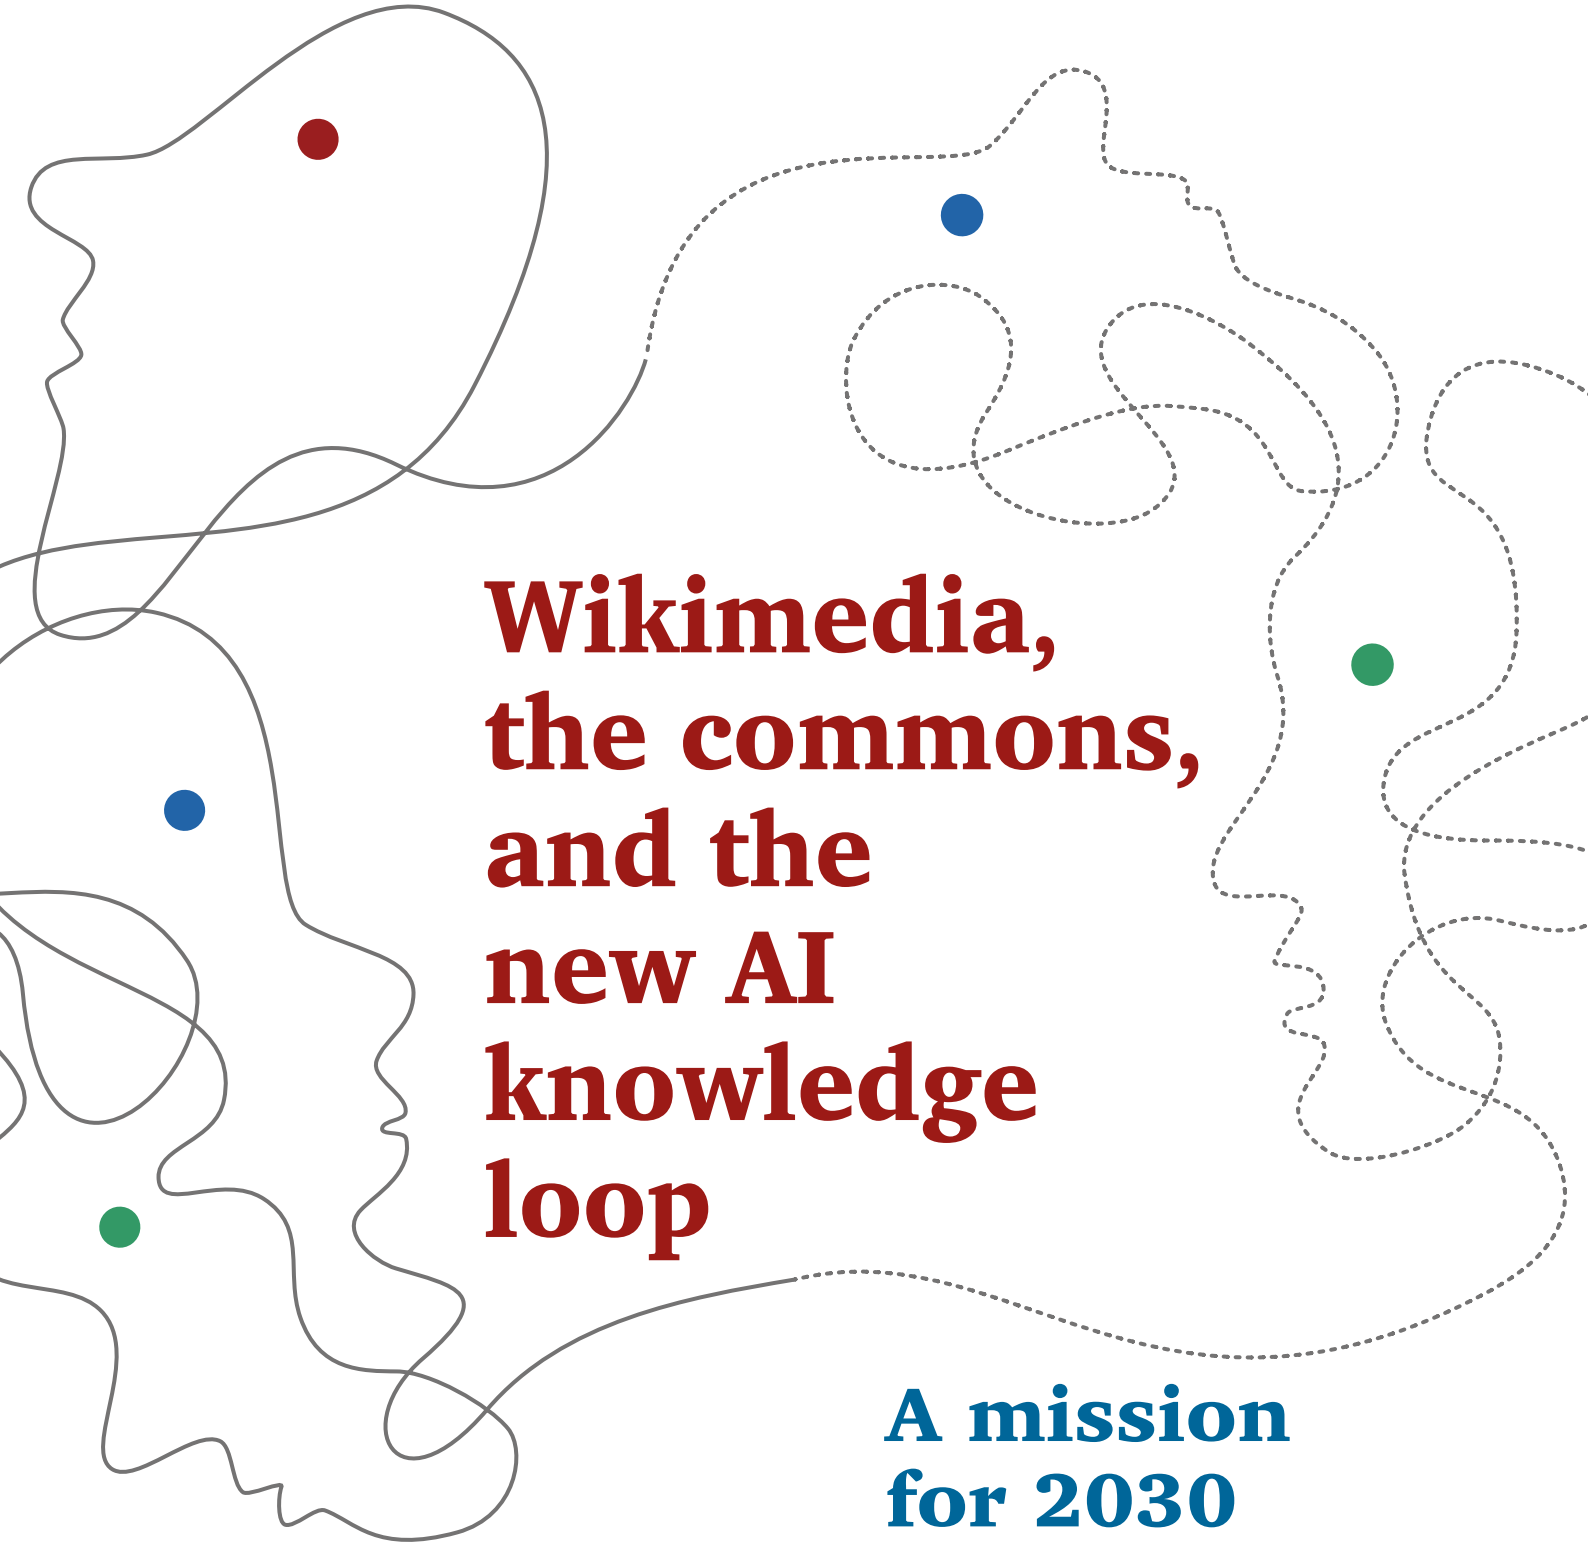

Wikimedia, the commons, and the new AI knowledge loop

A mission for 2030

WIKIMEDIA CH

OPEN _FUTURE

White paper

Version 1.0 - July 2026

by

Alek Tarkowski - Open Future

Ilario Valdelli - Wikimedia CH

Table of contents

Version 1.0 - July 2026

About this paper

This paper was developed by Alek Tarkowski, Strategy Director at Open Future, in collaboration with Ilario Valdelli, Innovation Programme Lead at Wikimedia CH. It was commissioned by Wikimedia CH and was initiated at a roundtable organized with IMD Business School in Lausanne in November 2025.

The position presented in this paper does not represent those taken by any single organization. It reflects a direction that emerged from conversations with Wikimedia community members, movement leaders, and external experts who see both the scale of the challenge and the potential of a coordinated response. The paper aims to give these emerging patterns a shared language, while acknowledging that different perspectives and open questions remain. It is intended as a contribution to a broader movement conversation – a starting point, not a conclusion.

It is the result of a sustained process of research, discussion, and iterative refinement that began in November 2025 with **the roundtable "Collective intelligence vs Artificial Intelligence"** . It is a deliberate and considered effort to understand the ongoing structural transformation of the knowledge ecosystem, and to articulate a direction for the Wikimedia Movement's response.

This is the first in a series of three documents. A governance brief, examining the specific norms and mechanisms needed to shape the new knowledge loop, and a technical direction paper, describing the AI

infrastructure and tools Wikimedia should be building or contributing to, will follow.

The goal is to build a coalition of movement organizations, community leaders, and aligned partners ready to pursue a shared Wikimedia and AI mission together. This paper is an invitation to that process.

Acknowledgements

We would like to thank John Andersson, Lane Becker, Ana Colom, Lucy Crompton-Reid, Jenny Ebermann, Jonathan Fraine, Remy Gerbet, André Golliez, Andrew Gray, Klaus Haensgen, Jack Hardinges, Franziska Heine, Christophe Henner, Tony Lai, Lodewijk Gelauff, Paul Keller, SJ Klein, Dominik Landwehr, Andrew Lih, Kasia Odrozek, Bhavesh Patel, João Alexandre Peschanski, Andrea Rizzoli, Andrea Scheller, Shani Evenstein Sigalov, Aditya Singh, Anastasia Stasenko, Alex Stinson, Joshua Tan, and other persons we have been talking to about Wikimedia, knowledge commons and the new knowledge loop.

Introduction

Dominik Landwehr

The rapid development of artificial intelligence in recent years will fundamentally change the way we produce, use, and share knowledge. And this also affects Wikipedia. AI is the greatest challenge the Free Knowledge movement has faced since its founding 25 years ago.

No one knows where this journey is headed, and neither do we. We are convinced that the Wikipedia movement must engage intensively with this issue in the coming years to determine what conclusions we should draw for Wikipedia and what changes need to be made. Time is of the essence: this discussion must begin now. It should be conducted openly and transparently and involve all affected parties. We now have the opportunity to set the course and ensure that the model of free knowledge will endure into the future.

This is the background to this white paper, which stems from an initial expert discussion that took place in 2025. It is a starting point and is intended to open the discussion; further events, both large and small, will follow.

Dominik Landwehr
President of Wikimedia Switzerland

Summary

The various conversations conducted for this project make the case for a mission-oriented response to the systemic transformation brought about by generative AI technologies. Without a coordinated and swift response, Wikimedia risks a scenario in which it remains widely used but is structurally invisible and lacks agency in a new knowledge loop: the emerging AI ecosystem.

This paper begins by tracing the history of Wikimedia's embedding in the broader knowledge ecosystems through three successive knowledge loops – the first tied to search engines, the second to social media, and the current third connected with emergent AI technologies – showing how each has shaped the flow of value to and from the commons.[1]

The emergence of the third loop, driven by large AI models, breaks the reciprocal dynamics[2] that sustained Wikimedia for over two decades, and raises questions about the sustainability of the knowledge commons and the agency of collective intelligence. **Today, AI systems consume the commons at massive scale without reciprocating the**

[1] Vincent, Nicholas, and Brent Hecht. "**A Deeper Investigation of the Importance of Wikipedia Links to the Success of Search Engines.**" arXiv.Org, 21 Apr. 2020, https://arxiv.org/abs/2004.10265.

[2] Khosravi, Mehrzad, and Hema Yoganarasimhan. "**Impact of AI Search Summaries on Website Traffic: Evidence from Google AI Overviews and Wikipedia.**" arXiv.Org, 5 Feb. 2026, https://arxiv.org/abs/2602.18455.

value that sustains it.[3] The result is not just a sustainability crisis caused by infrastructural strains – but a threat to Wikimedia's capacity to remain an active force in how knowledge is produced and governed.[4] Wikimedia is becoming a "dark commons": an invisible substrate feeding proprietary systems.

From this diagnosis, the paper develops a normative framework to address the challenge. It examines the **Paradox of Open**: the ways in which open sharing, under conditions of extreme power asymmetry, can enable extraction rather than equity.

Based on conversations with various stakeholders and analysis of the strategies of various organizations, the paper maps **three types of value flows (extractive, reciprocal, and generative), and three strategic postures available to stewards of the commons: defensive closure, passive acceptance, and an active stance.** The two most common responses today (closing access to protect the commons or maintaining openness and trust in the overall benefits of reuse) fail to address structural problems.

A recurring view among participants of the Lausanne roundtable was that only an active stance – focused on negotiating reciprocity norms and investing in commons-aligned AI development – can shift value flows from extraction toward generation.

Following this, three principles are introduced to guide Wikimedia's response: protecting collective intelligence, establishing reciprocity, and actively governing the knowledge ecosystem. The paper then proposes a Wikimedia and AI mission to reorient the value flows in the

[3] Keller, Paul. **Beyond AI and Copyright. Funding a Sustainable Information Ecosystem**. Open Future, 30 Jun. 2025, https://openfuture.eu/publication/beyond-ai-and-copyright/, Colom, Anna, Murray, Emma, & Poblet, Marta. Media, democracy and generative AI: a critical juncture. Data Tank, May 2026, https://www.datatank.org/news-insights/media-democracy-and-generative-ai-a-critical-juncture

[4] Southern, Matt G. "**Wikipedia Traffic Down As AI Answers Rise.**" Search Engine Journal, 20 Oct. 2025, https://www.searchenginejournal.com/wikipedia-traffic-down-as-ai-answers-rise/558803.

knowledge ecosystem towards the commons. The mission is the following:

By 2030, Wikimedia works in coalition with public AI developers and commons-aligned partners to build and govern AI interfaces and infrastructures through which the knowledge commons are produced and reach readers. As a result, the Wikimedia Movement is one of the primary user-facing knowledge interfaces and has agency over how knowledge is produced and used in the AI era.

It has two parts:

1. **Pro-commons governance**, focused on reshaping norms around attribution, reciprocity and ecosystem protocols; and
2. **Wiki AI technical stack**, supporting the development of commons-aligned data and application layers, and contributing to public AI model development.

Taken together, these actions will ensure Wikimedia's agency in the new knowledge loop. Supportive measures include a mission network that coordinates action across movement entities, volunteer editors and aligned partners like public AI developers; and mission investment to resource these efforts.

Wikimedia's knowledge loops

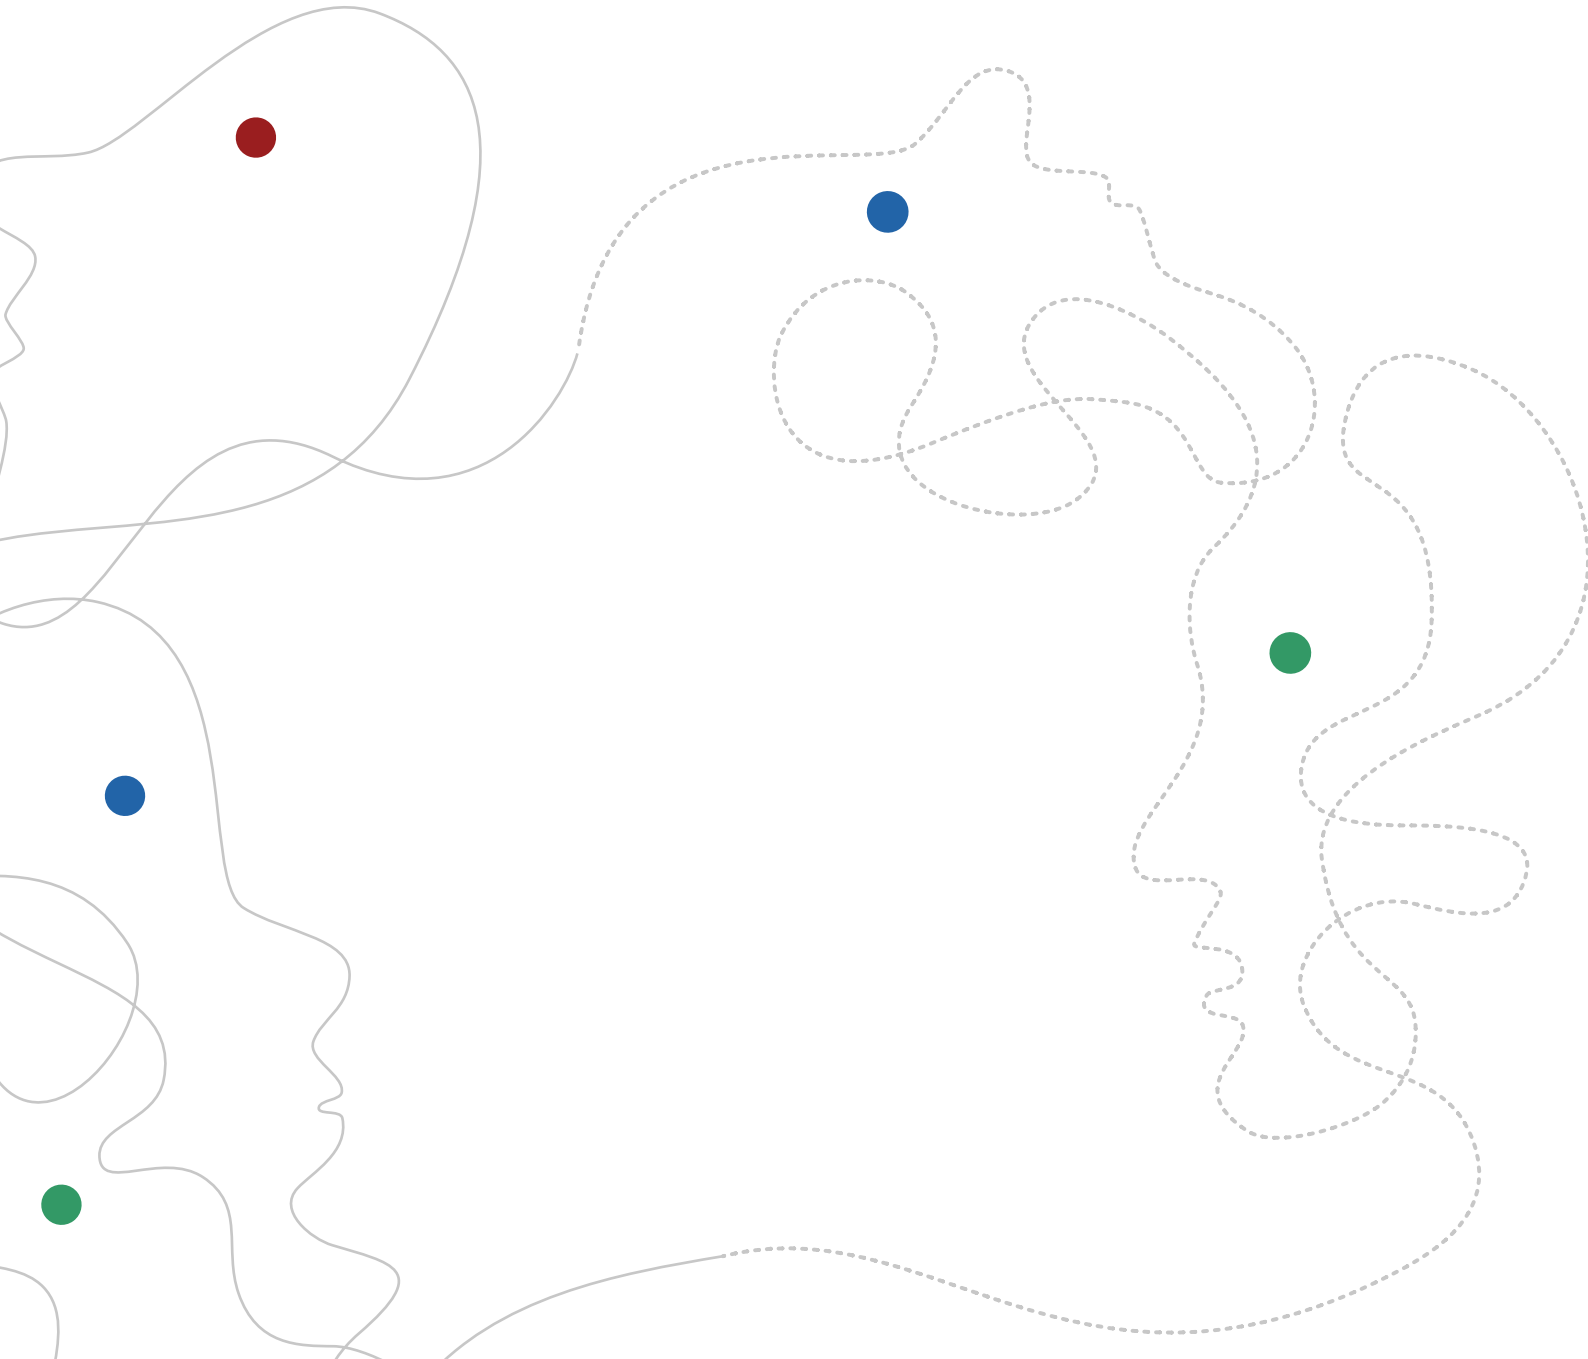

Wikimedia projects are not just platforms. They constitute a foundational layer of the global knowledge infrastructure, on which both human and machine knowledge systems increasingly depend. If this layer degrades, the quality of knowledge across the web – including AI systems that rely on it both as training data and verified reference – degrades with it. At the same time, Wikimedia was not built in a vacuum.

KNOWLEDGE LOOPS AND ECOSYSTEMS

Wikipedia and its sister projects have been made possible through the circulation of users, content and resources within the open web. This created three **knowledge loops, in which Wikimedia is embedded, and which sustain the knowledge ecosystem.** These loops consist of circulations of content and other resources from Wikimedia to various entities accessing and using them; and reciprocal flows of content, users or even funding back to Wikimedia.[5]

If Wikimedia's vision is "a world in which every single human being can freely share in the sum of all knowledge", then its success depends not only on internal factors – the successful creation of a body of such knowledge. It also requires external conditions that support the free sharing in this body of knowledge. To this end, the shape of the knowledge ecosystem, in which it operates, needs to be considered.

Knowledge ecosystems (also called information ecosystems – the terms can be used interchangeably) consist of institutions, infrastructures and actors that collectively produce, preserve, and distribute knowledge. Among these, commons-based peer production plays a key role, with Wikimedia as a prime example.

[5] Piccardi, Tiziano, et al. "**On the Value of Wikipedia as a Gateway to the Web.**" arXiv.Org, 15 Feb. 2021, https://arxiv.org/abs/2102.07385.

WHAT IS A KNOWLEDGE LOOP?

Knowledge loops are the key formations in a knowledge ecosystem, established through flows of resources and value generation between key nodes and actors. In an increasingly centralized web, these loops connect knowledge commons with the dominant platforms that determine the various flows of value, attention and resources on the web. These loops define specific structural relationships and value flows within digital ecosystems.

Positive loops are symbiotic and reciprocal, while negative ones are based on extraction of value, and uneven distribution of risks and harms. In other words, positive loops are equitable. They distribute knowledge, attention, and power in ways that sustain the commons and the communities that produce it.

The challenge posed by AI is therefore an equity challenge to the Wikimedia Movement as a whole. But it also intersects with longstanding inequities within the Wikimedia movement itself, between the Global North and the Global Majority, between dominant and under-resourced languages. Generative AI threatens to entrench them at planetary scale.

THREE LOOPS, THREE ALGORITHMS

There were three such major loops in the 25 years of Wikimedia's history.[6] Each has at its core algorithms that shape these loops and fundamentally change how people interact with information.

The first knowledge loop – which played a foundational role in Wikipedia's development over the last 25 years – was driven by the search algorithm. The second loop could have been tied to social media algorithms, but it did not develop into a fully reciprocal loop comparable

[6] The three loops do not cover all the key relationships established by Wikipedia – other loops can be identified. For example, the relationship with the Internet Archive continues to be of fundamental importance for Wikimedia.

to search. The third loop, emerging now, is tied to the spread and adoption of large AI models.

To be successful, today as in the past, Wikimedia needs to be aware of the ecosystem's structural conditions: the shape of these loops. Based on this understanding, it needs to be an active shaper of the conditions under which it operates. Establishing strong sharing norms through **Share Alike licensing**, or launching **Wikimedia Enterprise**, are two examples of this strategic posture.

The foundational loop: building an encyclopedia on the Open Web

Beneath the three downstream loops that structure this paper lies a fourth, foundational one. Before Wikimedia became embedded in a series of knowledge loops, it was built on top of the open web, and on the body of knowledge that the web made accessible. Wikipedia depends on journalism, scientific research, and heritage collections for the sources it cites. The early web made those sources easy to reach and to verify through hyperlinks – and thus made Wikipedia's verifiability policy practical at scale.

The open web of the early 2000s had commons-like properties: open protocols, low barriers to publication, norms of linking and reuse. This allowed Wikipedia to grow as a node in a wider ecosystem of public knowledge. This ecosystem was from the start geographically uneven and characterized by the digital divide that privileged the Global North. Wikipedia's verifiability policy meant from the start that Global North sources, made readily citable by the early web, were privileged. The foundational loop was generative, but not equitable.

A FOUNDATION NOW UNDER PRESSURE

The three loops that follow are characterized by increasing levels of extraction of value without reciprocity. And while the foundational loop of web-based upstream sources continues to exist, it has also been impacted: AI is degrading the sustainability of journalism, academic publishing, and other knowledge-producing institutions that Wikipedia depends on. This happens the most in areas where they are already underresourced, in Global Majority countries. A Wikipedia that survives the downstream loops but loses its source base will not be a sustainable encyclopedia.

Loop one: coopetition with search engines

For Wikipedia, the primary knowledge loop has, from the beginning, been tied to its position with search engines, and more specifically with Google.[7] Google generates over 90% of global search referrals, making it the dominant gateway to web content.[8]

THE ORIGINAL SYMBIOSIS

From the start, Google's PageRank algorithm elevated Wikipedia and suppressed clones that tried to free-ride on its resources. By 2008, most of Wikipedia's traffic came from Google search referrals – and Wikipe-

[7] Contributors to Wikimedia projects. "**Relationship between Google and Wikipedia.**" Wikipedia, 16 Feb. 2026, https://en.wikipedia.org/wiki/Relationship_between_Google_and_Wikipedia, McMahon, McMahon, Connor, et al.
"**The Substantial Interdependence of Wikipedia and Google: A Case Study on the Relationship Between Peer Production Communities and Information Technologies.**" Proceedings of the International AAAI Conference on Web and Social Media, vol. 11, no. 1, May 2017, pp. 142–51, https://doi.org/10.1609/icwsm.v11i1.14883.

[8] Bennett, James. "**The Search Engine Referral Report: Google 90% Dominance in Numbers.**" WebsearchAPI.Ai, 8 Feb. 2026, https://websearchapi.ai/blog/search-engine-referal-report-insights.

dia's content had a significant positive impact on Google's search quality.[9] Google also donated money to the Foundation, totalling over USD 10 million before the Wikimedia Enterprise partnership was established. These were discretionary gifts, not contractual terms: the value Wikimedia drew from the loop depended on Google's goodwill, not on any agreement it could enforce.

Wikipedia's fundamental knowledge loop was thus established: editors write articles, Google surfaces them, readers visit Wikipedia through search referrals, some of them become editors, and the cycle continues. In addition, and just as importantly, Google connected Wikimedia with people willing to support the network financially: readers visiting via search saw Wikipedia's fundraising banners, the key mechanism driving donations and ensuring Wikimedia's sustainability.

Nathalie Casemajor describes the relationship that emerged as a case of coopetition: simultaneous cooperation and competition.[10] And Wikimedia projects benefitted from this loop unevenly, as search algorithms are based on the "rich get richer principle": the large language editions captured the most traffic, while smaller Wikipedias remained peripheral.

ASYMMETRY SURFACES

This symbiotic relationship faced a challenge in 2012, when Google launched its Knowledge Graph, and began surfacing Wikipedia content directly in search results. The function of sharing knowledge was retained, and the knowledge panels proved a useful interface for many readers, but the value of the loop for Wikipedia decreased.

In 2018, YouTube (another company in Google's Alphabet group) star-

[9] Vincent, Nicholas, and Brent Hecht. op. cit.

[10] Casemajor, Nathalie. **L'entre-prise des communs Trajectoire des relations entre Wikimédia et Google**. Réseaux, 254(6), 113-146. 2025. https://doi.org/10.3917/res.254.0113

ted using Wikipedia content to combat disinformation – proving that Wikipedia constitutes essential infrastructure for the platform. In 2021, Wikimedia Enterprise was launched as a response to some of these developments, and Google was still the first paying customer.

Google also provided a reader-facing interface as Wikipedia itself remained an editor-facing platform. Yet even during the period of mutual benefit, the relationship was structurally asymmetric. Since it emerged, rather than being negotiated, there were no agreed terms that Wikimedia could invoke. It spelled a significant dependency of Wikimedia on a single, dominant commercial actor for its discoverability, and offered little meaningful leverage to negotiate the terms of the relationship.

No attempt has been made to precisely measure the overall benefit of this knowledge loop for Wikimedia. It is commonly acknowledged that it helped monetize traffic, generate donations, and turn readers into editors. The Foundation itself identifies changes to the traffic-editors-donations causal chain as a central risk.[11]

This is the lesson the foundational loop carries into the AI era: a beneficial loop that emerges on its own can also dissolve on its own. The next one will not configure itself in Wikimedia's favour – it has to be built.

THE LOOP UNRAVELS

Across the web, the ratio of AI crawler traffic to search referrals has shifted dramatically. Data from Cloudflare's network shows Google moving from crawling 2 pages for every visitor sent to publishers in 2020, to 6:1 in 2024, and 18:1 by mid-2025. This reflects a systemic shift in how search engines extract value from web content without returning

[11] Miller, Marshall. "**New User Trends on Wikipedia.**" Diff, 17 Oct. 2025, https://diff.wikimedia.org/2025/10/17/new-user-trends-on-wikipedia/

equivalent traffic.[12] For Wikipedia, the impact is measurable and specific. Between 2022 and 2025, the share of search referrals in human traffic fell from 48% to 36%. Google referrals dropped by 22%, and this decline alone explains virtually all the reduction in human traffic to Wikipedia over the same period.[13] A comprehensive analysis of Wikipedia traffic trends found that organic search decline accounts for 96% of total traffic loss – equivalent to over one billion monthly visits.[14]

Today, Google search continues to benefit from Wikipedia's knowledge, and readers benefit from Wikipedia's knowledge as mediated by Google, but Wikipedia benefits less and less from Google's referrals.[15]

[12] Belson, David, and Sam Rhea. "**The Crawl before the Fall... of Referrals: Understanding AI's Impact on Content Providers**." The Cloudflare Blog, 1 Jul. 2025, . https://blog.cloudflare.com/ai-search-crawl-referratio-on-radar/ .

[13] Leseberg, Nina, and Daniel Sigge. "**Wikipedia:WikiCon 2025/Programm/Reichweite, Relevanz, Resonanz: Wikipedia-Nutzung Im Wandel – Wikipedia.**" WikiCon 2025, 5 Oct. 2025, https://de.wikipedia.org/wiki/Wikipedia:WikiCon_2025/Programm/Reichweite,_Relevanz,_Resonanz:_Wikipedia-Nutzung_im_Wandel.

[14] Kemp, Simon. "**Digital 2025: Exploring Trends in Wikipedia Traffic – DataReportal – Global Digital Insights.**" DataReportal – Global Digital Insights, 23 Apr. 2025, https://datareportal.com/reports/digital-2025-exploring-trends-in-wikipedia-traffic.

[15] Khosravi, Mehrzad, and Hema Yoganarasimhan. op. cit.

Loop two: a missing connection with social media

In 2006, Facebook launched Newsfeed and around 2009-2011 shifted from a chronological feed to an algorithmic one. While Google's search algorithm aimed to direct users outward, to where they wanted to go and to information that they were seeking, social media algorithms were designed differently, with the intention of keeping users' attention on the platform. The business model was not to help people find knowledge, but to maximize engagement.

ENGAGEMENT VERSUS KNOWLEDGE

Social media platforms are not designed to point attention toward quality sources, but to capture it and hold it inside a continuous loop of content.Algorithmic amplification of emotional, sensational content has proven to be a major driver of distortions in public knowledge and discourse, and levels of trust in information from social media remains significantly lower than trust in other sources, and has declined over time.[16] Holding attention and degrading quality are not separate effects of this design; they are the same mechanism. At the same time, algorithmic feeds are personalized, while Wikipedia is built on the principle of a shared, negotiated and neutral point of view. A system built to retain users will not, as a rule, return them to the source. As a result, social media platforms have not enabled the emergence of a stable reciprocal knowledge loop with Wikimedia.

BURSTS, WITHOUT RECIPROCITY

[16] Eddy, Kirsten, and Elisa Shearer. "**How Trust in Info from News Outlets and Social Media Has Changed over Time.**" Pew Research Center, 29 Oct. 2025, https://www.pewresearch.org/short-reads/2025/10/29/how-americans-trust-in-information-from-news-organizations-and-social-media-sites-has-changed-over-time/.

Social media has nonetheless played some role in disseminating Wikipedia content, even if this has not translated into sustained reciprocal engagement. This is especially true of platforms focused on sourcing information, such as Reddit or Stack Overflow.[17] AI is not the first service to break attribution — social media platforms began this pattern by displaying Wikipedia content without a structural approach to sending users back to the source.

And when social media drives traffic to Wikipedia, they do it unevenly in sudden bursts of attention around specific articles. These spikes in traffic are often poorly absorbed. Community response to viral posts tends to be defensive: locking the article, reverting changes, rather than engaging or onboarding the people arriving.

THE TEXT-TO-VIDEO GAP

Finally, as social media shifted decisively from text to audiovisual content, Wikipedia did not follow. This principled stance came at a cost. As social media shifted decisively from text to audiovisual content, Wikipedia remained a textual medium – consistent with its mission, but increasingly out of step with how younger generations consume information. The cost was not only generational. Beyond younger users, a broad mass of readers of every age came to prefer information that is immediate and low-effort over information that is verified but demanding. This is not a failure of those readers: the demand for accessible, frictionless knowledge is legitimate and universal. The failure is that, so far, only systems willing to trade quality for ease have met it. Wikipedia protected quality at the cost of accessibility, and lost not just a generation but anyone who would choose an easy answer over a verified one.

Data on referral traffic reflects this structural gap: social media has not

[17] Vincent, Nicholas, et al. "**Examining Wikipedia With a Broader Lens.**" Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems [New York, NY, USA], 2018, pp. 1–13, https://doi.org/10.1145/3173574.3174140.

become a significant or stable driver of traffic or contributor growth for Wikipedia, and recent declines in referrals further compound the broader shifts caused by AI-mediated access to information.[18] The shift away from text-based knowledge consumption is most pronounced among younger users in the Global Majority. Wikipedia's mobile experience has been acknowledged as a barrier to growth in these regions.[19]

It's worth noting that, today, some of the strongest calls for innovative approaches to how Wikimedia presents its content involve changes that would bring Wikipedia closer to the logic of social media. This was well demonstrated at the **2026 Wikimedia Futures Lab**, where widely supported proposals concerned video content, personalized knowledge graphs and reader-facing UX, reflecting genuine community demand, but also the tension between Wikimedia's mission and the attention economy it has long resisted.

[18] Newman, Nic. "**Journalism, Media, and Technology Trends and Predictions 2026.**" Reuters Institute for the Study of Journalism, 12 Jan. 2026, https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2026.

[19] Wikimedia Foundation. "**Global market and audience research (2020)**". 2020. https://meta.wikimedia.org/wiki/Wikimedia_Foundation/Communications/Research/Global_market_and_audience_research_(2020)

The knowledge loop is breaking

Search-era loop

A symbiotic loop

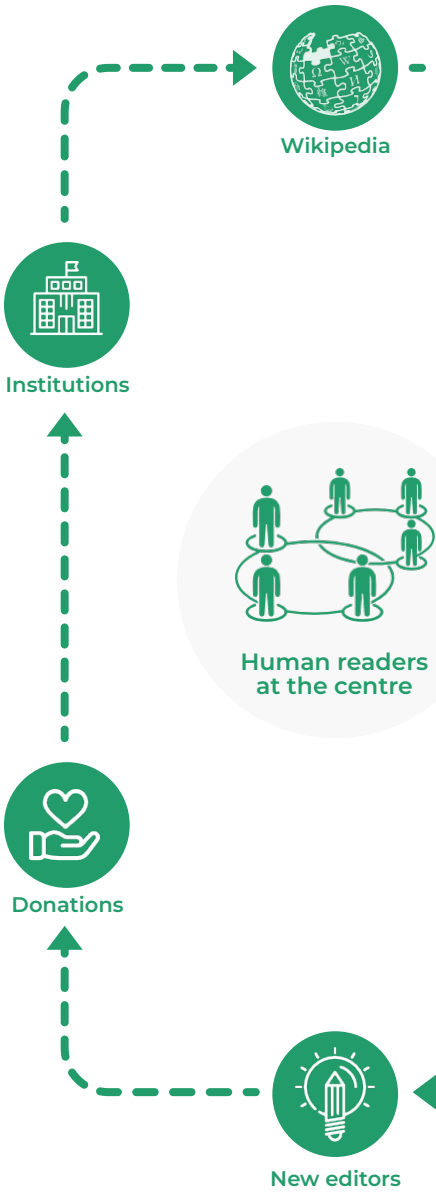

AI-era loop

An extractive loop

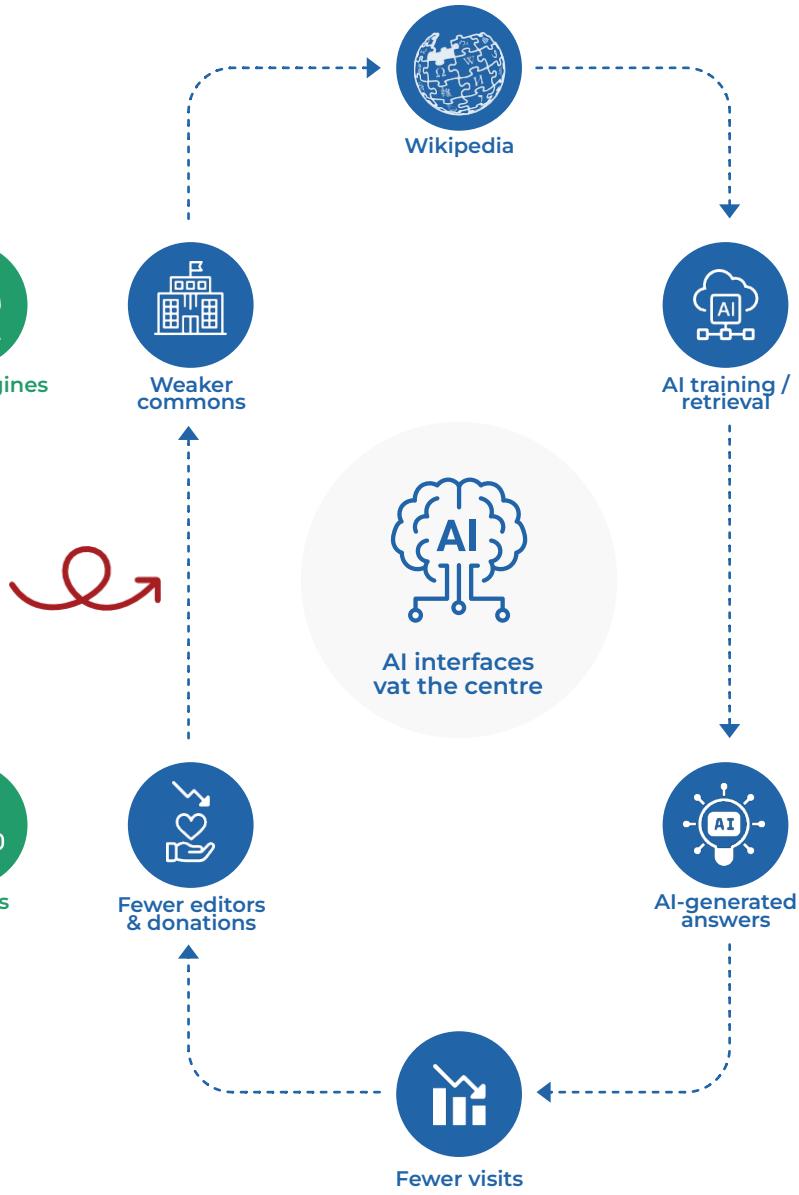

The problem is not only traffic loss.

It is the loss of Wikipedia's agency in the knowledge loop.

Loop three: Artificial Intelligence technologies and machine uses

This new knowledge loop is embedded in an emerging ecosystem in which access to knowledge is increasingly intermediated by AI systems. It is an ecosystem shaped from the very start by dramatic concentrations of power. While these have been typical of **an increasingly platformized web**, the pace of current change is unprecedented.[20]

A NEW LOOP, WITH A BROKEN SOCIAL CONTRACT

This leaves limited space for alternatives to what is emerging as the dominant loop, in which communication is mediated by a small range of AI services, owned by some of the biggest and richest companies in the world. In turn, this introduces new challenges for fulfilling the goal of democratizing access to knowledge.[21] In this knowledge loop, machines are becoming as significant users of Wikimedia content as humans, fundamentally altering the nature of the knowledge loop. It spells radical shifts for a knowledge commons built on the foundation of the first loop.

It is in this loop that challenges to knowledge equity become exacerbated. Since large language models are trained substantially on English Wikipedia, the new loop reproduces Wikimedia's existing structural biases at planetary scale. AI does not extract from a uniform commons but rather from a biased selection, caused by decades of digital inequalities and divides. This selection is then propagated by AI models as global default knowledge.

[20] Recap, Cecilia. "How Big Tech Turns Knowledge into Power." Critical Takes on Corporate Power, 16 Jun. 2025, https://criticaltakes.org/the-corporation/how-big-tech-turns-knowledge-into-power/.

[21] Gimpel, Lea, et al. "Democratizing AI for the Public Good – Open Future." G20 Brasil 2024, 2024, https://openfuture.eu/publication/democratizing-ai-for-the-public-good/.

In the new loop, the implicit social contract of the Open Web is broken: it is no longer true that you can "publish freely, get discovered via search, and monetize through attention".[22] Just as drops in search referral traffic are a clear sign of the first loop waning, so the emergence of the new loop can be observed through shifts in human and automated traffic.

MACHINES BECOME THE USERS

Before shifts in knowledge consumption began affecting human traffic to Wikipedia, the platform was being crawled as a unique source of training data for AI models. At least since 2019, AI labs have been using Wikipedia content as a key resource. Wikipedia provided them with structured data at the scale needed to build modern language models, and served as a reference database when checking the quality of other data.[23]

Between 2016 and 2022, traffic by self-identifying bots (mainly related to search services, but over time also AI crawlers, called "spiders" in Wikipedia nomenclature) grew by around 6% each year. Between 2022 and 2025, it grew by 10% each year and peaked at 5.9BN views in October 2023, when they accounted for 23.8% of all traffic. Spider share has since risen further to 25% in mid-2025 even as absolute volumes eased. By March 2026, spider traffic dropped to 3.5BN views, or 18.4% of overall traffic.[24]

In 2021, Wikipedia also started tracking unidentified automated traffic.

[22] Ojha, Saanya. "**The Internet Just Flipped Its Default.**" The Change Constant, 2 Jul. 2025, https://saanyaojha.substack.com/p/the-internet-just-flipped-its-default.

[23] Hao, Karen. **Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI**. Penguin Group, 2026, Brown, Tom B., et al. "Language Models Are Few-Shot Learners." arXiv.Org, 28 May 2020, https://arxiv.org/abs/2005.14165.

[24] See also: Henner, Christophe. "**Wikipedia at 25: A Wake-Up Call**", Wikimedia Meta-Wiki. 9 Jan. 2026. https://meta.wikimedia.org/wiki/User:Schiste/what-now;
Villa, Luis. "**Wikipedia's traffic drop: more on languages and freshness**". Luis Villa. 22 Apr. 2026. https://lu.is/2026/04/wikipedia-decline-by-topic/

This grew from an average of 0.9BN views monthly (4% of total traffic) to a peak in September 2025, when 4.5BN automated views constituted 20.4% of all traffic. Today, traffic from both identified and unidentified bots has lessened, but still constitutes 35% of all traffic, and continues to grow.

HUMAN TRAFFIC IS IN DECLINE

At the same time, human traffic to Wikipedia started declining. Over the years, there have been shifts in this traffic, and different language Wikipedias have observed different trends. Nevertheless, between January 2025 and January 2026, human traffic to all Wikipedias fell by 8%, as machine traffic continued to rise. According to Marshall Miller from the Wikimedia Foundation's Product Team, "these declines [in human traffic] reflect the impact of generative AI and social media on how people seek information, especially with search engines providing answers directly to searchers, often based on Wikipedia content".[25]

Machine traffic is also placing a significant strain on Wikimedia infrastructure. According to the Wikimedia Foundation, there has been a 50% growth since January 2024 in bandwidth used for downloading Wikimedia, driven largely by automated scraping of Wikimedia Commons.[26] Machine traffic, accounting for 35% of the total, is responsible for 65% of the most resource-consuming traffic, as automation complicates caching of content.[27] This data complements statistics on human and bot traffic by focusing on bandwidth and the infrastructural resources needed to service these two types of traffic.

[25] Miller, Marshall. Ibid.

[26] Mueller, Birgit, et al. "**Quo Vadis, Crawlers? Progress and What's next on Safeguarding Our Infrastructure.**" Diff, 26 Mar. 2026, https://diff.wikimedia.org/2026/03/26/quo-vadis-crawlers-progress-and-whats-next-on-safeguarding-our-infrastructure/.

[27] Mueller, Birgit, et al. "**How Crawlers Impact the Operations of the Wikimedia Projects.**" Diff, 1 Apr. 2025, https://diff.wikimedia.org/2025/04/01/how-crawlers-impact-the-operations-of-the-wikimedia-projects/.

Shifts in Wikipedia traffic statistics reveal the emerging shape of the new loop. It is characterized by:

- Radical growth in machine traffic, driven both by training-time data collection by crawlers, and by inference-time access from deployed AI systems. This includes AI search, retrieval-augmented generation, and other forms of real-time querying, also by AI agents. Machine traffic will continue growing and may eventually match or exceed human traffic.

- Intermediation of human activity by AI, mainly through AI search, which is often "zero click" and does not refer traffic to Wikipedia. These intermediating entities will control the flow of human activity and attention, which is crucial for Wikimedia's wellbeing. This is coupled with other factors that reduce human engagement on Wikimedia. Some are well understood, like reduced interest among the younger generations, due to the lack of innovation in reader-facing interface design.[28] Others are more hypothetical, such as reduced editor activity due to AI-related shifts in knowledge production. The overall outcome may lead to degradation of the commons.[21]

- Various second order effects that further degrade the knowledge ecosystem based on the Open Web. For example, as websites respond to AI with defensive closure mechanisms (including robots.txt files, paywalls, or Cloudflare's bot-blocking measures), Wikimedians's ability to reference the Web is constrained. **AI slop** published on the web degrades the quality of references, and growing dependence on AI tools may reduce overall trust in human knowledge work done by Wikipedians.

Together, these three factors represent a seismic shift for Wikimedia: a knowledge platform built for humans, but increasingly accessed and used by machines.

[28] Jemielniak, Dariusz. "**Wikipedia Faces a Generational Disconnect Crisis.**" IEEE Spectrum, 30 Jan. 2026, https://spectrum.ieee.org/wikipedia-at-25.

CONCENTRATION OF POWER DRIVES THE LOOP

These changes should be considered alongside other shifts taking place on the other side of the loop. While the search engine space was as concentrated as the emerging AI ecosystem, Wikimedia's symbiosis in the first knowledge loop was established before the full consequences of this dependency were felt. The AI loop, by contrast, is being shaped by concentrated actors from the outset. Generative AI has been centralized almost from the start by a small oligopoly with the capital and data infrastructure to shape the new knowledge loop on its own terms. This concentration of power is exacerbated by the dominant AI development paradigm. Scaling laws inherent to this paradigm place the largest actors in a position that others cannot compete with.

Extraction of value from the commons has characterized the web for a long time. The AI loop in many ways extrapolates trends visible in the two earlier loops, and all three loops are characterized by asymmetry of power that impacts Wikimedia. The AI loop just takes existing tensions to a level where they become completely unsustainable.

KEY FIGURES

- **Bot/spider traffic grew 10%** per year between 2022 and 2025, peaking at 23.8% of all traffic in October 2023; by March 2026 it had eased to 18.4% (3.5BN views)
- **Unidentified automated traffic peaked at 20.4%** of all traffic in September 2025; combined bot + unidentified automated traffic now sits at 35% of total traffic and continues to grow
- **Machine traffic** (35% of total) **causes 65% of the most resource-intensive load**
- **Bandwidth** for downloading Wikimedia content **grew 50%** since January 2024, driven largely by automated scraping of Wikimedia Commons
- **Human traffic to all Wikipedias fell on average 8%** in 2025 compared to the year before. In the first five months of 2026, the decrease was 11,7% on average. May 2026 recorded 12.4 billion human page views – the lowest

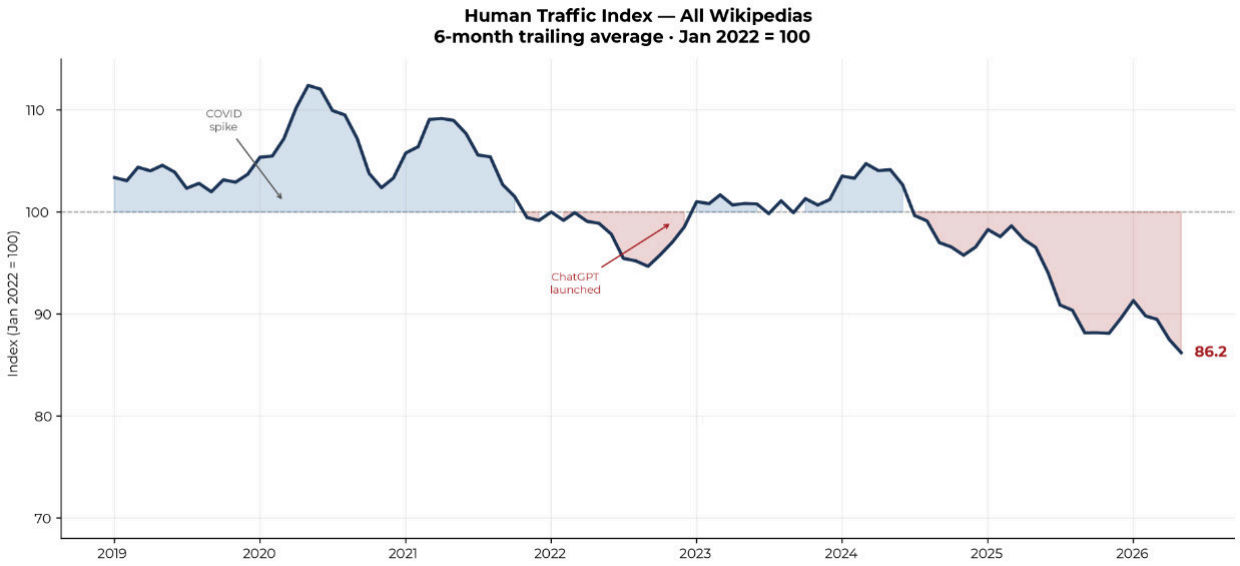

Facing the Paradox of Open

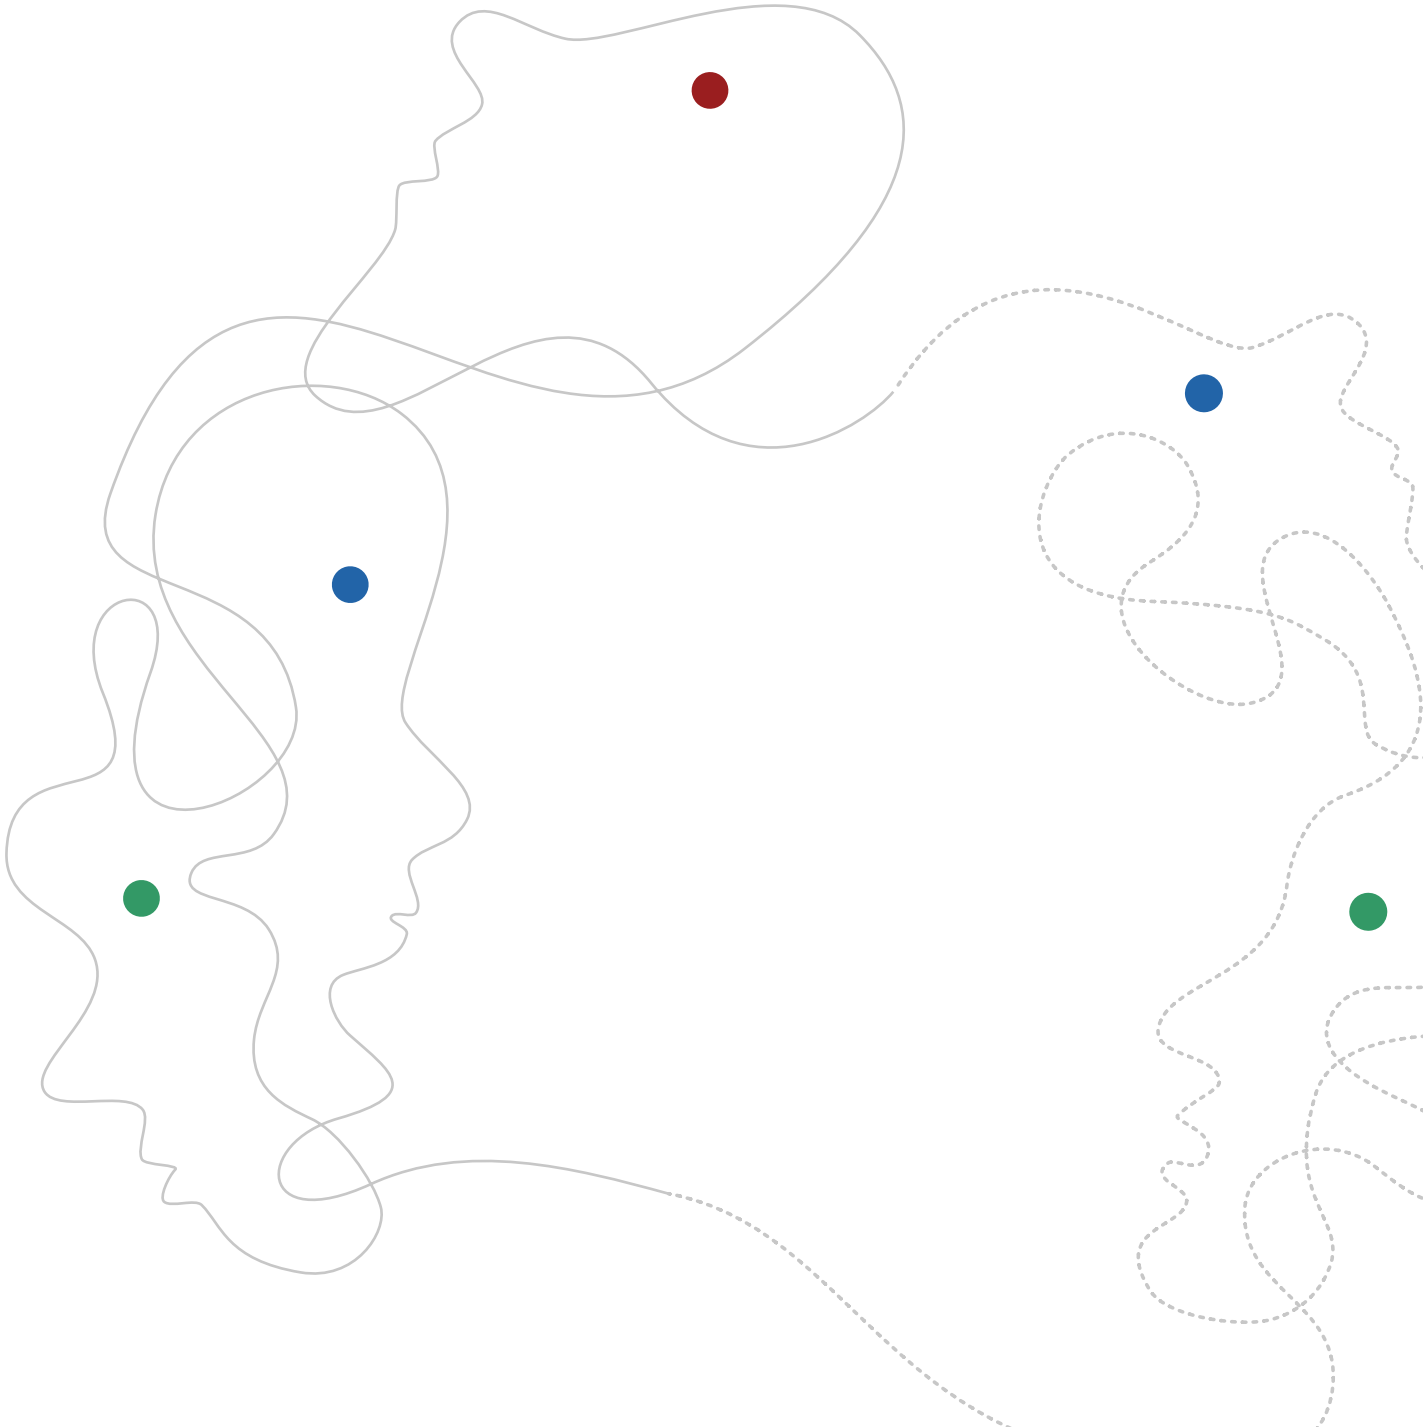

For any of Wikipedia's knowledge loops, the key question is whether it supports or hurts Wikimedia's sustainability. The new, third, knowledge loop, in its current form, largely extracts value from Wikipedia, without creating positive feedback mechanisms. These AI-related uses, happening at massive scale, are exactly the kind of broad reuse that the open licensing framework was designed to enable.

THE BROAD REUSE THE LICENSES INVITED

Wikipedia content has been used to train models that now serve hundreds of millions of users. By the original logic of free knowledge, this should count as a success story.

In 2014, 100 million Flickr photographs, shared with open licenses, were used in a training dataset for machine vision systems.[29] It was a rare case of massive reuse of free knowledge – which quickly turned out to be problematic. Concerns were raised over attribution, disregard for licensing conditions, but also lack of reciprocity and ethical use of the photographs. The Flickr case predated the generative AI boom, but already spelled out tensions that have intensified in recent years.

Free licensing was, in principle, open to new unexpected uses, and permission to use was granted beforehand. This category of uses – at massive scale, for commercial gain, without reciprocity – poses a challenge to open licensing frameworks. There is an often repeated sense that some form of social contract has been broken. According to Sarah Pearson, "It is unlikely that the incentives to share can outweigh all of the growing countervailing forces at play: economic, moral, safety, more. We cannot respond by accepting these risks and harms as inherent and inevitable costs of public sharing knowledge".[30] This tension is

[29] Tarkowski, Alek, and Zuzanna Warso. **AI Commons. Filling the Governance Vacuum on the Use of Information Commons for AI Training.** Open Future, https://openfuture.eu/publication/ai-commons/. Accessed 24 Apr. 2026.

[30] Hinchliff Pearson, Sarah. "**How to Keep the Internet Human.**" *Creative Commons*, 12 Feb. 2026, https://creativecommons.org/2026/02/12/how-to-keep-the-internet-human/.

A Dark Commons

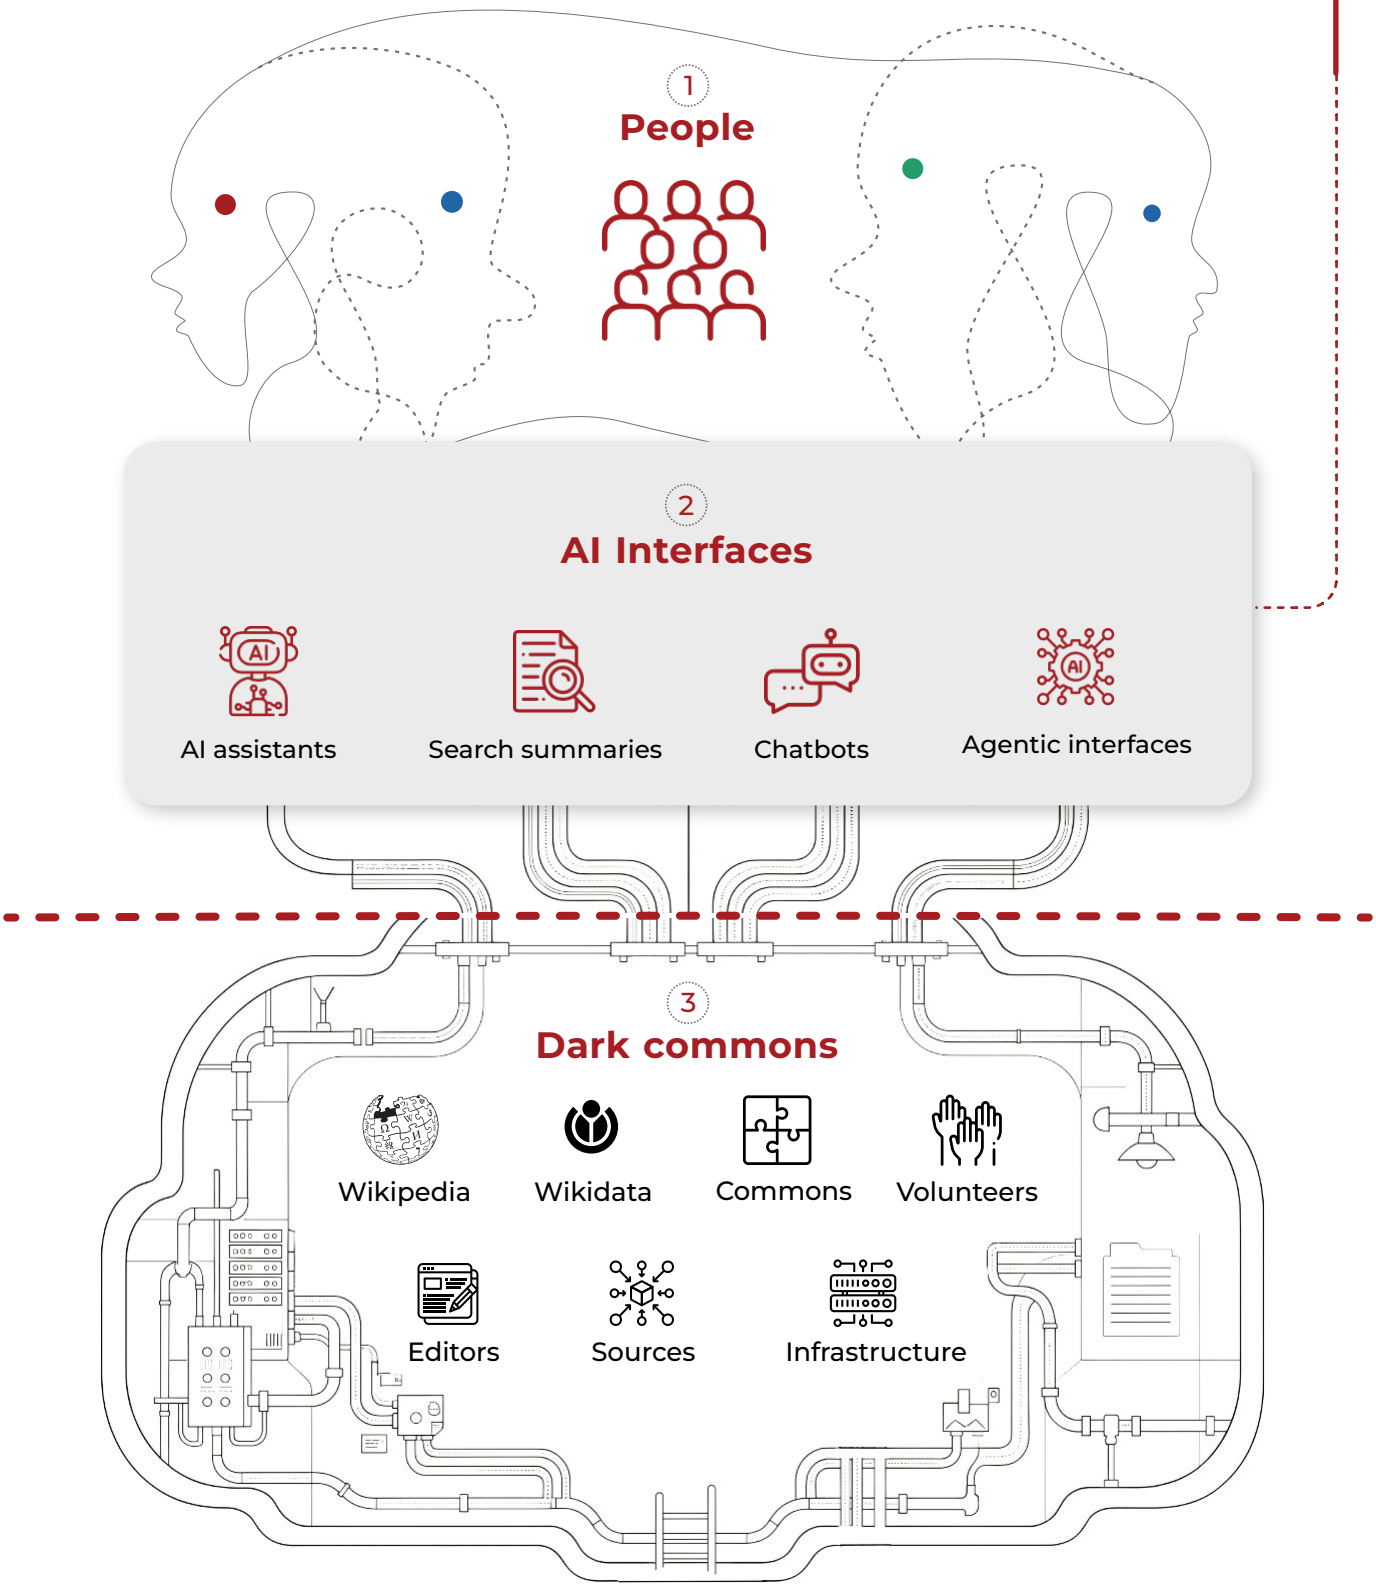

increasingly recognized within the open licensing community.[31]

THE RISK OF THE DARK COMMONS

Thus, a **paradox of open** occurs: sharing methods that were meant to challenge concentrations of power are prone to capture, and value extraction. In other words, power disparities have in recent years increasingly challenged knowledge equity.[32] This paradox has always been present, but now becomes acute. The knowledge commons are becoming a "dark commons": an invisible substrate that resources AI systems, which in turn do not reciprocate, or even acknowledge the role of these commons.

Historically, the knowledge commons have been built from a perspective of abundance, with a vision that successful reuse would ensure sustainability and further success. This vision has been expressed through copyleft mechanisms, in particular the Share Alike condition of Wikipedia's licensing model. The original knowledge loop also managed to balance the flow of value, by providing a flow of users back to Wikimedia.

[31] For other statements on the issue, see:
Ávila, Renata. "**Updating the Open Definition to meet the challenges of today.**" Open Knowledge Foundation. 16 Mar. 2023, https://blog.okfn.org/2023/03/16/updating-the-open-definition-to-meet-the-challenges-of-today/ ;
Krikorian, Raffi. "**Owners, not renters: Mozilla's open source AI strategy.**" Mozilla Foundation. 8 Jan. 2026, https://blog.mozilla.org/en/mozilla/mozilla-open-source-ai-strategy/ ;
Open Knowledge Foundation. Angela Oduor Lungati: "**When embracing open source, one must be intentional about being inclusive**". 25 Mar. 2024. https://blog.okfn.org/2024/03/25/angela-oduor-lungati-when-embracing-open-source-one-must-be-intentional-about-being-inclusive/ .
For a landscape scan of related issues see:
Collister, Lauren, Lippincott, Sarah and Skinner, Katherine E. "**Sustaining the Commons in the AI Economy: A Landscape Scan of Challenges and Strategies for Bridging AI Companies and Open Curated Collections**". Invest in Open, 21 Apr. 2026. https://investinopen.org/blog/sustaining-the-commons-in-the-ai-economy/

[32] See also:

FROM STRAIN TO EPISTEMIC CAPTURE

Today, the massive scale of reuse strains Wikimedia's resources and puts at risk its self-sustaining character. But this is the lesser risk – the one that is merely visible. The greater one is qualitative: a shift in who controls the tools through which knowledge is produced and accessed, and a narrowing of the very plurality of ways in which societies establish what is true..[33]

Unlike previous technologies, AI systems are increasingly used as tools for generating and mediating knowledge itself. Large AI models are closed, proprietary structures built on top of the knowledge commons. They are also becoming a key – and for some already the primary – interface through which people and organizations produce and consume knowledge.[34] This gives AI companies increasing influence over how knowledge is generated, validated and distributed. In the long run, the logic of proprietary AI platforms aims to displace alternative approaches, creating a dependency not just on their infrastructure, but on the single epistemological model encased in their AI technologies.[35]

This has direct consequences for Wikimedia. On one hand, it will remain a valuable source of knowledge for AI systems – but value as a source is not the same as agency over how knowledge is made. On the other hand, as AI generated content proliferates, the real risk will be not growing volumes of AI slop, but dependence on a single epistemological source of knowledge. On one proprietary way of generating, validating and ranking knowledge, displacing the plurality of methods – many sources, many forms of verification, the negotiated point of view – through which knowledge has been contested and established.

[33] Mazzucato, Mariana, et al. **Governing Artificial Intelligence in the Public Interest.** UCL Institute for Innovation and Public Purpose, 28 Jul. 2022, https://www.ucl.ac.uk/bartlett/publications/2022/jul/governing-artificial-intelligence-public-interest.

[34] Keller, Paul. op. cit.

[35] Stasenko, Anastasia. **"The Will to Generate."** *Le Substack de Anastasia*, 17 Apr. 2026, https://anastasiastasenko.substack.com/p/the-will-to-generate.

This is where the loss is most fundamental. Wikipedia's real contribution was never only the free knowledge, but also the open and plural process through which that knowledge is produced.

A world that routes knowledge through one epistemic interface not only extracts value built through this process, but also makes it invisible, and eventually unnecessary. In the most acute scenarios, the platforms producing trusted knowledge can collapse when their loops with human contributors and readers break — a pattern already visible elsewhere on the web.

STACK OVERFLOW AS A WARNING

The case of Stack Overflow illustrates what such platform collapse looks like. For fifteen years, it operated as a robust knowledge commons for software development, produced and curated by a community of human developers. At its peak, the platform processed over 200,000 questions per month. Then, beginning in late 2022, AI coding tools offered a faster, more personalized alternative. Monthly question volume collapsed by over 75%, returning by early 2025 to levels not seen since the platform's 2008 launch. Notably, Stack Overflow was an important source of training data for AI models.[36]

The dynamics of the new knowledge loop are affecting not just Wikimedia, but the entire ecosystem of knowledge and cultural production: public media, news outlets, libraries, open access repositories, open-source infrastructure, or cultural heritage institutions. AI systems "increasingly substitute for direct engagement with original sources, systematically eroding the visibility, traffic, and revenue streams that sustain human information production".[37] What happened to Stack Overflow is the trajectory awaiting Wikimedia and other knowledge commons that do not act.

[36] Anderson, Tim. "**Stack Overflow Seeks Rebrand as Traffic Continues to Plummet – Which Is Bad News for Developers.**" *Devclass*, 13 May 2025, https://www.devclass.com/ai-ml/2025/05/13/stack-overflow-seeks-rebrand-as-traffic-continues-to-plummet-which-is-bad-news-for-developers/1623624.

[37] Keller, Paul. op. cit.

The Paradox of Open

- Open licensing
- Free reuse
- Knowledge equity
- Public benefit

Openness without reciprocity becomes extraction.

- Massive scraping
- Commercial AI
- No attribution
- No return
- Infrastructure pressure

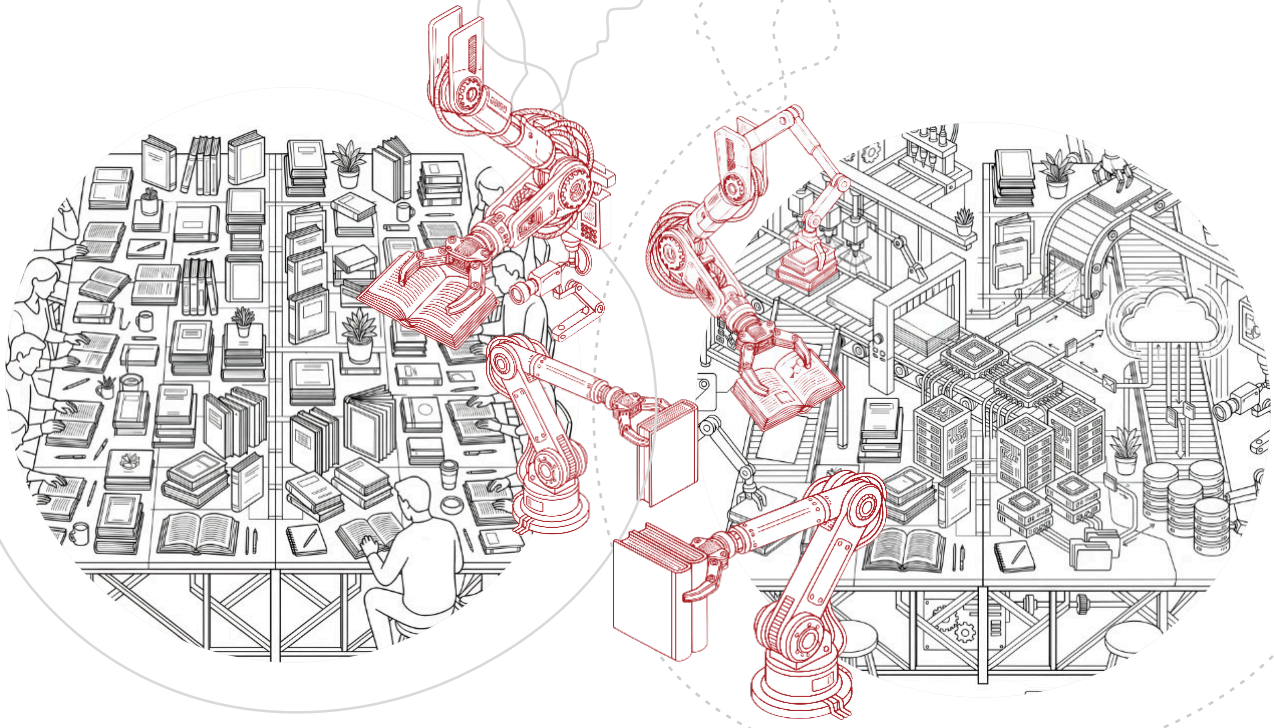

A shared challenge across the knowledge commons

Wikimedia is unusual in having the institutional scale, global recognition, and normative legitimacy to respond to these pressures. Most other knowledge commons (open-source projects, heritage institutions, open access repositories) have fewer resources to address the same structural dynamics.

A SYSTEM-WIDE PRESSURE

Strains on infrastructure are common across the knowledge commons ecosystem. OpenStreetMap reports being "**at constant war**" with scrapers, with its senior engineer seeing coordinated scraping from 100,000 IPs in a single week. Open source projects including KDE, GNOME, and Fedora have experienced **repeated outages caused by AI crawlers**. A 2025 report from Michael Weinberg at GLAM-E Lab documented that of 43 surveyed heritage institutions, 27 attributed recent traffic spikes to AI training bots.[38]

In addition, many platforms are also losing human traffic to their sites along with key sources of revenue and engagement. **Pew Research study** shows that click-through rates fall by 50% when users encounter AI search summaries, and only 1% of them click on links within the summaries. Chartbeat data **cited by Reuters Institute** shows a 33% drop, year-on-year in Google traffic to news sites. This is compounded by significant drops in referral traffic from major social media sites.

[38] Weinberg, Michael. **Are AI Bots Knocking Cultural Heritage Offline?** GLAM-E Lab, 2025, . https://www.glamelab.org/products/are-ai-bots-knocking-cultural-heritage-offline/

WIKIMEDIA AS AN ANCHOR

The governance frameworks, technical standards, and negotiation strategies that Wikimedia develops are not only relevant to its own sustainability. Wikimedia is well positioned to contribute frameworks, tools, and even shared infrastructure for other – usually smaller – actors that face the same pressures but lack the scale and institutional capacity to respond on their own.

Actions taken by Wikimedia have ripple effects across the knowledge commons. Wikimedia's strategy should be firmly based on principles of partnership and solidarity with this broader ecosystem, and in particular with sources that constitute Wikipedia's foundational loop.[39]

At the same time, Wikimedia does not need to build everything from scratch. Smaller commons-based actors are also experimenting with responses to AI-driven challenges. Examples include work done by Pleias on using open data (and synthetic data derived from it) for **building AI training datasets**; the **Transkribus project**, which has built hundreds of open models for working with historical documents – and governed them as a cooperative. These initiatives are too small to reshape the knowledge loop on their own, but they form a development ecosystem that can be combined with Wikimedia's own efforts to develop machine learning technologies, such as **the LiftWing project**. Therefore, a deliberate strategy of mutual learning would strengthen the resilience of the entire ecosystem.

[39] Trebor, Scholz, and Esposito Mark. "**Building a Solidarity Ecosystem for AI.**" *Stanford Social Innovation Review*, 5 Feb. 2026, https://ssir.org/articles/entry/artificial-intelligence-solidarity-ecosystem.

The uneven geography of the new loop

Wikimedia itself is not a monolith, and Wikipedia itself has over 300 active language versions. While trends across these language projects are broadly similar, and the dynamics of the largest projects shape what happens to Wikimedia as a whole, treating the impact of AI as uniform would be a mistake.

UNEVEN IMPACTS WITHIN WIKIMEDIA

The pressures fall hardest on the communities whose knowledge is already under-represented. The challenge of the new AI loop is directly tied to Wikimedia's long running problems with knowledge equity.

Machine traffic impacts various Wikipedia projects differently and for small ones can constitute a much larger share of overall traffic. These distinctions demand a detailed analysis. Especially important are those Wikipedias that – even if small – provide structured knowledge in low resourced languages. They have become key resources for AI development, precisely because they are often the only structured sources in those languages.

Yet the communities that maintain them are the least equipped to respond to the pressures described above. Editors of four African-language Wikipedias estimated that 40–60% of articles consist of uncorrected machine translations. In September 2025, the Greenlandic Wikipedia was closed by its administrators, citing AI-generated content as actively harmful to the integrity of the language. These indicators suggest that power imbalances in the new knowledge loop are felt most acutely by smaller Wikipedias, and that the resulting harms are unevenly distributed globally, falling disproportionately on Global Majority communities.[40]

[40] Judah, Jacob. "**How AI and Wikipedia Have Sent Vulnerable Languages into a Doom Spiral.**" *MIT Technology Review*, 25 Sep. 2025, https://www.technologyreview.com/2025/09/25/1124005/ai-wikipedia-vulnerable-languages-doom-spiral/.

THE MOVEMENT'S OWN EQUITY GAP

The Wikimedia Movement is not innocent as it faces the challenge. English Wikipedia's structural dominance, the under-resourcing of smaller-language projects, and content partnerships skewed toward Global North institutions lead to equity gaps. The challenge has been recognized in the movement since the 2010s, when the Foundation acknowledged that underrepresentation of women and the Global Majority is a structural issue. **The Strategic Direction**, adopted in 2017, sets Knowledge Equity as one of the two directions for the Movement. The same AI systems that extract value from the commons without reciprocity also propagate the commons' existing blind spots at scale.

The new knowledge loop, with its extractive dynamics, compounds the challenge with knowledge equity. Since language models are trained on Wikimedia, its epistemic biases become the epistemic biases of AI itself. AI-related trends do not merely inherit Wikimedia's existing equity gaps – they freeze them in place and amplify them, propagating them through a new cultural technology at massive scale.

Knowledge equity therefore means the continuing need for greater representation of Global Majority knowledge in Wikimedia, and as a result in generative AI models – but this measure is not sufficient. A deeper challenge is that Wikimedia's knowledge infrastructures themselves embed assumptions in which non-Western knowledge may not fit. The issue was best illustrated by a failed attempt, made already in 2010, to allow oral **citations** in Wikipedia.

FACING ALGORITHMIC COLONISATION

There is also a global dimension to the paradox of open. The challenges to open sharing, from a Global Majority perspective, map onto a long history of colonial extraction. Since actors with the capacity to extract at scale are concentrated in the Global North, the Global Majority now

faces data colonialism.[41] Abeba Birhane writes about algorithmic colonisation: the export of Western AI systems, assumptions, and extractive practices into communities that have no agency over how the resulting technologies are built or governed.[42] This puts into question the basic assumption of free knowledge – that removing legal barriers to reuse leads to equity.

A response to the challenge of the new knowledge loop needs to build on the Knowledge Equity direction, and shift resources toward knowledge commons activists and institutions in the Global Majority. The concentrations of power in the Wikimedia Movement itself have a geographical dimension that needs to be acknowledged, as investments are made into technical capacity and the Wiki AI stack is developed. Global majority partners thus need to play a key role in the mission that we propose. Similarly, technological solutions can be built in a way that address the needs and match the capacities of Global Majority communities. For example, the proposal for a federated AI Commons ecosystem describes an alternative AI loop built through decentralized development of AI technologies, based on data commons and with public control of technology.[43]

The federated AI commons proposal shows that Global Majority actors can play a key role in Wikimedia's AI mission, even if constrained by compute scarcity and other structural inequalities. The overall assumption should be that Global Majority solutions can still benefit the Global North, but not the other way around. Solutions designed for the most developed regions will propagate inequalities by default.

[41] Couldry, Nick and Mejias, Ulises A. **The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism.** 2019. Stanford University Press.

[42] Birhane, Abeba. "**Algorithmic Colonization of Africa**". SCRIPT-Ed 17 (2): 389–409. 2020. https://doi.org/10.2966/scrip.170220.389

[43] Gimpel, Lea, et al. Ibid.

Three principles of Wikimedia and AI work

As described in the previous section, the Wikimedia Movement needs to adopt an active stance in the new knowledge loop and aim not just to defend the commons, but to strengthen its generative capacity. This effort should be driven by three principles:

Principle 1: Protect and strengthen collective intelligence

Collective intelligence was a term used by **Pierre Lévy** to describe a knowledge space that emerges thanks to networked media.[44] In this sense, Wikipedia is not just a collection of articles – it is also a living infrastructure of collective intelligence, which emerges from the work and interactions of Wikimedia editors and stewards. It is human-driven, decentralised, and self-correcting through diversity (although the quality of that self-correction depends on the breadth of participation, which still remains uneven across the movement).

Both automation through AI tools and the declining health of editor and reader communities are weakening the conditions that enable this collective process to emerge and sustain itself. Without this collective process, Wikimedia would not function as a knowledge system, regardless of how much data it contains.

Collective intelligence and the knowledge commons have been strongly intertwined, as the latter depends on the former for its quality and trustworthiness. Historically, Wikimedia relied on a virtuous cycle in which collective intelligence produced high-quality knowledge, which in turn attracted users, contributors, and resources.

[44] Lévy, P., **Collective Intelligence: Mankind's Emerging World in Cyberspace**, Perseus Books, 1997

PROTECTING COLLECTIVE INTELLIGENCE

Today, approaches aimed at ensuring sustainability of the new knowledge loop need to recognize and strengthen this human process, directly tied to Wikipedia's function as a source of ground truth for AI and search systems, which treat it as the canonical reference layer. While Wikipedia rejects being a primary source, its ability to provide a citation-backed, continuously revised consensus is invaluable.

Active strategies to mitigate risks in the new knowledge loop cannot just protect Wikimedia as a data resource –they need to support collective processes that ensure its sustainability and health. **Wikimedia Foundation's AI strategy** reflects elements of this approach: it prioritizes human agency, while assuming that knowledge should still be generated by humans, supported by AI.

Wikimedia's strategy for the new knowledge loop must include mechanisms to counteract the erosion of the editor pipeline, to ensure that AI tools augment rather than replace human editorial judgment, and to protect epistemic diversity. While Wikimedia has not yet achieved the goal of knowledge equity, AI raises the risk of flattening this diversity further.

Principle 2: Reciprocity and symbiotic partnerships

Reciprocity is not only a normative principle, but also a structural condition for sustaining the knowledge commons. The original knowledge loop between Wikipedia and search engines was symbiotic: both sides generated value and returned it. The importance of reciprocal loops that help the commons create value and sustain itself was expressed in **the Share Alike principle**, which is a core element of Wikimedia's sharing framework and ensures that reuse contributes back to the com-

mons. The new AI-driven loop disrupts this balance.

Wikimedia Enterprise marked a crucial first step by establishing that large-scale commercial machine use of the commons should include mechanisms of value return, including financial and non-financial forms. **A sustainable knowledge loop depends on the development of broader norms: attribution that makes the commons visible rather than a hidden substrate; transparency about how commons-derived knowledge is used and transformed; and contributions back that strengthen the shared resource**.

BEYOND FINANCIAL RETURNS

Reciprocity can take other forms than financial returns. For example, AI platforms could share data on information usage (what users ask for, what answers receive low ratings or contradict ground truth, where coverage is thin, etc.) that would be highly valuable for Wikimedia communities as they try to allocate editorial attention.

Without reciprocity, the knowledge commons becomes a resource that is endlessly consumed and never replenished. Over time, this would hollow out the ecosystem that produces trusted knowledge and the AI systems that depend on it.

This requires engaging with major AI actors, to ensure that large-scale reuse of the commons aligns with principles of attribution, transparency, and returns value to the commons.

Principle 3: Deliberate shaping of the knowledge ecosystem

The Wikimedia Movement's response to the growing impact of AI technologies has, so far, been largely defensive and focused on "on-wiki"

solutions: either seeking to **control the impact of AI on Wikimedia infrastructure**; or framing **rules for the use of AI tools** on Wikipedias. This leaves the architecture of the new knowledge loop: who mediates access to knowledge, and on whose terms, in the hands of a handful of commercial AI companies. Wikimedia must acknowledge that this is not a tooling or infrastructural question, but an ecosystem issue – and act at three escalating levels.

STEP ONE: GOVERNANCE OF THE NEW LOOP

The first step requires engaging in governance of the new loop, ensuring that Wikimedia and other organizations building the commons have agency in the broader ecosystem. The new knowledge loop must be open and undistorted by concentrations of power. Wikimedia is one of the few actors with the scale and legitimacy to shape these new rules.

This means engaging with the bodies where the rules of the new knowledge loop are being written – such as the Agentic AI Foundation, a Linux Foundation consortium whose founding members are exclusively large technology companies, and which is tasked with the governance of key AI ecosystem protocols like MCP and A2A. Still, having a say in how extractive infrastructures are governed will not, on its own, shift value flows toward the commons.

STEP TWO: BACKING COMMONS-ALIGNED ALTERNATIVES

The second step is about backing commons-aligned alternatives. Wikimedia must also use its credibility and position to shift the field toward responsible actors and practices, by contributing to the development of public AI infrastructure. Public AI describes AI systems treated as public digital infrastructure. The framing has been proposed by the Public AI Network and elaborated in reports from Mozilla Foundation, and Bertelsmann Stiftung with Open Future. It rests on three pillars that distin-

guish it from the dominant model of AI development: open and non-discriminatory access, mission-driven public goals, and public control through democratic accountability. Public AI is not a single model or institution but an ecosystem of efforts demonstrating that AI infrastructure can be built differently when it is treated as a public good.

This means endorsing and amplifying AI development projects such as **Apertus** in Switzerland, **OLMo** in the United States, **Pleias** in France, **Bielik** in Poland, **SEA-LION** in Singapore, **Masakhane** across Africa, and many other initiatives around the world. It means putting Wikimedia's weight also behind early-stage efforts to develop similar solutions in other regions, like **Latin America**. All of these initiatives demonstrate that the technical path for public, reciprocity-based AI is viable. Wikimedia's brand and institutional legitimacy, deployed behind these projects, can tilt the field toward responsible actors. This also means developing shared standards for training on commons data with reciprocity built in, and building AI competencies and capacity in communities.

STEP THREE: BUILDING PUBLIC AI

The third and ultimate step, needed to bring this principle to life, means building Wikimedia's own public AI. **This kind of public AI system, created in partnership with aligned AI developers and other organizations willing to provide necessary resources, is to the new knowledge loop what the wiki was to the first: the core piece of infrastructure through which a commons meets its public.**

The only way to secure a user-facing platform in the new knowledge loop is to participate in building and shaping it. Wikimedia is structurally positioned to do so, as both a key source of training data and the most established example of commons-based knowledge production. This means contributing to the development of public AI models trained on commons data, developed with Wikimedia's collective-intelligence ethos, designed with reciprocity in mind, and governed accord-

ing to commons values; and building the interfaces through which these models meet the public.

Commons health, value flows and strategic postures

The three principles can form the basis for various strategic postures. This section introduces a framework that begins with three indicators of commons health, followed by three types of strategies that can be adopted in response to various types of value flows in the commons.

THREE INDICATORS OF COMMONS HEALTH

Interventions meant to secure the sustainability and resilience of Wikimedia in the knowledge loop can be reviewed against the following indicators for the health and sustainability of the knowledge commons:

1. **Openness**. This has traditionally been the key indicator, and the strength of the commons has been measured by the volume of openly shared resources – for example, the size of a Wikipedia. This supply-side metric takes sustainability for granted, and often ignores demand side-metrics. Even access metrics, collected by Wikimedia, fail to capture uses and their outcomes. This indicator alone is no longer sufficient, as strategies focused on maximizing openness fail to address challenges to the sustainability of the commons.
2. **Sustainability**. A healthy digital commons has mechanisms that ensure its long term institutional, technical, and financial sustainability. This is an indicator increasingly considered in current discussions on healthy open-source and commons ecosystems.
3. **Value generation**. This assumes that the real value of digital commons lies not just in their availability. The value of the commons lies not just in usage, but in the ability to produce new knowledge, contributors, and ecosystem capacity. Openness on its own is not necessarily generative, and additional measures are needed to ensure a productive commons.

In order to establish a strategic position in the new knowledge loop, Wikimedia needs to take all three indicators into account. Optimizing for openness alone leaves aside issues related to the paradox of open and the extraction of value from the commons. Ensuring just sustainability can lead to strategies that limit access to the commons. Optimal strategies also require placing value generation in scope.

THREE TYPES OF USES AND VALUE FLOWS

These strategies also need to consider three types of value flows – three ways, in which the commons can be used for value generation:

1. **Extractive uses:** Value flows out of the commons without return. Commons resources are consumed by external actors who capture the resulting value entirely. Example: commercial platforms scraping Wikipedia content to train proprietary models, with no contribution back to the Wikimedia ecosystem.

2. **Reciprocal uses:** Value is exchanged between the commons and external actors on negotiated terms. Use of commons resources is paired with a return that supports the commons' continued operation. Example: Wikimedia Enterprise provides structured data access to commercial AI developers in exchange for revenue that funds Wikimedia infrastructure.

3. **Generative use:** Value circulates back into the commons and expands it. Use of commons resources produces outputs that are themselves open and commons-aligned, strengthening the ecosystem as a whole. Example: commons-sourced training data is used to build public AI models that are governed as a commons.

From Extraction to Generation

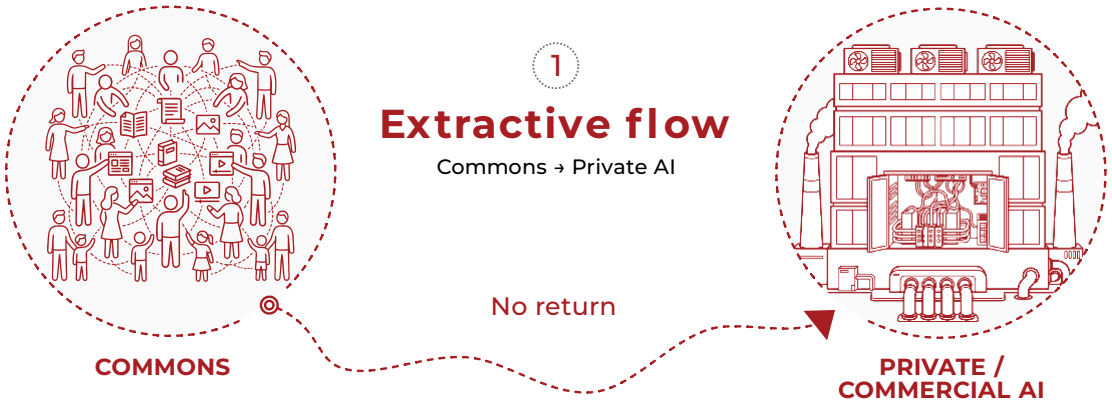

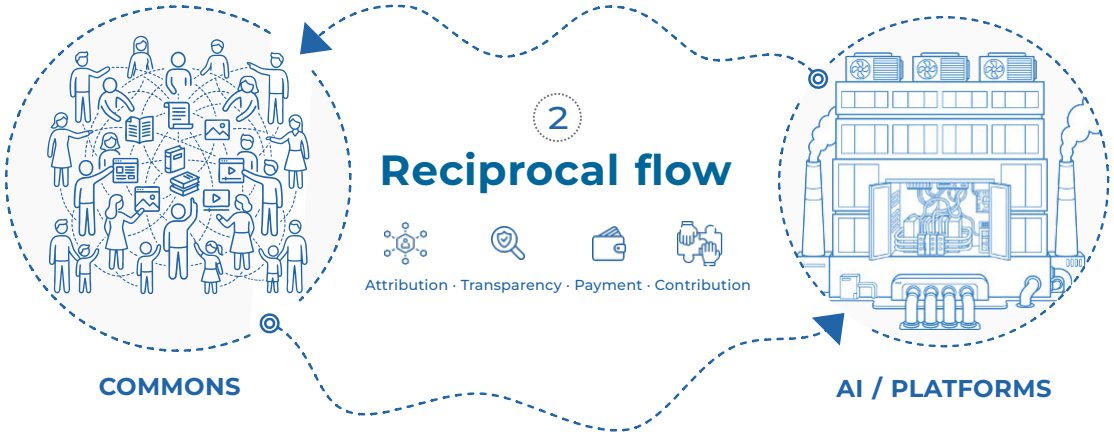

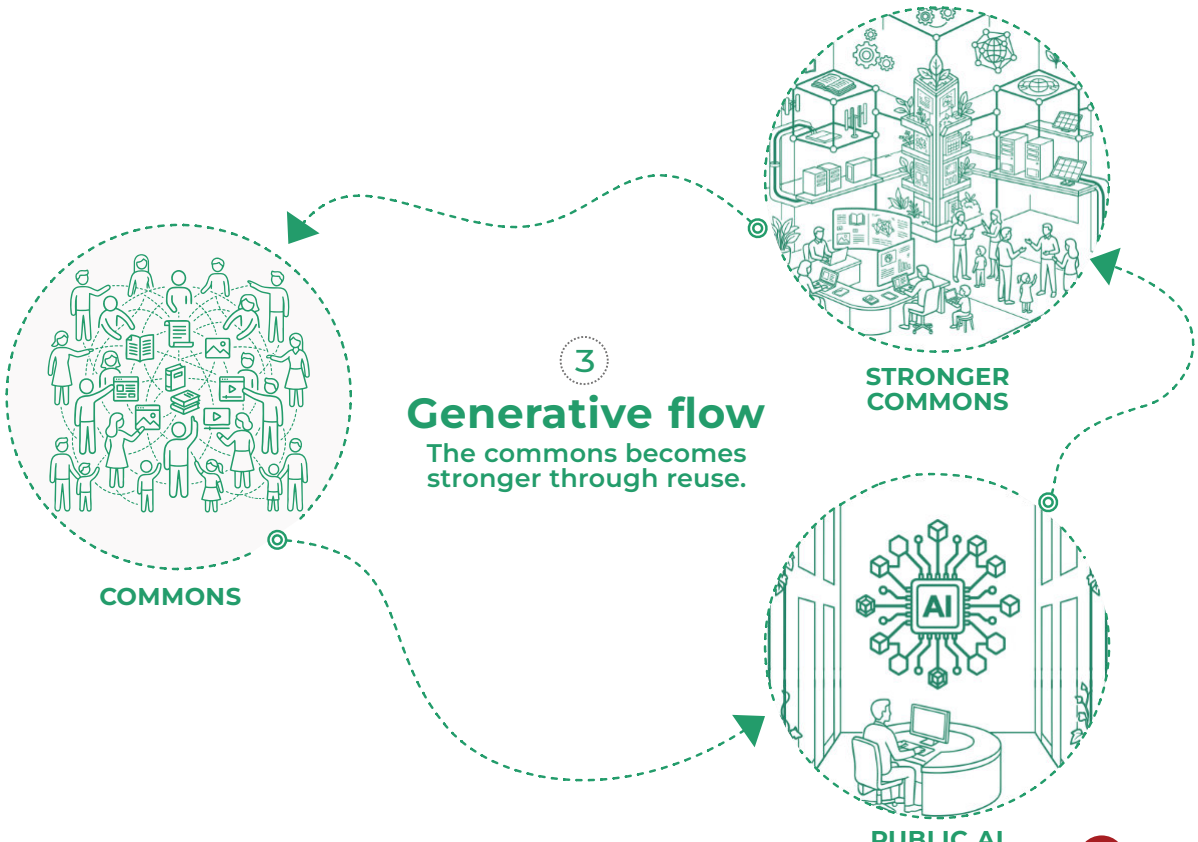

THREE STRATEGIC POSTURES

Finally, as stewards of various digital commons navigate the changing online environments, three strategies emerge:[45]

1. **Defensive closure:** Restricting access to commons resources in response to extraction, and in the name of sustainability. This reduces harms but also limits the generative potential of the commons itself.
2. **Passive acceptance:** Maintaining openness without actively governing the terms of reuse. Until recently, this was the default posture for most commons. It treats all uses as equally legitimate and valuable, and assumes sustainability will follow from sharing.
3. **Active stance:** Proactively shaping the terms of engagement with AI systems, aimed at securing access and control over commons-aligned solutions. This is the only posture likely to shift the value flow toward generation.

Mapping the three types of value flows with the three strategies available to the stewards of the commons results in the following matrix:

	1. Extraction	2. Reciprocity	3. Generation
1. Defensive closure	Rate limiting and targeted blocking of AI scrapers	Transactional closure through exclusive licensing deals	Not possible.
2. Passive acceptance	Wikimedia in the AI-era knowledge loop, where the predominant downstream mode is extraction.	2021: Wikimedia Enterprise introduces basic reciprocity for Wikimedia, without aiming at structural changes to the loop.	Wikimedia in the pre-AI knowledge loop, where the predominant downstream mode is generation.
3. Active stance	Not feasible.	Active enforcement of norms: attribution, Share Alike, transparency, AI sustainability commitments.	Public AI: trained on commons data, aligned with commons through governance measures, in service of human collective intelligence.

[45] Nathalie Casemajor has proposed a similar typology of strategic actions that Wikimedia can take in face of LLMs, by distinguishing defensive, adaptive and offensive strategies.
See: https://praxis.encommun.io/media/notes/note_20490/casemajor_diff_sat6918.pdf

Defensive closure and passive acceptance are surprisingly similar in their outcomes: both fail to address the structural problems the knowledge commons face today. Defensive closure responds by protecting the commons, but at the cost of reducing its value. In its transactional variant, it can slide into an outright enclosure, as the Reddit and Stack Overflow cases show. There, closure was not a defense of the commons but the mechanism that made it monetizable. Passive acceptance looks different on the surface, but ends in the same place: maintaining the appearance of openness while the commons is hollowed out through extraction. Both are reactive postures that leave the terms of the loop defined by others.

Passive acceptance was not always a failure. For roughly two decades, keeping things open did produce generative reuse: remixes, derivative commons, and a civic knowledge infrastructure. This posture worked because in the previous knowledge loop value was to a greater extent recirculated. In the AI era, the dominant downstream use of open commons is extraction by actors who do not return value.

Wikimedia, to its credit, has already taken steps to ensure some forms of reciprocity, with Wikimedia Enterprise being a prime example of this posture. But on the whole the movement has operated much closer to passive acceptance of the ongoing changes. In the face of the emerging new knowledge loop, the Wikimedia movement needs to move beyond passive acceptance of the changing environment, avoid both defensive closure and its transactional variant, and build an active posture focused on reciprocity and value generation.

The importance of Wikimedia Enterprise

To date, Wikimedia Enterprise stands as one of the most important governance innovations pointing in this direction. Launched in 2021 with the primary aim of ensuring Wikimedia's financial sustainability, it predates the generative AI boom. By establishing a paid access tier for a structured data API, it has been successful in defining a norm that large-scale commercial machine use of the commons can involve value return, rather than remaining entirely unstructured.

In the financial year 2024-2025, the initiative generated $8.3M revenue, 4% of the Foundation's total revenue.[46] While still a relatively small share, this marks a meaningful evolution from the traditional open default and contributes to diversifying Wikimedia's funding model. The partner ecosystem is notable and includes major AI model developers (save for Anthropic and OpenAI), traditional and AI-native search companies, and smaller actors building alternative, open-source AI solutions.[47]

More importantly, Wikimedia Enterprise establishes a governance precedent: that reuse of the commons at scale can be shaped through negotiated relationships, rather than relying on broad permissions granted in advance through Open Access policies.

Yet financial reciprocity is only a first step. The next move is to go beyond financial remuneration and establish broader pro-commons norms –including attribution that preserves visibility, transparency in how knowledge is reused and transformed, and contributions back that strengthen the commons itself. These norms would extend beyond bilateral agreements and begin to shape the ecosystem as a whole.

[46] Wikimedia Foundation. "**Wikimedia Enterprise Financial Report: Fiscal Year 2024 – 2025.**" *Diff*, 24 Nov. 2025, https://diff.wikimedia.org/2025/11/24/wikimedia-enterprise-financial-report-fiscal-year-2024-2025/

[47] Reynolds, Chuck. "**Announcing New Wikimedia Enterprise Partners for Wikipedia's 25th Birthday.**" Wikimedia Enterprise, 15 Jan. 2026, https://enterprise.wikimedia.com/blog/wikipedia-25-enterprise-partners/

The Wikimedia and AI mission

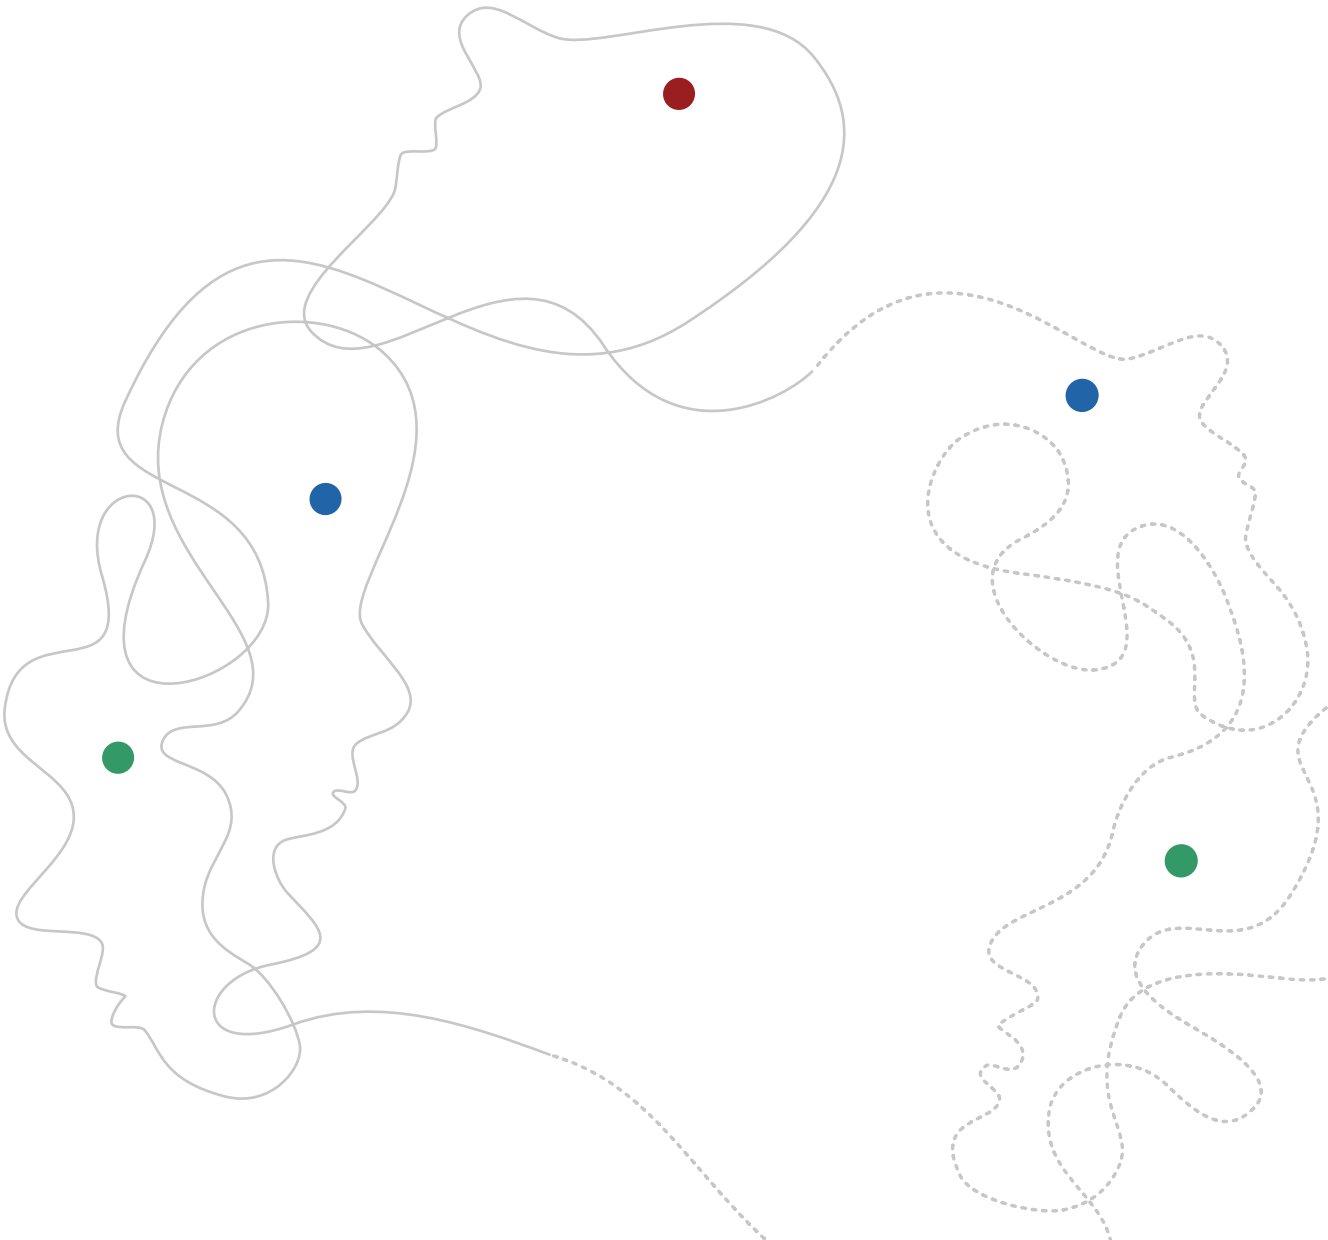

Facing a systemic transformation of how knowledge is produced and circulated, the Wikimedia movement needs a mission: a time-bound, cross-sectoral commitment to shape the ecosystem, not just respond to it. The challenges are systemic in nature. Isolated interventions will therefore not achieve an ecosystem-level shift.

WHY A MISSION, NOT JUST A STRATEGY

The mission-oriented approach developed by Mariana Mazzucato,[48] and since refined by the OECD and a wider literature, offers the most coherent playbook for coordinated action across a network of diverse actors.[49]

At the heart of this approach is an idea that public and commons actors can become shapers of markets and online spaces: setting direction, and using available tools to steer whole ecosystems. A mission is thus more than a way to fix problems – it offers a chance to reorient an ecosystem and support value creation while interrupting its extraction. The Wikimedia Movement is positioned to apply this approach to the new knowledge loop.

Mazzucato herself has extended this argument directly to AI, arguing that the central question is not whether to regulate AI but how to actively steer its development toward public value creation rather than value extraction".[50]

[48] Mazzucato, Mariana. "**Mission-Oriented Innovation Policies: Challenges and Opportunities.**" *Industrial and Corporate Change*, vol. 27, no. 5, Oct. 2018, pp. 803–15, https://doi.org/10.1093/icc/dty034.

[49] OECD. **Forging New Frontiers in Mission-Oriented Innovation Policies.** OECD Publishing, 2026, https://www.oecd.org/en/publications/forging-new-frontiers-in-mission-oriented-innovation-policies_d13d0142-en.html.

[50] Mazzucato, Mariana. "**AI for What? Public Value Creation versus Extractive Rents.**" Mission Economics by Mariana Mazzucato, 3 Feb. 2025, https://marianamazzucato.substack.com/p/ai-for-what.

WHAT MAKES A WELL-DESIGNED MISSION

Not every mission is the same. The Wikimedia and AI mission is a complex one, as both the problem and responses to it are contested. This means that the mission must make its framing choices explicit rather than assume consensus; it must build legitimacy through participatory formulation, especially with editor communities and Global Majority Wikipedias whose exposure is greatest; and it must be evaluated developmentally.

A well-designed mission must also be bold and socially legible (a non-expert can repeat it and say why it matters); targeted, measurable, time-bound; ambitious but plausible on a 5–10 year horizon; cross-disciplinary, cross-sectoral, and cross-actor by necessity rather than by invitation; and open to a portfolio of competing bottom-up solutions rather than a single predetermined approach.[51]

COLLECTIVE INVESTMENT LEADS TO COLLECTIVE CLAIM

One of Mazzucato's most persistent arguments is that collective investment in innovation entails a collective claim on its returns. Applied to Wikimedia, this clarifies that a platform that supports humanity's collective knowledge should not only claim fair compensation, but also establish structural mechanisms that ensure the new knowledge loop is symbiotic rather than parasitic with the commons.

In this way, the Wikimedia and AI mission will have an impact far beyond Wikimedia itself. If successful, it should make the broader knowledge commons more resilient as they adapt to the changing shape of the web, and shift the norms of the new knowledge loop and the agentic web emerging today.

[51] Mazzucato, Mariana. "**Governing the Economics of the Common Good: From Correcting Market Failures to Shaping Collective Goals.**" *Journal of Economic Policy Reform*, vol. 27, no. 1, Dec. 2023, pp. 1–24, https://doi.org/10.1080/17487870.2023.2280969.

Achieving this requires treating the challenge not as a tool question but as an ecosystem question — investing in infrastructures and governance, not only applications. And if the mission succeeds, and Wikimedia becomes a contributor and co-operator of public AI infrastructure, the dichotomy between AI tools for Wikimedia and Wikimedia's contribution to the AI ecosystem stops being relevant.

Without a coordinated, mission-oriented approach, Wikimedia risks remaining a passive provider of raw material for systems it does not influence. In such a scenario, Wikimedia would continue to be widely used, but with diminishing influence over how knowledge is produced, validated, and governed.

Elements of the Wikimedia and AI mission

We propose the following mission for the Wikimedia movement and its aligned partners:

By 2030, Wikimedia works in coalition with public AI developers and commons-aligned partners to build and govern AI interfaces and infrastructures through which the knowledge commons are produced and reach readers. As a result, the Wikimedia Movement is one of the primary user-facing knowledge interfaces and has agency over how knowledge is produced and used in the AI era.

For twenty-five years Wikimedia has understood its role as the production of the commons. The reader-facing interfaces have been, to a large extent, delegated to search engines. With intermediation by AI services, delegation is no longer survivable. Interface agency is today crucial for the free sharing of knowledge.

The mission is not to build "Wikimedia's own ChatGPT". It is to build, together with commons-aligned AI partners, a category of interfaces that make the knowledge commons directly available to readers on terms the commons itself sets. Interface agency must be pursued without undermining the conditions of collective intelligence that produce Wikimedia. Therefore, these interfaces must direct readers back to Wikipedia itself, which is how the editor and donor pipeline is sustained. And secondly, the interfaces must be plural, and accept various approaches, both in terms of technologies used, and local contexts and perspectives with which they are aligned.

This is underwritten by (i) a public, commons-governed Wiki AI stack, first deployed in Switzerland and scaled through a mission network of aligned public AI developers, and (ii) access conditions that require any AI system relying on the commons at scale to contribute to its attribution, reciprocity, and renewal.

This mission rests on two pillars:

1. **Pro-commons Governance:** shaping the rules of the new knowledge loop. The aim is to ensure, by 2030, that access and use of the commons (both at AI training and inference scale) is conditional on systemic commitments that are aligned with the sustainability of the commons.

2. **Wiki AI Development:** the direction for technology development, with two parallel targets. Firstly, a Wiki ML platform is needed to augment collective intelligence on Wikimedia. Second, this includes the development of AI knowledge interfaces, including AI agents.

The two pillars are mutually reinforcing. Governance work has more force when backed by credible alternatives; and reciprocity norms are easier to negotiate when Wikimedia is an active participant in the AI ecosystem, and not merely a data source for it. And, in turn, governance frameworks provide clarity for Wikimedia's partners and allow norms to spread more broadly in the ecosystem.

PRO-COMMONS GOVERNANCE

Attribution, Share Alike for AI, transparency, reciprocity, standards.

TECHNICAL DIRECTION

Data layer, model/tool layer, application layer.

Central Mission 2030

By 2030, Wikimedia works in coalition with public AI developers and commons-aligned partners to build and govern AI interfaces and infrastructures through which the knowledge commons are produced and reach readers. As a result, the Wikimedia Movement is one of the primary user-facing knowledge interfaces and has agency over how knowledge is produced and used in the AI era.

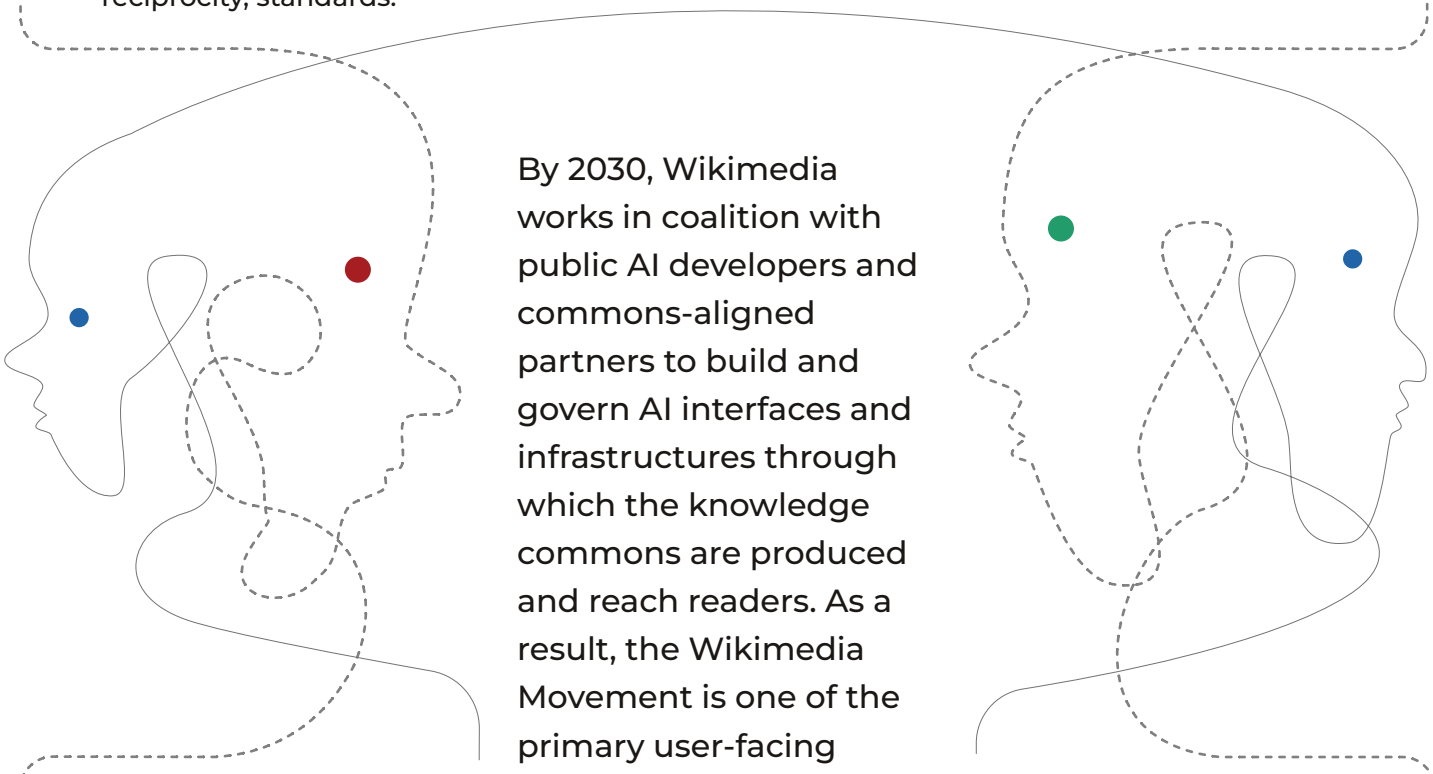

MISSION NETWORK

WMF, WMDE, affiliates, volunteers, researchers, public AI partners, other commons.

MISSION INVESTMENT

Funding, staff capacity, legal expertise, political engagement.

Pillar 2

Wiki AI Development

Actively shaping the norms that govern knowledge flows in the AI era:

- Rethinking attribution and Share Alike norms for AI contexts
- Developing reciprocity mechanisms beyond financial payment
- Securing representation in bodies developing AI ecosystem protocols

Pillar 1

Pro-Commons Governance

Building the technical capacity to remain an active participant in the new knowledge loop:

- **Data and knowledge layer**: Wikipedia as ground-truth corpus, Wikidata as structured knowledge graph, Wikimedia Enterprise as access API
- **Model and tool layer:** partnerships with open AI developers, experimentation infrastructure, task-specific models
- **Application layer:** tools that support editors, and new interfaces for knowledge delivery

The goal is not to replicate proprietary AI.
It is to ensure Wikimedia has agency in how the new knowledge loop is built, and that commons-aligned AI development is a credible alternative.

PRO-COMMONS GOVERNANCE: SHAPING THE RULES OF THE NEW KNOWLEDGE LOOP

The Wikimedia Movement needs to play an active role in shaping the norms that govern how knowledge flows in the new knowledge loop. This includes rules for the reuse of Wikimedia content and infrastructure, building on the foundation laid by Wikimedia Enterprise and other ongoing work. It then extends to norms developed and shared with aligned partners, including organizations from the digital commons and developers of public and open source AI. It also engages in broader ecosystem governance, where Wikimedia can use its position as a globally recognized and trusted source of knowledge to push responsible approaches.

Three governance issues are key, and should be the starting point for governance work:

Terms of downstream use. Attribution and Share Alike are Wikimedia's fundamental norms for open access and reuse that need to be re-designed. This also includes API governance rules.

Reciprocity mechanisms. Longstanding challenges around financial and structural sustainability must be addressed in the context of the new knowledge loop as well.

Ecosystem protocols. A number of standards and protocols are beginning to structure how AI technologies operate and impact the commons (MCP, A2A, etc.).

Ethical data collection and stewardship. Rules for how datasets are collected, documented, licensed, quality-assessed, and maintained. For some types of data, this also requires dealing with issues of consent, representativeness, and access conditions.

While most of these governance measures concern the ecosystem external to Wikimedia, and the various downstream uses, the on-wiki governance of AI is relevant as well. The draft AI Guidelines, which describe recommended practices for developing and deploying tools

that use Artificial Intelligence (AI) within Wikimedia projects, are an important starting point for this.

WIKI AI DEVELOPMENT: THE DIRECTION FOR TECHNOLOGY DEVELOPMENT

The Wiki AI Development element of the mission is delivered through a collective technical capability set that is owned and governed as a commons: the Wiki AI Stack. It is the technical architecture through which Wikimedia retains agency in the new knowledge loop. The Wiki AI stack consists of three layers:

- **Data and knowledge layer**, including Wikipedia as a ground truth corpus, Wikidata as structured knowledge graph, Wikimedia Enterprise as APIs for automated access;
- **Model and tool layer**, including both the state-of-the-art public AI model and AI tools for Wikimedia-specific tasks and interfaces for real-time knowledge retrieval;
- **Application layer**, built on the first two layers and includes both augmentation applications (inward facing) and interface applications (outward facing).

Some of these elements are already in place, especially at the first layer. Various machine learning tools have also been developed in the past, including those for multilingual content translation. Still, many elements of the Wiki AI Stack need to be modernized, or exist currently only as proposals.

At the Data Layer, there is an urgent need to ensure that accurate, well-resourced information reaches AI systems at scale. This requires not just maintaining and improving current projects and resources. New content partnerships are needed, as well as targeted data collection through crowdsourcing (for example, Wikispeech datasets). Wikimedia can offer content partners a pipeline where data is structured, contextualized and made available through Enterprise – something

that is hard for other actors to replicate. Developing quality signalling infrastructure that provides machine readable data on the editorial state of content should also be a priority.

This pillar also includes the development of a space for experimentation and product development. Wikimedia's production-grade infrastructure is poorly suited for rapid AI experimentation. A complementary environment is needed: one that allows faster deployment, community-built tools, and access to adequate hardware.

A state-of-the-art public AI model is a prerequisite for the mission to be complete. Without it, commons-aligned interfaces have no commons-aligned model to route to, and the mission will become dependent on resources that are not fully aligned in terms of values, and that lack necessary transparency. Wikimedia cannot build such a model on its own, as frontier model development requires significant compute, engineering capacity, and capital at a scale beyond any movement entity. But Wikimedia is positioned to be a defining contributor to a collective public AI effort, particularly on the data side: shaping training corpora, contributing structured knowledge, setting reciprocity norms, and lending the institutional legitimacy that distinguishes public AI from its commercial counterparts. Even for a model built by partners; Wikimedia can ensure that it is commons-aligned. These public AI infrastructures could then be deployed to augment Wikimedia.

Enabling Condition 2
Mission Investment

The mission cannot be carried by any one organisation. A Knowledge Commons and AI Cluster is needed: a coordinated network of Wikimedia entities, volunteer communities, research institutions, and public AI development organisations. The Wikimedia Foundation and Wikimedia Deutschland bring core infrastructure. Smaller affiliates bring agility and policy connections. Volunteers bring legitimacy. External partners — open AI labs, researchers, other knowledge commons— bring technical capacity and resources.

Enabling Condition 1
Mission Network

The mission requires sustained, shared resourcing — financial, institutional, and political. Funding should follow a co-financing model involving the Wikimedia Foundation, key affiliates, and aligned partners from the public AI ecosystem, so that no single organisation bears the full cost or controls the direction. Institutional investment means committing staff time, legal expertise, and policy engagement as core priorities, not side projects. Political investment means movement leaders taking visible, independent positions in the external governance processes where the future of AI is being decided.

Enabling conditions

In addition to its two pillars, a mission approach requires coordination measures and investments that resource the mission. In this part both are outlined, showing how the Movement can align on achieving the goals of the mission.

MISSION NETWORK: COORDINATION THROUGH A WIKI AI CLUSTER

The Wikimedia Movement is not a single organization – it is a distributed system of interdependent actors, each with a specific role, capacity, and mandate. A strategy that relies on any one actor – however large or well-resourced – to carry the mission alone will fail to address the challenges that Wikimedia is facing.

Therefore, the Wiki AI cluster needs to be established as a coordination mechanism for a group of movement entities, external partners and volunteers. This cluster should be aligned on a positive vision of augmenting human collective intelligence with AI tools, and on the principles outlined above.

The cluster should build on what already exists, and on the diversity of capacities and resources that various stakeholders can bring. WMF and WMDE, the two largest movement organizations, hold the core infrastructure and engineering capacity.

Smaller movement entities can bring specialization, external partnerships, and the ability to move faster on experimentation. They can also contribute to governance work, community engagement, or connections to national AI ecosystems. Even if they lack the capacity to develop technical solutions, they can help shape ecosystem norms around responsible AI practices. These entities, many of them based in the Global Majority regions, also ensure geographical diversity of the cluster. And the technology roadmap should build on the **Global Majority Wikimedia Technology Priorities**, established in 2024 through a

process facilitated by Wikimedia Movimento Brasil.

Volunteer editors are the third core group in the cluster, which needs to develop capacity to support and scale volunteer work, with organizational support. The challenge is designing participation that is meaningful without being tedious. More generally, shifting to a more active, mission-oriented posture will require support from editors and communities. The cluster, driven by institutional actors, should closely collaborate with, and support those volunteers who want to engage in the mission.

Finally, the cluster should cooperate with aligned organizations. On one hand, it should include institutions that steward other knowledge commons. On the other, it should cooperate with AI development labs (for example, **Apertus**, **Pleias**, **Eleuther**) and individual researchers that can bring technical capacity and additional resources (for example, compute infrastructure).

The governance of the Wiki AI cluster – including questions of decision-making authority, resource allocation, and accountability between movement entities – will require dedicated attention.

MISSION INVESTMENT: RESOURCING THE MISSION

The Wikimedia and AI mission cannot succeed without deliberate investment. Governance work, technical development, and network coordination all require resources – financial, human, and institutional – that do not currently exist at the scale needed.

Mission investment has three dimensions. First, financial investment: sustained funding for the governance, technical, and coordination activities described above. This requires a co-financing model involving WMF, key affiliates, and where possible external partners from the public AI ecosystem. No single organization should bear the full cost, and no single funder should have sole influence over the direction of the mission.

The Wikimedia Endowment is best positioned to provide seed investment for the mission. The Endowment's funding strategy focuses on investments in areas where technological innovation intersects with access to information, and aims to strengthen the knowledge ecosystem and help Wikimedia remain resilient in face of rapid changes in the digital landscape. The Endowment should therefore support this mission, focused on the greatest challenge in Wikimedia's history.

Second, institutional investment: the commitment of organizational capacity – staff time, legal expertise, policy engagement – from movement entities. The mission cannot be resourced through grants alone; it requires that organizations treat it as a core strategic priority, not a side project.

Third, political investment: the willingness of movement leaders to take visible positions in external governance processes, including the bodies developing AI ecosystem protocols and global debates on AI (including the Geneva 2027 AI Summit process). The scope and mechanisms of mission investment will be further developed as the Wiki AI cluster takes shape.

Space to act: Switzerland and the 2027 Geneva AI Summit

The governance of the new knowledge loop will not be decided in a single place or moment – but some places and moments offer more opportunities than others. Switzerland is one of those places, and the next few years are one of those moments.

Switzerland has developed a distinctive approach to AI development: it is a country that is innovating in AI while simultaneously investing in the governance frameworks and public infrastructure needed to ensure that innovation serves the common good.

This is expressed in concrete institutional choices: the Swiss government's engagement with international AI governance and global standard-setting processes; and the Swiss Confederation's decision to fund **Apertus** (a fully open, multilingual large language model built by ETH Zurich, EPFL, and CSCS) as public infrastructure.

This creates an opportunity for Wikimedia CH to act as an early mover. Its role is not to lead the mission, but to open the space in which the movement can act.

Switzerland will host the **Geneva AI Summit in 2027** – which is an opportunity for Wikimedia to present a bold position on AI governance norms, and lead in proposing an AI development roadmap in service of collective intelligence and the knowledge commons. Wikimedia CH, as a Swiss organization embedded in a national ecosystem that combines strong research capacity with a balanced approach to AI governance, can participate in preparatory processes as a recognized institutional actor.

Launching the mission

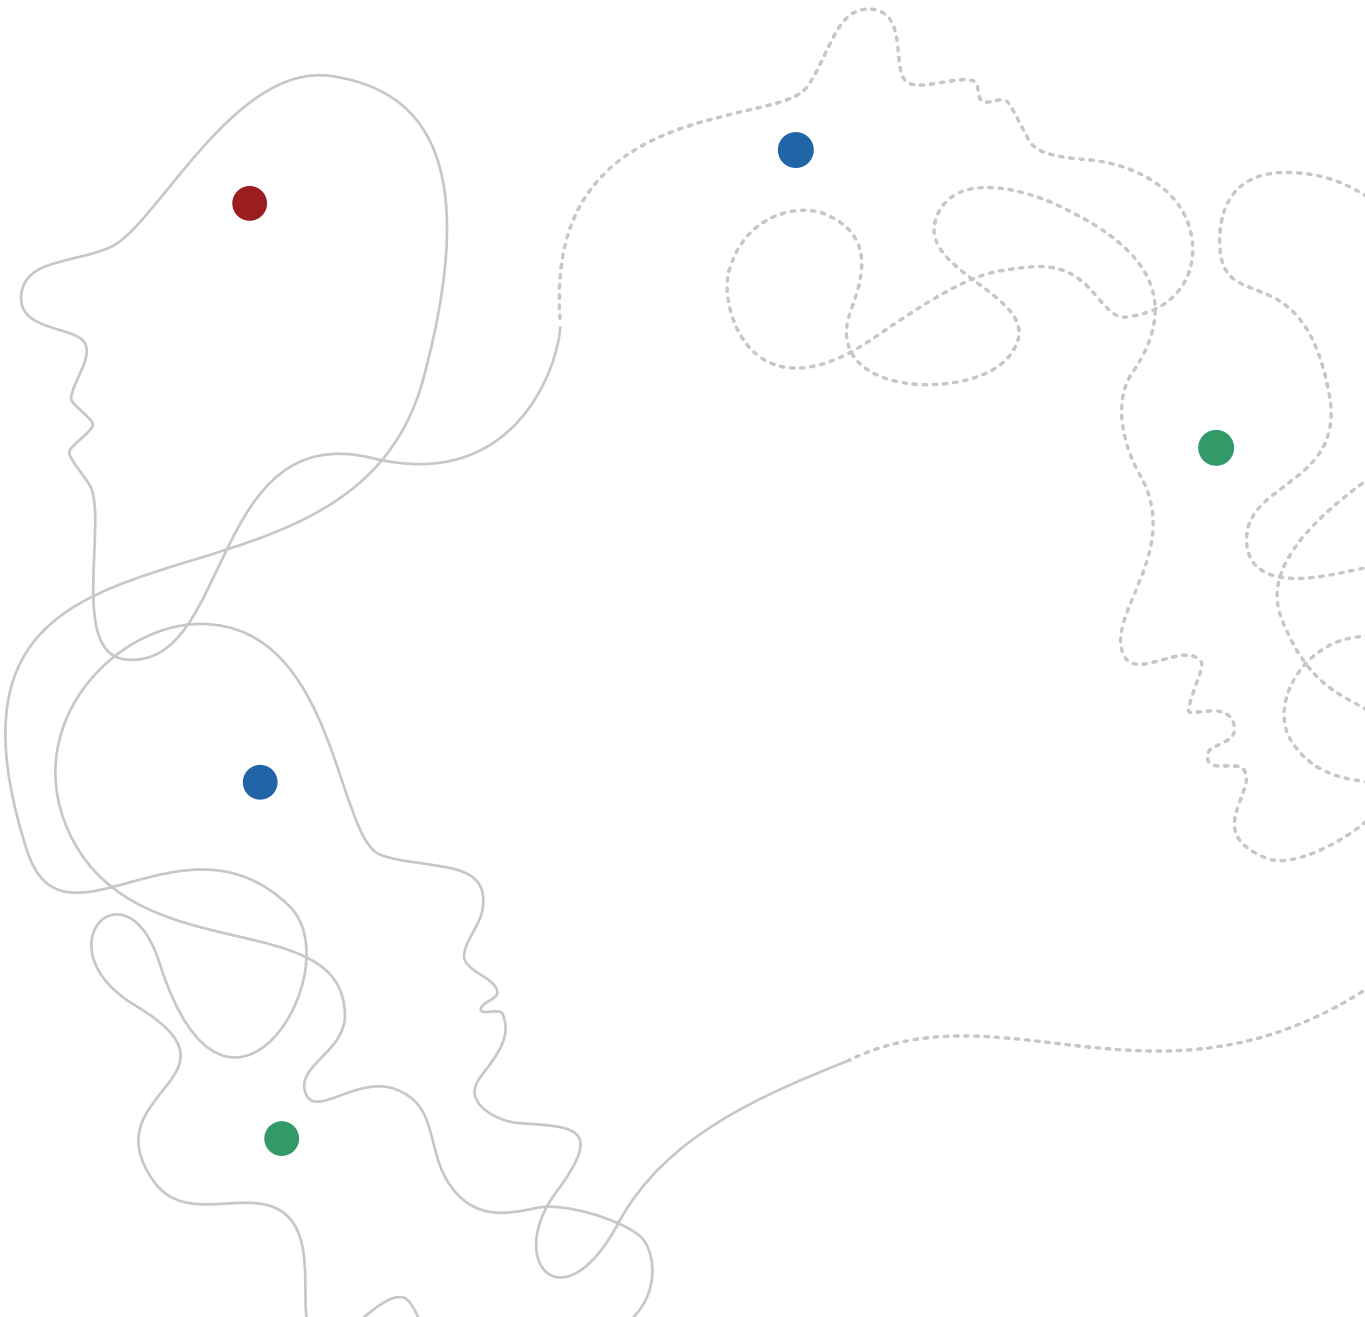

The framework and principles set out in this paper point toward a single proposition: **by 2030, Wikimedia should co-build and co-operate the AI systems through which the knowledge commons is produced and reaches its readers – so that the movement retains agency over how its knowledge is produced, presented, attributed, and renewed.**

No single Wikimedia entity or group of volunteers can carry it out alone. The immediate next step, therefore, is to form the Wiki AI cluster to serve as the mission network, and to secure seed funding. The priority in the first phase is to engage a critical mass of aligned actors, build a shared understanding of the challenge, and agree on what coordinated action to take. This paper should be read as the start of that process, not its conclusion. It proposes a direction; the coalition that forms around it will specify the detail.

If this does not happen, the third loop settles into its current trajectory: zero-click AI answers, declining editor recruitment, attribution and reciprocity stripped, the dark commons sustaining an information economy that no longer sustains it.

The window for establishing a sustainable position for Wikimedia in the new knowledge loop is currently open. The interface layer of AI-mediated knowledge is consolidating, and the compute and partnerships that make commons-governed alternatives possible are being set now. The window will not remain open indefinitely.

If the mission succeeds, Wikimedia will end this decade not as a dark commons feeding systems that it does not influence, but as a commons in the full sense. Visible in the key interfaces of the new web, having control over terms of its reuse, more equitably present around the world, and generating new knowledge.

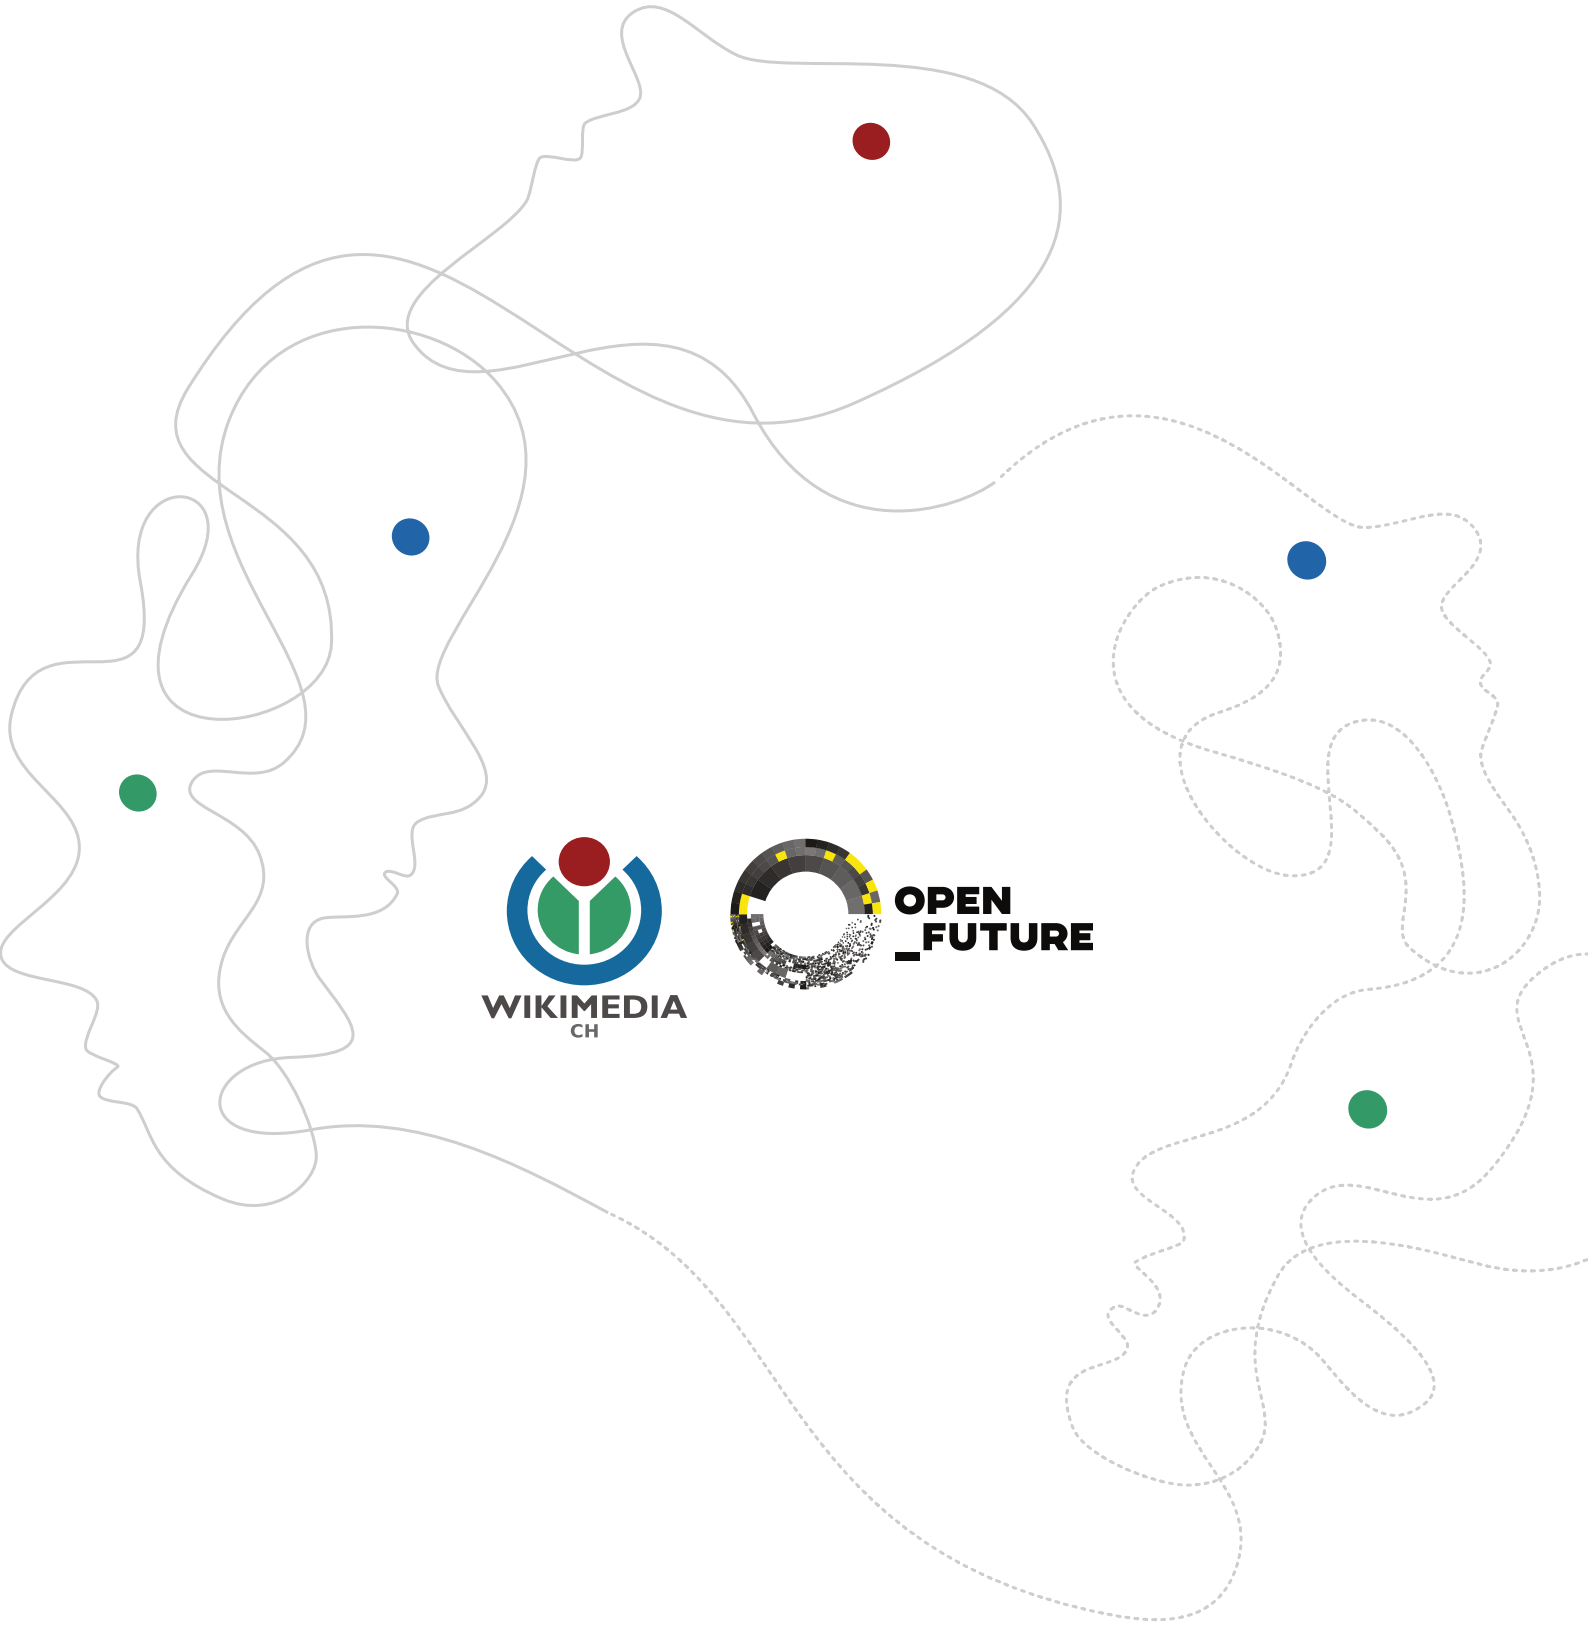

Design: zerogravita.com