

Loved. Feared. And Here to Stay.

How People in German- and French-Speaking Switzerland Think, Use, and Talk About Generative AI

By Sandra Cortesi, Alexa Hasse,
Kirsten Müller-Daubermann, and
Urs Gasser

A cinematic close-up portrait of a young person with red hair and freckles sitting by a window on a modern tram or train. Sunlight streams through the window, illuminating their face as they gaze thoughtfully outside. Green scenery passes by in the background. Calm, introspective mood, shallow depth of field, natural light, realistic skin texture, modern European public transport. Optional style modifiers: golden hour light, 50mm lens, quiet reflection, cinematic realism

ai-identities



Universität
Zürich^{UZH}

Loved. Feared. And Here to Stay.

How People in German- and French-Speaking Switzerland Think, Use, and Talk About Generative AI

OVERVIEW

Host:

The project is led by an interdisciplinary team at the University of Zurich, anchored at the Institute of Communication and Media Research (IKMZ).

Supporters:

The UZH Foundation and YouMedia support the project through non-financial contributions such as fundraising support, communication, and stakeholder engagement. The Technical University of Munich, through the TUM Think Tank, supports the project as an academic partner for research exchange and translation of findings.

Funding and Oversight:

The project is funded by an unrestricted gift from Sunrise GmbH. Any views, opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and should not be interpreted as reflecting the views, policies, or position, either expressed or implied, by Sunrise.

Acknowledgements:

We are deeply grateful for the research assistance and editorial support provided by members of the Youth and Media Lab and the Creativity Lab teams, including Gabriel Garbers, Joelle Grewe, Noha Halim, Camila Hidalgo, Anyu Jiang, Annabel Jones, Ulrike Liebert, and Andras Molnar. We also extend our sincere thanks to Lionel Brossi, Christian Fieseler, Amanda Lenhart, and Aurelia Tamò-Larrieux for reviewing an earlier draft of this report and for their thoughtful comments and feedback. We are likewise grateful to Benjamin Fretwurst and Lisa Vogel at Livo Mergen for their support in analysing the survey data. Finally, we thank Camila Romero Otárola (cnromero.com) and Angela Mötteli for their artistic vision, their development of the visualizations, and the creative layout and design of this report.

Contact:

Sandra Cortesi
Department of Communications and
Media Research, University of Zurich
Andreasstrasse 15
8050 Zürich, Switzerland
s.cortesi@ikmz.uzh.ch

SUGGESTED CITATION

APA:

Cortesi, S., Hasse, A., Müller-Daubermann, K., & Gasser, U. (2026). *Loved. Feared. And here to stay: How people in German- and French-Speaking Switzerland think, use, and talk about generative AI*. Department of Communications and Media Research, University of Zurich. aidentities.ch

KEYWORDS

Switzerland, artificial intelligence, AI, generative AI, metaphors, mental models, public opinion, excitement, concerns, technology adoption, digital divide, risk perception, participation gap, learning, education, lifelong learning.

PUBLICATION DATE

May 2026

LICENSE

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International ([CC BY-SA 4.0](https://creativecommons.org/licenses/by-sa/4.0/)) license.

TABLE OF CONTENTS

↘ SUMMARY	04
↘ KEY FINDINGS	05
I. ATTITUDES TOWARDS AI	06
II. THE USE OF GENERATIVE AI TOOLS	12
III. THE RELATIONSHIP BETWEEN ATTITUDES AND USE	18
21 What the Association Does and Does Not Tell Us	
21 The “Used a Few Times” Group	
22 Implications for the Public Conversation	
IV. WHAT GENERATIVE AI METAPHORS REVEAL	24
27 Why Metaphors Matter: A Conceptual Note	
27 About the Metaphor Exercise	
29 The Difference Between Attitudes and Frame	
30 The Language of Concern	
33 The Language of Excitement	
35 The Language of Ambivalence	
V. WHAT THIS MEANS	36
38 The Structural Imbalance in Public Debate	
39 Mental Models Matter for Communication and Governance	
39 Ambivalence Is an Asset, Not a Problem to Be Solved	
VI. ABOUT THIS STUDY	41
42 Research Questions	
44 Participants	
44 Research Team	
↘ REFERENCES	45



SUMMARY

This report presents findings from the first ai-identities (aidentities.ch) study, a survey of 1,000 people aged 15–79 in German- and French-speaking Switzerland conducted in October 2025. It examines how people feel about the growing presence of artificial intelligence (AI) in daily life, how frequently they use generative AI tools, and how the two relate. A supplementary metaphor exercise – asking respondents to complete the sentence "For me, generative artificial intelligence is like __ because __" – provides a qualitative window into the conceptual frames that shape people's understanding.

The data reveals that the largest single group is ambivalent: respondents who are neither dismissive nor swept up in enthusiasm, but who are navigating real tradeoffs and forming views that resist easy categorization. That nuance is the study's most important finding, and the one with the clearest implications for how organisations, policymakers, and communicators engage with these questions.

THREE BROADER TAKEAWAYS ARE WORTH EMPHASISING

The debate is structurally imbalanced. Those most likely to raise strong concerns are also those with the least direct experience. Their concerns are legitimate and do not require hands-on familiarity to be well-founded. However, governance conversations risk being dominated by those already comfortable with the technology.

Framing matters. Respondents who frame generative AI as an extension of human capability likely engage very differently with oversight and risk discussions than those who frame it as a replacement. Effective communication must first acknowledge the mental model a given person or community holds, then speak to it directly; abstract reassurances about responsible development will not move the needle for someone whose primary concern is loss of human agency.

Ambivalence deserves engagement. The largest single group holds both excitement and concern simultaneously, which is a reflection of genuine complexity in a rapidly evolving landscape. Organisations that create space for open, two-sided dialogue will likely build more durable trust than those that communicate only excitement.

KEY FINDINGS

1.

Concern outweighs excitement.

40% are more concerned than excited about AI in daily life, compared with 21% who are more excited than concerned. The dominant public mood is cautious, not enthusiastic.

2.

A substantial group has not taken a side.

39% describe themselves as equally concerned and excited about AI's growing role in daily life. Rather than polarised, the picture is a public with mixed and sometimes contradictory attitudes.

3.

Age is the sharpest dividing line.

Among under-30s, 61% use generative AI tools often and only 6% have never used them. Among over-60s, 52% have never used them and only 13% use them often. The same age gradient appears in attitudes: 60% of over-60s are more concerned than excited versus 31% of under-30s.

4.

Women in this sample are more frequent users than men, and somewhat less concerned.

Women report more frequent generative AI tool use than men (40% vs. 31% reporting "often") and are somewhat less "more concerned than excited" (37% vs. 44%). This finding should be interpreted cautiously: the mean attitude difference is modest (2.8 vs. 2.6 on a five-point scale), and online panels tend to over-represent digitally active populations.



5.

Attitude scores and conceptual language often diverge.

Even respondents who score maximum excitement on the attitude scale frequently reach for cautious, ironic, or ambivalent metaphors to describe generative AI. What people say they feel and how they conceptualise generative AI are not the same thing.

6.

Concerns are more social in character.

When concerned respondents describe what worries them, the dominant themes are loss of human agency and control, or more specific themes such as surveillance, power imbalance, and cognitive displacement.





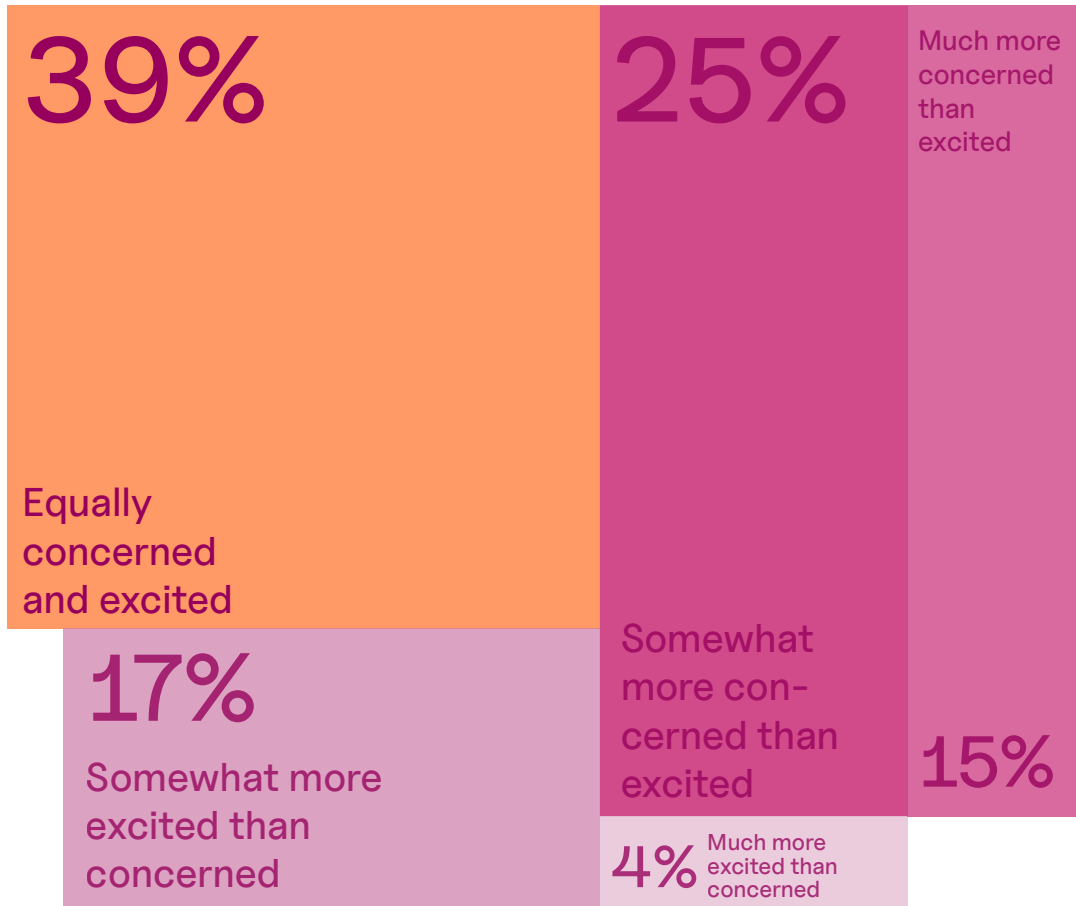
/Prompt:A realistic cinematic image of a man standing in a supermarket aisle, holding a shopping basket and turning his head to look toward the camera. Behind him are neatly stacked shelves filled with packaged food products. Warm indoor lighting, sharp focus on the subject with slight background blur. Casual clothing, everyday moment, candid documentary style, high realism. Style modifiers: lifestyle photography, 35mm lens, commercial realism, natural expressions

ATTITUDES TOWARDS AI

ATTITUDES TOWARDS AI

The most common single response to the question “Thinking about the increasing use of artificial intelligence in daily life, would you say it makes you feel more excited or more concerned?” is neither much more nor somewhat more excited nor concerned. It is both. Across the five response options, the midpoint – “equally concerned and excited” – draws the largest single share, at 39%.

**FIGURE 1:
CONCERN-EXCITEMENT DISTRIBUTION**



The concerned responses combine to a “more concerned than excited” total of 40% (25% somewhat more concerned + 15% much more concerned). “More excited than concerned,” at 21%, is the smallest group. This reframes the public debate: this survey finds that the largest combined group is the concerned one, at 40% and a substantial middle that is not ignorant or indifferent, but holding two attitudes simultaneously.

FIGURE 2: "MORE CONCERNED THAN EXCITED" AND "MORE EXCITED THAN CONCERNED"

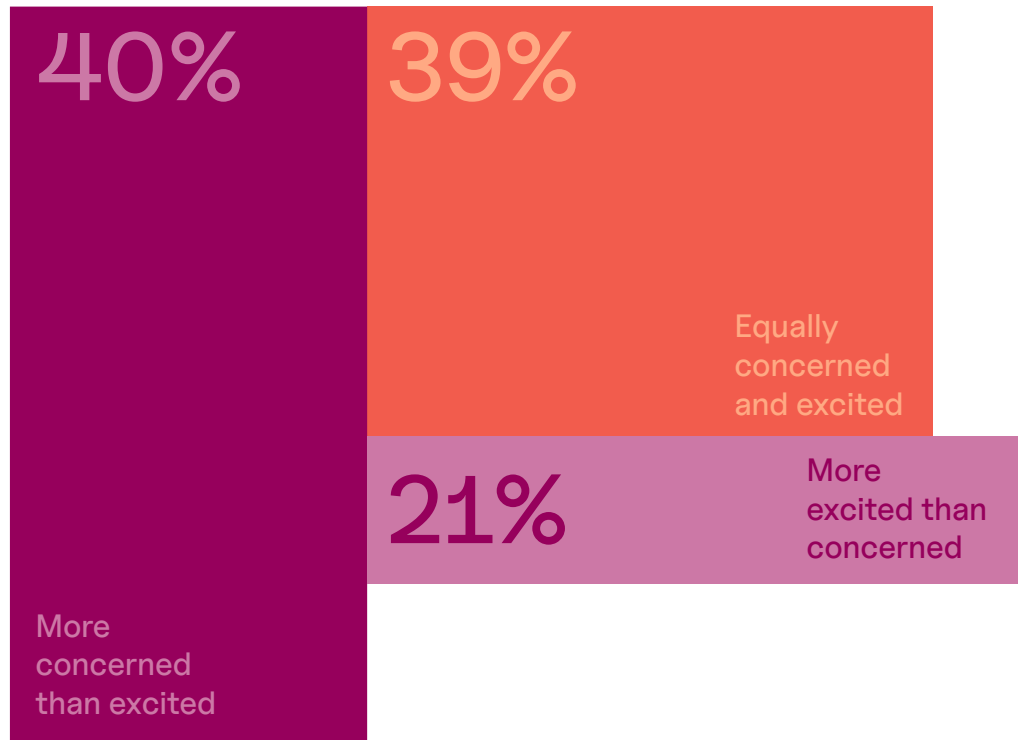


FIGURE 3: EXCITEMENT VS. CONCERN DISTRIBUTION ALONG DEMOGRAPHIC VARIABLES

The demographic fault lines are clear. In regard to **age**, concern is substantially higher among older adults versus younger adults: 60% of over-60s are more concerned than excited, compared with 31% of under-30s. Men are somewhat more likely to be more concerned than excited than women (44% vs. 37%), though the **gender** gap is modest compared to the age divide. **Education** runs in the same direction: among those who did not complete secondary school, 49% are more concerned than excited, falling to 31% among those with tertiary education, though this difference is driven more by lower concern among those with higher education than by higher excitement, which varies less across education groups. **Regional** differences are also visible, with French-speaking respondents being more likely to be more concerned than excited than their German-speaking counterparts (49% vs. 37%).

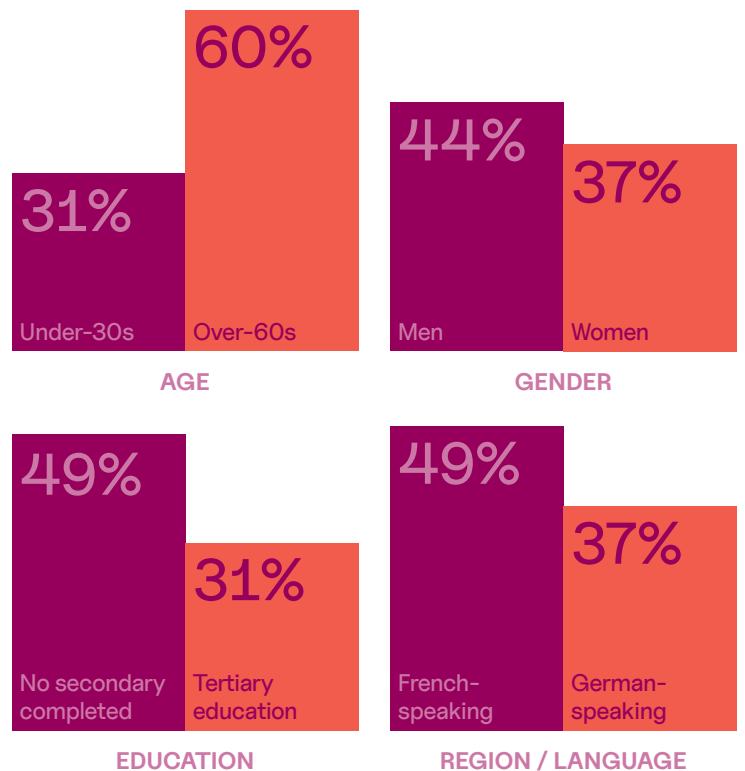


TABLE 1: AI EXCITEMENT VS CONCERN

Answers	AI Use			Attitudes				Age				Education (ISCED)			Region		Gender	
	Total	Never	Few	Often	Concerned	Both	Excited	<30	30-44	45-59	60+	1-2	3	5-8	GSS	FSS	F	M
Much more concerned	15%	39%	13%	2%	37%	0%	0%	9%	16%	11%	25%	25%	15%	10%	12%	23%	13%	17%
Somewhat more concerned	25%	42%	26%	14%	63%	0%	0%	22%	21%	24%	35%	24%	28%	21%	25%	26%	24%	27%
Equally	39%	19%	50%	39%	0%	100%	0%	43%	34%	45%	30%	32%	37%	44%	41%	32%	38%	40%
Somewhat more excited	17%	1%	10%	34%	0%	0%	79%	22%	21%	16%	8%	16%	15%	20%	17%	15%	19%	14%
Much more excited	4%	0%	1%	11%	0%	0%	21%	4%	9%	3%	2%	2%	5%	5%	5%	3%	6%	2%
Total	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Mean	2.7	1.8	2.6	3.4	1.6	3.0	4.2	2.9	2.9	2.7	2.3	2.5	2.6	2.9	2.8	2.5	2.8	2.6
95% CI	±0.1	±0.1	±0.1	±0.1	±0.0	±0.0	±0.1	±0.1	±0.1	±0.1	±0.1	±0.2	±0.1	±0.1	±0.1	±0.1	±0.1	±0.1
Valid n	1000	241	402	357	405	387	209	225	243	264	252	161	504	335	733	267	515	482

Question F1: Thinking about the increasing use of artificial intelligence in daily life, would you say it makes you feel more excited or more concerned?
 Scale: 1 «Much more concerned than excited», 2 «Somewhat more concerned than excited», 3 «Equally concerned and excited», 4 «Somewhat more excited than concerned», 5 «Much more excited than concerned»

"For me, generative AI is like a chameleon because it adapts perfectly, but you can never be completely sure what is really real."



HOW DOES SWITZERLAND COMPARE INTERNATIONALLY?

In this survey, 40% of respondents are more concerned than excited about AI's growing use in daily life, 39% feel equally concerned and excited, and 21% are more excited than concerned. How does that sit within the broader global picture?

- A useful reference point is the 25-country Pew Research Center study conducted in 2025, which found a global median of 34% more concerned than excited, 42% equally concerned and excited, and 16% more excited than concerned (Poushter et al., 2025). Switzerland's profile sits modestly above the global median on concern – but also slightly above it on excitement. Rather than pulling strongly in either direction, Switzerland occupies a position of heightened but balanced ambivalence.
- Within Europe, the picture is more varied. Italy and Greece emerge as among the most concerned publics (Poushter et al., 2025). Germany and Sweden, by contrast, come in considerably lower, at 29% and 31% respectively, making participants in this study (40%) closer to the Southern European pattern than the Northern one.
- At the other end of the global spectrum, countries such as China, Indonesia, and Thailand report excitement levels that far surpass those in Europe, with upwards of 76–80% of respondents describing themselves as excited about AI (IPSOS, 2024; OECD, 2025; Poushter et al., 2025).

A note on comparability: the Pew study draws on nationally representative samples across 25 countries, whereas this survey covers German- and French-speaking regions of Switzerland via a weighted online panel. The figures are broadly comparable in structure but not directly equivalent, and should be read as indicative rather than precisely matched.



/Prompt: A warm, photorealistic scene of a father standing inside a pharmacy, gently holding a sleeping newborn baby in a soft fabric baby carrier against his chest. He looks down at the baby with a quiet smile. Wooden shelves filled with medicine boxes line the background. Soft indoor lighting, shallow depth of field, intimate and tender mood, documentary-style realism. Style modifiers: soft focus, warm tones, human connection, natural parenting moment

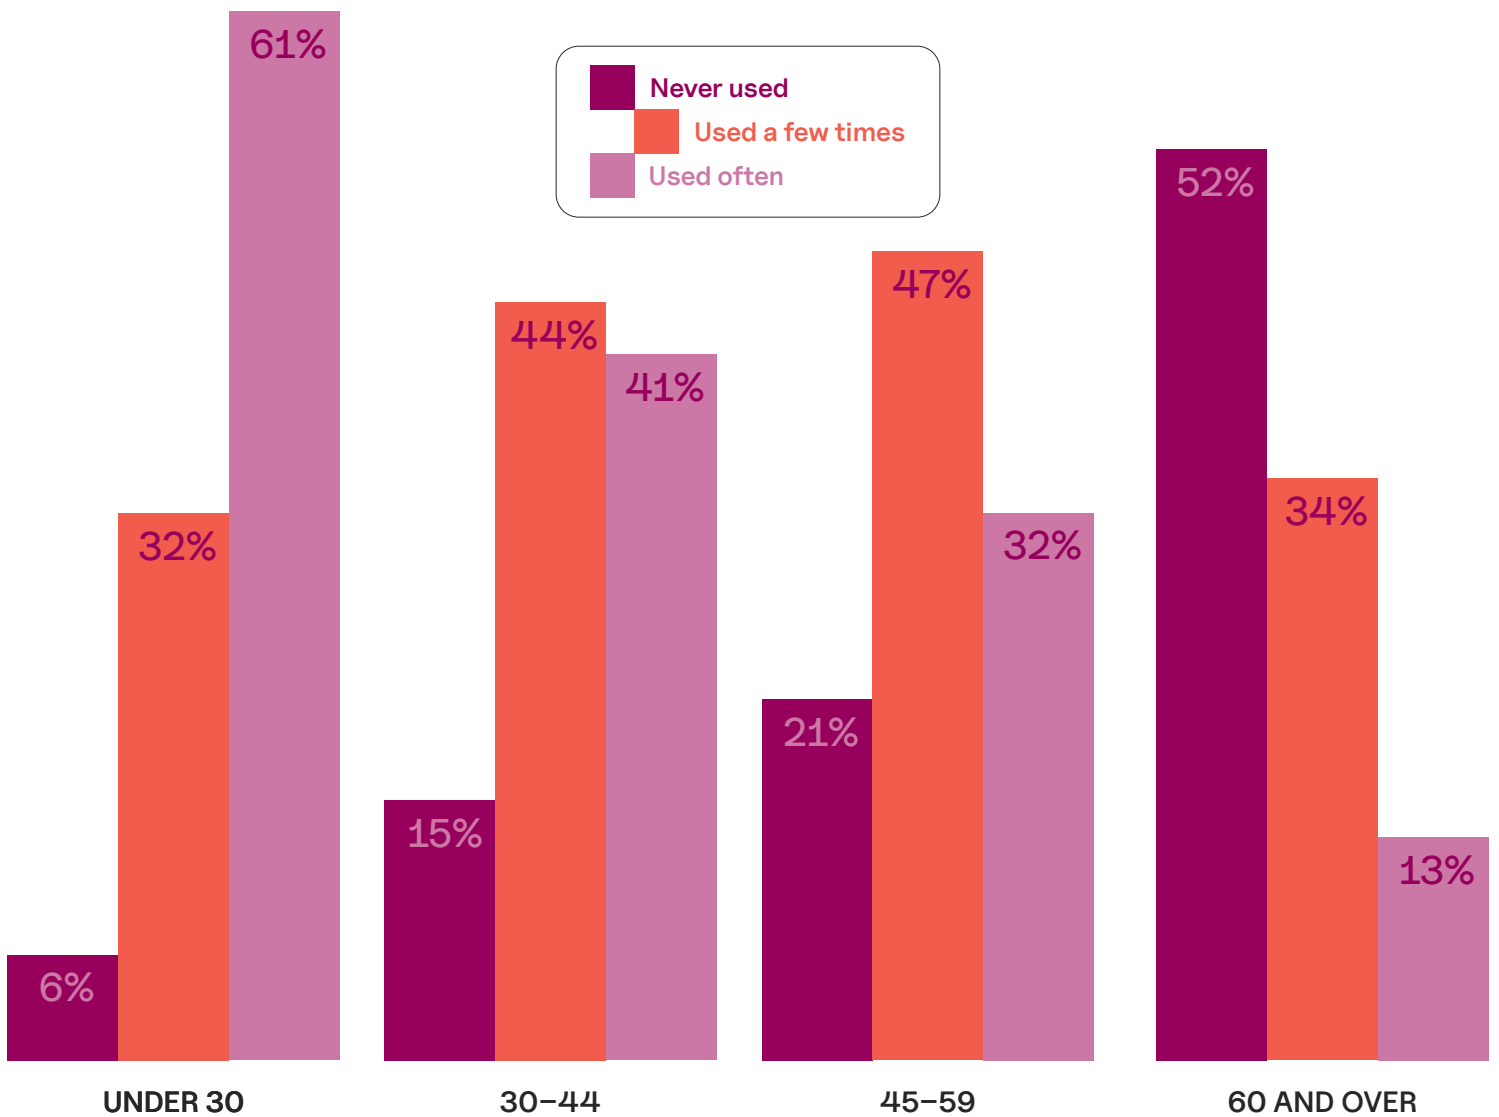
THE USE OF GENERATIVE AI TOOLS

THE USE OF GENERATIVE AI TOOLS

Most respondents (76%) have used generative AI tools at least once. Just over a third (36%) use them often, whereas nearly a quarter (24%) have never used them at all. This snapshot exists within a rapidly shifting landscape: Swiss tracking data from 2024 found that 19% of internet users reported using generative AI tools often or regularly, a figure that had more than doubled by 2025 (Latzer & Festic, 2024; Latzer et al., 2025). The present survey finds 36% reporting frequent use, consistent with this trajectory.

Aggregate adoption figures, however, conceal as much as they reveal. The distribution across demographic groups is uneven and follows a clear pattern.

FIGURE 4: FREQUENCY OF GENERATIVE AI USE, BY AGE



Source: ai-identities Pilot Survey, October 2025 (n=1,000, weighted).
Question F3, German- and French-speaking Switzerland.

Table 2: Frequency of Generative AI Tools Use

Answers	AI Use				Attitudes			Age				Education (ISCED)			Region		Gender	
	Total	Never	Few	Often	Concerned	Both	Excited	<30	30-44	45-59	60+	1-2	3	5-8	GSS	FSS	F	M
Never	24%	100%	0%	0%	48%	12%	1%	6%	15%	21%	52%	30%	31%	12%	21%	33%	23%	26%
Few	40%	0%	100%	0%	38%	52%	21%	32%	44%	47%	34%	36%	38%	46%	40%	39%	37%	43%
Often	36%	0%	0%	100%	14%	36%	78%	61%	41%	32%	13%	34%	32%	43%	39%	28%	40%	31%
Total	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Valid n	1000	241	402	357	405	387	209	225	243	264	252	161	504	335	733	267	515	482

F3: How often have you used generative artificial intelligence tools?

I have never used AI tools = 1, I have used them a few times = 2, I have used them often = 3

Age is the dominant variable. Among under-30s, 61% use generative AI tools often and only 6% have never used them. For the over-60s, those numbers invert: 52% have never used generative AI tools and only 13% use them often. The middle-aged groups (30-44 and 45-59) sit much closer to each other than to either extreme. **Education** follows a consistent but less extreme pattern. Among those with tertiary education, 43% use generative AI tools often and 12% have never used them. Among those with lower educational attainment, the never-used share rises to 30%, whereas frequent use falls to around 34%. The difference is real and consistent, but age remains the stronger variable. The **regional** gap is worth specific attention. In French-speaking Switzerland, 33% have never used generative AI tools, compared with 21% in German-speaking Switzerland. Frequent use also differs: 28% versus 39%. This roughly eleven percentage point gap appears consistently across both use and attitude measures throughout the data, making it one of the more robust regional findings in the survey. The reasons for this gap warrant investigation. Possible explanations include differential availability of high-quality French-language AI interfaces, occupational and sectoral differences between the regions, and cultural differences in how new technologies are adopted. Italian-speaking Switzerland is not included in this wave of data collection, and represents a clear limitation of the current data set. The French-language subsample ($n = 267$) is substantially smaller than the German-speaking one ($n = 733$); therefore subgroup comparisons should be treated with appropriate caution, and findings here should be regarded as warranting targeted research rather than as definitive characterisations of the regions.



"For me, generative AI is like a hammer because it is very useful as a tool, but dangerous as a weapon."

GENDER INSIGHTS

A long-standing body of research documents persistent gender differences in technology adoption, with men typically engaging more frequently and across a wider range of online activities (Fallows, 2005; Helsper, 2010). International evidence on generative AI broadly supports this pattern (Carvajal et al., 2024; Humlum & Vestergaard, 2025; International Telecommunication Union, 2025; Otis et al., 2025; Tang et al., 2025; World Economic Forum & LinkedIn, 2025), though estimates of the gap vary considerably – from women being around 25% less likely to use generative AI tools than men (Otis et al., 2025) to near-parity (Deloitte, 2025).

This survey's data tell an interesting story: taken together, women report slightly higher levels of generative AI tools use and slightly less concerned attitudes toward AI than men. Specifically, women are somewhat less likely to have never used generative AI tools (23% vs. 26%) and report more frequent (= often) generative AI tool use than men (40% vs. 31%). Women also are slightly less concerned than men, with an average score of 2.8 on a 1–5 scale (compared to 2.6 for men, where 1 = much more concerned and 5 = much more excited). This pattern is also reflected in their attitudes: women are less likely than men to say they are more concerned than excited (37% vs. 44%) and more likely to say they are more excited than concerned (25% vs. 16%). The difference is modest but worth noting.

The most plausible reading of this survey data is not that the gender gap has reversed, but that it has largely closed. Swiss longitudinal data show extremely rapid diffusion of generative AI tools among Swiss internet users. In 2023, men were much more likely than women to have used generative AI at least once (47% vs. 27%), but by 2025, the gap is minimal (74% vs. 71%) (Latzer et al., 2025). This is consistent with earlier research suggesting that once adoption barriers fall, women's use patterns can converge with or surpass men's (Fallows, 2005; Gottfried & Park, 2025).

Two caveats are important. First, online panel surveys tend to over-represent digitally active populations, and any skew toward digitally active women would affect these numbers. Second, the gender differences are modest. This finding warrants replication in nationally representative data before strong conclusions are drawn.

Women report slightly higher levels of generative AI tools use and slightly less concerned attitudes toward AI than men. Specifically, women are somewhat less likely to have never used generative AI tools (23% vs. 26%) and report more frequent (= often) generative AI tool use than men (40% vs. 31%).



/Prompt: A cinematic close-up portrait of a young person with red hair and freckles sitting by a window on a modern tram or train. Sunlight streams through the window, illuminating their face as they gaze thoughtfully outside. Green scenery passes by in the background. Calm, introspective mood, shallow depth of field, natural light, realistic skin texture, modern European public transport. Optional style modifiers: golden hour light, 50mm lens, quiet reflection, cinematic realism


THE RELATIONSHIP BETWEEN ATTITUDES AND USE

THE RELATIONSHIP BETWEEN ATTITUDES AND USE

One of the most consistent patterns in the data concerns the relationship between generative AI tool use and attitudes toward AI. Among those who have never used generative AI tools, concern overwhelmingly dominates: 81% report being more concerned than excited, while only 1% lean toward excitement. Among people who use generative AI tools often, this pattern reverses. Only 16% report being more concerned than excited, while 45% say they are more excited than concerned.

At the same time, a substantial portion of respondents occupy a middle ground. Among non-users, 19% feel equally concerned and excited, a share that rises to 50% among those who have used AI a few times. Even among frequent users, 39% continue to hold this view. Taken together, this suggests that while experience with AI is associated with a shift from concern toward excitement, a significant share of people – especially those with some experience – maintain an equally concerned and excited position.

FIGURE 5: AI ATTITUDES BY FREQUENCY OF USE



Source: ai-identities Pilot Survey, October 2025 (n=1,000, weighted).
Questions F1 x F3, German- and French-speaking Switzerland.

At the same time, a substantial portion of respondents occupy a middle ground. Among non-users, 19% feel equally concerned and excited, a share that rises to 50% among those who have used AI a few times. Even among frequent users, 39% continue to hold this view.

What the Association Does and Does Not Tell Us

It is important to be precise about what this pattern shows. The data demonstrates an association between use and more positive attitudes. It does not establish the direction of causation, and several interpretations are plausible.

The first interpretation is that direct experience with generative AI tools reduces concern – a pattern consistent with the “mere exposure effect” documented in cognitive psychology (Zajonc, 1968). On this reading, higher concern among non-users reflects the distance from the technology, and would likely shift with greater engagement.

The second interpretation is that more AI-positive individuals self-select into use: people who are already open to new technologies are both more likely to try generative AI tools and more likely to feel excited about AI. Demographic variables are known to shape AI attitudes before any direct exposure occurs (Venkatesh et al., 2003), and usage may therefore reflect a prior disposition.

The third interpretation is that habitual use produces familiarity and social normalisation without necessarily constituting better-calibrated judgment about AI. UNESCO’s guidance on AI ethics explicitly cautions that informed public concern is a feature of responsible AI governance rather than an obstacle to it (UNESCO, 2021), and frequent use does not guarantee that the concerns that remain are the only ones worth holding.

Tracking the same individuals over time would be needed to separate these effects. The present findings should be read with that limitation clearly in mind. The report refrains from asserting which of these interpretations is most plausible, because the data does not permit that judgment.

The “Used a Few Times” Group

The group that has engaged with generative AI tools a few times but not adopted them more frequently is analytically interesting for reasons distinct from the causal question. Half of this group (50%) describes itself as equally concerned and excited – a substantially higher proportion than the 39% in the total sample. This group appears to sit at a threshold: sufficiently exposed to have formed a differentiated affective response, but not so immersed that the tension between concern and excitement has resolved in either direction.

Research on technology adoption identifies this threshold as a critical moment: people who have tried something but not yet made it a habit are more likely than others to simply stop using it, particularly when their feelings are mixed (Rogers, 2003). Which way this group tips is likely to depend on perceived usefulness, one’s social environment (e.g., peer, family, school, and workplace attitudes around AI), and the quality of early experiences with AI. This is a question the study intends to examine in subsequent longitudinal waves.

Implications for the Public Conversation

The data suggests that the public conversation about AI has a structural feature worth acknowledging. The voices most likely to express strong concern are also those with the least direct experience of what they are concerned about. This is not equivalent to saying their concerns are uninformed or invalid. Loss of human agency and control are real concerns that do not require having used generative AI tools to be well-founded. This feature of the conversation is relevant for how debate is structured and who is heard, not for whether non-users' concerns deserve to be taken seriously.

This feature of the conversation is relevant for how debate is structured and who is heard, not for whether non-users' concerns deserve to be taken seriously.

IV.



/Prompt: A photorealistic indoor scene of a person sitting at a desk, working on a laptop while holding a smartphone to their ear. Soft natural window light fills a minimal, bright room with neutral walls. The person has long hair and wears a simple grey sweater. Their expression is thoughtful and focused. Clean composition, modern lifestyle photography, shallow depth of field, calm and professional atmosphere. Style modifiers: soft daylight, editorial photography, Scandinavian interior, realistic skin texture

WHAT GENERATIVE AI METAPHORS REVEAL

**FOR ME,
GENERATIVE
ARTIFICIAL
INTELLIGENCE
IS LIKE _____
BECAUSE _____.**

Metaphors → Metaphors are a way of speaking as well as thinking. When we describe something unfamiliar by comparing it to something we already know, we are reaching for a vivid image in our mind's eye, and revealing how we actually understand it. The comparison we choose carries assumptions about what the thing is, how it works, how it might affect us, and how we engage with it. Someone who describes generative AI as a tool will think about and will respond to it differently than someone who describes it as a force of nature.

Why Metaphors Matter: A Conceptual Note

Metaphors are more than just colourful ways of saying things. Research in cognitive linguistics demonstrates that they shape how people organise abstract concepts, including technologies they do not fully understand (Lakoff & Johnson, 1980; Maas, 2023). The difference between generative AI is like “a tool” and generative AI is like “a force of nature” is not rhetorical: the first implies human control and responsibility, the second implies something operating beyond human intentionality. These framings have practical consequences for governance. For example, someone whose mental model is an intern – capable but requiring supervision – will likely respond differently to regulatory proposals than someone whose model is a Big Brother or the Sorcerer’s Apprentice.


Standard attitude surveys capture what people think they feel. Metaphors can offer a partial window into the conceptual frames that shape which risks feel salient, which interventions feel intuitive, and how open someone is to changing their mind. A person who reaches for a tool metaphor and a person who reaches for a monster metaphor may score identically on the five-point concerns vs. excitement scale while inhabiting meaningfully different conceptual worlds. This is why listening to how people talk about generative AI, not just what they score on a scale, matters for anyone trying to communicate about or regulate it.

About the Metaphor Exercise

As an exploratory complement to the quantitative measures, this survey asked respondents to complete the sentence: “For me, generative artificial intelligence is like _____ because _____.” The open responses were coded by the research team into broad emotional categories – positive, negative, mixed (that is, both positive and negative), and neutral – and analysed for patterns across attitudes, use, and demographic variables.

Several methodological cautions are important. The coding was developed inductively from the response material and should be understood as exploratory rather than confirmatory. Formal inter-rater reliability assessment is planned for subsequent waves; the valence coding presented here has not yet been subjected to systematic reliability testing across independent coders. Where percentage figures are reported from this coding, they should be read as indicative rather than precise numbers. The qualitative dimensions of the analysis – the specific framings and their governance implications – are more reliable than the aggregate quantitative breakdowns. The examples quoted below are selected to represent the range of framings observed; they are not representative of the overall sample in a statistical sense.

"For me, generative AI is like a soap bubble because it will burst at some point."



The Difference Between Attitudes and Frame

The most striking finding is not visible in attitude scores at all. Consider three respondents who each score maximum excitement on the attitude scale:

«eine Katze am Keyboard – weil sie so tut, als wüsste sie alles, und manchmal stimmt's sogar» [“a cat at a keyboard – because it acts as if it knows everything, and sometimes it's even right”]

Respondent, 46, negative, very excited

«moderne Sklaven – weil sie das Leben einfacher machen wird» [“modern slaves – because they will make life easier”]

Respondent, 55, positive and negative, very excited

«Sauser¹ – weil man nicht immer genau weiss wie der geschmack ist, aber es trotzdem fein ist» [“cider (“Sauser”) – because you don't always know exactly what the taste is like, but it is still tasty”]

Respondent, 23, positive, very excited

All three are very excited about AI. None reaches for an enthusiastic metaphor. Their attitudes and their conceptual language point in different directions. It suggests that sorting people into 'pro-AI' and 'anti-AI' camps based on attitude scores alone misses something important about how they actually think. Someone can perceive AI to be exciting while holding a fundamentally humorous or sceptical frame for understanding it. Such contradictions and coexisting orientations warrant further exploration.

At the aggregate level, metaphor valence broadly mirrors attitude. Respondents coded as more concerned than excited use predominantly negative metaphors (approximately 61%), while those coded as more excited than concerned use predominantly positive ones (approximately 69%). The equally concerned and excited group has the highest share of mixed responses, consistent with their intermediate position. These aggregate patterns, however, are less interesting than the individual-level mismatches they obscure, and the percentage figures carry the interpretive nature of exploratory coding.

¹ A traditional, seasonal autumn drink in Switzerland, Germany, and Austria, consisting of fresh grape juice that has begun to ferment.

The Language of Concern

Among more concerned than excited respondents, a few themes and subthemes emerge from the responses. First, loss of human agency and control:

«etwas Unheimliches – weil es sich unkontrolliert weiterentwickelt. Die Geister, die ich rief, die werd ich nimmer los» [“something uncanny – because it keeps developing uncontrolled. The spirits that I called, I can never get rid of them”]

Respondent, 73, negative, more concerned [echoing Goethe's Zauberlehrling, 1797]

A subtheme is surveillance and power imbalance:

«Big Brother – weil ich denke, es wird ein Profil von mir erstellt» [“Big Brother – because I think a profile is being created of me”]

Respondent, 56, negative, very concerned

Another subtheme is cognitive displacement – AI as something that does your thinking for you:

«die Auslagerung des menschlichen Gehirns – weil wir daran ver dummen werden» [“the outsourcing of the human brain – because we will become stupid from it”]

Respondent, 53, negative, very concerned

The cognitive displacement concern deserves particular attention. It connects to longstanding anxieties about what capacities humans lose when technologies routinise previously human cognitive tasks (Carr, 2010; Dellermann et al., 2019). These are substantive concerns with an empirical literature behind them.

A separate analysis of the metaphor responses against an established AI risk taxonomy (Slattery et al., 2024) reveals that the dominant themes are human-computer interaction (concerns related to problematic relationships between humans and AI, including overreliance and loss of human agency) and AI system safety failures (concerns related to AI behaving unsafely, doesn't do what it's intended to do, is unreliable, or could cause harm). Other categories that dominate formal AI policy debates – such as socioeconomic and environmental impacts, misinformation, privacy and security, malicious use, and discrimination and toxicity – are registered far lower in our data.

This is a meaningful gap, but one that deserves careful framing. When people in this survey reach for a risk frame, they reach predominantly for human agency and control concerns. This is partly a question of how people tend to articulate concerns about complex systems: in social and relational language rather than technical language. It does not follow that their underlying concerns are non-technical; rather, that the same concern (about opacity, for example) can be expressed as a technical worry about model interpretability or as a social worry about uncanny, uncontrollable behaviour. What the metaphor data shows is the language through which concerns are held, which has direct implications for how communication about AI should be framed. Technical reassurances will often fail to address more social concerns.

Current AI regulation, particularly in the European context, is substantially oriented around technical risk categories (Veale & Borgesius, 2021): model audits, bias testing, transparency requirements, and safety benchmarks. The EU AI Act (European Parliament and Council of the European Union, 2024), in its final form, does include provisions addressing some social concerns – fundamental rights impact assessments, transparency obligations, AI literacy, human oversight requirements, and prohibitions on certain social scoring systems. Whether these provisions are sufficient to address the kinds of social and political concerns dominant in this survey’s data is a legitimate open question. This report cannot resolve it, but the gap between how publics frame AI risk and how governance frameworks categorise it is a gap worth examining as the Act is implemented.

"For me,
generative AI
is like Formula
1 because the
car does a lot,
but you still
need a driver."



The Language of Excitement

Among excited respondents, the dominant framings are expansive rather than instrumental – generative AI as something that extends human reach rather than replacing human capacity:

«eine Welt der Erfüllung und Kreativität – weil es völlig in anderen Dimensionen denkt» [“a world of fulfillment and creativity – because it thinks in completely different dimensions”]

Respondent, 37, positive, very excited

«eine Horizonterweiterung – weil sie mich auf Gedanken bringt, die ich allein nicht gehabt hätte» [“a broadening of horizons – because it brings me thoughts I wouldn't have had on my own”]

Respondent, 38, positive, very excited

The distinction between extension and replacement metaphors matters. Respondents who frame AI as extending human capacity – a horizon-broadener, a creative partner, a different kind of mind – tend to position themselves as agents using a powerful tool. Those who frame it as replacing human capacity – outsourcing the brain, doing the thinking for you – tend toward concern, even when they score high on excitement. This parallels findings in the human-computer interaction literature on how “augmentation” versus “automation” framings shape technology adoption and perceived threat (Brynjolfsson & McAfee, 2014; Hancock et al., 2011). The metaphors people reach for may, in other words, be predictive (as discussed by Lakoff and Johnson [1980] though not in the context of AI specifically) of how they engage with AI in practice – a question worth pursuing in future research.

"Necessary evil" is more resigned in tone, but it too resists simple categorisation: the respondent is not against AI or for it, but has thought through a position on it – one that many attitude scales would designate as a neutral stance.

The Language of Ambivalence

The responses from the “equally excited and concerned” group most directly capture a sense of mixed feelings, and are often the most inventive:

«Lebkuchen – weil es lecker aber klebrig ist»
[“gingerbread – because it's tasty but sticky”]
Respondent, 16, positive and negative, equally concerned and excited

«ein notwendiges Übel – weil wir früher oder später die KI brauchen werden» [“a necessary evil – because sooner or later we will need AI”]
Respondent, 68, positive and negative, equally concerned and excited

Although ambivalence toward new technologies is well-documented (Mick & Fournier, 1998), it is notable how articulately respondents hold two frames at once without resolving the tension between them. Gingerbread is genuinely appetising and genuinely messy, and the respondent seems comfortable with both at once. “Necessary evil” is more resigned in tone, but it too resists simple categorisation: the respondent is not against AI or for it, but has thought through a position on it – one that many attitude scales would designate as a neutral stance.

V.



/Prompt: A photorealistic cinematic portrait of an older woman sitting on a wooden bench beside a calm alpine lake, surrounded by lush green trees and dramatic mountains in the background. She is turned slightly away but looks back over her shoulder toward the camera with a gentle, warm smile. Natural daylight, soft sunlight reflecting on the water, shallow depth of field. She wears a dark outdoor jacket. Peaceful, serene mood. Shot at eye level, documentary-style photography, high realism, natural colors. Style modifiers: 35mm lens, f/2.8, Swiss alpine landscape, lifestyle photography, soft haze

WHAT THIS
MEANS

Three insights emerge from this data as genuinely consequential for researchers, but more importantly, for anyone thinking about how Switzerland navigates the next phase of AI integration.

The Structural Imbalance in Public Debate

Scholarship on public participation in technology governance has long identified a “participation gap” – the structural underrepresentation of certain publics in the conversations that shape technology policy (Jenkins, 2006; Robinson et al., 2020; van Dijk, 2020). The AI data in this survey illustrates a specific form of this gap: the people most likely to express strong concern are systematically the people with the least direct experience of what they are concerned about.

Non-users – predominantly older respondents – are forming views about AI without hands-on exposure to current tools. What this does and does not imply requires careful reasoning.

It does not imply that their concerns are less valid. The science and technology studies (STS) literature offers an important corrective here. Wynne’s (1992, 2006) work on lay expertise demonstrates that non-expert concern is often grounded in legitimate knowledge about social context, power relationships, and institutional trustworthiness that technical experts may underweight. Jasanoff’s (2005) concept of civic epistemologies argues that nationally and socially specific ways of knowing and trusting are not deficits to be overcome but rather features of democratic legitimacy. Loss of human agency and control, or more specific themes such as surveillance, power imbalance, and cognitive displacement are not concerns that require direct experience with generative AI tools to be well-founded. They are concerns about social issues, and they can be held intelligently by people who have never used such tools.

What the structural imbalance does imply is something more specific: it means the public conversation risks treating as broadly representative a debate that is substantially conducted by and for those already familiar with the technology. Non-users’

concerns are real, but non-users may not be equipped to distinguish as easily between concerns that are empirically well-supported now, concerns that may become well-supported as the technology develops, and concerns that are speculative or based on inaccurate information. Frequent users, conversely, may have developed familiarity that normalises concerns that still warrant attention.

There are, in other words, two distinct problems worth distinguishing. The first is the possible absence of well-calibrated concern: non-users may hold valid concerns in an undifferentiated form that more engagement with the technology could sharpen and refine. The second is the potential absence of legitimate voices from governance conversations: non-users’ structural concerns about loss of human agency and control may be underweighted in policy contexts dominated by those with stronger incentives toward excitement.

A better question than “how do we reassure the concerned?” – which assumes that concern is a deficit to be corrected – is: “how do we ensure that Switzerland’s public conversation about AI is inclusive, valuing voices across the use and attitude spectrum?” A separate but equally important question is: “what would it mean to take the social and political content of non-users’ concerns seriously as governance inputs, not merely as attitudes to be shifted?”

These are questions for schools, workplaces, public institutions, and governance bodies, as much as for technology companies. This report is diagnostic rather than prescriptive: translating these findings into specific policy interventions would require stakeholder engagement, institutional context, and deliberative processes that go beyond what a survey-based study can deliver.

Mental Models Matter for Communication and Governance

The metaphor exercise, however exploratory, points toward something the quantitative data cannot show on its own: people come to any conversation about AI with a pre-existing conceptual frame, and that frame likely shapes what they hear and what they want.

An abstract assurance about “ethical AI development” lands differently for someone whose mental model is a hammer than for someone whose mental model is Big Brother. The message may be formally identical, but the frame determines whether it registers as reassuring or evasive.

For policymakers and communicators, this has a concrete implication: effective AI communication needs to start by acknowledging the frames people already hold and speak to those specifically, rather than leading with abstract principles. An oversight-centred framing – one that describes AI as requiring active human supervision, capable of error, and subject to accountability – is more likely to address the Sorcerer’s Apprentice concern than a reassurance that developers are following ethical guidelines. Whether such framings are also more effective in generating durable public trust is an empirical question that subsequent research – including experimental work with competing framings – could usefully pursue.

Ambivalence Is an Asset, Not a Problem to Be Solved

A large group in this survey (39%) refuses to take a side and holds both concern and excitement simultaneously. They sit disproportionately among occasional users – people with enough experience to see both opportunity and risk, but not enough to have resolved the tension into a settled position. They use complex and creative language to describe generative AI: necessary evils, gingerbread, colleagues who talk a lot but don’t really know anything.

The temptation, from a communications or policy perspective, is to treat this group as a problem: people who need to be moved to a clearer position, resolved into either support or opposition. Yet, a public that holds complexity might be better equipped for nuanced governance conversations than one that has sorted itself into clear camps. The equally concerned and excited group is paying close attention to both sides of an evolving argument, which continues to expand and diversify as the technology advances.

Rather than attempting to resolve ambivalence prematurely, the goal of public communication about AI should be to give it something substantive to work with: honest accounts of what is known and unknown about AI’s effects, realistic descriptions of oversight mechanisms and their limits, and deliberative processes that include rather than work around the voices most worried about the technology’s consequences.

“how do we ensure that Switzerland’s public conversation about AI is inclusive, valuing voices across the use and attitude spectrum?”

VI.



/Prompt: A realistic outdoor scene of a man in his early 60s walking a medium-sized light-colored dog on a leash along a muddy forest path in late autumn. Tall bare trees line the path, fallen leaves cover the ground. The man wears a practical outdoor jacket, jeans, and rubber boots. The dog walks happily beside him. Overcast natural light, muted earthy tones, centered composition. Calm, everyday documentary feeling, photorealistic. Style modifiers: cinematic realism, 50mm lens, natural textures, European woodland

ABOUT THIS STUDY

Research Questions

The survey posed three questions.

First, an attitude measure: “Thinking about the increasing use of artificial intelligence in daily life, would you say it makes you feel more excited or more concerned?” Respondents answered on a five-point scale from much more concerned than excited to much more excited than concerned, with equally concerned and excited at the midpoint. This item mirrors the wording used in Pew Research Center international surveys, enabling contextualised comparison.

English	German	French
<p>Q1 Thinking about the increasing use of artificial intelligence in daily life, would you say it makes you feel more excited or more concerned?</p> <ul style="list-style-type: none"> • 1 = Much more concerned than excited • 2 = Somewhat more concerned than excited • 3 = Equally concerned and excited • 4 = Somewhat more excited than concerned • 5 = Much more excited than concerned 	<p>Q1 Wenn Sie an die zunehmende Nutzung von künstlicher Intelligenz im Alltag denken – fühlen Sie sich dadurch eher begeistert oder eher besorgt?</p> <ul style="list-style-type: none"> • 1 = Viel besorgter als begeistert • 2 = Eher besorgter als begeistert • 3 = Gleichermassen besorgt und begeistert • 4 = Eher begeisterter als besorgt • 5 = Viel begeisterter als besorgt 	<p>Q1 Lorsque vous pensez à l'utilisation croissante de l'intelligence artificielle dans la vie quotidienne, diriez-vous que cela vous rend plutôt enthousiaste ou plutôt inquiet/inquiète?</p> <ul style="list-style-type: none"> • Beaucoup plus inquiet/inquiète qu'enthousiaste • Plutôt plus inquiet/inquiète qu'enthousiaste • Autant inquiet/inquiète qu'enthousiaste • Plutôt plus enthousiaste qu'inquiet/inquiète • Beaucoup plus enthousiaste qu'inquiet/ inquiète

Second, a use frequency measure: “How often have you used generative artificial intelligence tools?” with three response options: never, a few times, and often. This item mirrors the wording used by Latzer and Festic (2024).

English	German	French
<p>Q2 How often have you used generative artificial intelligence tools?</p> <ul style="list-style-type: none"> • 1 = I have never used AI tools • 2 = I have used them a few times • 3 = I have used them often 	<p>Q2 Wie häufig haben Sie bereits generative künstliche Intelligenz Tools verwendet?</p> <ul style="list-style-type: none"> • 1 = Habe ich noch nie verwendet • 2 = Habe ich schon ein paar wenige Male verwendet • 3 = Habe ich schon oft verwendet 	<p>Q2 À quelle fréquence avez-vous déjà utilisé des outils d'intelligence artificielle générative?</p> <ul style="list-style-type: none"> • Je ne les ai jamais utilisés • Je les ai utilisés quelques fois • Je les ai utilisés souvent

Third, a metaphor prompt: “Complete this sentence: ‘For me, generative artificial intelligence is like __ because __.’”

English	German	French
<p>Q3 Complete this sentence: For me, generative artificial intelligence is like _____ because _____.</p>	<p>Q3 Vervollständigen Sie diesen Satz: Für mich ist generative künstliche Intelligenz wie _____, weil _____.</p>	<p>Q3 Compléte cette phrase: Pour moi, l'intelligence artificielle générative c'est comme _____, parce que _____.</p>

Open responses were coded by the research team into four valence categories (positive, negative, neutral, mixed) and a richer set of analytical dimensions covering ontological status, role and control relationships, and temporal framing, drawing on the tradition of systematic metaphor analysis in qualitative research (Schmitt, 2005). This coding was developed inductively from the response material and should be understood as exploratory: a first-pass framework for listening to how people describe AI in their own words. It is not a fully validated content analysis, and the qualitative dimensions beyond valence have not been quantified in this report. Coding decisions for ambiguous cases were discussed within the research team; formal inter-rater reliability assessment is planned for subsequent iterations of the present survey. Quantitative figures derived from this coding should be treated as indicative rather than precisely calibrated.

We also conducted an exploratory, theory-informed analysis of the negative metaphor responses by mapping them onto an established AI risk taxonomy (Slattery et al., 2024). For more information, visit <https://airisk.mit.edu>.

Risk Category	Description
Discrimination & Toxicity	Risks related to unfair treatment, harmful content exposure, and unequal AI performance across different groups and individuals.
Privacy & Security	Risks related to unauthorized access to sensitive information and vulnerabilities in AI systems that can be exploited by malicious actors.
Misinformation	Risks related to AI systems generating or spreading false information that can mislead users and undermine shared understanding of reality.
Malicious Actors	Risks related to intentional misuse of AI systems by bad actors for harmful purposes including disinformation, cyberattacks, and fraud.
Human-Computer Interaction	Risks related to problematic relationships between humans and AI systems, including overreliance and loss of human agency.
Socioeconomic & Environmental	Risks related to AI's impact on society, economy, governance, and the environment, including inequality and resource concentration.
AI System Safety, Failures & Limitations	Risks related to AI systems that fail to operate safely, pursue misaligned goals, lack robustness, or possess dangerous capabilities.

Rather than applying a strict coding scheme, responses were interpretively reviewed and situated within the taxonomy's categories, allowing us to identify broad patterns in how respondents framed potential risks.

Participants

The analytic sample consists of 1,000 participants aged 15–79 recruited via a commercial online panel in October 2025 in German- and French-speaking Switzerland. The sample covers approximately 92% of the Swiss population by linguistic region; Italian-speaking Switzerland is not included in this wave and represents a clear limitation of the current data.

The raw sample was weighted to Swiss Federal Statistical Office (FSO) population statistics for 2024, using joint distributions of age, education, and gender – a methodologically conservative approach that accounts for interactions between these variables. The residual differences between weighted proportions and FSO targets are small (typically 0.1–0.3 percentage points).

As with all online panel surveys, people with limited or no internet access are not captured. In a study about AI use specifically, this means the non-user group is almost certainly undercounted relative to the broader population. The 24% who report never having used AI should be read as a floor, not a ceiling. The true extent of the participation gap described in this report is likely larger in the full Swiss population than the data shows.

No significance tests or effect sizes are reported in the body of this report. For a sample of $n = 1,000$, modest differences across demographic subgroups may achieve statistical significance; conversely, differences described in qualitative terms (“substantially higher,” “modestly above”) should not be read as implying specific significance thresholds. Full data tables including subgroup confidence intervals and relevant test statistics are available at aidentities.ch.

Research Team

The ai-identities study is conducted by researchers at the Department of Communication and Media Research at the University of Zurich, in partnership with researchers at the Technical University of Munich. For more information please visit aidentities.ch.

REFERENCES



- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.
- Carr, N. (2010). *The shallows: What the internet is doing to our brains*. W. W. Norton.
- Carvajal, D., Franco, C., & Isaksson, S. (2024). *Will artificial intelligence get in the way of achieving gender equality?* (NHH Department of Economics Discussion Paper No. 03). Norwegian School of Economics. <https://ssrn.com/abstract=4759218>
- Deloitte. (2025). Women and generative AI: The adoption gap is closing fast, but a trust gap persists. Deloitte Insights. <https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2025/women-and-generative-ai.html>
- Dellermann, D., Calma, A., Lipusch, N., Weber, T., Weigel, S., & Ebel, P. (2019). The future of human-AI collaboration: A taxonomy of design knowledge for hybrid intelligence systems. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*. <https://doi.org/10.48550/arXiv.2105.03354>
- European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union.
- Fallows, D. (2005, December 28). *How women and men use the internet*. Pew Internet & American Life Project. <https://www.pewresearch.org/internet/2005/12/28/how-women-and-men-use-the-internet/>
- Gottfried, J., & Park, E. (2025, November 20). *Americans' social media use 2025*. Pew Research Center. <https://www.pewresearch.org/internet/2025/11/20/americans-social-media-use-2025/>
- Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*, 53(5), 517–527. <https://doi.org/10.1177/0018720811417254>
- Helsper, E. J. (2010). Gendered internet use across generations and life stages. *Communication Research*, 37(3), 352–374. <https://doi.org/10.1177/0093650209356439>
- Humlum, A., & Vestergaard, E. (2025). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1). <https://doi.org/10.1073/pnas.2414972121>

- International Telecommunication Union. (2025). *WTISD-25: Gender equality in digital transformation*. <https://www.itu.int/hub/2025/05/wtisd-25-gender-equality-in-digital-transformation>
- IPSOS. (2024). *The Ipsos AI Monitor 2024: Changing attitudes and feelings about AI and the future it will bring*. Ipsos. <https://www.ipsos.com/en/ipsos-ai-monitor-2024-changing-attitudes-and-feelings-about-ai-and-future-it-will-bring>
- Jasanoff, S. (2005). *Designs on nature: Science and democracy in Europe and the United States*. Princeton University Press.
- Jenkins, H. (2006). *Confronting the challenges of participatory culture: Media education for the 21st century*. MacArthur Foundation.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.
- Latzer, M., & Festic, N. (2024). *Generative artificial intelligence tools: Knowledge, use and attitudes in Switzerland*. University of Zurich. <https://mediachange.ch/research/artificial-intelligence-applications>
- Latzer, M., Festic, N., Odermatt, C., & Birrer, A. (2025). *World Internet Project – Switzerland 2025*. University of Zurich. <https://mediachange.ch/research/wip-ch-2025/>
- Maas, M. M. (2023). *AI is like... A literature review of AI metaphors and why they matter for policy*. Institute for Law & AI. <https://law-ai.org/ai-policy-metaphors/>
- Mick, D. G., & Fournier, S. (1998). Paradoxes of technology: Consumer cognizance, emotions, and coping strategies. *Journal of Consumer Research*, 25(2), 123–143. <https://doi.org/10.1086/209531>
- OECD. (2025). *How do people experience new technologies and generative AI?: Insights from a few countries worldwide* (OECD Policy Insights on Well-being, Inclusion and Equal Opportunity, No. 23). OECD Publishing. <https://doi.org/10.1787/49b8d10e-en>
- Otis, N. G., Delecourt, S., Cranney, K., & Koning, R. (2025). *Global evidence on gender gaps and generative AI* (HBS Working Paper No. 25-023). Harvard Business School. <https://www.hbs.edu/faculty/Pages/item.aspx?num=66548>
- Poushter, J., Fagan, M., & Corichi, M. (2025). *How people around the world view AI*. Pew Research Center. <https://pewresearch.org/global/2025/10/15/how-people-around-the-world-view-ai>

- Robinson, L., Schulz, J., Blank, G., Ragnedda, M., Ono, H., Hogan, B., ... Khilnani, A. (2020). Digital inequalities 2.0: Legacy inequalities in the information age. *First Monday*, 25(7). <https://doi.org/10.5210/fm.v25i7.10842>
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Schmitt, R. (2005). Systematic metaphor analysis as a method of qualitative research. *The Qualitative Report*, 10(2), 358–394. <https://doi.org/10.46743/2160-3715/2005.1854>
- Slattery, P., Saeri, A. K., Grundy, E. A. C., Graham, J., Noetel, M., Uuk, R., Dao, J., Pour, S., Casper, S., & Thompson, N. (2024). The AI Risk Repository: A comprehensive meta-review, database, and taxonomy of risks from artificial intelligence. <https://doi.org/10.48550/arXiv.2408.12622>. Visit also: <https://airisk.mit.edu>
- Tang, C., Li, S. K., Hu, S., Zeng, F., & Du, Q. (2025). Gender disparities in the impact of generative artificial intelligence: Evidence from academia. *PNAS Nexus*, 4(2), pgae591. <https://doi.org/10.1093/pnasnexus/pgae591>
- UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- van Dijk, J. A. G. M. (2020). *The digital divide*. Polity Press.
- Veale, M., & Borgesius, F. Z. (2021). Demystifying the Draft EU Artificial Intelligence Act. *Computer Law Review International*, 22(4), 97–112. <https://doi.org/10.9785/cr-2021-220402>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- World Economic Forum & LinkedIn. (2025). *Gender parity in the intelligent age*. World Economic Forum. https://reports.weforum.org/docs/WEF_Gender_Parity_in_the_Intelligent_Age_2025.pdf
- Wynne, B. (1992). Misunderstood misunderstanding: Social identities and public uptake of science. *Public Understanding of Science*, 1(3), 281–304. <https://doi.org/10.1088/0963-6625/1/3/004>
- Wynne, B. (2006). Public engagement as a means of restoring public trust in science – hitting the notes, but missing the music? *Community Genetics*, 9(3), 211–220. <https://doi.org/10.1159/000092659>
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2, Pt.2), 1–27. <https://doi.org/10.1037/h0025848>