrieval pipelines that serve as baselines for comparison.
- ◇ **Develop two advanced experiments:** Design and test **two** sophisticated retrieval approaches, such as neural ranking/re-ranking, and your own idea!

Your **final submission** must be a *.zip* file containing:

- ◇ The source code implementing your retrieval experiments (this may include scripts or Google Colab notebooks).
- ◇ The final project report, prepared using the provided template. You can download the template in word format form [here](#), or the latex version from [here](#). The report must be up to 5 pages of content.
- ◇ A draft or the final version of your presentation.

6 Assessment Criteria

Credits. The entire project carries a total of 3 credits. The initial part, which includes understanding the task and its challenges, proposing your ideas, and completing the development described in Phase I, will account for **1.5 credits**. The subsequent Phase II is allocated the remaining **1.5 credits**. In Phase I, the focus is on evaluating your understanding and application of the fundamental concepts covered in the labs. Phase II presents an opportunity for each team to focus on an area of their personal interests. It encourages **exploration** and **innovation**, as teams have the freedom to choose and develop their own ideas.

Bonus: The team that develops the most complete, innovative, and high-performing retrieval system within each project will receive an additional **1 credit**.

The projects will be evaluated based on their:

- ◇ **Functionality:** The system should be able to retrieve relevant documents; the absolute performance is not the main focus, as long as the retrieval process functions correctly.
- ◇ **Innovation:** Originality of the proposed solutions.
- ◇ **Technical Implementation:** Quality of coding, algorithm design, and data handling.
- ◇ **Team Collaboration:** Communication and teamwork.
- ◇ **Presentation:** Written and oral presentations.

7 Additional Information and Support

For any additional information or support, please send an email at georgios.peikos@unimib.it and david.labarbera@unimib.it.