Using an AI interviewer for the first round of Outcome Harvesting

12 Nov 2026

> We have a chapter (Rassmann et al. 2026) in *Demokratie im Zeitalter Künstlicher Intelligenz: Praxis und Ethik KI-gestützter Evaluation* (Democracy in the Age of Artificial Intelligence: Practice and Ethics of AI-assisted Evaluation), a new book from the German Centre for Integration and Migration Research (DeZIM), edited by Olaf Kleist, Mirjam Weiberg and Simon Müller and published by transcript Verlag. The whole book is open access, and although the title is in German, our chapter is in English, along with several others. This page is a short summary of what we did and what we found, and the full chapter has the detail.

What we tried

Together with Kornelia Rassmann and Birgit Müllegger, we ran a pilot in Austria's KLAR! climate adaptation programme in Bad Ischl and Ebensee, where we built a "Harvest Assistant" in QualiaInterviews to do the first round of Outcome Harvesting interviews.

Anyone who has done Outcome Harvesting will know how much work it takes to get from someone's story to a proper outcome statement, with who changed what, when and where, why it matters and how the programme contributed. We wanted to see how much of that first step an AI interviewer could take on.

Who took part and how it went

We invited 39 people and 19 took part, ranging from steering committee members to residents who knew very little about the programme. The interviewer adjusted its questions to who it was talking to, so, for example, it didn't press residents on internal programme details they wouldn't have known about.

Interviews took around half an hour on average, and some people paused for a few hours and came back later to finish. Most found it easy and pleasant, with 13 of the 17 who answered that question describing it that way (one wrote "Thank you, it was really fun writing with you"), while a few found it a bit repetitive or long, which was useful to hear.

What came out of it

We used a custom GPT to go through each transcript and draft outcome statements, which gave us 56. A trained harvester then checked each one, merged duplicates where several people had described the same change, and removed those that weren't behavioural changes or weren't linked to KLAR!, which left 38. Only 6 of these met the strict SMART criteria, which sounds low until you remember that a first interview with a human harvester rarely produces finished outcomes either. The other 32 became leads for follow-up interviews, either with the AI or with a person. Sending out the interview link and running the AI analysis took around 15 minutes of our time, and we couldn't have interviewed that many people so quickly with the team we had available.

It also meant we heard from people outside the programme much earlier than usual, since the first round of Outcome Harvesting normally sticks to the people closest to it.

Letting people check their own outcome

The part we liked most was the summary at the end of each interview, where the Harvest Assistant wrote up the outcome in Outcome Harvesting format and asked the person to confirm it. Everyone agreed theirs was accurate (one said "almost correct"), although in two interviews the conversation ended before the summary appeared. It meant people could see, and correct, how their words were being recorded, which matters a lot in participatory evaluation and also saves a round of emails afterwards.

Getting the interview guide right

Writing the interview guide took *a lot* of trial and error. When the instructions were too polite or too general, the interviewer stopped probing early, and in one interview it moved on to a second outcome before the first had been properly explored. If you're planning something similar in QualiaInterviews, a few things helped us:

- say how many follow-up questions the interviewer should ask before moving on
- test the guide on a few sample interviews before going live
- check that it works well in your respondents' language
- stress test the interview and even play different personas (including difficult respondents) to make sure the AI-interviewer knows how to behave in different situations

Where the AI fits, and where it doesn't

In the end the chapter argues for a mix, with the AI interviewer helping to reach more people and humans doing the follow-up, the judgement and the sense-making with stakeholders. Reaching more people only helps if the analysis keeps up and if people trust the process, so the chapter also has a section on data protection, consent and who might still be left out.

If you'd like to try a Harvest Assistant of your own, or adapt the approach for another method, get in touch and we'll happily help you set one up.

We sometimes use AI to help draft these pages, but the final text is ours.

References

Rassmann, Powell, Caldas Cabral, & Müllegger (2026). *The Role of AI-Driven Interviewers in Evaluating and Promoting Democratic Processes*. In *Demokratie Im Zeitalter Künstlicher Intelligenz: Praxis Und Ethik KI-Gestützter Evaluation*. https://doi.org/10.14361/9783839465424.