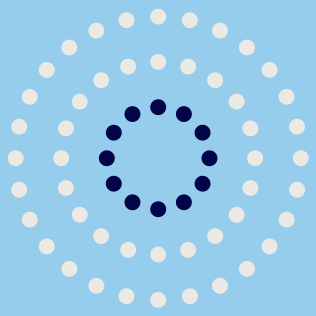


CONVERSATIONAL CAPITAL

Does your assistant handle conversations well?

How the Conversational Capital Rating (CCR) judges how well an assistant handled what the customer came for, from getting the answer right to leaving them glad they asked.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Craft · Conversation quality · Evaluation · AI assistants

Two assistants can both close a conversation with “Is there anything else I can help with?” One has told the customer something that is not true. The other has fixed the problem, confirmed what happens next and remembered that the customer called about the same thing yesterday. Most dashboards score them the same, because both conversations ended and neither customer complained in the chat.

Craft is how the Conversational Capital Rating (CCR) tells them apart. It does not ask whether the conversation ended, or how quickly. It asks how far up the ladder of good service the assistant got: whether it was right, whether it could be trusted, whether it actually helped, whether it used what it knew about this customer, and whether it did something the customer would remember.

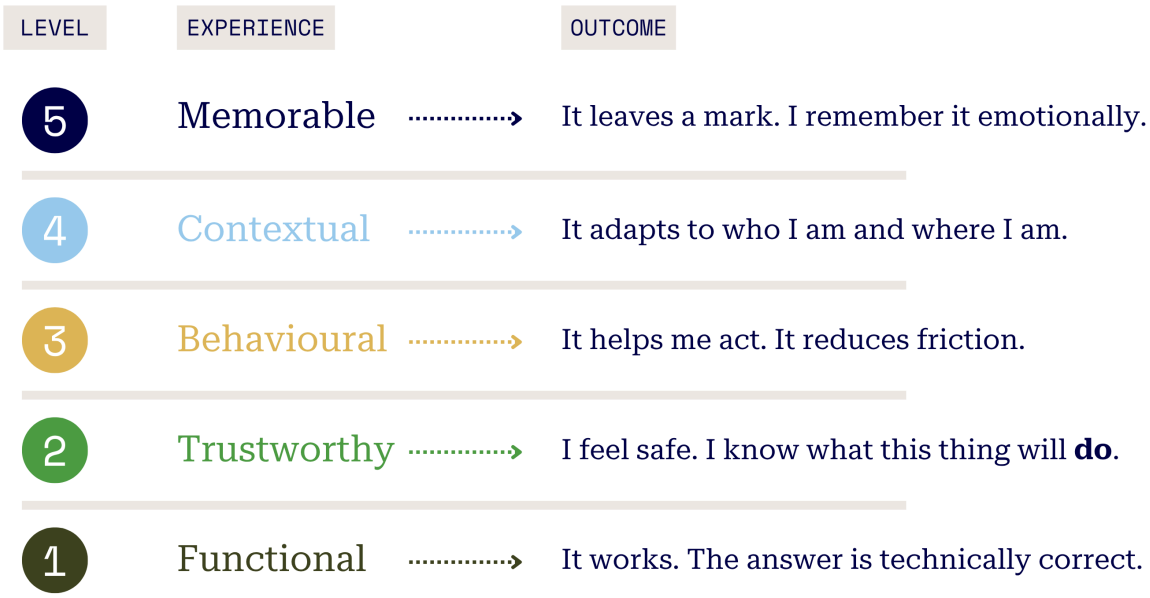
01 Summary

WHAT IT IS	A score out of 100 for how well the assistant handled each conversation, placed on five levels of service quality, from simply working to genuinely memorable.
WHY IT MATTERS	Customers do not average their experiences. One confident wrong answer, one dismissive redirect or one problem left with them to solve outweighs many smooth exchanges, and it is what they remember about the brand.
HOW IT WORKS	A language model reads each conversation and answers a short set of yes or no questions for each level, in order, quoting the turn that decides each one. The conversation stops at the first level it does not reach.
HOW WE KNOW IT WORKS	The same conversation judged twice lands on the same level almost every time, the readings agree closely with expert review on the hardest cases, and assistants that completed the customer's task score clearly higher than those that did not.[1]

02 How it works

The five levels come from the Conversational Capital framework, and each level includes everything below it.

- Functional: the assistant answers correctly and does what it claims.
- Trustworthy: it knows its limits, admits what it does not know, keeps the customer's dignity and hands off cleanly when it should.
- Behavioural: it guides the customer through and gets the thing done, rather than just reading out information.
- Contextual: it uses what matters about this customer, what they said earlier, their history or their situation, so that the help fits them.
- Memorable: it anticipates a need the customer had not raised, or, when it cannot help, closes with clear expectations of what happens next.



The five levels of the Conversational Capital framework, as the customer experiences them.

A language model[2] reads each conversation with the reason for contact already identified, together with a short note of what must be said or done for that kind of request. A refund must say when the money will arrive; a payment change must confirm what was done. It then answers yes or no questions for each level in turn, citing the turn and the words that decide each answer. The conversation is placed at the last level where every answer was yes, with partial credit for progress towards the next. A set of conversations ruled on by an expert shows the model where the borderlines fall.

Take a telecoms customer whose mobile data has stopped working. They mention that they called yesterday about the same problem. The assistant checks the account, explains honestly what it can and cannot see, walks the customer through two settings on their phone and confirms the data is back. That conversation is Functional, Trustworthy and Behavioural, and scores in the 60s. It does not reach Contextual, because it started from scratch: it never used yesterday's call, and the customer repeated the checks they had already been through. Had it picked up where yesterday left off, it would have scored in the 80s. Had it told the customer the fault was fixed when the phone still showed no signal, it would have scored below 20, however polite the rest of the conversation was.

03 What it means for clients

Craft shows where your conversations stop on the ladder and why. Because every reading names the question that stopped it and quotes the conversation, the evidence points at specific fixes: answers given without checking, redirects that say nothing about the customer's problem, tasks left for the customer to finish, details asked for twice. These are the fixes conversation design makes, and the next rating shows whether they worked. Because each kind of request carries its own note of what good handling must include, the assistant is judged on what that customer needed.

1 Reliability is reported as two figures with 95% confidence intervals. For everyday conversations the figure is consistency rather than expert agreement: the same conversation judged twice, independently, lands on the same level 94% of the time (89 to 96%). For complex conversations, the hardest cases in development, the model's level matched an expert reviewer's ruling 84% of the time (65 to 94%), and was within one level 96% of the time (80 to 99%). 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.

