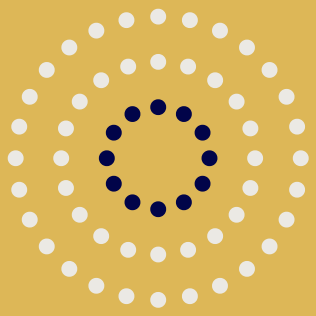


CONVERSATIONAL CAPITAL

How little data do you need to get a Conversational Capital Rating?

How the Conversational Capital Rating (CCR) is built from data your organisation already holds, why so little of it is needed, and what it gives back.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Data · Privacy · Value · Evaluation · AI assistants

When an organisation considers exploring a CCR for its assistant, the first question usually is ‘what do you need from us?’ Rating an assistant sounds as if it should need a large extract of customer records, a new integration and months of work.

The Conversational Capital Rating (CCR) needs none of that. It is built from three things your organisation already has: a sample of your assistant's conversations, anonymised, a short spreadsheet of counts your contact centre already reports, and one cost figure from your budget. From this data we return a rating, the evidence behind it, and what your assistant delivers in financial terms.

01 Summary

WHAT IT IS	The data behind a rating: a sample of conversations, a few counts by topic, and the cost of a contact with a person.
WHY IT MATTERS	A rating that needs little and identifies no one can start in weeks rather than months, without stalling in a data project.
HOW IT WORKS	Conversations show how well the assistant handles what customers come for. The counts show whether customers then went round it to a person. By understanding the cost of a contact, we provide an analysis of the financial value your assistant is creating or losing.

02 How it works

Each of the three things you send answers a different question, and each comes from a system you already run.

WHAT YOU SEND	WHERE IT COMES FROM	WHAT IT TELLS YOU
A sample of conversations, a few hundred per channel, masked[1]	The assistant platform's standard export	How well the assistant handles what customers come for, at the right pace, and where conversations break down or leave customers worse off
Counts by topic and month: conversations the assistant handled, contacts that reached a person, and, if you have it, how many customers came back[2]	The contact centre's routine reporting	Whether customers go round the assistant, topic by topic, and how that has moved over six months
The cost of one contact with a person	The contact-centre budget	What the assistant saves, and what it costs when customers come back

If you also have them, your churn rate and yearly revenue per customer let the CCR also put a floor under what the assistant is costing you in customers at risk.

A few hundred conversations are enough to show an assistant's habits. Take an insurer that sends 400 conversations from the last quarter, a spreadsheet of about thirty topics over six months, and its cost of a call to its contact centre. From the conversations, the CCR finds that the assistant handles policy changes politely and quickly. From the counts, it finds that a third of the customers whose policy change it completed rang the call centre afterwards to check. That is around 50,000 policy changes in six months, 16,700 of them followed by a repeat contact, and the cost of those calls set beside the savings on the rest.

03 What it means for clients

The ask is small enough to start quickly. There is no integration to build and no new tracking to switch on. What we need is an export of conversations and a simple spreadsheet.

What comes back is large by comparison: a rating out of 100 with its six parts, the share of customers going round the assistant by topic, the value it builds, the repeat contacts it causes and the customers at risk, in financial terms, and the fixes in the order that pays back first.

They give you something most organisations lack: a return on the assistant measured from what customers actually did, rather than from containment figures. Linking customers across the assistant and the contact centre turns an estimate of who came

back into a count, and your own churn figures put a value on the customers at risk.

You can use this once or keep using it. A Baseline Assessment rates the assistant at a point in time: one sample of conversations, six months of counts, one report and a readout. If you continue with continuous measurement, you send the same small set each month, so a change in the assistant, or in how customers respond to it, shows while there is still time to act.

1 Personal details should be masked before data leaves you, and you remain its controller; CDI processes it under a data processing agreement, uses it for your rating only and deletes the conversations once the report they were used for is done.

2 The share of customers who came back to a person is measured when you count it from your own linked records, estimated when only a sample of conversations carries the answer.

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.

