

Collective intelligence vs artificial intelligence

Report from the roundtable
on Wikimedia and AI

+++



IMD / Real learning
Real impact



Report Index

Objective	02
Context	04
Why Collective Intelligence?	07
Roundtable Insights	10
Wikipedia's knowledge production at a crossroads	11
Wikipedia, democracy and information integrity	13
Challenges	15
Dealing with the paradoxes of openness	16
Wikipedia as a data commons	18
Public and private knowledge	20
The new knowledge loop	22
Tensions	26
The time for change is now	30
Next Steps	32
Participants	34

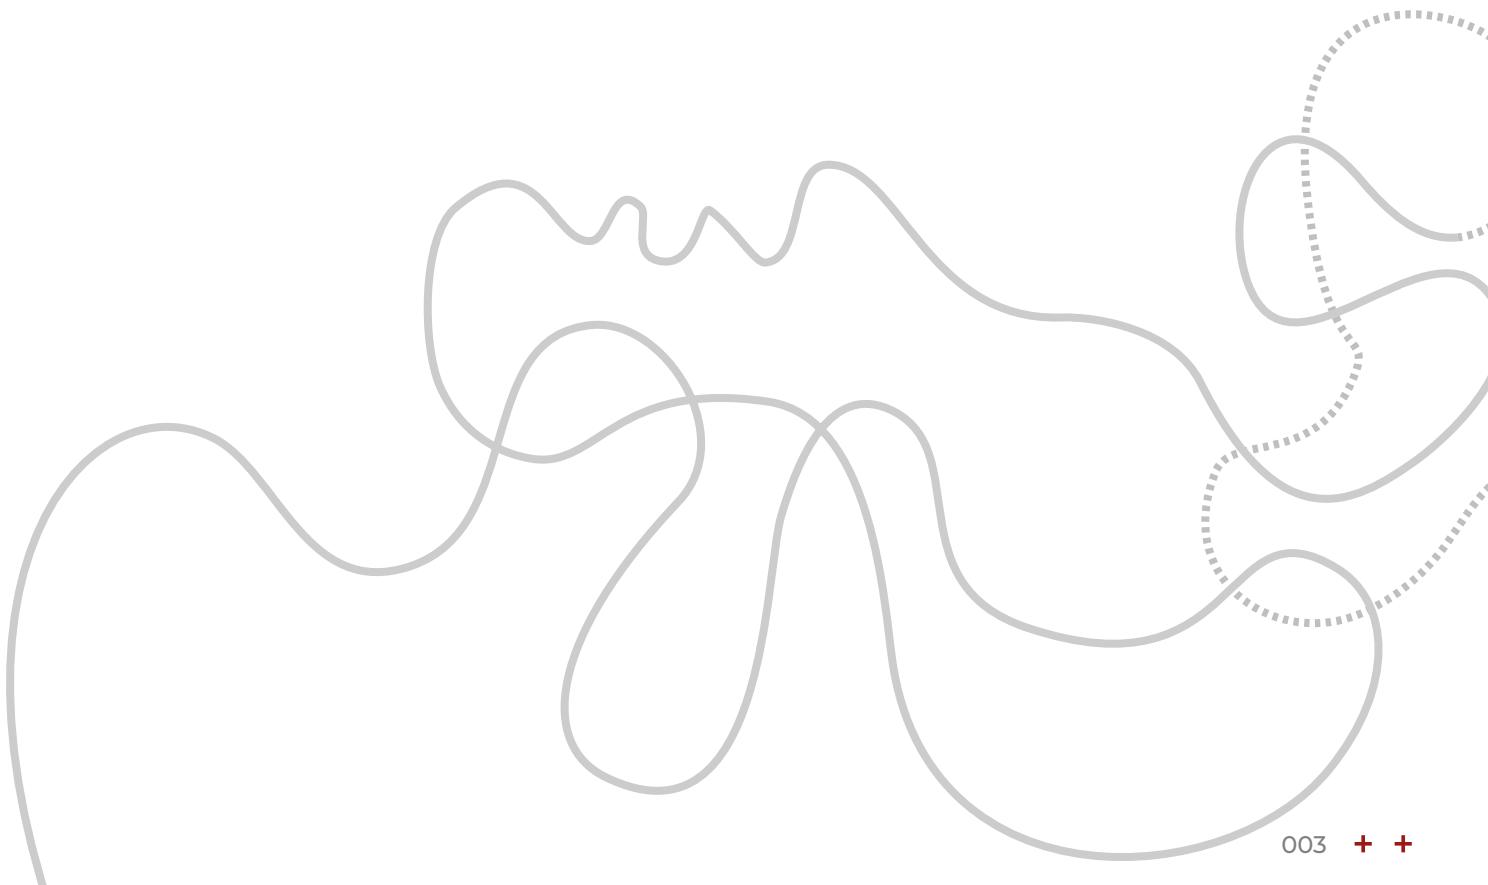
Objective



At the heart of the roundtable “Collective Intelligence vs Artificial Intelligence” was a simple but urgent question: *What happens to the Wikimedia Movement when AI stops merely reading Wikipedia and starts replacing it as a key source of knowledge?* A new “knowledge loop” is emerging in which access to knowledge is increasingly intermediated by AI tools and machines are becoming as important as humans as users of knowledge. This creates a serious risk that knowledge commons like Wikipedia and other Wikimedia platforms will be used for the benefit of AI development, without companies giving back to them.

Our core thesis is that Wikimedia must redefine its role in the age of AI as the backbone of a public, human-governed knowledge infrastructure.

To this end, Wikimedia CH, Open Future and IMD Business school organised a roundtable on 4 November 2025 in Lausanne, Switzerland. The “Collective Intelligence vs Artificial Intelligence” roundtable brought together 20 Wikimedians, AI developers, data scientists, data governance experts, journalists and researchers. The insights from the event will inform the development of a white paper that will shape the Wikimedia Movement's shared position on generative AI technologies.



Context

The Peak Wikipedia
moment is now

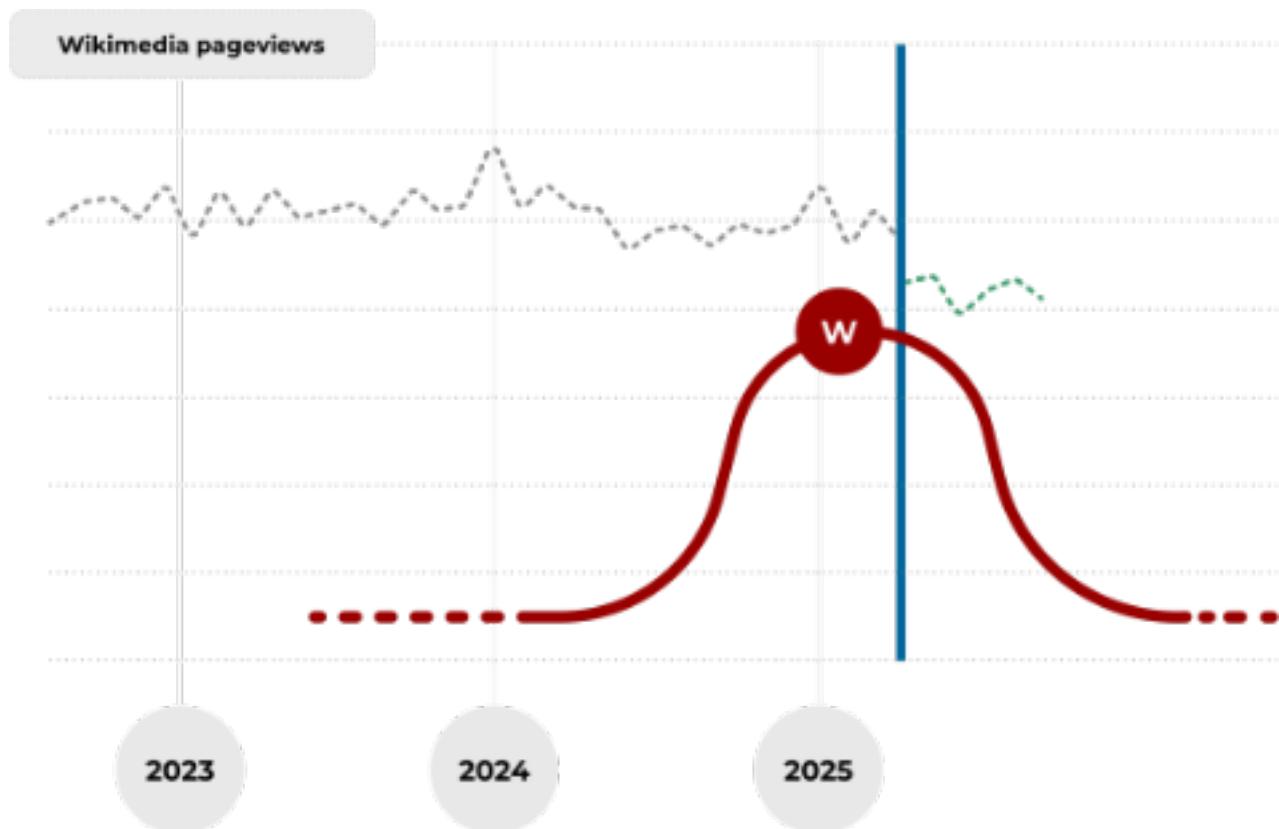


The AI ecosystem stands at a pivotal moment. The existing ecosystem of the Open Web, with search engines playing the key role, is waning. And the rules, economic models, and power structures of the new knowledge loop — which will shape the knowledge ecosystem of the future — are currently being defined.

Wikipedia, at 25 years, is at a tipping point. On one hand, this is the online encyclopaedia's peak moment as a major global platform with a validated model of developing a knowledge commons. Wikipedia has never been so widely used, not only by billions of

humans, but also at massive scale by machines.

But despite its peak strength and use, there is a growing sense that the influence of Wikipedia— and therefore the entire Wikimedia Movement— might be in decline. Wikipedia could increasingly become an invisible layer of the AI ecosystem: heavily used as training data and as a live API endpoint, but no longer visible to, or directly visited by, human users. Wikipedia would continue to feed AI systems, but risk losing editors, resources and ultimately its power to shape how knowledge is produced and governed.



This moment is both a challenge and a strategic opportunity. There is an unprecedented — and limited — window of opportunity for Wikimedia to adjust to the changing landscape, establish its position in the new knowledge loop, and influence its evolution before it crystallises around proprietary models and profit-driven logics.

Our goal through the roundtable and the forthcoming white paper is to propose ways in which AI solutions can support — not replace — collaborative knowledge creation, while safeguarding Wikimedia's role as a source of original, human knowledge. We also aim to develop policies and norms that will ensure sustainability of the new knowledge loop.



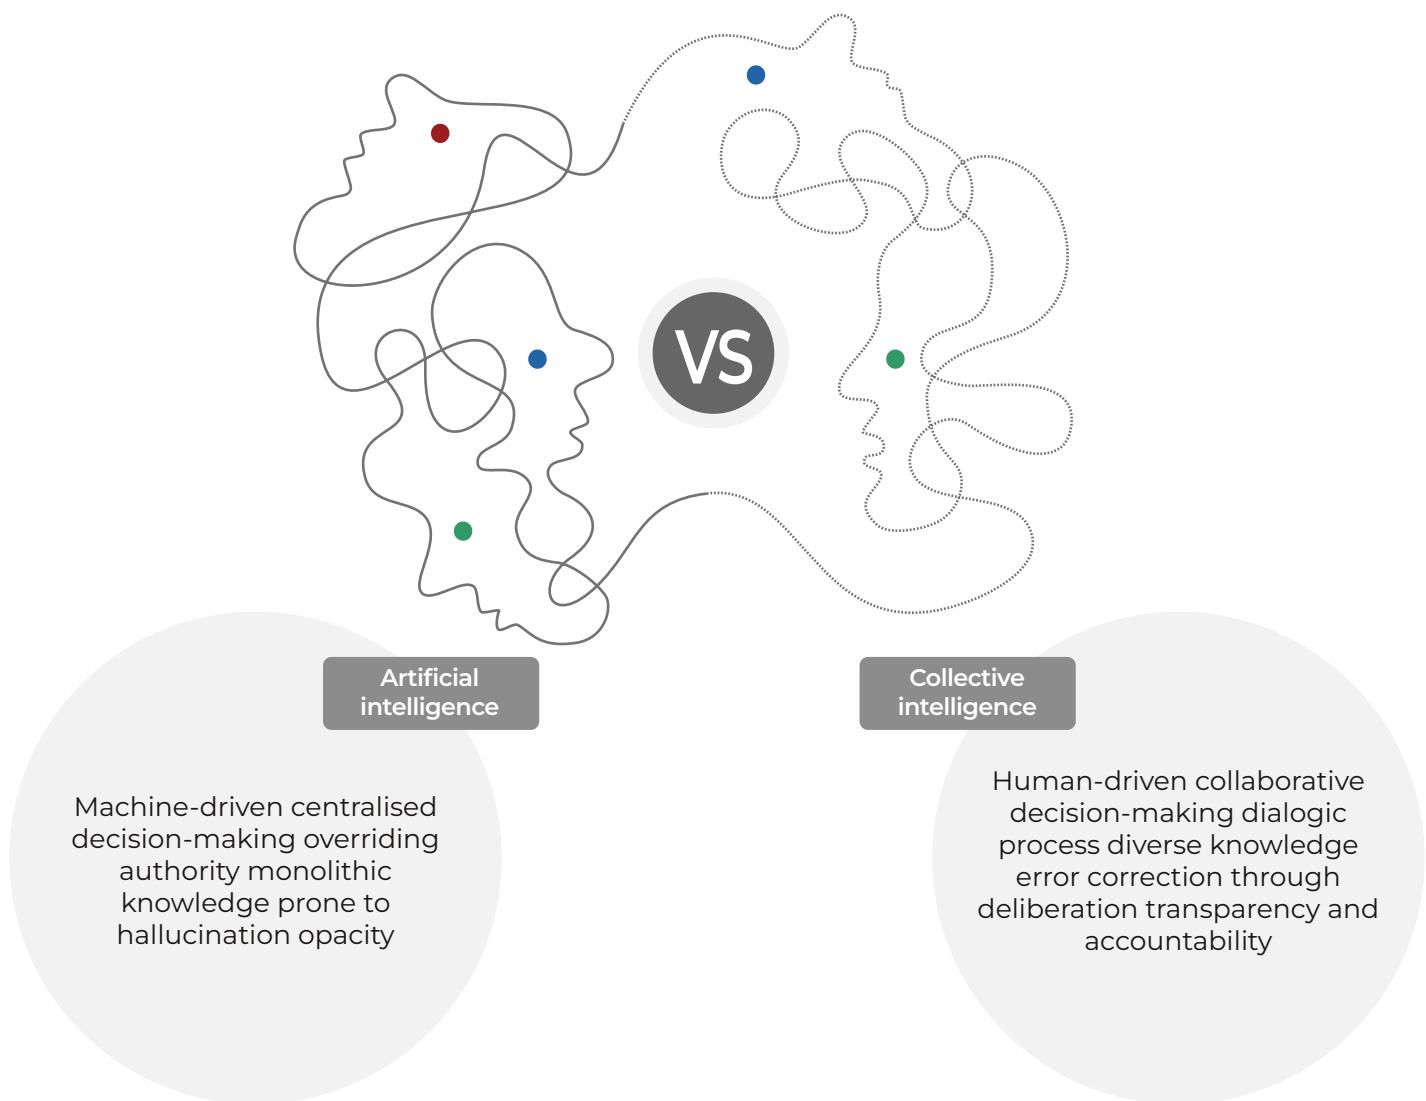
Why Collective Intelligence?

Wikimedia platforms as a product
of collective intelligence



Collective Intelligence is the shared cognitive capacity that emerges when individuals collaborate, pool knowledge, or coordinate efforts to solve problems or make decisions. It is human-driven, relies on collaborative decision-making, draws from diverse knowledge sources, operates through decentralised processing, and benefits from error correction through diversity. Collective intelligence is what emerges on Wikimedia

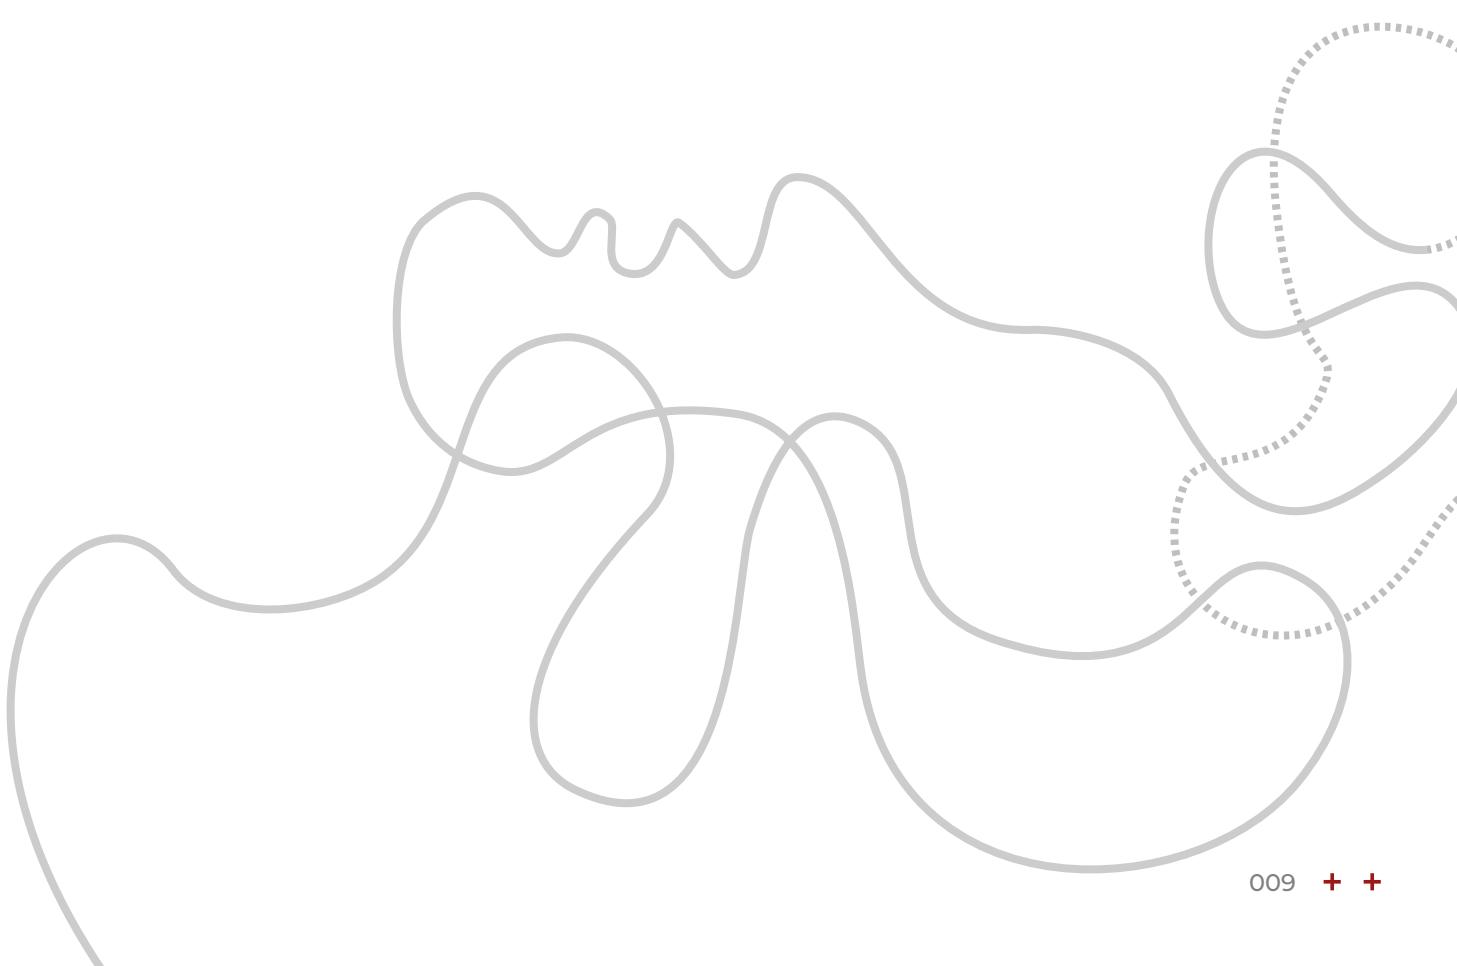
platforms, as editors collaborate and seek consensus on a body of trusted knowledge. In this sense, Wikipedia is not just a website or a collection of articles, but a living infrastructure of collective intelligence. Given that generative AI uses the outputs of this collective work for training and inference data, any sustainable approach to AI must recognise, protect, and strengthen this human process.



Anna Jobin, Vice-President of the Board of Wikimedia CH, in her opening statement said that collective intelligence emerged on Wikipedia from the very beginning. Collective intelligence not only fosters the quality in a way that Wikipedia and the Wikiverse demonstrate, it also fosters trust. This reliability and trust cannot be delegated to machines.

Thinking of Wikipedia as an outcome of collective intelligence allows us to compare its creation to that of generative AI technologies. There is a tension between the collective effort of building a knowledge commons and the central-

ised effort of developing closed, proprietary AI systems. In an era where generative AI is rapidly reshaping how knowledge is produced, validated, and distributed, it is essential to critically examine how collective intelligence can interact with, and benefit from, the use of AI tools.



Roundtable Insights



Wikipedia's knowledge production at a crossroads

At the roundtable, Prof. Misiek Piskorski, Dean of IMD Business School, spoke about emerging trends related to generative AI and their impact on Wikipedia. In recent years, the development of LLMs has enabled a strange new content production loop to emerge. LLMs are trained on Wikipedia articles and can therefore produce similar content — only not as trustworthy. Disintermediation of Wikipedia by AI technologies means that users no longer need to visit Wikipedia directly to obtain information from the platform.

There are signs that Wikimedia's position—established over more than two decades— as one of the most popular web platforms and a unique source of trusted knowledge, is beginning to erode.

The clearest sign of this shift is recent data showing an 8% decrease in user traffic, combined with 50% growth in overall traffic attributed to bots.

Looking forward, the balance between human and machine uses of Wikipedia is the key strategic challenge. For the first twenty-five years, Wikipedia grew to be a site used and read by over a billion humans every year. Recent years have seen increased machine use of Wikipedia, as bots download its content in bulk, and agentic tools like AI search access Wikipedia in real time. In the coming years, various futures are possible. A positive scenario would be that controlled use of generative AI tools supports human editors and readers, and machine uses are properly governed, so that they stay sustainable. Already, various bots have been created to support the work of human editors — a balance of human and AI that can serve as a reference

point for future directions. However, two negative scenarios are possible as well. In one, the only “visitors” to Wikipedia are AI agents, which fully replace human users. In another, AI increasingly writes Wikipedia.

The current moment creates an opportunity for Wikipedia to become a verification layer—the source of “ground truth” for technologies that still hallucinate. As Wikipedia’s main assets are its neutrality and verifiability thanks to human oversight, its community of editors is key to ensuring its future. A strategic posture towards generative AI therefore requires established rules for machine uses of Wikipedia and other Wikimedia platforms.

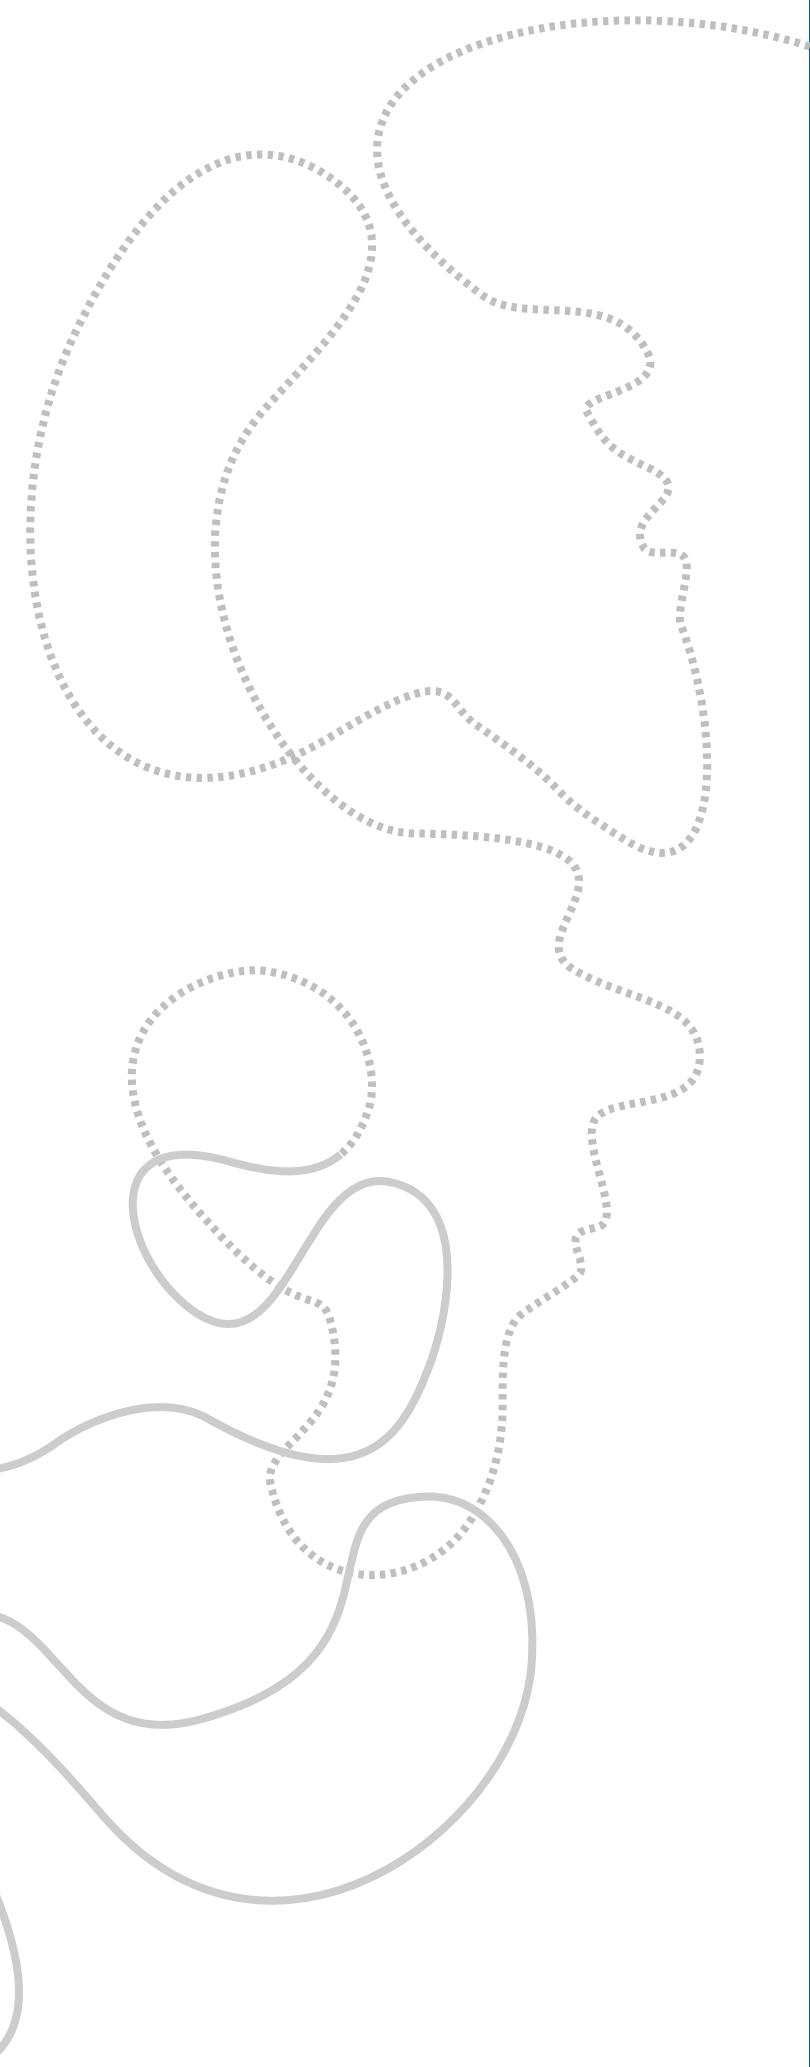
Wikipedia, democracy and information integrity

Lucy Crompton-Reid, Executive Director of Wikimedia UK framed the issue based on her organisation's vision of a more informed, democratic and equitable society through open knowledge. In her view, the challenge goes beyond the well-being of an online encyclopaedia— it's that access to reliable, representative and free knowledge is one of foundations of democratic society. The threat of increased bot traffic and decreased human visitors is less fundamental than the threat caused by a transformation of how knowledge and information are produced, consumed, and trusted. This has implications for civic life, equity and democracy, and applies well beyond the boundaries of Wikimedia projects, across all knowledge ecosystems.

The strategic challenge then becomes one of upholding the integrity of free knowledge in an information ecosystem that is increasingly shaped by generative AI technologies. The much-critiqued opacity of most AI systems is the antithesis of Wikimedia's commitment to transparency. Meanwhile, Wikipedia's collective knowledge production is the opposite of generative AI solutions that are starting to be seen as a single, authoritative source. Wikimedia's strategy needs to support contributors as they navigate this new space of knowledge production, while also modelling a values-driven and inclusive approach to AI and promoting it among Wikimedia's audiences and partners.

In other domains— from journalism to public-interest research— organisations are already describing this situation as a “collapse of the commons.” Wikimedia has the opportunity, and

arguably the responsibility, to become the clearest and most principled voice articulating what is at stake for public knowledge in the age of AI: not only access to information, but the very possibility of a shared, publicly governed understanding of truth.



Challenges



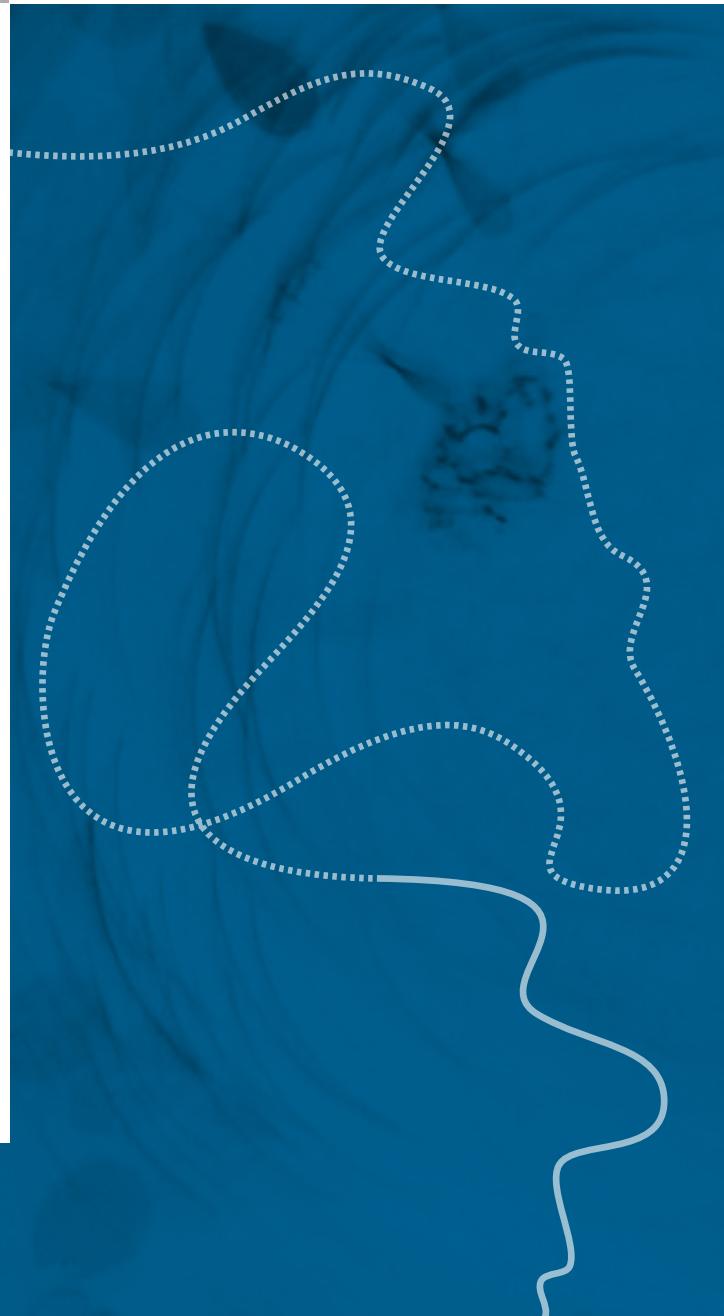
Roundtable participants identified various challenges related to AI development and deployment, but also to the broader changes in online information and knowledge ecosystems. We were particularly interested in building a more expansive understanding of the issues, going beyond a focus on technical strain on Wikimedia's infrastructure or the relatively well-identified challenges related to automation of knowledge production on Wikimedia.

1

Dealing with the paradoxes of openness

Use of Wikimedia content by the largest AI companies is a prime example of how knowledge equity is hindered by disparities in power. Traditional forms of openness — described as “radical open access” — can lead to a sense that free knowledge is being exploited or even stolen. Therefore, Wikimedia needs to acknowledge these limitations to open licensing and consider evolving the concept and the mechanisms of ensuring openness. The Wikimedia Enterprise project is a first step in this direction.

Licensing frameworks should be re-examined if AI use is compliant with open licensing terms, but still harmful





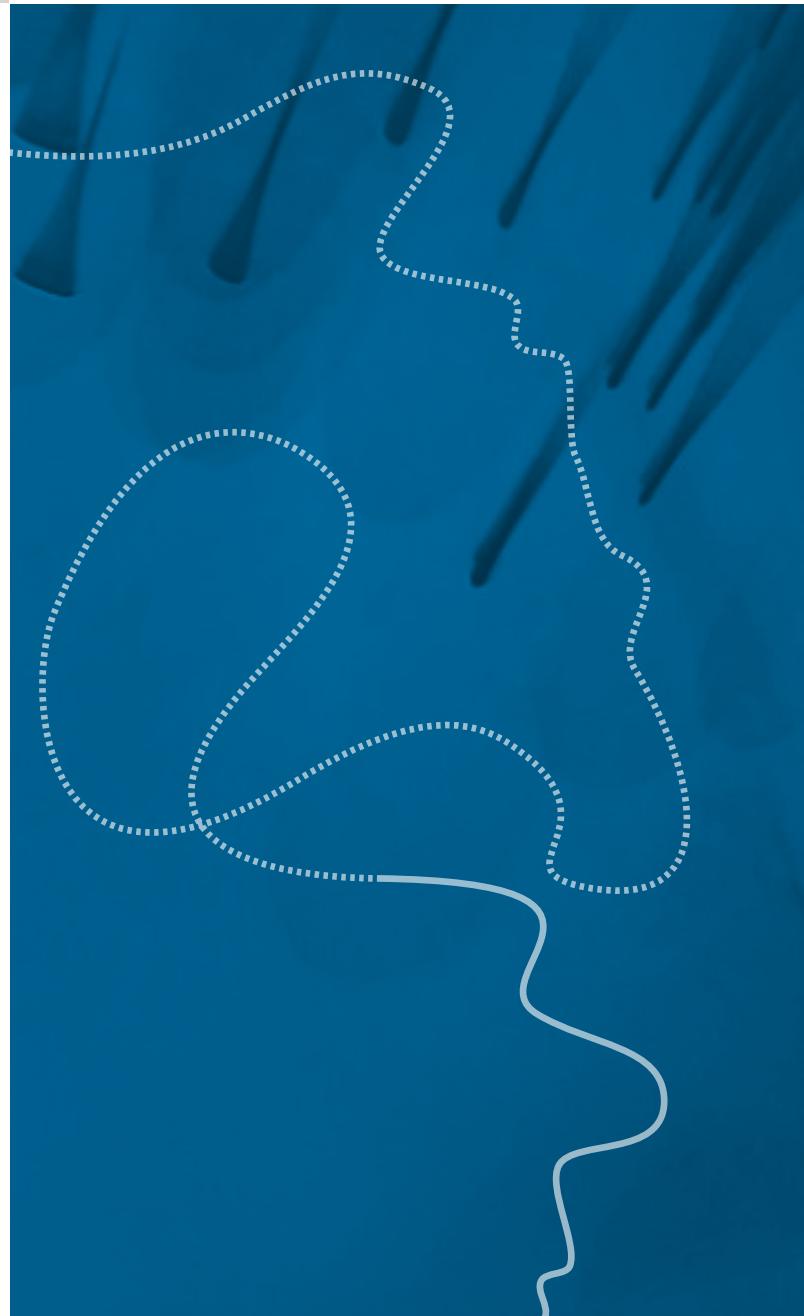
to the knowledge commons, or the community that stewards these commons. For example, the Wikimedia Enterprise project is already establishing new boundaries and leverage to enforce conditions on commercial users, for example through API and robot policies. The overall challenge is to develop a model that ensures sustainability without abandoning, or limiting, openness. Wikimedia also needs to address the sustainability of the broader ecosystem of knowledge organisations, media and other information providers. An ecosystem approach should use a combination of solutions such as governance standards, technical solutions, or regulation, including taxation schemes.

2

Wikipedia as a data commons

Wikipedia is a “mountain of gold” for AI training and development, because its data is verified by humans. There are challenges to the sustainability of producing human knowledge which is then transformed into data for AI. But this is also an opportunity for Wikimedia to develop and promote new approaches to data creation and governance in AI development, based on collective intelligence. Work done by Wikimedia Deutschland on the Wikidata Embedding project is a first step in this direction.

There is rising demand for data created by humans, largely due to the importance of model post-training processes, where such data can play a key role. Existing means of creating this data — often by startups touting



“human data” — are very different from collective intelligence efforts like Wikimedia platforms. They are individual, outsourced, top-down, and aimed at minimising costs, rather than collaborative knowledge development. There is space to offer an alternative data pipeline, based on a vision of collective intelligence.

While synthetic data might seem like the antithesis of human data, it is an increasingly important resource for AI

development that helps ensure quality or protect privacy. Just as with human knowledge work, Wikimedia could offer a model synthetic data pipeline that is collectively developed and governed.

Finally, public debate often paints Wikipedia as a complete resource for AI training, and AI models as trained on “all of world’s data.” In fact, there are large knowledge gaps both in Wikipedia and in AI training datasets. Also, there are big differences between the different language Wikipedias and how they are impacted by AI. While the public debate focuses on the largest Wikipedias, smaller ones can play a key role in developing truly diverse and multilingual AI tools. Knowledge work needs to include ways of mapping and representing gaps, as well as degrees of certainty. Wikipedia is not a monolithic body of knowledge. The differences between Wikipedias provide an opportunity to create more nuanced and dynamic representations of knowledge.

3

Public and private knowledge

Generative AI development puts in sharp contrast the divergence between proprietary and public knowledge. The former — data held by any commercial platform or AI company — is growing rapidly and sees much innovation but remains inaccessible to the public. Private organisations are actively collecting and managing bodies of “private truth” that contribute to concentrations of power, and a sense that public knowledge is being exploited for private gain.

While public knowledge has value, its modes of production and consumption are outdated, and lack of access to private data limits the possibilities of growing the body of public knowledge.



Patterns of knowledge consumption, and of reaching consensus as to the validity of knowledge, are not keeping up with rapid technological change, while rise of disinformation challenges integrity of the information ecosystem. Generative AI models are opaque sources of knowledge that ultimately are shaped not through public debate and with public oversight, but by choices sometimes made by individuals in positions of power.

As a result, there is a risk of the public knowledge ecosystem collapsing, and fragmenting into bodies of private knowledge, and fractured views of the world and of what truth is. Amongst these challenges, Wikipedia remains the prime — and proven — example of how knowledge can be collectively developed, verified and stewarded. The erosion of trust in public knowledge and institutions that create them is a major challenge for Wikimedia. In this environment, Wikipedia can focus on its value as a source of “ground truth”: a body of neutral, verifiable knowledge built through consensus.

To remain relevant, Wikimedia also needs to be more innovative to adapt to shifting patterns of knowledge consumption, and to the risk of disintermediation by AI systems. There is a need to consider new interfaces and new modes of interacting with Wikipedia. Ideas include video formats or machine learning solutions such as chatbots, MCP servers, or editing interfaces.

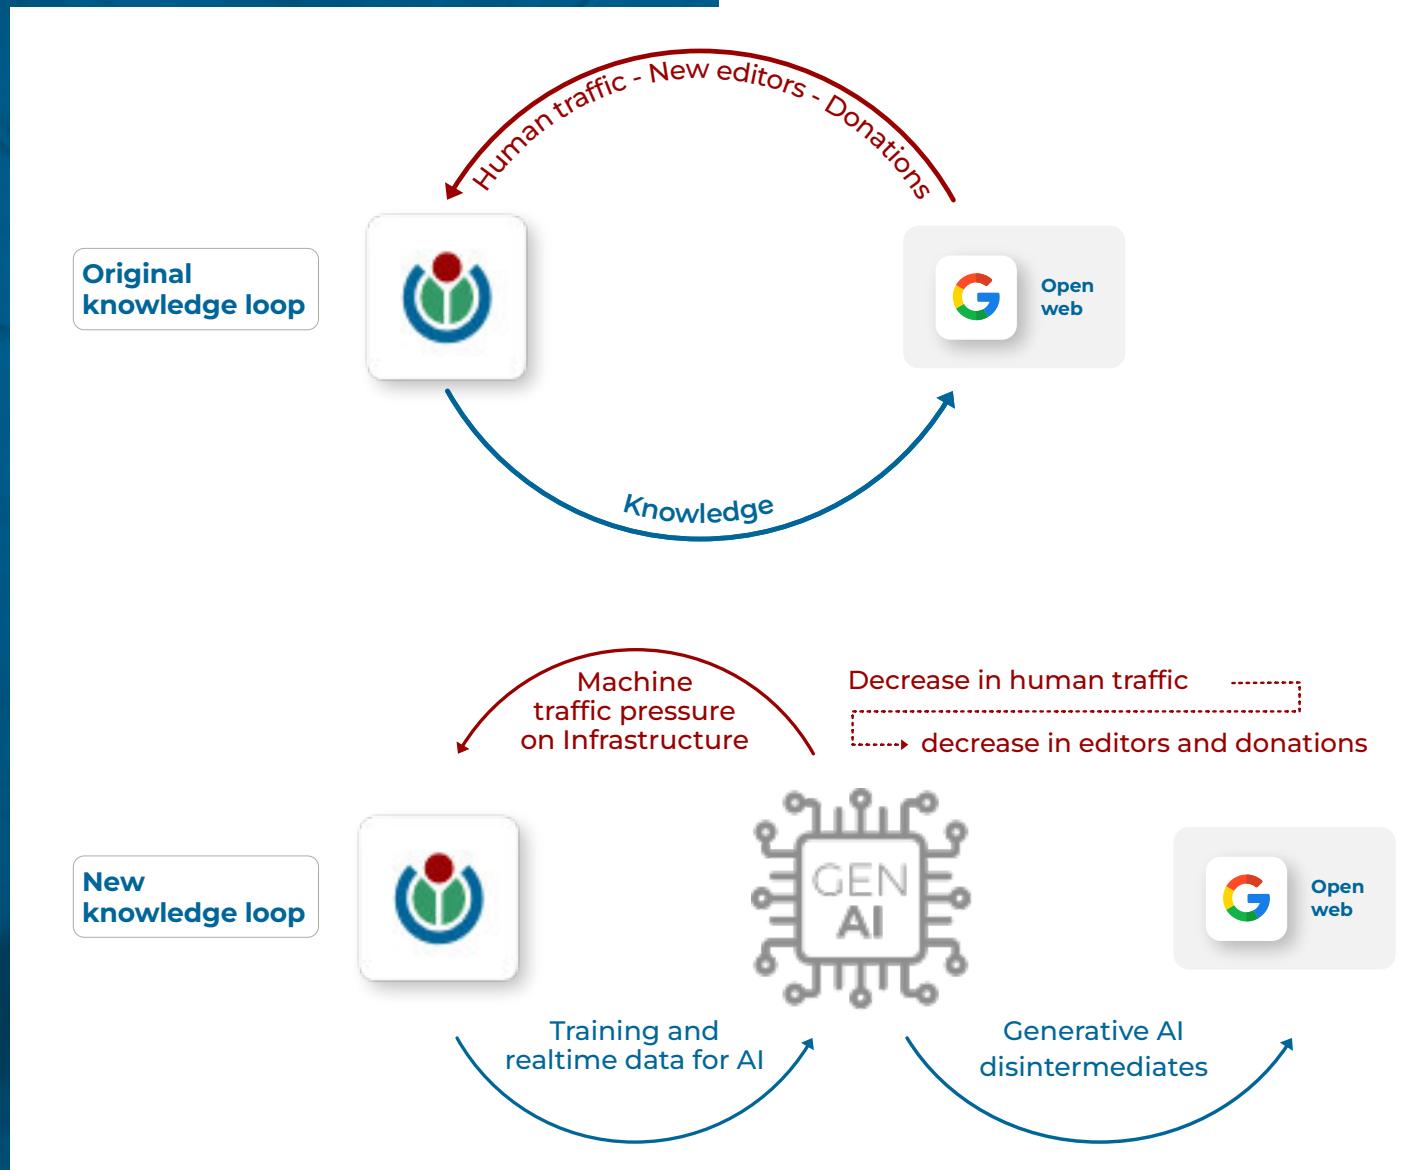
4

The new knowledge loop

Wikipedia is facing a challenge of disintermediation by AI services, with users increasingly relying on AI tools to obtain information that they would previously seek on Wikipedia. As a result, the original feedback loop at the heart of Wikipedia's sustainability as a project of collective intelligence is breaking with decreasing overall visits, fewer users are inclined to become editors. The fundraising model, reliant on views of the fundraising banners, is also at risk. The recently reported 8% yearly drop in visits by human users to Wikipedia is a sign of this challenge. In recent years, Wikipedia has largely relied on Google search to drive traffic



to its website. This symbiotic relationship is now fracturing. Furthermore, as the knowledge ecosystem reacts to the challenge of generative AI, many organisations are limiting or closing access to their resources. This can have a negative effect on Wikipedia, as editors rely on access to these primary sources.



To address this challenge, Wikimedia needs to understand the new knowledge loop that is emerging, with AI services as key actors that intermediate access to knowledge. Understanding the shape of the new, emerging online ecosystem is a prerequisite to action. Dealing with this challenge also requires an ecosystem perspective, including partnerships with other knowledge organisations facing the same challenge: media and journalists, libraries, and other collective intelligence efforts.

Wikimedia needs to decide where it wants Wikipedia to sit in this emerging architecture: as a passive data source, or as an active shaper of standards, interfaces, and governance norms for how AI systems access and represent knowledge. It also needs to understand who its partners are, which intermediaries can support Wikimedia's mission, and under what conditions.





Tensions



As the Wikimedia Movement seeks to address these challenges, it will face various tensions that it needs to navigate.



1

vs

Wikimedia's mission of knowledge equity

Wikimedia is committed to knowledge equity and providing free content.

Strain on Wikimedia's infrastructure

Infrastructural costs put a strain on efforts to deliver free knowledge as a service. The principle of openness conflicts with power imbalances, as shared knowledge is exploited.

2

vs

Reuse of Wikimedia content

AI development leads to the reuse of Wikimedia content at an unprecedented scale — seemingly in line with its mission.

Incentives to contribute to Wikimedia platforms

Pervasive reuse, at mass scale, can threaten motivations and incentives for contributing to Wikimedia projects. Disintermediation by AI can lead to fewer users reaching Wikipedia and going on to become editors.

3

Wikimedia engaging with AI development

Wikimedia needs to actively engage with AI development, if it wants to influence the emerging ecosystem and its norms.

vs

Wikimedia not giving in to AI hype

Many in the Wikimedia Movement are anxious about giving in to AI hype. Dealing with this tension means finding the right mix of human activity and automated functions.

4

Wikipedia as free knowledge

For AI developers, Wikipedia is valuable as a source of high-quality data for AI training. Over time, its value could decrease if knowledge can be produced cheaply by LLMs.

vs

Wikimedia as a collaborative community

Wikimedia platforms' continuing advantage lies in their collective intelligence: a process of human-centric deliberation and content production.

5

Wikipedia as a destination

For the last 25 years, Wikipedia has been the source of first resorts for users seeking information on the Web. Its sustainability depends on being a destination site for human users.

vs

Wikipedia as ground truth

Wikipedia might become a “source of last resort,” a reference providing ground truth to AI systems. A new mode of sustainability would be required, which accounts for Wikipedia's value being distributed across the new information ecosystem.

6

Centralised action by Wikimedia Foundation

As Wikimedia seeks solutions to AI-related challenges, there is a tendency to centralise since many of the potential solutions are technical or infrastructural and require expert knowledge.

vs

Collective action by Wikimedia Movement

Decentralisation remains a key advantage of the Wikimedia Movement, which should seek multiple, local approaches to these challenges. Uncovering how various Wikimedia entities are complementary is key to dealing with this tension.



The time for
change is
now



It was noted in the roundtable that the discussion — and the Wikimedia Movement in general — lacked sense of urgency. It is possible that in five years, no one will visit Wikipedia, or that in 10 years, the platform might no longer exist. This sense of urgency is clear in other organisations working in the public knowledge ecosystem, for example in newsrooms.

The Movement's lack of urgency might be because Wikipedia itself was a force of disruptive innovation that changed the existing knowledge ecosystem — much like the LLMs today. Unlike more established organisations that had to navigate the transition from analogue to digital, Wikimedia projects have never experienced massive changes in ways of working with knowledge. There is a lack of experience of change and, possibly, a sense of invulnerability. There is an urgent need to move beyond describing the challenges of AI and the crisis of sense-making, towards a shared strategy and coordinated action. Without this shift, Wikimedia risks becoming a highly used but politically weak infrastructure: indispensable to AI systems, but increasingly invisible, underfunded, and unable to defend the public interest.

Wikimedia is big enough to help shape the new knowledge ecosystem — but it needs to act now. It needs to be the voice of the open web and the knowledge commons, presenting a different vision of AI based on core principles and collective intelligence. To achieve this, Wikimedia needs to align closely both with other stewards of the knowledge commons and developers of open source and public AI, such as the Swiss Apertus. Together, they should propose an approach covering the full life cycle of AI development and deployment, ensuring that:

- Wikimedia content is used in a sustainable way
- The collaborative vision of Wikimedia is present in various stages of AI development
- Mechanisms ensuring sustainability of the knowledge commons are in place
- There are normative alternatives to the dominant generative AI models that can augment the collective labour of building Wikimedia

Next Steps



Building on the outcomes of the roundtable, Wikimedia CH will publish a white paper in 2026 outlining its vision for Wikimedia's role in AI development and in the new knowledge loop. It will offer the Wikimedia Movement and its partners a proposal for a "Wikimedia in the era of AI" mission and strategy, including:

- a shared policy position on AI governance
- a funding proposal for sustaining the mission
- and a vision for developing technical elements and tools for a Wiki AI stack

The challenges that Wikimedia faces are significant. Yet the Wikimedia Movement, acting together in the spirit of collective intelligence, can successfully face this challenge. Wikipedia's 25th Anniversary is a great opportunity to launch a new mission.

The challenges that Wikimedia faces are significant. Yet the Wikimedia Movement, acting together in the spirit of collective intelligence, can successfully face this challenge. Wikipedia's 25th Anniversary is a great opportunity to launch a new mission.

✚ Authors

Alek Tarkowski
Open Future Foundation

Ilario Valdelli
Wikimedia CH

✚ Organizers

Faith Chilombo Avolio
IMD Business School

Delia Fischer
IMD Business School

Misiek Piskorski
IMD Business School

Andrea Scheller
Wikimedia CH

✚ Participants

Ricardo Chavarriaga
ZHAW

Anna Colom
Data Tank

Lucy Crompton-Reid
Wikimedia UK

Jennifer Ding
The Alan Turing Institute

Jenny Ebermann
Wikimedia CH

Martin Jaggi
EPFL / Swiss AI Initiative

Pierre-Carl Langlais
Pleias

Joel Niklaus
Hugging Face

Kasia Odrozek
Independent Advisor / UNESCO

Nicholas Perry
Wikimedia Enterprise

Andrea Rizzoli
IDSIA / Dalle Molle Institute for Artificial Intelligence

Giovanni Pennisi
Wikimedia Tech Community

Elsa Reichling
Ringier Group

Ferdinando Traversa
Wikimedia Italia

Tendik Tynystanov
Thomson Reuters Foundation

Jens Ohlig
Wikimedia Deutschland

Tomoko Yokoi
IMD Business School

Jonathan Zimmermann
Google / OKA.Wiki



This report is published under the terms of the
[Creative Commons Attribution License](#)

