

Citizens' Assembly on **artificial intelligence**

FRENCH-SPEAKING SWITZERLAND, 2025

Final report

Survey results and citizens' proposals

February 2026

Table of contents

Executive summary	3
Preamble	5
Introduction	7
Survey on artificial intelligence in French-speaking Switzerland	8
Composition of the Citizens' Assembly	21
Methodology and proceedings of the Citizens' Assembly	23
Final wording of issues, objectives and proposals	27
Conclusion	31
Acknowledgements	33

Executive summary

This report presents the results of the Citizens' Assembly on artificial intelligence, organized by the EPFL AI Center in collaboration with the Université de Genève in November 2025. It draws on two complementary sources: a survey of 734 residents of French-speaking Switzerland and the deliberations of 40 citizens gathered over four days.

MAIN FINDINGS FROM THE SURVEY

A divided but engaged population. Nearly half of respondents (47%) say they are familiar with AI, and 87% have already used at least one AI tool. ChatGPT is the most widely used tool (70%). Opinions range between moderate optimism and concern, with no dominant trend.

High trust in academic institutions. Universities and research institutes are trusted by 57% of respondents to develop AI responsibly. Swiss authorities also enjoy a notable level of trust (42%). By contrast, technology companies (10%) and especially social media companies (less than 2%) generate strong mistrust.

Concerns centered on security and disinformation. Malicious uses of AI (81%) and deepfakes (77%) are the main sources of concern, ahead of privacy (65%) and job losses (59%).

A clear demand for national governance. The Swiss government is identified as the primary actor for AI governance (31%), and the protection of personal data is identified as the top political priority (68%).

THE PROPOSALS OF THE CITIZENS' ASSEMBLY

After four days of deliberation, the 40 participants formulated **20 concrete proposals** structured around five major issues:

Issue 1 — The role of the state: creation of a federal AI office, long-term funding for research.

Issue 2 — Access and education: awareness-raising for all audiences, promotion of human interactions in the face of an "all-AI" world.

Issue 3 — The world of work: anticipating socio-professional impacts, supporting career transitions.

Issue 4 — Traceability: labeling human creation, strengthening copyright.

Issue 5 — Responsible practices: ethical legislation, fighting cybercrime.

The detailed wording of these proposals, drafted and approved by the citizens themselves, appears in the section "Final wording of issues, objectives and proposals" of this report.

Preamble

General message from the participants in the Citizens' Assembly

The 40 participants in the Citizens' Assembly on artificial intelligence welcome the process carried out during this collective experience and proudly present the 20 proposals that emerged from it, in the hope that they will be well received.

The participants defined five major issues, each paired with two objectives and corresponding proposals, some addressed to society at large and others to political actors.

This preamble stems from the spontaneous remarks made by participants at the end of the Citizens' Assembly. It aims to convey, in summary form, the spirit and convictions that emerged from this collective deliberation, and to serve as a compass for reading this report.

Artificial intelligence is developing very rapidly, perhaps too rapidly for the pace of politics. The participants stress the need to recognize this urgency and to rise to the challenge. They call on political authorities, institutions, civil society and the wider public to act without delay, anticipate future transformations and show foresight. Artificial intelligence is a powerful technology. It is crucial to seize the opportunities offered by its immense potential while recognizing the risks associated with this digital revolution.

As a cardinal principle of their reflections, the participants agree that everything possible must be done to ensure that AI serves human beings. They call for defining appropriate uses of AI and adapting it to the specific needs of different groups. In this respect, particular attention must be paid to young people, who should be guided toward a conscious and measured use of AI in order to protect them from risks such as dependency.

Although all participants recognize AI's immense potential to improve society, they nevertheless ask that its use be framed. An "all-AI" world appears neither desirable nor appealing to them. The participants therefore call for artificial intelligence not to be used systematically in every area of human life, and they urge that everyone be left the choice of whether or not to use AI.

→ "AI, good servant, bad master" — The participants do not support AI making decisions in their place, and therefore call for limiting the autonomous decision-making capacity of AI systems and for developing a relationship of control rather than dependence.

After deliberating as humans about AI, and doing so without using AI, the participants emphasize the fundamental importance of human ties. They urge everyone to keep human friends, value human relationships and avoid humanizing AI, even though it can sometimes imitate certain human characteristics. In conclusion, the participants in the Citizens' Assembly on artificial intelligence stress the need to preserve and promote critical thinking about AI, and to foster its expression within spaces of collective human intelligence, as this citizens' assembly sought to do.

Introduction

Artificial intelligence represents one of the most significant technological transformations of our time. Its rapid development raises fundamental questions that affect every aspect of life in society: work, health, education, privacy and democracy itself. Given the scale of these challenges, it is essential that citizens be able to speak out and help shape the future of this technology rather than passively undergo its effects.

In this spirit, the EPFL AI Center, in collaboration with the Pôle de recherche en innovations démocratiques at the Université de Genève, decided to organize a Citizens' Assembly on artificial intelligence in French-speaking Switzerland. This deliberative process sought to give the floor to a panel that was as diverse and representative as possible given the practical constraints (in-person format, four full days), in order to gather its perceptions, concerns and expectations and then formulate concrete proposals for political authorities and the broader community. Ahead of the Assembly, a survey was conducted among people living in French-speaking Switzerland in order to establish an initial snapshot of opinions on AI. The results of this survey, together with the proposals arising from the deliberations of the Citizens' Assembly, are presented in this report.

The EPFL AI Center, as a producer of independent scientific knowledge on artificial intelligence, has set itself the mission of informing and educating the public about the issues raised by this technology, and of serving as a trusted platform to facilitate informed public debate.


Ahead of the Citizens' Assembly, the Federal Statistical Office provided 7,000 randomly selected addresses from all French-speaking cantons or the French-speaking regions of multilingual cantons, while ensuring an appropriate balance across age groups (from age 16 onward) and sexes. These individuals received a letter inviting them to complete a short questionnaire on their uses of and opinions about artificial intelligence, and could also indicate their willingness to take part in the four days of the Citizens' Assembly. Of all the people contacted, 734 responded to the questionnaire. Among them, about one quarter (24%) indicated that they wished to participate in the Assembly, which took place over two weekends at the EPFL AI Center in November 2025. From these volunteers, 40 citizens were selected, with the aim, as far as possible, of ensuring balanced representation across cantons, sexes, education levels, ages and political orientations.

Survey on artificial intelligence in French-speaking Switzerland

The survey conducted ahead of the Citizens' Assembly made it possible to gather the perceptions, attitudes and concerns of 734 respondents from French-speaking Switzerland regarding artificial intelligence. The results presented below provide an overview of respondents' views before the deliberative process.

Methodological note: The questionnaire was sent by post to 7,000 addresses randomly selected by the Federal Statistical Office. Respondents completed the survey online over a period of several weeks. The response rate of around 10% (734 responses) is explained in particular by relatively short deadlines and by the technical nature of the topic, which may have discouraged some people who did not feel sufficiently familiar with AI. Unless otherwise stated, all percentages are calculated on the basis of all 734 respondents (N=734).


"How familiar are you with artificial intelligence (AI)?"



Respondents are split relatively evenly between those who say they are familiar with AI and those who are less so. About 47% consider themselves rather or very familiar (37% rather familiar, 10% very familiar), while 53% consider themselves not very or not at all familiar (37% not very familiar, 16% not at all familiar).

→ Key message: Survey respondents remain divided in their familiarity with AI, with a slight majority saying they are not very or not at all familiar.


"Overall, how do you feel about AI and its impact on society?"



Overall feelings about AI are nuanced and do not clearly lean in one direction. A little more than a third of respondents (37%) express a positive opinion, about 30% adopt a neutral position, and nearly a third (31%) express a negative view. Strongly negative opinions ("very negative") are nevertheless more frequent than strongly positive opinions ("very positive"), suggesting a degree of asymmetry in the intensity of perceptions.

→ Key message: Opinions are balanced between optimism and pessimism, but very negative attitudes are more frequent than very positive ones.

"Please indicate your level of agreement with the following statements"




Respondents recognize AI's potential to improve everyday life: 63% say they agree or strongly agree with the statement that AI makes life simpler and more practical, compared with 15% who disagree. Likewise, 56% believe that AI will significantly improve healthcare and medical diagnosis (14% disagree).

At the same time, concerns about risks are very strong. Nearly 69% of respondents consider AI a serious threat to privacy and data security (8% disagree). A large majority (70%) want the development of AI to be strictly framed by public authorities (12% disagree).

The strongest consensus concerns transparency: 90% of respondents believe that AI systems should always clearly inform users when they are interacting with a machine rather than a human. Only 2% express disagreement.

→ Key message: While the benefits of AI are acknowledged, the demand for oversight and transparency commands very broad consensus.


"To what extent do you trust the following organizations to develop and use AI responsibly?"



Levels of trust vary depending on the actor. Universities and research institutes enjoy the highest level of trust: 57% of respondents give them high or very high trust, and only 9% low or no trust. Swiss public authorities also receive a notable level of trust (42% high trust, 16% distrust). European Union regulatory bodies (27%) and international organizations such as the UN or WHO (27%) occupy an intermediate position. By contrast, technology companies (Google, Microsoft, OpenAI) inspire only limited trust (10%), while 50% express low or no trust. Social media companies (Meta, TikTok) are viewed with pronounced mistrust: 83% of respondents give them low or no trust, and barely 2% high trust.

→ Key message: Academic institutions and Swiss authorities inspire trust, whereas technology companies, and especially social media companies, generate marked mistrust.


"How concerned are you about each of the following potential AI-related problems?"



The strongest concerns relate to malicious uses of AI (hacking, cyberattacks): 81% of respondents say they are very or extremely concerned, and only 3% slightly or not at all concerned. The difficulty of distinguishing real information from AI-generated content (deepfakes, disinformation) is a major concern for 77% of respondents. Loss of privacy and misuse of personal data worry 65% of respondents. The risk that AI systems may become too powerful or uncontrollable concerns 63% of respondents. Job losses linked to automation worry 59% of them. Finally, biased or unfair decisions made by AI systems concern 50% of respondents.

→ Key message: Risks related to cybersecurity, disinformation and privacy dominate concerns, ahead of employment issues and algorithmic bias.


"How enthusiastic or excited are you about each of the following potential benefits of AI?"



Respondents are most enthusiastic about reducing the time spent on repetitive and administrative tasks: 53% say they are very or extremely enthusiastic, compared with 15% who are only slightly or not at all enthusiastic. Improvements in healthcare and medical diagnosis also generate notable enthusiasm (49%). Innovations aimed at fighting climate change and protecting the environment generate moderate enthusiasm (39%), as does the development of new products and services (37%). Improvements in education and training divide respondents more strongly (36% enthusiastic, 26% slightly or not enthusiastic). Finally, improving Switzerland's economic competitiveness generates the most restrained enthusiasm (28%).

→ Key message: Productivity gains and medical applications generate the strongest enthusiasm, while economic and educational benefits prompt more mixed reactions.


"Do you think AI and automation will create more jobs than they eliminate in Switzerland?"



A relative majority of respondents are skeptical about net job creation from AI. In total, 50% believe that AI will probably not or certainly not create more jobs than it eliminates (31% probably not, 20% certainly not). Only 14% anticipate a positive balance (11% probably yes, 3% certainly yes). One third of respondents (33%) say they are uncertain.

→ Key message: Half of respondents doubt that AI will create more jobs than it eliminates, revealing significant concern about the future of work.


"Do you think your own job (or professional field) will be replaced or deeply transformed by AI in the next 5 years?"



Views are divided regarding the impact of AI on respondents' own jobs. About 34% consider it rather or very likely that their job will be replaced or deeply transformed by AI (22% rather likely, 12% very likely). Conversely, 39% consider this prospect unlikely (16% rather unlikely, 23% very unlikely). A notable share (12%) say they are uncertain, and 12% are not concerned (currently without employment).

→ Key message: One third of working respondents anticipate a profound transformation of their job due to AI, but a slightly larger share feels sheltered from these changes.


"Which of the following AI applications or services have you personally used over the past year?"



A large majority of respondents have already experimented with AI tools. ChatGPT or similar conversational assistants have been used by 70% of respondents, making it the most widespread tool. Translation and language-correction tools have been used by 56%, and AI-based search engines (Bing Chat, Google Gemini) by 51%. More specialized uses remain in the minority: 19% have used image or video generation tools (DALL-E, Midjourney), 12% AI tools at work (coding assistants, analysis tools), and 8% AI-based health or fitness applications. Only 13% of respondents say they have not used any AI tool.

→ Key message: Adoption of AI tools is already widespread, with nearly 9 out of 10 respondents having used at least one tool, led by ChatGPT.


"Overall, how satisfied are you with your experience using these AI tools?"



Satisfaction with AI tools is largely positive. Across all respondents, 62% say they are satisfied or very satisfied with their experience (45% satisfied, 17% very satisfied), while only 5% express dissatisfaction. It should be noted that 10% of respondents say they have never used these tools.

→ Key message: A large majority of respondents say they are satisfied with their experience using AI tools, while dissatisfaction remains marginal (5%).


"Who should bear the primary responsibility for the governance of AI and its impacts?"



When asked which actor should bear the primary responsibility for AI governance, respondents place the Swiss government first (31%), followed by a large number of people saying they do not know (27%). Independent oversight bodies (NGOs, expert committees) are mentioned by 17% of respondents. The European Union (8%), international organizations (7%), and technology companies themselves (6%) receive more modest scores.

→ Key message: The Swiss government is the preferred actor for AI governance, but the high level of uncertainty points to a need for greater institutional clarity.

"Switzerland should prioritize the following aspects in its AI policy (choose up to two)"



Because this was a multiple-choice question (up to two answers), the results are presented here as a share of all responses. Protection of privacy and personal data comes out clearly on top (36.4% of responses), followed by ethical guidelines and AI transparency (21.8%), and the prevention of uncontrolled or dangerous AI (21.3%). Preventing unemployment and supporting workers account for 9.1% of responses, while economic considerations (6.2%) and equitable access (5.2%) are mentioned less often.

→ Key message: Data protection and the ethical framing of AI are the priorities cited most often by respondents, far ahead of economic considerations.

Composition of the Citizens' Assembly

The EPFL Citizens' Assembly on artificial intelligence included 40 participants. Among them were 19 women and 21 men, ensuring an overall balanced gender distribution. The group consisted mainly of Swiss nationals, most of them born in Switzerland, while also including people with diverse backgrounds.

How the 40 participants were selected. The selection was made from among the people who responded to the questionnaire and indicated their willingness (and availability) to take part in the Citizens' Assembly across all four days. The objective was to form a group that was as diverse as possible and close to the population of French-speaking Switzerland on key criteria (notably age, sex and level of education).

In practical terms, the information declared in the questionnaire was grouped into categories (age groups, sex, education level, and political self-placement). Distribution targets were then defined (based on population statistics and the available responses). A weighted random draw was then carried out: each volunteer was assigned a selection probability adjusted so as to meet these targets as closely as possible across all criteria at the same time. This approach preserves the random nature of the selection while ensuring better representativeness than a simple random draw.

After this initial draw, the selected individuals were contacted and invited. In the event of a refusal or lack of availability, replacement participants were randomly drawn from the other volunteers, according to the same principle, until a final group of 40 participants was obtained.

Three groups should be distinguished. It is useful to distinguish (A) all survey respondents (n=734), (B) the people who said they were interested and available for four days (n=177), and (C) the participants who were actually present (n=40). A "volunteer effect" can be observed between (A) and (B): people willing to participate for four days are more often from the canton of Vaud (26.6% in (B) versus 17.3% in (A)), more often have tertiary education (university/UAS: 49.2% in (B) versus 44.0% in (A)), and are slightly less right-leaning on the political scale (8-10: 6.8% in (B) versus 11.2% in (A)). The final group (C) retains substantial diversity while partially correcting some of these gaps, especially in education (university/UAS: 37.5% in (C), versus 49.2% in (B)) and sex distribution (21 men, 19 women). On political self-placement (scale 0 = far left, 10 = far right), all three groups are overall centered (median = 5 in (A), (B) and (C)). In the overall sample (A), 25.8% place themselves between 0-3, 63.0% between 4-7 and 11.2% between 8-10 (n = 663; 71 missing responses). Among interested and available people (B), the distribution remains centered but the 8-10 share is lower (27.3% / 65.8% / 6.8%; n = 161; 16 missing responses). The final group (C) remains centrist and pluralistic, with positions ranging from 0 to 10 (mean 4.95; 29.7% / 62.2% / 8.1%; n = 37; 3 missing responses).

The final group covers all age brackets (18-29: 20.0%; 30-44: 27.5%; 45-59: 30.0%; 60+: 22.5%) and all French-speaking cantons (Vaud: 22.5%; Jura: 20.0%; Valais: 15.0%; Genève: 15.0%; Fribourg: 10.0%; Berne, partie francophone: 10.0%; Neuchâtel: 7.5%). The higher share from Jura and the lower share from Neuchâtel should be interpreted cautiously: in a group of 40 people, a few refusals and replacements (as well as the randomness of the selection) can materially affect the cantonal distribution.

The majority were professionally active, mainly full-time or part-time, and very few were unemployed. Declared economic situations were generally stable, with most households indicating that they had sufficient resources, or even a surplus. Finally, the assembly was characterized by strong civic engagement, illustrated by regular participation in popular votes, which points to a strong social and democratic anchoring.

Methodology and proceedings of the Citizens' Assembly

The Citizens' Assembly on artificial intelligence took place over four days, spread across two weekends, on 1 and 2 November and 15 and 16 November 2025. The methodology adopted followed a multi-step progression designed to enable participants to acquire knowledge about AI, deliberate collectively in order to identify and select priority issues, and then formulate structured proposals. The process was designed by Demoscan, a Swiss association dedicated to democratic innovations, especially deliberative and sortition-based approaches, in order to ensure progressive capacity-building, informed deliberation and the collective formulation of issues and recommendations.

DAY 1 – INFORMATION AND FAMILIARIZATION WITH THE ISSUES

The first day of the Citizens' Assembly was devoted to knowledge acquisition and to establishing a shared framework for deliberation. It aimed to help participants better understand the issues related to artificial intelligence while laying the methodological foundations needed for the later collective work.

To this end, the day began with activities designed to foster mutual acquaintance and group dynamics. Participants exchanged in small groups about their motivations for taking part in the assembly and about their current uses of artificial intelligence. This phase aimed to map the diversity of experiences and expectations while establishing a climate of trust conducive to deliberation. The participants collectively defined the rules governing the exchanges. These included respect for everyone's voice, active listening, balance in speaking time, and the confidentiality of internal group discussions.

The day continued with an expert presentation given by Professor Marcel Salathé (EPFL AI Center), whose purpose was to provide a factual and accessible introduction to artificial intelligence. This presentation addressed the historical development of AI, its operating principles, the resources required for its development (data, computing power and energy), as well as the issues related to open models and Swiss initiatives.

Finally, the day included a discussion with people from political and institutional circles, namely National Councillors Isabelle Chappuis (VD / Le Centre) and Gerhard Andrey (FR / Les Verts), shedding light on issues of digital sovereignty, governance, data security and regulation.

Overall, this first day sought to establish a shared base of knowledge and methods, essential for the collective identification of issues and the formulation of proposals over the following days.

Finally, the day concluded with specific work devoted to the introduction of the report. Participants were invited to formulate, in groups, the key messages and intentions they wanted to see reflected in the report's preamble.

DAY 2 — DELIBERATION AND IDENTIFICATION OF THE ISSUES

The second day of the Citizens' Assembly was devoted to deepening knowledge, taking inclusion issues into account in the deliberation, and identifying and structuring the priority issues that would guide the rest of the process.

The morning began with a group-dynamics activity centered on inclusion and equality of participation. Through a specific exercise ("the power walk"), participants were made aware of the importance of moral equality among assembly members and of possible asymmetries in speaking time.

The day continued with a series of expert hearings, in the form of presentations followed by question-and-answer sessions. These interventions, given notably by Professor Marcel Salathé and Olga Baranova (CH++), aimed to deepen participants' knowledge of the transition linked to artificial intelligence.

After these hearings, group reflection phases enabled participants to identify the key information emerging from the exchanges with the experts and to formulate the main lessons they retained. This step was intended to foster collective appropriation of the content and prepare the problem-framing phase.

The afternoon was devoted to continuing the workshop sessions. Participants were invited to enrich the discussions in light of the experts' contributions, and then to begin a second round of workshops aimed at identifying and formulating major issues. The groups exchanged views about their expectations and concerns in order to formulate issues that were clear, complete and understandable to people who had not participated in the process.

At the end of this work, the workshop results were presented in plenary and then submitted to a preference vote. By the end of the second day and the first weekend, eight clusters of issues had been identified.

DAY 3 — DEEPENING AND FORMULATING THE OBJECTIVES

The third day of the Citizens' Assembly was devoted to selecting, prioritizing and clarifying the issues identified during the previous days. It constituted a pivotal stage of the deliberative process, aimed at narrowing down and structuring the issues so as to enable in-depth and operational work in the next phase.

Participants worked within their respective groups on the basis of the eight clusters of issues identified at the end of the second day. The task was to deliberate collectively in order to determine which five issues were considered the most important and highest priority, while formulating the arguments justifying those choices.

It was explicitly envisaged that groups could propose merging certain issues when these were considered closely linked, provided that the logic was clearly explained and validated by the group. The plenary reports and collective discussions led to the following decisions: some issues were integrated into others when they reflected convergent problem areas, while one issue was set aside as an autonomous theme, though its elements could be taken up transversally in other frameworks. At the end of this collective deliberation, five main issues were retained for the rest of the process, and they appear later in this report.

Following the deliberation that led to the selection of the five priority issues, participants once again split into five working groups, each assigned one of the selected issues. The objective of this sequence was to clarify the directions to be favored in response to the assigned issue by identifying and formulating one or two main objectives. The groups were invited to focus on clear objectives, formulated in operational terms and able to serve as the basis for developing concrete proposals.

DAY 4 — FORMULATION AND FINALIZATION OF THE PROPOSALS

The fourth and final day of the Citizens' Assembly was devoted to drafting, discussing and finally approving the proposals associated with the objectives of the five selected issues.

Participants worked on formulating concrete proposals for each issue, based on the objectives defined the day before. Each issue gave rise to distinct proposals addressed to society and to political actors.

The proposals were then discussed in plenary session. This phase of collective deliberation aimed to share all the proposals that had been formulated, debate them and, where appropriate, carry out trade-offs by vote in order to retain, for each issue, two proposals in their final wording. Editorial adjustments were made in real time on a large screen, under the participants' supervision, in order to guarantee fidelity to the decisions taken collectively.



Final wording of issues, objectives and proposals

The issues presented below are the result of the deliberative process, combining knowledge inputs, group discussions and collective arbitrations. Their wording, including the choice of terms and verbs, has not been modified and corresponds strictly to the version approved by the participants in the Citizens' Assembly.

Translation note: The participants in the Citizens' Assembly formally approved only the original French wording of the issues, objectives and proposals. The texts below are English translations. The French version in the original report (*Rapport de l'Assemblée citoyenne sur l'intelligence artificielle*) remains the authoritative version approved by the participants.

ISSUE 1: Determine the degree of state involvement and responsibility in AI, in its development and implementation in Switzerland.

Objective 1: Adopt a national AI strategy.

Objective 2: Promote the development of AI in accordance with human rights, enriched by Swiss values, in order to shape the future of AI.

FOR SOCIETY	FOR POLICYMAKERS
1. Implement the strategies defined by political bodies by putting in place collaborations between politicians, experts and citizens (for example: commissions, mandates, meetings, audits, citizens' assemblies) to guide and assess the proper implementation of these strategies. 2. Strengthen national and international research in the field of AI, notably on technological, social and ethical aspects (calls for proposals, competitions, conferences, universities, etc.).	1. Create a new federal office for AI and new technologies to define and implement a political strategy that respects human rights and Swiss values (neutrality, humanitarianism, democracy, peace, diplomacy, etc.). 2. Ensure sustainable funding dedicated to the needs of AI research and development (for example EPFL, ETH Zurich, universities of applied sciences and universities, the Apertus project).

ISSUE 2: Democratize access to AI and promote equality, encourage healthy, responsible and ethical use, notably to limit dependency and foster social interactions, including in the real world.

Objective 1: Raise public awareness of AI, train people in its responsible use, and facilitate access for everyone.

Objective 2: Promote social, cultural, sporting and artistic gatherings so as not to fall into an "all-AI" world.

FOR SOCIETY	FOR POLICYMAKERS
1. Encourage and support associations, media outlets, media libraries, schools, adult education institutions, etc., to raise awareness and inform/help the public understand AI, its uses and its ecological, geopolitical, social, economic and public-health implications, and to address cases of dependency. 2. Create/promote/inform people about specific local events for a broad public (festivals, open days, interactive simulations, forum theatre, etc.), or integrate into existing events, in order to raise awareness of AI and prioritize human rather than digital contact.	1. Create and support organizations / pools of experts that can serve as reference points for training / awareness-raising / risk prevention for the population, cantons and regions, notably in cooperation with education and health services. 2. Increase funding for social, cultural, sporting and artistic activities (local sports clubs, scouts, theatre, conservatory, dance, etc.) to strengthen social interactions.

ISSUE 3: Prepare the world of work for the AI transition with regard to social and financial protection and to technological changes in jobs.

Objective 1: Anticipate the socio-professional impacts linked to the advent of AI.

Objective 2: Guarantee dignity, and human and financial support, for people who may lose their job or experience a change in their role due to AI.

FOR SOCIETY	FOR POLICYMAKERS
1. Develop and adopt protective measures and adapt training to anticipate the evolution of the world of work linked to AI. 2. Dedicate a share of the funds provided by Swiss umbrella organizations and other professional unions at the federal and/or cantonal level to training and retraining linked to job losses due to AI.	1. Tax or levy providers that generate lasting job losses because of AI. 2. Encourage and rethink career transitions, notably toward new AI-related professions (for example by adapting unemployment insurance to the challenges posed by AI).

ISSUE 4: Guarantee the origin and traceability of outputs and distinguish what comes from humans from what comes from AI.

Objective 1: Label human creation.

Objective 2: Adapt and strengthen copyright protection.

FOR SOCIETY	FOR POLICYMAKERS
1. Create a label that guarantees the origin of outputs (example: "Swiss valid ®"): • proof of human creation • tamper-proof / encrypted digital signature (e.g. blockchain) 2. Allow the author to refuse the use and modification of their work by AI.	1. Federal recognition of the label, awareness-raising and support (financial, guidance, etc.) for its dissemination and use in Switzerland and abroad. 2. Adapt and strengthen copyright protection within a basic national and international legal framework linked to AI.

ISSUE 5: Define responsible and ethical practices for AI users (legal entities and natural persons) that guarantee data protection and the reliability of tools shared online.

Objective 1: Define criteria for good practices for users and providers and create instruments (national labeling body, legislation, etc.) to implement and assess these good practices.

Objective 2: Fight cybercrime.

FOR SOCIETY	FOR POLICYMAKERS
1. Put in place educational tools and awareness courses on data protection, scam prevention, the dangers linked to the use of AI and the internet, etc., from childhood and for all ages, in school and out-of-school, associative and professional settings.	1. Put in place legislation and a specific independent audit body that verifies ethical compliance and good-practice criteria (label). 2. Create a federal "task force" composed of members of government, police and independent experts to strengthen, with human and financial resources, the fight against cybercrime intensified by AI.

Conclusion

The Citizens' Assembly on artificial intelligence was organized by the EPFL AI Center in collaboration with the Pôle de recherche en innovations démocratiques at the Université de Genève (UNIGE). The deliberative process was designed and conducted by the association Demoscan, which was responsible for the methodology, facilitation and smooth running of the work. The entire arrangement was based on a central principle of neutrality, both in the design of the process and in the selection and presentation of the content, in order to guarantee a balanced, non-prescriptive framework for reflection conducive to the formation of informed opinions.

This project is part of the EPFL AI Center's broader mission: to produce independent scientific knowledge, contribute to the development of critical and informed citizens, and foster informed public debate in the face of a technological revolution with major societal implications.

The results and proposals arising from this Assembly constitute a working basis intended to inform academic, institutional and political reflection. The next steps aim to highlight these contributions, ensure their dissemination and explore the ways in which they can be integrated into future debates and initiatives related to the development and governance of artificial intelligence.

Participant feedback (exit survey)

At the end of the four days, a final questionnaire was completed by 39 participants. It sheds light on how the process was experienced and on its effect on perceptions.

- **Very high perceived neutrality:** participants rate the neutrality of the teams very positively (EPFL: 9.8/11; Demoscan/UNIGE: 9.4/11; facilitation: 10.1/11).
- **Deliberative quality:** scores are high for respect (10.1/11) and for the facilitation team's ability to create a constructive climate (10.1/11) and to take opposing arguments into account (9.7/11).
- **Final report strongly endorsed:** satisfaction with the final report is very high (average 10.4/11), with no negative rating.
- **Change in attitudes:** on average, participants report a more favorable perception of AI at the end of the process, with a statistically significant shift.

The full report on this exit survey is available on the EPFL AI Center website.

Learn more:

- The EPFL AI Center: <https://ai.epfl.ch/>
- Pôle de recherche en innovations démocratiques (Université de Genève):
<https://www.unige.ch/pidem/>
- Demoscan: <https://demoscan.ch/fr/>

Acknowledgements

The organizers warmly thank all the people who contributed to the design, implementation and smooth running of the EPFL Citizens' Assembly on artificial intelligence. Their commitment was crucial in ensuring the scientific, methodological and deliberative quality of the process.

The people involved in organizing, supporting and facilitating the Assembly included:

- Prof. Marcel Salathé, EPFL AI Center
- Prof. Nenad Stojanović, Université de Genève / Demoscan
- Dr. Victor Sanchez-Mazas, Université de Genève / Demoscan
- Christiane Amici-Raboud, CYC partenaires / Demoscan
- Dr. Francesco Veri, Zentrum für Demokratie Aarau / Demoscan
- Nicolas Machado, EPFL AI Center
- Maroussia Schaffner-Portillo, EPFL AI Center
- Nathalie Likhite, Facilitator
- Kim Ingold, Facilitator

The organizers finally extend their sincere thanks to all the participants in the Citizens' Assembly for their commitment, the time devoted to the process and the quality of the exchanges. Their active and thoughtful participation constituted the core strength and richness of this deliberative process.

This project was funded by the EPFL AI Center and Stiftung Mercator Schweiz.