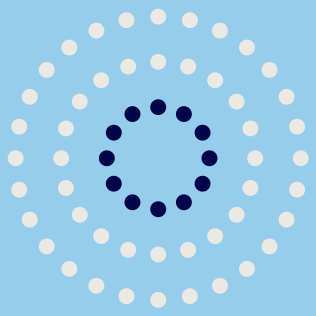


CONVERSATIONAL CAPITAL

Do your conversations move at the right pace?

How the Conversational Capital Rating (CCR) judges whether a conversation with an assistant moved at the pace the customer's need called for.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Velocity · Pace · Evaluation · AI assistants

A customer who has lost their card wants it stopped, now, in as few words as possible. A customer telling their insurer that their partner has died needs room to say what happened. Both conversations can go well, but they go well at very different speeds. Measure them by the same clock and you reward the assistant that hurries the grieving customer and penalise the one that gives them time.

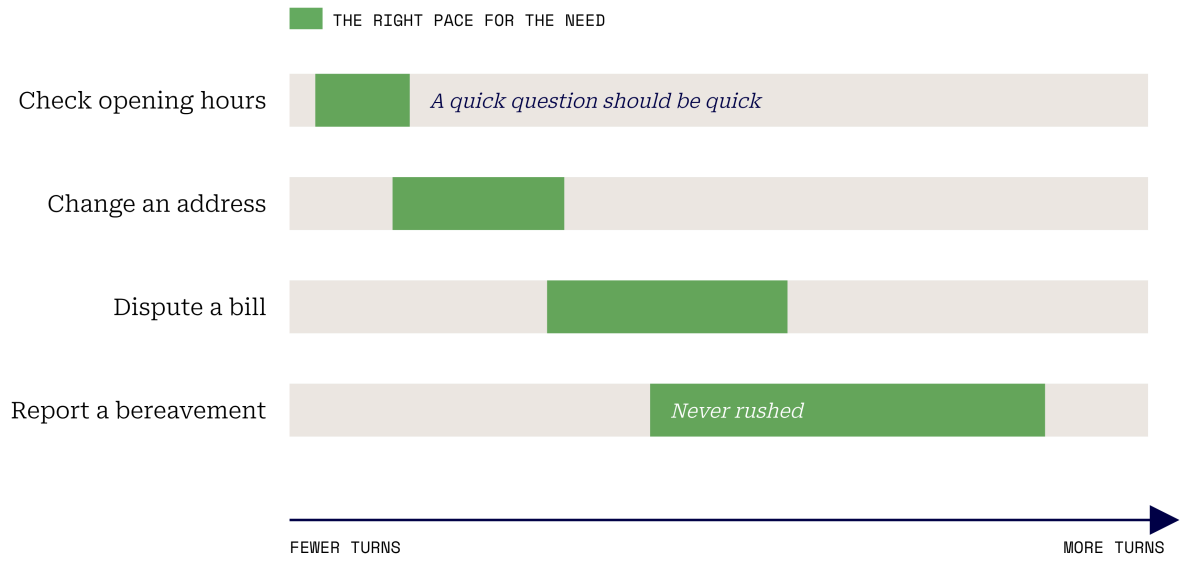
Velocity is how the CCR measures pace. It does not ask whether a conversation was fast. It asks how much of the customer's effort moved their request forward, judged against the pace that kind of conversation should take.

01 Summary

WHAT IT IS	A score out of 100 for how much of the customer's effort in a conversation moved their request forward, against the pace that kind of request should take.
WHY IT MATTERS	A conversation can end politely and still waste the customer's time. Repeated questions, details asked for twice and steps that could have been one are what customers remember.
HOW IT WORKS	A language model reads each conversation against the pace that kind of request should take, and must quote the customer's words for every step it counts as wasted.
HOW WE KNOW IT WORKS	The model's readings agree with expert review on the hardest conversations, and the same conversation judged twice gives much the same score.[1]

02 How it works

Every reason for contact in the CCR's taxonomy carries an expected pace. A lost card should be dealt with in a few turns; a disputed bill needs some diagnosis; a complaint or a new claim needs room for the customer to explain what happened. If a customer discloses a bereavement, serious illness or vulnerability, the conversation is judged as needing room, whatever it was about.



The right pace depends on the need: a quick question should be quick, a bereavement should not be rushed.

A language model[2] reads each conversation against that expected pace and identifies the customer effort that did not move the request forward: repeating something already said, supplying details a second time, chasing an answer, or steps the assistant could have avoided. Each must quote the customer's words, and a separate check confirms the quote is really there before it counts. Conversations that were rushed, or that left the customer stranded, are flagged too.

Velocity is reported out of 100 across the conversations that can be judged. One that moves to a private channel or a phone call before the request is dealt with is left out, because what happens next cannot be seen.

Take a telecoms customer whose broadband keeps dropping. The assistant asks for their account number, then, in a separate message, their postcode, then their date of birth, before it starts on the fault. The fault is then fixed in three more turns. Nothing went wrong in a way the customer would complain about, and the conversation ends well. But three of the customer's six turns were spent on identity checks that one question could have covered, so most of that effort counts as waste and the conversation scores well below what its outcome suggests. That is the pattern Velocity exists to find: conversations that resolve, politely, in more steps than they needed.

03 What it means for clients

Velocity points to the effort your customers spend that they should not have to. Because every wasted turn comes with its reason and the customer's own words, the

evidence shows where the time goes: identity checks split across several messages, questions the customer has already answered, answers that send them round again. Those are design fixes, and the score shows whether they worked. And because pace is judged by topic, an assistant is never rewarded for hurrying a customer who needed time. As with every part of the CCR, the first run on your conversations is treated as calibration, with a sample of readings checked by an expert before any CCR is reported.

1 Reliability is reported as two figures with 95% confidence intervals. For complex conversations, the model's readings agreed with the expert on all of them: 100% (84 to 100%). For everyday conversations the figure is consistency rather than expert agreement: the same conversation judged twice, independently, lands within 10 points of its first score 86% of the time (74 to 93%). These figures describe 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.

