




Artificial intelligence

By Bernard Reber


Can artificial intelligence (AI) replace human intelligence in the most complex cognitive activities? To test this hypothesis, Bernard Reber reflects on an experiment he took part in during the Citizens' Convention on the End of Life. In dealing with such a sensitive subject, he wonders about the ability of AI to summarise and faithfully convey the abundance of debates and the diversity of opinions expressed, and queries whether it can even provide useful advice simply on the organisation of a conference.



and intelligent choices



Bernard Reber is a moral and political philosopher and a research director of the Centre National de la Recherche Scientifique (CNRS) (National Centre for Scientific Research) at the Centre de Recherches Politiques (CEVIPOF) (Centre for Political Research). His work focuses on deliberation, pluralism, responsibility, evaluation, governance, argumentation, institutional design, environmental ethics and controversial technologies. He was a member of the college of guarantors of the Convention citoyenne sur la fin de vie (Citizens' Convention on the End of Life).



Deliberation is a complex, individual and collective cognitive activity that goes far beyond all the calculations and simulations that the tools produced by artificial intelligence (AI) researchers and engineers are capable of. I use the term AI for simplicity's sake. These tools, as varied as their fields of application, can 'reproduce human-related behaviour such as reasoning, planning and creativity', to use the European Parliament's definition. This definition, according to which they are capable of simulating, synthesising and providing the appearance of intelligence, is a generous one. It will still be a long time before they can achieve the level of human deliberation, and in particular that of ethics and politics, involving beliefs, emotions and responsibilities.

This should not distract us from what they can already offer or – above all – from what we would like them to 'learn'. An exercise such as democratic deliberation can indeed benefit from tools derived from AI research, provided we are intelligent about the choices we make. It is important, for example, to know at what stage of a deliberation to use them, for what purpose, and what can be promised. I will consider two options here – one real and the other imaginary. The first is an experiment I participated in that sought to improve the way in which the partial results of the Convention Citoyenne sur la Fin de Vie (CCFV) (Citizens' Convention on the End of Life) were communicated to the general public. The second relates to a conference I was to give on the ethics of AI, for which I asked ChatGPT, the conversational agent using generative AI, for advice on organising a citizens' convention. I'd also like to emphasise that any technology requires you to explain what you're trying to do with it. This is

also true if we want the technology to do what our brains do differently, and often much better. Our brains are still a largely unknown landscape – a constant source of amazement for the researchers who specialise in them. Technological simulations therefore take detours. The Greek word for technology is 'cunning' (*mètis*).

Where to position AI and what to promise?

Let's take the case of the CCFV, organised by the Economic, Social and Environmental Council of France from December 2022 to April 2023. I was one of the four guarantors responsible for ensuring compliance with the principles of the Convention. Because it was largely made up of exchanges and conversations among the 184 citizens drawn by lottery, experts and vetted witnesses, the Convention produced a mass of language data that could have been fed into and trained by an AI tool. AI processing could also have been applied to the exchanges of the governance committee, which decided on the procedures it considered most conducive to collective deliberation, framing the exchanges within the Convention and then implemented by facilitators.

Could these moments of exchange (organisation, debates and implementation) have benefited from AI and in what way? The possibility of using this tool arose on two occasions during the Convention. First, when a question arose about the form to give to the final version of his sub-group's manifesto, one of the participants suggested asking ChatGPT. As a salesman, he had already used this application to make advertisements. 'Don't waste your time with that,'



Closing session of the Citizens' Convention on the End of Life, in the debating chamber of the Economic, Social and Environmental Council, Paris, 2 April 2023.

replied his colleagues, who burst out laughing at the suggestion, which they considered inappropriate for such a delicate and serious task. As guarantor, I can attest to the great care that went into the collective drafting of this text, where the search for the right word sometimes involved bitter negotiations. This type of drafting is eminently political, in the noblest sense of the term.

The opportunity to test AI arose again at the end of April 2024, during the Convention's reporting weekend, when the President of the Republic and several ministers explained how they were dealing with the Convention's proposals. During this weekend, the director of a company came to present an experiment carried out by his team, with the help of staff from the Economic, Social

and Environmental Council, which he described as a 'global first'. The experiment, carried out in the months following the Convention, was designed to show that AI processing could improve the way debates are presented to the general public. It focused on a limited body corresponding to two specific sequences of the Convention, which the AI tool was asked to summarise.

'Thanks to artificial intelligence, we will now be able to make the work of the Citizens' Convention on the End of Life fully accessible to the general public and to the members of parliament who will be working on the text,' said the company director at the presentation. 'We are convinced that this revolution will change our relationship with institutions for years to come.' This promoter



Hearings held at the Citizens' Convention on the End of Life in 2022 included, from left to right: Haïm Korsia, Chief Rabbi of France; Antony Boussemart, Joint President of the Buddhist Union of France; Mgr Marc Alric, Bishop of Neamt and representative of the Assembly of Orthodox Bishops of France. Christian Krieger, president of the Protestant Federation of France, was also a contributor.

of AI and of broad citizen participation in public debate even went so far as to argue that making the documents from the Convention available to the public (texts and more than thirty hours of video streaming), not to mention the numerous press articles, could lead to true transparency, because nobody in the general public would have the time to immerse themselves in the verbatim account of this twenty-seven-day experience.

As a guarantor myself, it was my job, along with my three colleagues, to ensure that a well-established idea of transparency was respected. Total transparency is undesirable. It can even be detrimental to deliberation, as is well documented in the relevant scientific literature. We do not deliberate in the same way in public as behind the doors of a committee or working group. I was somewhat surprised that the members of the Convention had not been consulted about this project, as they had been for the research observers who had presented their research intentions to the Convention members. I learned that AI had taken a long time to integrate the proposed section. For example, it had been unable to read the voting tables. Yet there were many of these (over 500) in the section in question that constituted organised material which was easy to process. Second, it appeared that the material extracted from the two sequences chosen for the experiment was the easiest to analyse. These included the final report, which was very well written, organised and readable by any curious and somewhat virtuous citizen, as well as all the debates in the plenary sessions, moderated by talented facilitators. It would have been more interesting to see AI take up the challenge of analysing the much longer conversations conducted in subgroups, which are more chaotic and where, as in any conversation, participants lose track and forget what has been said. It was this type of help that one of the participants, who was also very knowledgeable about AI languages, had in mind when he bombarded the promoter of the 'global first' with questions. If it had taken the AI tool at least two months *ex post* to process these two limited sections, how much more time would have been needed to analyse the conversations supporting the process, to produce simultaneous summaries? And what would be its quality?

Any technology requires you to explain what you're trying to do with it.

Rating the AI version of transparency: 'can do better'

The qualitative dimension of the work carried out by AI became central when the leader of this experiment shared an example that was very important to him. The following question had been asked during a session: 'How did the Convention take religious arguments into account?' Using extracts from the speeches of various religious representatives heard in plenary, the AI tool replied in two lines by denomination and by speaker, without bothering to specify their function. One of the speakers was recorded as a Freemason, but this was a mistaken identity, AI having confused him with the Orthodox representative, Monsignor Alric, who can be seen on the video. ChatGPT users are used to these factual errors. However, in this case, given the stakes and the small number of speakers involved, such carelessness is cause for concern. As for the sources, they referred to excerpts from the hearings of just four officials. Why four when there were five? One of the extracts also contained comments made by the facilitator. Needless to say, the material was the transcript of the audio and had not taken into consideration the video image, otherwise it would not have substituted an Orthodox bishop with a Freemason.

In another example, to present the Catholic Church's position on the end of life, AI produced just two statements made at a hearing of the Archbishop of Paris, Monsignor Ulrich: 1) 'The position of the Catholic Church is based on the belief in an option for life that must not be interrupted by man'; and 2) 'Accompanying those who are dying is essential, but the decision to end a life is not acceptable.' Readers would be left wanting more rather than these simple assertions. Otherwise the word argument has no meaning. And if, for the sake of transparency,



Rather than talking about a revolution for democracy and new transparency, we should be cautious.

they wanted to delve into all the speeches, they would be even more disconcerted. They would find arguments put forward by the archbishop, but that were not captured by AI, namely that life is a gift from God and that God does not take it away at the moment of death. They would then discover that the prelate referred to the prohibition on killing and provided justifications for this prohibition. AI clearly missed the essential point of the prelate's speech and lost his argument. The same occurred for other speakers whose arguments were omitted by AI.

Not only did AI fail to answer the initial question of how the Convention took religious arguments about the end of life into account, not only did it mix people up, but its results also said nothing about the far more interesting exchanges that took place between some Convention members and the religious leaders who were heard for over an hour and a half. To really answer these questions, we would have had to start by looking at the way the report records those discussions (bearing in mind that the term 'religion' only appears five times) and go back to where the religious leaders interviewed are mentioned. Most importantly, a much more detailed analysis was needed of certain Convention members' references to their religious convictions. This excellent subject is more within the reach of researcher-observers taking notes throughout the process than that of an algorithm.

Rather than talking about a revolution for democracy and new transparency, we should be cautious. The experiment carried out on this fragment of the proceedings added screens rather than transparency, and paradoxically took a very long time. Its promoter praised the time saved for an outside audience, without saying anything about the fastidious working methods and the choices (technical or interpretative)

that were made to produce such a meagre, partially erroneous and generally misleading result. If argumentation is one of the essential elements of deliberation, it needs to be identified and presented in its entirety, so that we can review it and know whether it is valid or not, and not just fish out a few assertions. The AI used to achieve this form of transparency therefore deserves a 'can do better' rating.

A ChatGPT convention

Invited to give a talk on the ethics of AI, I decided to put ChatGPT to the test by asking it for advice on organising a citizens' convention. It's impossible to reproduce here the whole of the edifying conversation with the application that went on for twelve pages. It was supposed to give me ideas, but all it did was list the things I needed to worry about, which are well known to any organiser or researcher in the field, namely: identifying the participants (and here it was thinking of experts and a target audience, not randomly selected citizens); choosing a venue; setting an agenda; informing the participants and engaging them in respectful discussions; taking note of the results; and keeping the participants informed about how their results would be used.

In the ensuing conversation, AI reinvented, in the course of my questions, problems familiar to philosophy, such as the difficulties of judgement. I did not find the famous formulations of the American philosopher John Rawls on this subject, nor his points of disagreement with his German counterpart Jürgen Habermas. On the other hand, AI has suggested some unexpected ways for dealing with the major difficulties of moral pluralism that lie at the heart of such experiments. In this case, it surpassed the usual practice, which is more concerned with respecting and including participants.

Our wisdom must therefore be exercised with regard to the relevance and appropriateness of the prostheses that AI tools are.

Meanwhile, its definition of deliberation was limited to a ‘shared understanding of the subject, considering several perspectives in order to make decisions based on collective wisdom’. This irenic approach pays little heed to disagreements and political uncertainties.

The wings of desire

Bearing in mind what was at stake in a convention about life and death, I asked the AI tool this question: ‘As an AI agent, do you have any moral intuitions?’ It replied: ‘As a linguistic model of AI, I have no personal beliefs, emotions or moral intuitions. I am programmed to provide information and answer questions [...] and I do not have the capacity [...] to form personal opinions. However, it is important to bear in mind that I am not capable of having personal experiences [...], and that my answers are limited by the information and knowledge contained in the textual data on which I have been trained. Although I have no personal emotions or moral intuitions, I can provide neutral and impartial information and advice on these and other aspects of the deliberative process, based [...] on the principles of deliberative democracy. My goal is to support participants in their efforts to engage in fair, inclusive and productive deliberation, and to help facilitate the process of collective decision-making.’ Like one of the two angels portrayed by Wim Wenders in the film *Wings of Desire*, in this respect AI might envy humans, who are capable of desires, emotions and intuitions, and give up its immortality to enjoy the pleasures of the senses. And, we might add, the wisdom of choices and the aptness of feelings, in a nod to the fine title of moral philosopher Alan Gibbard’s book, *Wise Choices, Apt Feelings*.

Our wisdom must therefore be exercised with regard to the relevance and appropriateness of the prostheses that AI tools are, like any technique, particularly in the context of political deliberations, where the moral dimension is fundamental. The point is not to save time or make promises of transparency, but to take the time to ask the right questions. The first is to know what deliberation is. Artificial intelligence will always be based on the intelligence of the choices we make and the questions we ask.

■ REFERENCES

→ Reber, B.,
*La délibération
responsable entre
conversation et
considération. Conditions
d’un grand débat
démocratique*, London:
ISTE Éditions, 2023.