

## **ChatGPT on top of your data: Implementing intelligent enterprise search in combination with an LLM**

**Dr. Moritz Mueller – moritz.mueller@squirro.com**

### **1. About the Case Organization**

This is an insurance company from Switzerland. They have about 200,000 private clients as well as about 15,000 institutional clients. Their main business is sale of insurance policies (health, accident, travel, etc.).

### **2. About the Challenge**

The main challenge for the insurance is that they need to provide customer support. Especially for health insurance, there is typically a major peak before the end of November every year when people in Switzerland can change their health insurance provider for the following year. Additionally, there are numerous support requests via email as well as via the telephone support. Most of the customer requests are in German, French and Italian. The insurance company has a call centre team of 20 people which is supported by additional resources during peak times. As a knowledge base, they have numerous FAQ documents, insurance descriptions and any kind of PDF documents that the customer service staff needs to consult to answer the questions.

Typical support questions are:

- Which services are covered by insurance package ABC.
- Which insurance limits exist for insurance package ABC.
- Etc.

### **3. What We Did**

We provided the insurance company with our SquirroGPT solution<sup>1</sup> which combines the information retrieval system of Squirro with the capabilities of a Large Language Model (LLM), in this case ChatGPT. We simply loaded the selected internal knowledge base, i.e., the internally available FAQ and service & product description documents into Squirro. Then we provided the customer support staff with a chat functionality which answers their service specific questions based on the available documents.

Technically, the approach can be split into 2 main steps: First we conduct a semantic search with the question on the loaded documents. Then once we have identified the pages that are relevant to answer the question, we pass those selected pages together with the question to ChatGPT to get an answer to the question. Since we control the content which is passed to the large language model and we validate that the answer is based on those documents, we can also provide the respective evidence to the question. That means the support staff do not only get an answer to their question, but we also provide a direct link to the page inside the document on which the answer is based on. This speeds up the answering process significantly and thus has of course also very significant cost-saving effects.

The main innovation here is that we combine the Search Engine Squirro with ChatGPT. We need the Search Engine to first identify the relevant document which contains the

---

<sup>1</sup> <https://squirro.com/enterprise-generative-ai-and-large-language-models/>

information to answer the question. Only then we can pass on the question with this evidence to the LLM.

#### **4. Challenges and Lessons Learned**

The initial hurdles were concerns in terms of data privacy and potential “hallucination” of the LLM. Data privacy is an issue if sensitive content is fed to the LLM, where it then has the potential to re-emerge in replies to other queries. “Hallucination” can occur because LLMs work by predicting likely next words in a given sequence, based on their learning set. This produces highly plausible output, but it is not necessarily true.

The data privacy issue is addressed by using the Azure OpenAI API which guarantees that no data is used for training purposes of the LLM and any content passed to the API stays private.

The “hallucination” component of the LLM is handled by always providing the evidence, that means the actual document on which the answer is based on. Customer service staff can then still look at the underlying FAQ and product factsheet to substantiate their answer. We also conducted some training on how to best formulate the question to get relevant answers.

It is also important to make sure that the underlying FAQ documents are kept updated. With the Squirro Search Engine we can use the search ranking functionality to make sure that the latest version is used.

#### **5. Impact and Benefits**

The practical business benefits of our solution are massive timesavings for handling individual customer requests. We can already now in the initial phase measure that customer calls are answered faster as the customer support employees find the information faster. We also have direct feedback from the customer service teams that our solution helps them enormously during their daily work. We are now also getting suggestions of which additional documents to add to the knowledge base and are working on a functionality to flag outdated knowledge base content for review & correction.

One of the main reasons for our success was the easy and fast setup of our solution. We had a head start, as the customer had an existing knowledge base and we simply needed to upload those documents into SquirroGPT. There was no further customization needed.

#### **6. Next Steps**

As a next step, we intend to further automate the customer service process by using our SquirroGPT solution to pre-formulate answers to customer emails. For this we are passing the customer email to the chat functionality and letting the LLM formulate an automatic answer to the customer request based on the provided knowledge base. The service agent will then still have to read through it and can make changes, but it creates very significant timesaving in terms of answering customer requests by email.

**Note:** There is a free trial available by Squirro to test the described functionality on your own documents: <https://squirro.com/apps/generative-ai-trial/>