Credit Mutuel: Using AI to get the right information to customer advisors

Summary

The drive for efficiency in customer service is often at odds with a desire to deepen the customer relationship. Specialists are better able to address questions related to specific products or services but can lack the context necessary to serve customers whose relationship with the company is broader than one domain alone. At Credit Mutuel, AI allowed them to scale dedicated points-of-contact for customers across many products by providing curated information to their customer reps.

Opportunity: Enhance the service provided by human advisors

At Credit Mutuel, each customer has a dedicated advisor. The advisor acts as a first point of contact, helping customers to navigate their relationship with Credit Mutuel across various products (e.g., checking, savings, mortgages, investments). The quicker and easier it is for an advisor to access relevant information, the quicker they can respond to customer requests, and the more time available to serve other customers. With around 3 million calls and 7 million emails received per month, any improvement in resolution time can have a significant impact.

Challenge: Inconsistent documentation across products and groups

The challenge of having a single advisor across many products is the burden that it places on that advisor to have the necessary information at their fingertips. To resolve a customer query, advisers (who are typically generalists) use internal search engines or phone calls to source answers about specific products. But individual banks in the Credit Mutual network organize their information differently, which complicates the search. Moreover, the language and terminology can also differ. This means that typical, off-the-shelf language models are insufficient for prioritizing the information presented to these advisors.

Solution: Custom word embeddings and hierarchical classification

To create a language search customized for their products, Credit Mutuel first collected all the questions faced by their customer advisers over a 3–4-month period, and then curated answers to those questions (which took an additional 4 months), repeating this effort for each of 11 business domains currently in production. Then they trained a Deep Learning model for custom word embeddings and used this to train an individual Support Vector Machine (SVM) model for each domain, to select the answers most likely to address each question. They also built tens of thousands of dialog steps to support the collection of any missing information from the initial question. The initial domain classification (which in this setup could focus on only short, simple opening questions) was developed using a FastText[1] model which performed as well as the next best attempted (i.e., BERT), but was much quicker, yielding an F1 score of 90% with only 10-15s weekly training time and 20-30ms classification time. Splitting in this way helped to minimize number of classes in each domain specific SVM model.

Outcomes: Improved answer quality and quicker call resolution

The improved language models enhanced the quality and speed of answers. The Virtual Assistant is now able to provide good answers to 85% of customer cases (and 2 million additional answers to customers each year), while also reducing the time to resolution from 3 minutes to 1 minute on average. This overall time saving (for customers and advisers) was in the order of tens of thousands of hours each month.

In this case, we see a great example of AI being used not to solve a problem (i.e., provide an answer) directly, but as an integral part of a human-led workflow, generating a smaller and more targeted set of proposed outputs where humans can apply their subjective judgement.

[1] An open source NLP library developed by Facebook AI. https://fasttext.cc