CONVERSATIONAL CAPITAL

What did your customer come for?

How the Conversational Capital Rating (CCR) sorts every conversation with an assistant by the reason the customer came.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Taxonomy · Contact reasons · Evaluation · AI assistants

Before you can say whether a conversation went well, you have to know what the customer came for. A customer checking a balance and a customer whose claim has stalled for a month both deserve a good conversation, but good looks different in each. The first should be over in seconds. The second needs time, care and a real outcome. Judge both by the same yardstick and the result tells you very little.

To do this, we use a taxonomy – a list of reasons customers contact an organisation. Every conversation is given one of these reasons, and that determines what counts as success, how quickly the conversation should move and how much is at stake. Using this taxonomy means different organisations, or the same organisation a year apart, are measured against the same standard.

01 Summary

WHAT IT IS	Each conversation is given one reason for contact from a controlled taxonomy: a core every organisation shares, plus extensions for different industry sectors.
WHY IT MATTERS	The reason for contact sets the standard the conversation is judged against.
HOW IT WORKS	A language model reads each conversation against written definitions of every reason and says how sure it is. Uncertain cases go to a stronger model that also sees how experts ruled on similar cases.
HOW WE KNOW IT WORKS	Checked against conversations ruled by experts, the system's label matches the expert ruling reliably, on everyday and complex conversations alike.[1]

02 How it works

The taxonomy is a list of contact reasons. The core covers the kind of things every organisation sees, such as requests for account access, a disputed charge, a fault or a complaint. Industry sector extensions add reasons such as a claim for an insurer or moving a number to a new network for a telecoms provider, meaning that

conversations are measured in their industry context. A language model[2] is used as a closed-set classifier that reads each conversation alongside the full list of reasons, each with a definition and notes on where it ends and its neighbour begins. It works through a fixed sequence of questions, assigns one reason and states how confident it is. Labels it is unsure of go to a second, stronger model, which also sees worked examples of how experts ruled on similar cases.

The contact reason helps decide how quickly the conversation should move and how much it was worth to the customer. Take an insurance customer who writes: “Three weeks and still no word on my claim.” Read as a complaint, the conversation is judged on how well it repairs goodwill. Read as a claim-status check, it is judged on whether the customer got a clear answer quickly. The taxonomy's definitions settle which reading is correct.

The taxonomy currently covers banking, insurance, telecoms and retail, with public services and higher education in draft and other sectors in development.

03 What it means for clients

Results are comparable with others in your sector and with your own over time, because every conversation is sorted the same way before it is scored. And they are fitted to you: the first run on your conversations is treated as calibration, with a sample of labels checked by an expert before any CCR is reported.

1 Reliability is the share of conversations where the system's reason matches an expert's ruling, reported as two figures with 95% confidence intervals. For everyday conversations this is 88% (80 to 92%), and for complex conversations it is 82% (78 to 86%), measured on a standing set of hard cases ruled by human experts. These figures describe how reliably conversations are sorted by reason for contact, one step in producing a CCR. They are not a measure of the accuracy of the CCR as a whole, which depends on every step in the process.

2 Anthropic is our model provider. Where conversations are processed is agreed with each client before any data is sent, every model call is logged, and client data is not used to train models.

Conversation Design Institute

About Conversational Capital and the CCR

Conversational Capital is the trust, goodwill and willingness to engage that an organisation builds up when people feel respected, understood and helped in their conversations with it, and loses when they do not.

The Conversational Capital Rating (CCR) is how we measure it. It weighs two sides of every conversation with an assistant: the value it creates, by handling what the customer came for well and at the right pace, and the trust it spends, where the conversation breaks down or the customer leaves worse off than they arrived. The outcome is a single rating out of 100, with the evidence behind it, and a priority list of how to raise it, and so build the trust your brand depends on. It is built for assistants that use large language models and increasingly do things for the customer, not just answer their questions.

The CCR turns the quality of your AI conversations into a measure you can track and act on. It shows which kinds of conversation are building trust and which are spending it, and where you should invest first. It also shows what the assistant really saves, and what it is really costing you in repeat calls, lost customers and trust, months before those costs reach your accounts, so that further investment can be judged on its return.

