Making medical science understandable
Large Language Model platform for patient-friendly content | Case study

Dominik Ślęczka, Krzysztof Odrzywołek, Lizaveta Sakovich, Patrizia Marciniak, Tomasz Drapała, Aleksander Siwek, Michał Czekański, Dominika Glanowska, Dawid Rymarczyk, Łukasz Nowak

ABSTRACT & BACKGROUND

The client aimed to develop a tool that will help patients to bridge the gap of understanding scientific language. The goal was not only to simplify complex content and streamline the workflow for specialists, but also to enhance the clarity and accessibility of health and medical knowledge. The project is designed to help patients better understand medical materials while enabling the organization to communicate in an open, clear, and modern way.

Our team has developed an AI-powered system that translates complex scientific documents into clear, patient-friendly language. The solution combines multiple approaches and utilizes Large Language Models (LLMs) to ensure accurate and understandable outputs.

PRODUCT DESCRIPTION

This project resulted in two separate cloud hosted web applications. The first application allows users to refine, transform and summarize scientific documents into an output that is understandable by patients. Users start by uploading documents and choosing some of the predefined options such as templates. Then the content is refined by users by requesting LLM to transform the text or create a summary.

A side-by-side editor displays original and modified text with highlighted changes, with version control that allows tracking, comparing, and restoring previous edits. AI can generate new versions of the text, simplify language, or adapt tone according to the user's instructions, while allowing manual review before the AI changes are applied. The system also supports document summarization and maintains formatting. In a later phase, the system was extended with capabilities to generate visual elements and infographics tailored to simplified documents, further supporting clarity and patient comprehension.

To enable customization of generated text, a concept of a Lexicon was introduced. A Lexicon which maps specific scientific phrases to their easier to understand counterparts, preferred by users. The second application allows users to create such Lexicons, by combining Social Media Posts and LLMs.

COMPUTATIONAL FLOW

The original scientific document is processed by the LLM with the assumption that apart from the source document, additional context is provided, that help to accomplish desired output requirements.

The first type of enrichment is a lexicon. It is a domain specific dictionary of scientific terms and their simple counterparts. In case of a company having particular preferences towards the vocabulary used by target audience, using a lexicon allows for easy management of such dictionary. It is useful to adhere to specific narration of the output document.

To enable creating reproducible results, without requiring the user to type long and specific prompts, a concept of Strategic Documents has been introduced. After establishing proper and precise requirements, this type of additional context can contain guidelines and templates toward the expected output document. This approach enables creating output for multiple scientific documents, while making sure that all these documents have similar layout, narration and target the same target audience. In case when apart from the source file, additional document contain more contextual information, a user can upload more documents using the "Additional Documents" feature. The uploaded documents are attached to the context provided to LLM.

Apart from working on text documents, the application supports creating image output. The intended type of image to be generated is an infographic. This combines well with the workflow intended for text creation. The proposed workflow allows the user to work and refine the text visible on the image as a first step. Afterward, working image on the a second step. To create an infographic, the user uploads all desired context, and using existing textual workflow refines the description that will appear on the resulting infographic. After being defined, the description is used as context for the result image.

METHODS / TECHNOLOGIES

Both applications are implemented using Python's Streamlit framework. The underlying connectivity with LLMs is handled using Python's langgraph and langchain libraries.

Handling large textual documents imposes potential problems when combined wit LLMs. The most probable issues that can happen are either running out of context windows, or low quality results. An approach used in the both applications is Retrieval Augmented Generation (RAG). This allows the application to extract only relevant pieces of information from the source documents. This allows for smaller amount of information sent to the LLM without sacrificing quality of output.

One of the most time-consuming parts of the project was prompt engineering in the sense of creating an optimal system prompt, creating guidelines toward specific requirements that the output documents need to meet.

RESULTS AND DISCUSSION

Our efforts resulted in an application which satisfied the needs of our client. It is both universal and extensible, as it allows for fulfilling various requests expressed by user, while enabling creation of reproducible content. It is easy to extend the application capabilities by adding new strategic documents containing guidelines for new types of documents.

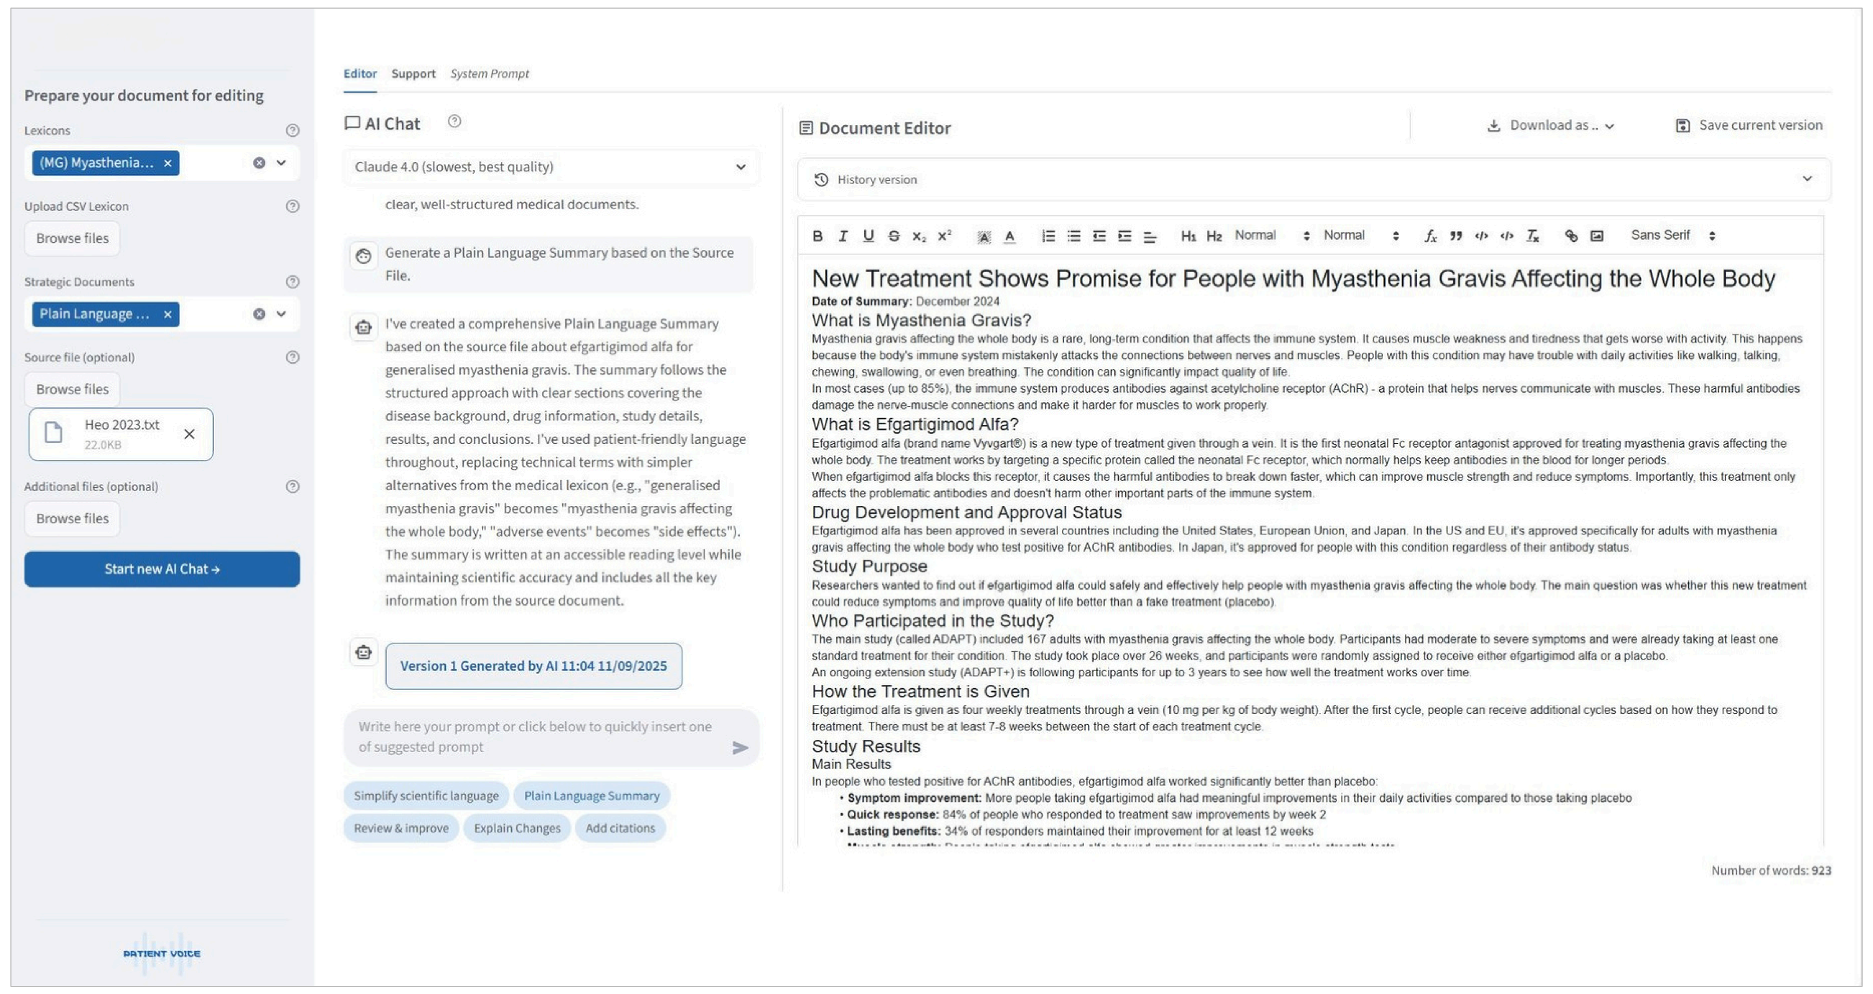

Figure 1: Main view of the Phase 1 Application. The LLM has generated the document, the user can edit the document manually or request the LLM for additional changes.

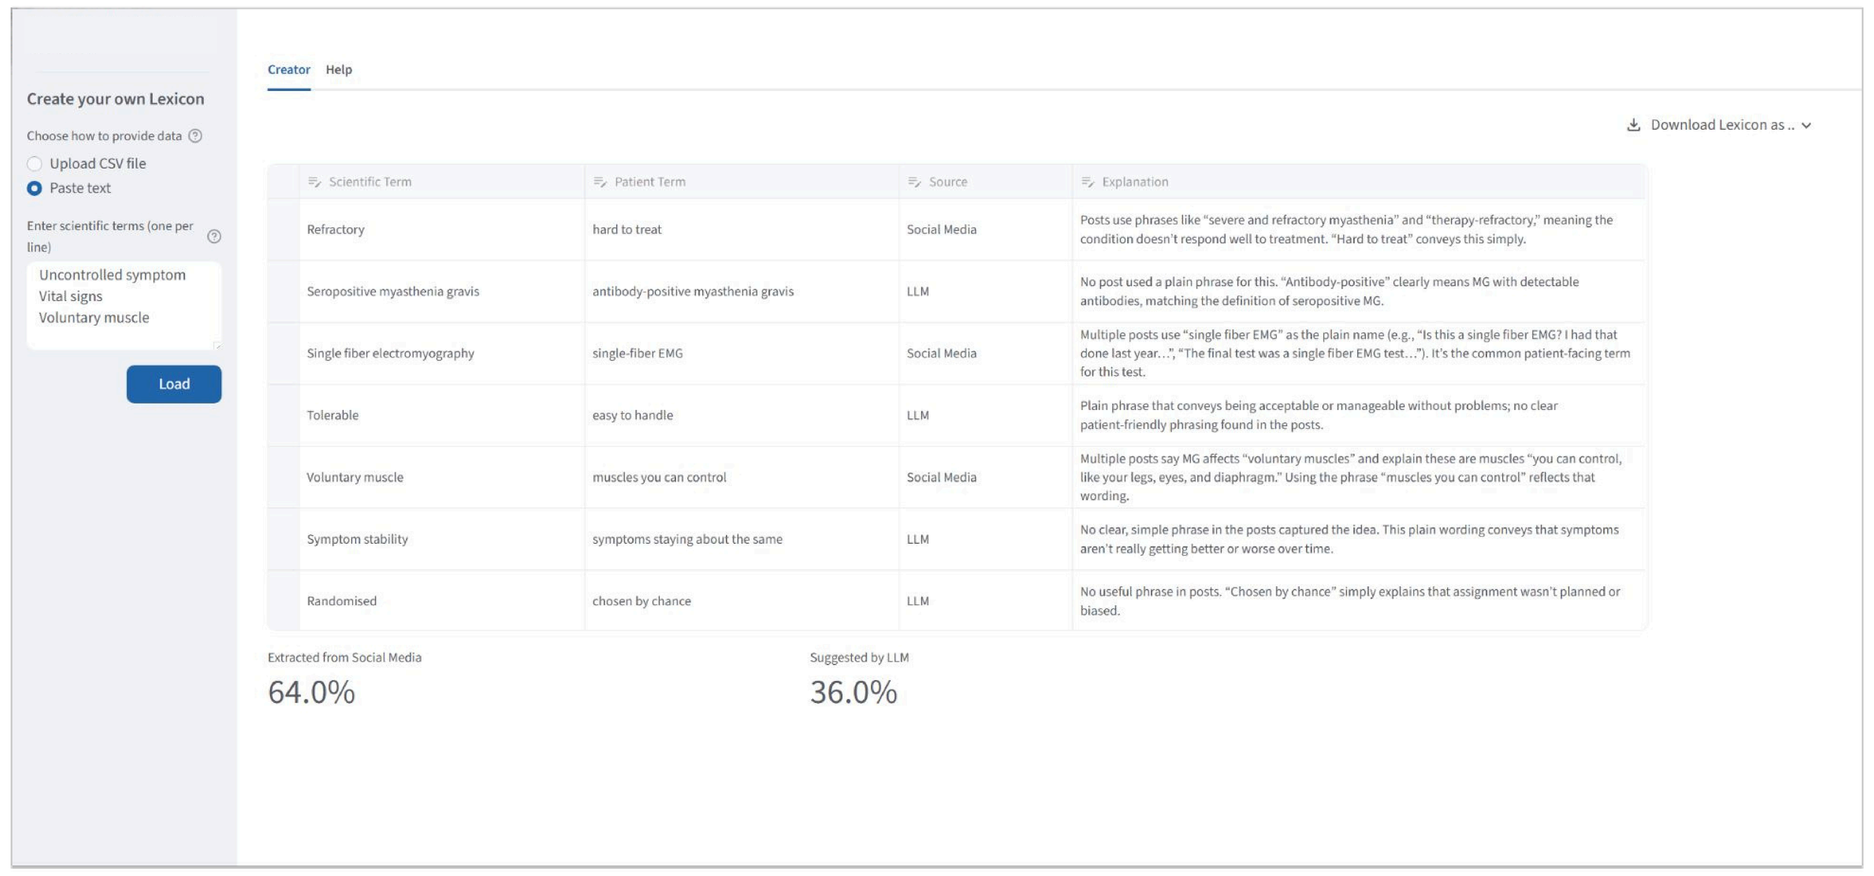

Figure 3: Main view of the Phase 2 application - Lexicon Creator. After providing a list of scientific phrases, the user can review suggested plain language counterparts. Suggestions are provided by the LLM which has access to Social Media Posts relevant to the analysed domain of knowledge.

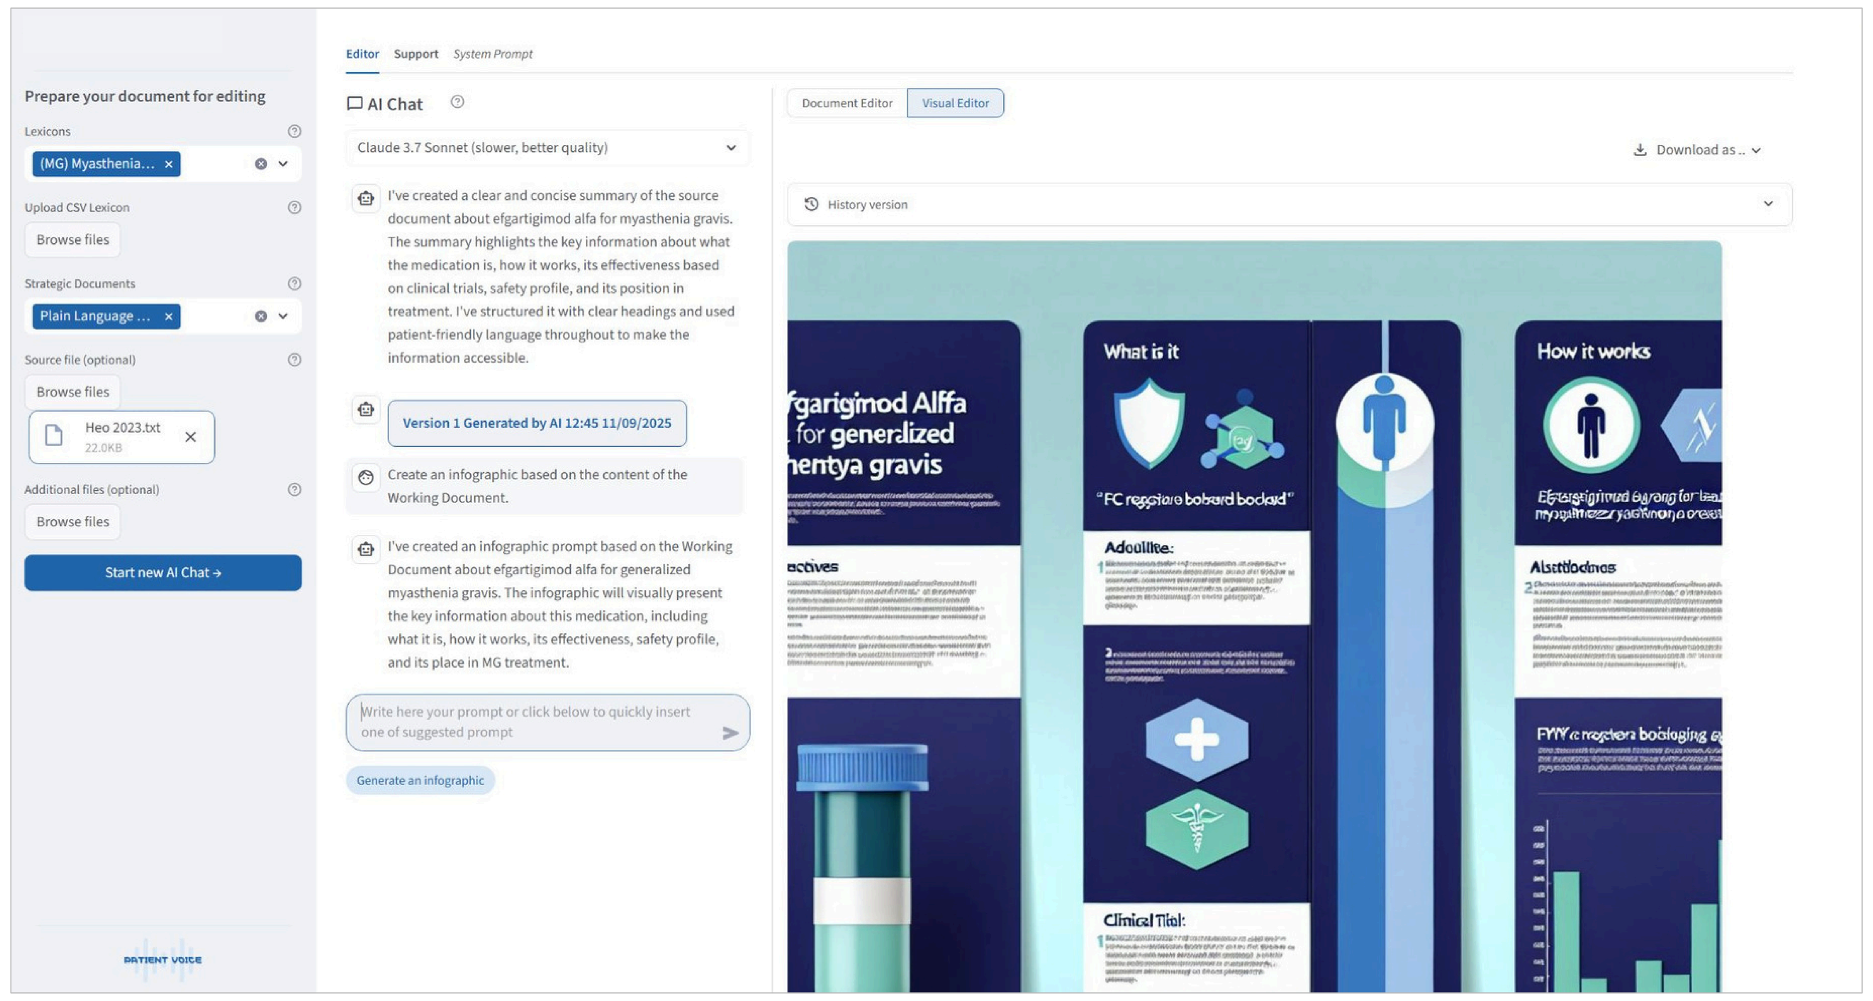

Figure 4: The view of generated image created using Phase 1 application. The initial version of image generation uses Dall-E 3 model, which will be replaced by another model in the future, due to low quality of image output.

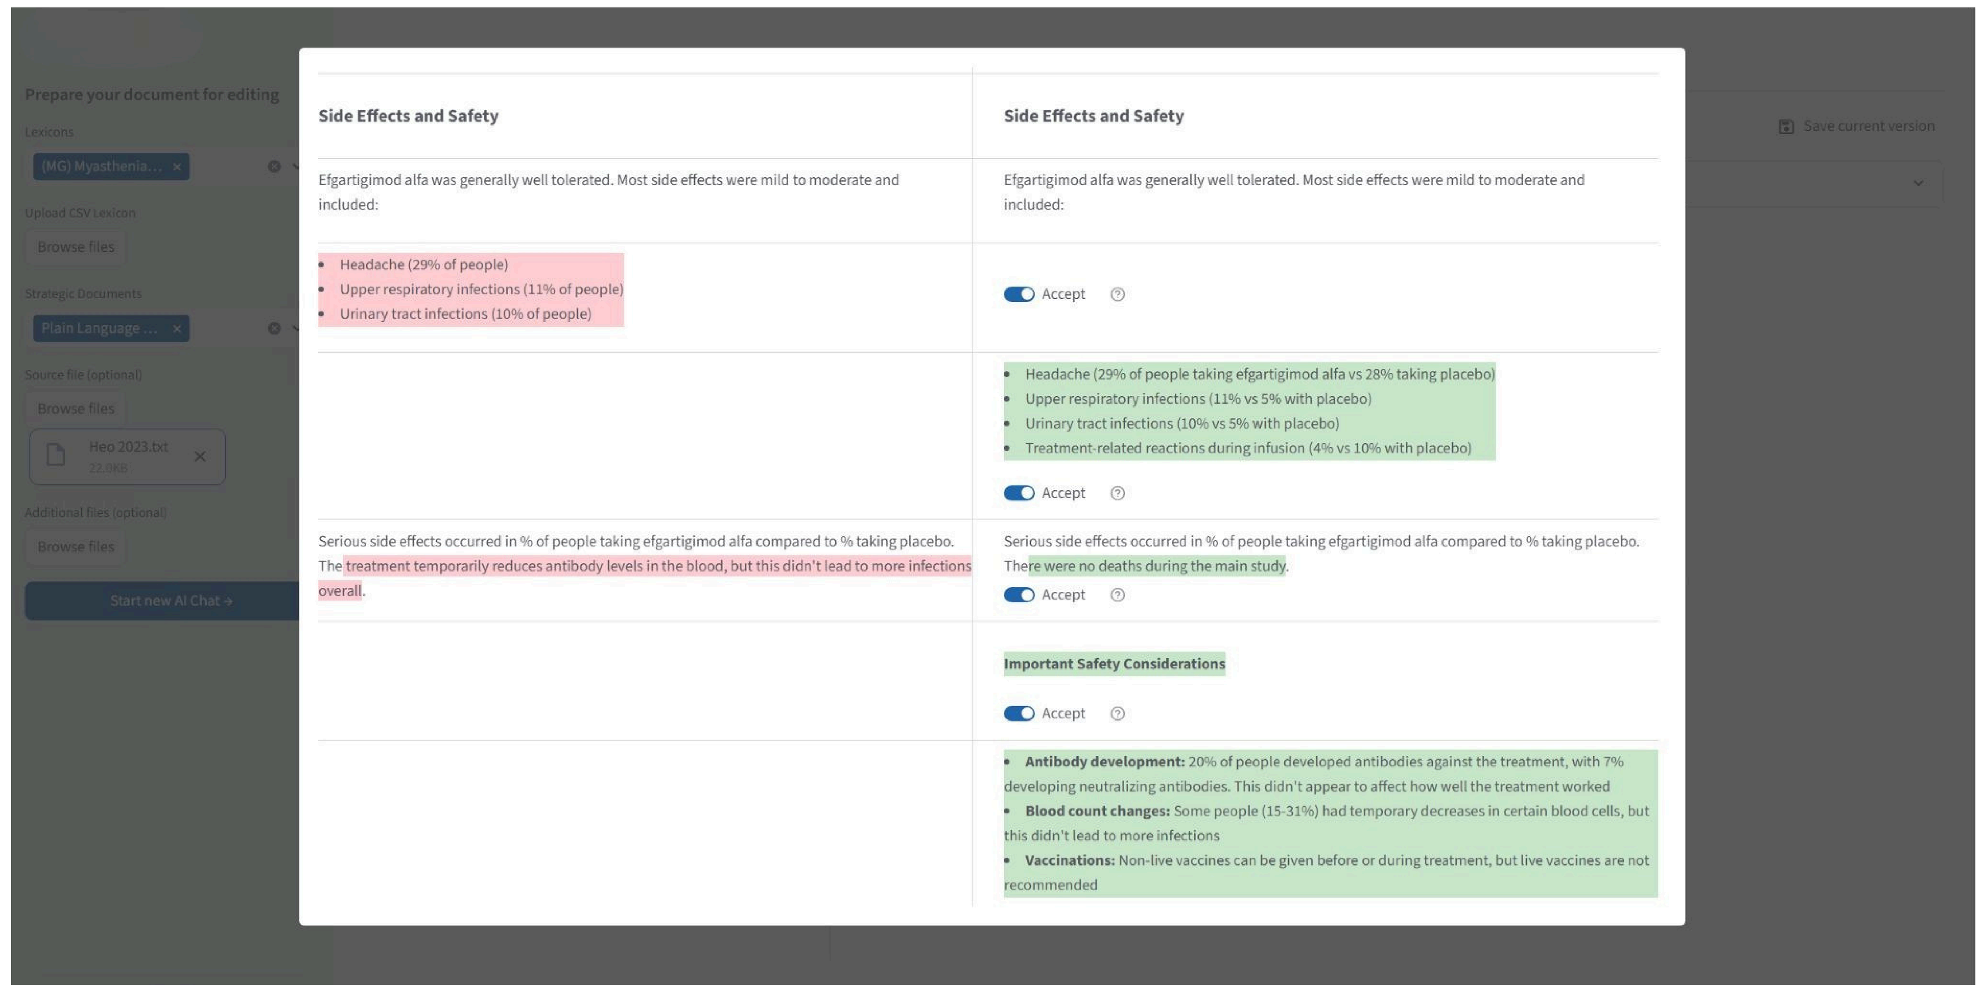

Figure 2: Change comparison view, with ability of accepting or rejecting changes. The previous version of the document is on the left. The newly created version is on the right. Green and red highlight additions and deletions when compared to the previous version. The user can accept or reject all changes at once using the buttons at the top of the popup, or select only the changes that are desired.

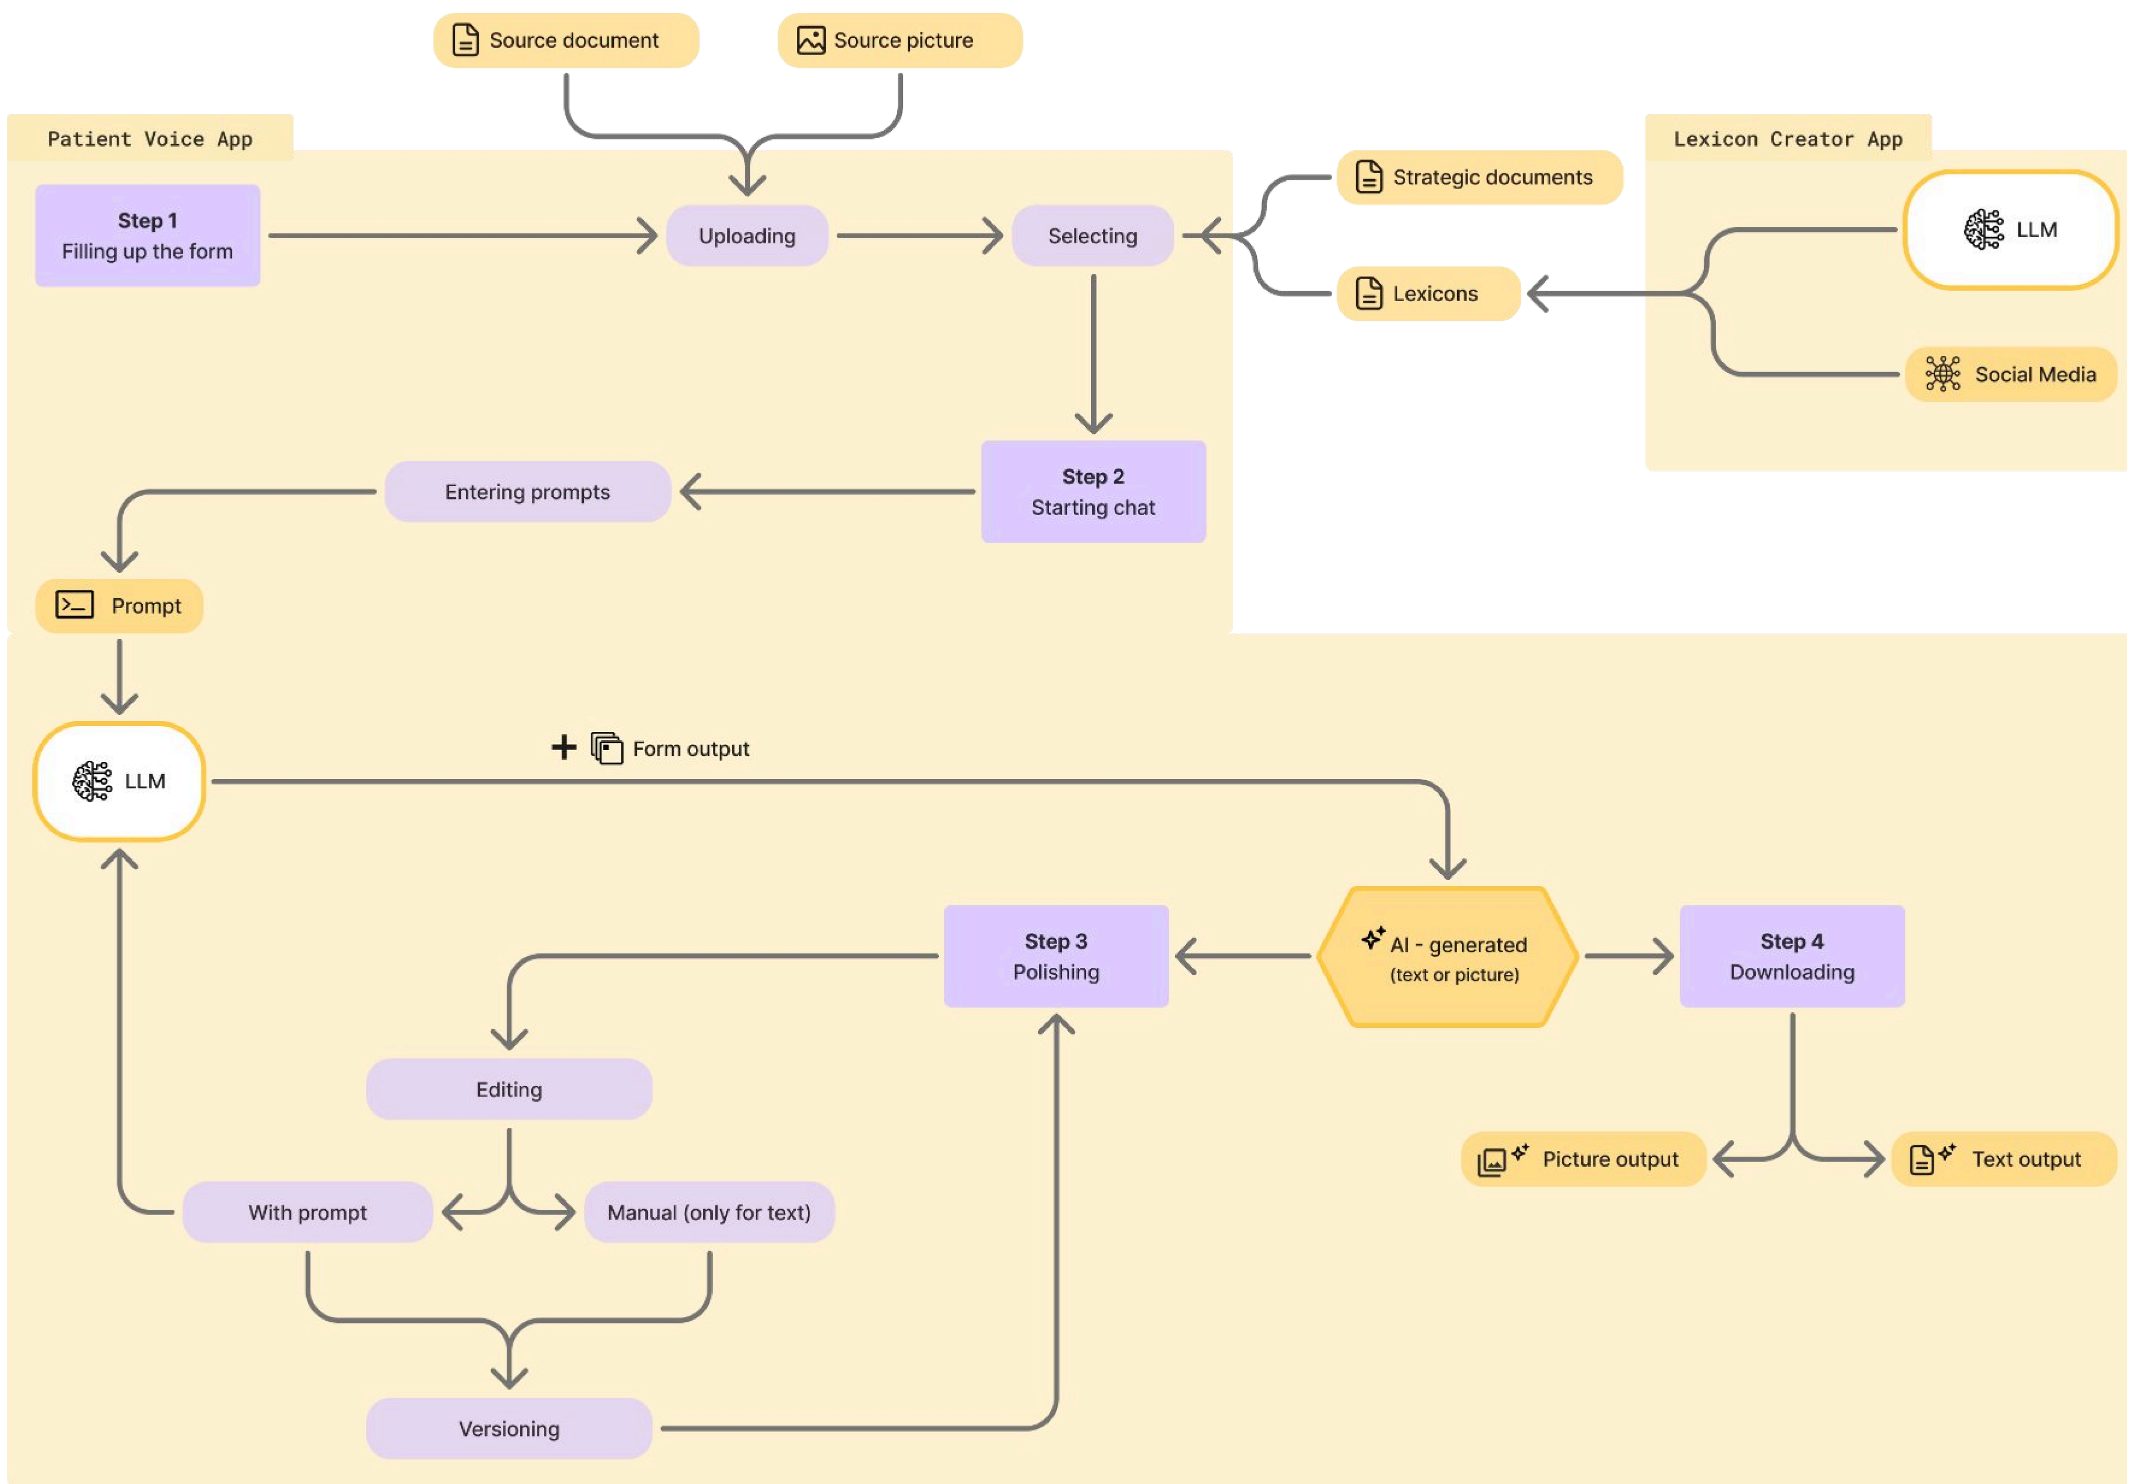

Figure 5: A diagram illustrating the user interactions and data flow within the system, showing how information is uploaded, processed, transformed, and presented back to the user.

IMPACT

This project equipped the client with an AI-powered tool that transforms complex scientific materials into clear, patient-friendly content while ensuring consistency and efficiency in the editing process. The solution helps specialists save time, maintain a unified communication style, and deliver health information that patients can more easily understand and engage with. By improving clarity and accessibility, the project supports stronger patient trust and empowerment, while also positioning the organization as an innovator in transparent and modern healthcare communication.

Download the poster here: