The Challenge of Generative AI
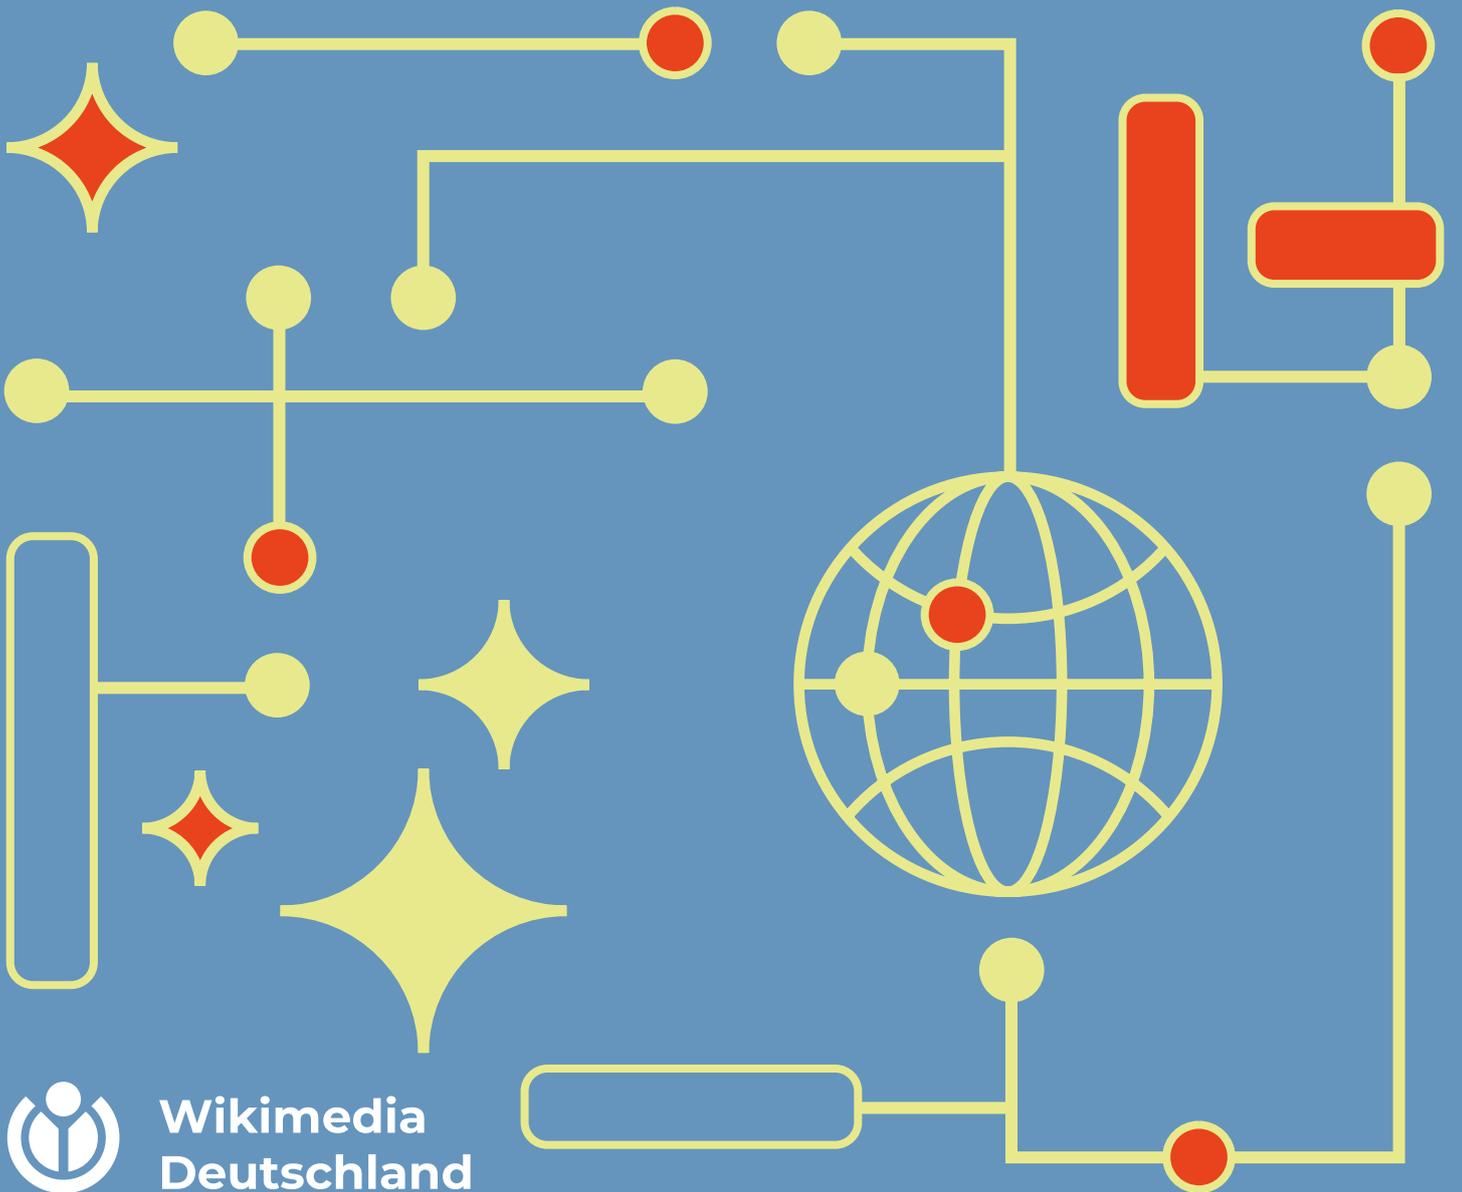
Opportunities and Risks for Free Knowledge

**Wikimedia Deutschland**

Contents

Foreword

Generative AI is transforming knowledge and information ecosystems as well as the way we find, consume, and develop knowledge. This does affect wiki projects, but the situation is complex. While the threat seems clear in some respects, such as the abbreviated presentation of content by chatbots or the hallucinating of entire discourses and their sources, AI chatbots can also facilitate translation within Wikipedia or research work for articles. Discussions are in full swing within the Wikipedia community on these and other AI-related aspects.

Open-source communities are also addressing the topic of generative AI. Since 2024, Wikimedia Deutschland has been running the Embedding Project, which aims to enable and facilitate the development of more ethical AI models using Wikidata. The goal of the Embedding Project is for us to shape technological change together with our communities, rather than just watching generative AI transform the knowledge landscape.

With this in mind, this brochure not only highlights the risks for wiki projects but also presents solutions from the communities and Wikimedia Deutschland, as well as our recommendations for educators. We intend to actively shape this technology and not just be shaped by it.
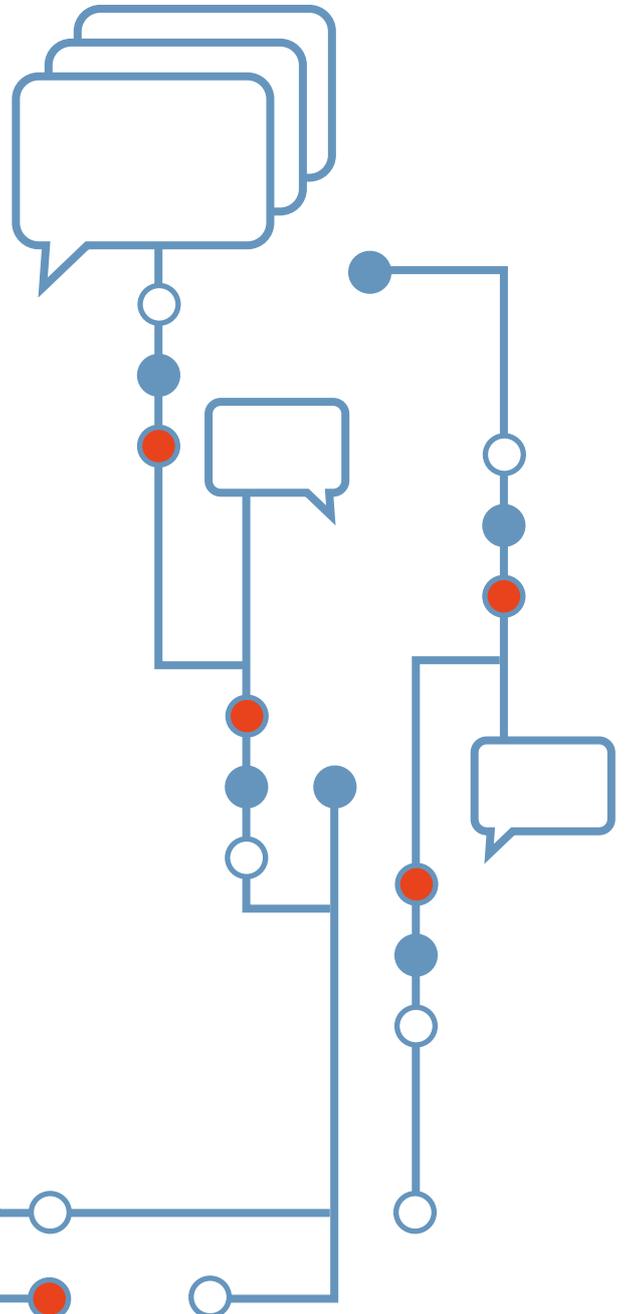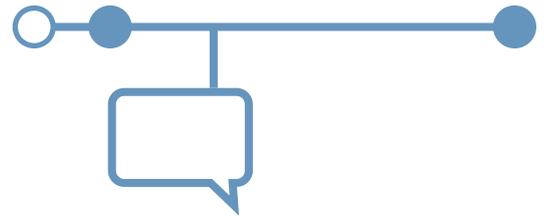
1. Wikipedia and generative AI

Artificial intelligence is changing how and where people search for and find information online—and how we share knowledge about the world. What does this mean for the future of Wikipedia? Is the AI hype a passing fad or a lasting challenge? Should the use of chatbots be permitted or prohibited on Wikipedia? This is a controversial topic in the German-speaking community, just as it is in all other wiki communities worldwide. At Wikimedia Deutschland, we are accompanying this process and helping to shape it.

Knowledge infrastructure
Wikipedia as a data provider — for fair conditions

Familiar large language models (LLMs) like ChatGPT or Google Gemini are largely trained using data from Wikipedia. This is not least due to the fact that the online encyclopedia offers ideal conditions for these LLMs. The structure of the articles is clear: they contain links and citations, provide context, and are available in over 300 language versions.[1] The collected knowledge is moderated by humans, accessible worldwide, and, above all, free of charge. This is a major advantage from the perspective of the corporations. However, massive web scraping poses a risk to Wikimedia organizations. Automated access by bots and crawlers places an enormous load on servers, consumes a lot of resources, and thus compromises the stability of the platform for human users. While Wikipedia can easily cope with short-term spikes in traffic—such as during current events—the constant, massive number of requests from bots poses a structural problem. This results in a lack of resources which should actually be benefiting volunteers and users of Wikimedia projects.[2]

For this reason, the Wikimedia Foundation (WMF) launched a paid service for large commercial users of Wikimedia content several years ago, namely the Wikimedia Enterprise API. [3] This paid interface bundles existing, publicly available data from Wikimedia projects in a way that makes it easier for commercial companies to reuse it. At Wikimedia Deutschland, we also believe Free Knowledge requires fair conditions.

Knowledge consumption
Changing reading habits

Wikipedia is not just a data pool used to train AI. It is also one of the most frequently cited sources in generative AI responses[4]. ChatGPT, OpenAI's chatbot, in particular, prefers to use it as a reference.[5] The problem is that an increasing number of people are satisfied with the summaries provided by ChatGPT or the Google AI assistant when searching for information. They visit the websites listed as sources just half as often if they have read the AI overview beforehand.[6] Between mid-July and mid-August 2025 alone, ChatGPT saw a 52 percent drop in referral traffic.[7] Many websites including those of NGOs, for example, are at risk of losing visibility entirely as a result of this development. Critics see this as a threat to diversity in civil society.[8]

The development also costs Wikipedia readers. Current studies show that young people in particular are now decreasingly reading Wikipedia directly. The proportion of users in the 18 to 24 age group worldwide has recently fallen significantly.[9] A secondary aspect is that, when measuring page views, it is now necessary to make a complex distinction between bots and human readers.[10] However, the trend is clear: the number of human visits is declining.[11] In addition to AI, there are other reasons for this, such as increased use of social media for information purposes as well as shorter attention spans.

New approaches are required to maintain Wikipedia as the most important source of freely available and verified knowledge, while also continuing to inspire young people to read and participate. Developing these is one of the main tasks for the future. Wikimedia Deutschland is tackling this challenge together with dedicated volunteers—among other things, by creating discussion forums, such as the Wikimedia Futures Lab in January 2026.[12]
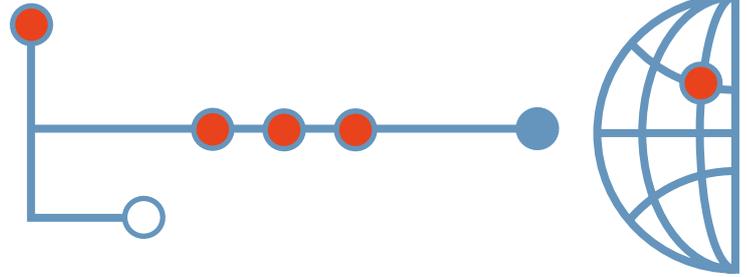
Knowledge production
AI discussions in the German-speaking community

If using chatbots is now a matter of course, especially for the younger generation, shouldn't their use also be legitimate when authoring articles? Few other topics related to AI and Wikipedia are as controversial as the use of AI tools for editing. The advancement of generative AI models is making it increasingly easy to generate text that at least seems to replicate the quality of Wikipedia text.

The risks are obvious: AI models are still unreliable. They often create hallucinations—inaccurate, unsubstantiated or even fabricated statements. Including this type of text in Wikipedia without careful review could damage the credibility of the encyclopedia as a whole. If language models were then to be trained using incorrect Wikipedia texts, it would set in motion a spiral of distorted knowledge, ultimately leading to a complete loss of trust.

The community's cooperation and reciprocal monitoring will continue to be necessary. As things stand today, generative AI is capable of summarizing existing human knowledge. It cannot, however, participate in the process of debate and consensus building. Neither does it check its sources, discover objects buried in archives, or take photos of insufficiently documented locations. These are all things that volunteers working on Wikimedia projects do every day.

In a community vote in February 2026 (called an "opinion poll" on Wikipedia), a clear majority of participants voted in favor of strict limits and clear rules for AI-generated content in the online encyclopedia. 208 Wikipedians voted in favor of a complete ban on texts generated or edited with large language models such as ChatGPT, while 108 voted against it. "All contributors are expected to write their own encyclopedic texts," is what was eventually agreed upon by the community.[13]

However, the use of AI tools for error detection remains permissible. Machine translations (for example, of articles from other language versions of Wikipedia) will also be allowed, provided that the text and references are carefully checked for accuracy.[14]

A solution similar to the „AI Cleanup" project[15], which has been active on the English Wikipedia since December 2023, is being launched on the German-language Wikipedia. It consists of over 100 volunteers systematically tracking down and deleting erroneous posts made by LLMs.

Against this backdrop, the Wikimedia Foundation, which operates Wikipedia, has developed an AI strategy entitled „Humans First"[16], which Wikimedia Deutschland also supports. No AI can replace the dedication that volunteers have shown for 25 years in their meticulous efforts to provide reliable encyclopedic knowledge. Accordingly, LLMs should facilitate the work of volunteers rather than replace them. The motto of the „Humans First" initiative is „Making sure AI serves people and knowledge stays human."
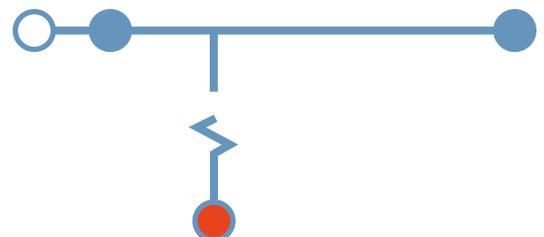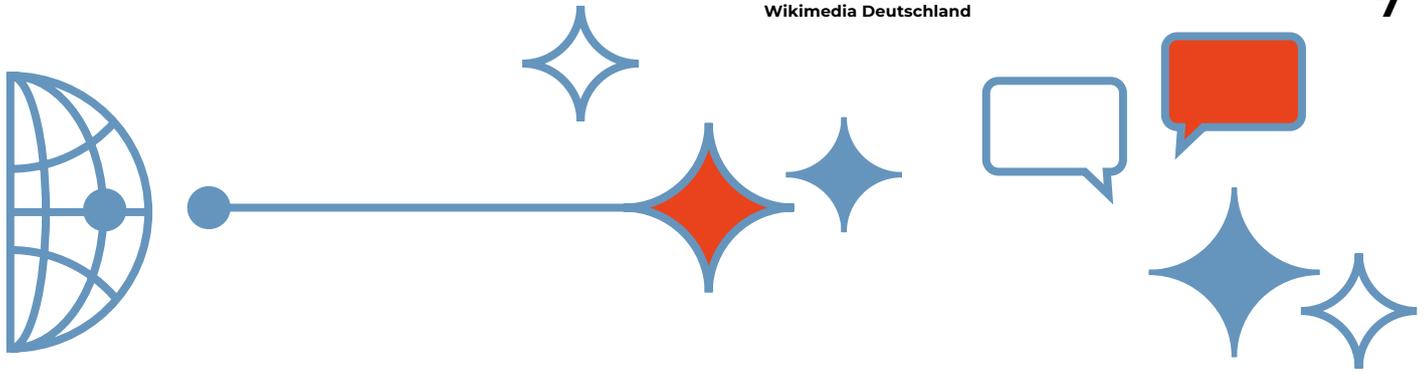
Knowledge manipulation
Risks posed by generative AI

One possible threat scenario involves attacks on Wikipedia content using generative AI. An example presented at the Wikimania conference in 2025 showed how ChatGPT rewrote the English Wikipedia article on the Russian invasion of Ukraine using a targeted prompt in such a way that the attacked country was blamed for the war. Performing this type of manipulation at scale would jeopardize the credibility of Wikipedia as a whole. AI is fast and inexpensive. Human corrections take time.

For this very reason, the German Wikipedia now also has a so-called quick deletion policy.[17] Articles can be removed immediately without lengthy discussions about deletion if they are recognizable as having been created by generative AI without human review, or if the evidence is implausible, clearly hallucinatory or misattributed.

A number of effective protective mechanisms against attempted manipulation—whether by humans or language models—were already in place. These include pages with unreviewed versions, where Wikipedians who specialize in this area can specifically search through new changes.[18] Another example is the watchlist candidates page.[19] It contains a long list of articles about people and topics that are particularly vulnerable to manipulation. AI tools are now also helping identify edits that were most likely created with the help of chatbots like ChatGPT.[20] Among other things, the community project AI and Wikipedia[21] lists such tools for checking sources or combating vandalism. Several volunteers from the German-speaking community have taken on the task of identifying and deleting incorrect AI-generated articles. They meticulously check sources and proof such as ISBN numbers and use abuse filters to detect suspected AI edits.[22] In this way, they ensure that Wikipedia remains a reliable, human-curated source of knowledge.

Knowledge equity
Potential for AI use on Wikipedia

The potentials and risks of human-AI cooperation are viewed differently depending on the country and Wikipedia community. For example, a tool called WikiVault[23] has been in use on the Korean Wikipedia since April 2025. It enables the translation and creation of articles that are optimized for the Wikipedia style using the support of AI. The tool is very popular in the community. This does not mean that the use of WikiVault is automatically recommended for other Wikipedia language versions. WikiVault inventor Ykhwong even emphasizes, „WikiVault is not a tool that replaces editors. Rather, I envision this tool as a starting assistant—something that makes beginning easier, a guide for taking that important first step in writing a Wikipedia article."[24]

In this context, there is also discussion as to whether smaller language versions of Wikipedia, which contain only a few articles due to the limited number of native-speaking volunteers, could benefit from AI contributions. Wouldn't it contribute to greater knowledge equity if language models were responsible for the growth of these editions? The case of the Greenlandic Wikipedia provides a counterargument. Its administrator, Kenneth Wehr—himself not a native speaker—had to put the entire project up for discussion because the output was increasingly flooded with incorrect AI-supported translations.[25] The Greenlandic Wikipedia has since been shut down.

A member of Kannada (a language in South India) Wikipedia has, however, clearly recommended AI-assisted article creation. The author, Pavanaja, does point out the limitations of generative AI and the time-consuming process of fact-checking. On the other hand, it is deemed a „game changer."[26] However, Pavanaja does not want to leave the Kannada Wikipedia solely to AI. Its growth must be guided by human expertise, „ensuring that every article is not just generated, but meticulously curated and referenced. The future of knowledge is collaborative, and AI is simply our newest collaborator."

2. Open education and AI

The „humans first" strategy also includes the demand to prioritize open-source or open-weight models when using generative AI—i.e., open or at least partially open systems, whose algorithms are transparent. This requirement is particularly important in education policy. The use of AI applications has long been part of everyday life in German schools and universities. However, the AI systems used in this context are predominantly proprietary systems from large commercial providers—language models, the databases, functionality, and algorithms of which cannot be independently verified. This poses a problem especially with regard to data protection, which is particularly important in the sensitive context of education. Wikimedia Deutschland therefore supports open AI in education. We are committed to an education system that focuses on participation, transparency and the common good. This is particularly important when it comes to the use of generative AI.

Education policy
Recommended actions to increase control of data

Open AI solutions fundamentally differ from the models used by large tech companies. They are transparent, compliant with data protection regulations and can be adapted to the needs of schools and universities. Educational institutions retain control over data, content and applications. The education policy team at Wikimedia Deutschland is working on setting the course for AI systems that are open and oriented toward the common good. Among other things, it intends to achieve this by raising awareness among political decision-makers about the importance of this issue. To this end, we have launched a series of events with the title „Forum Open AI in Education".[27] Questions were discussed in workshops, expertise was pooled and a wide range of perspectives
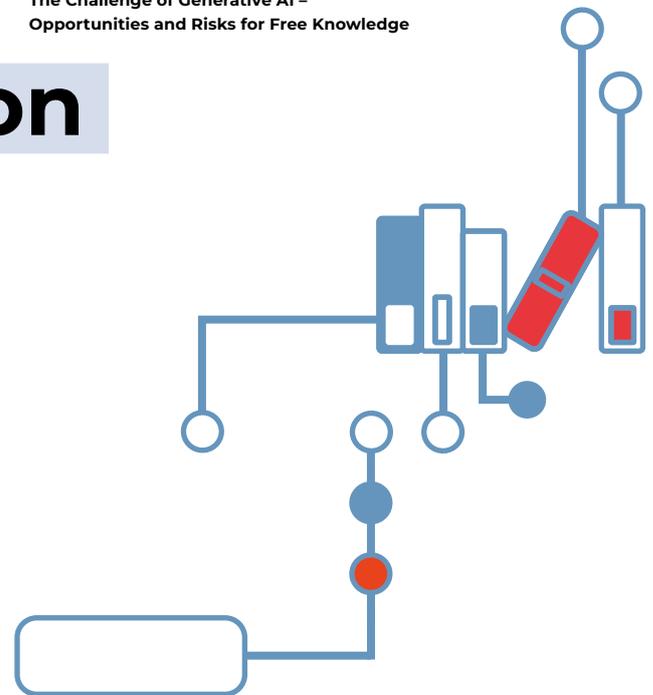
was taken into account. The result is ten concrete recommendations for action in education policy for AI solutions that are open and oriented toward the common good. We have summarized these recommendations in the publication „Offene KI für alle!" (Open AI for All!).[28] They are also available as wiki content.[29] The recommendations primarily relate to infrastructure and access, open educational practices and fundamental rights in the digital space.

Educational practice
How the recommendations can be implemented

To address the question of how these recommendations can be implemented in schools, we have launched the nationwide conference series „Offene KI in der Schule" (Open AI in Schools)[30] in collaboration with the Lower Saxony State Institute for School Quality Development. Over the 2025-2026 period, experts from the fields of educational practice, policy, administration, business and civil society are developing concrete measures to integrate open AI solutions into the school system in a practical, sustainable and nationally compatible manner. Emphasis is being placed on three key areas: the legal framework (particularly with regard to the EU AI Act and the General Data Protection Regulation), technical implementation (required infrastructure and systems) and sustainable training formats for teacher qualification.
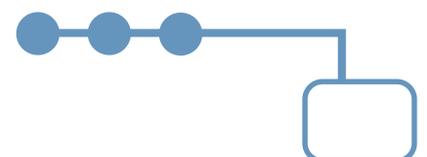
3. The Wikidata Embedding Project – for ethical AI
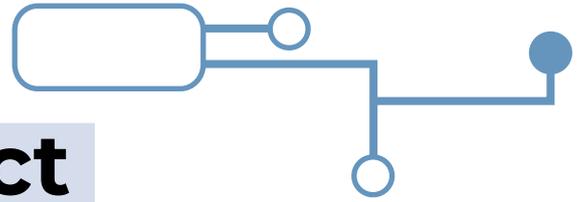
Our commitment to fair and open AI solutions is not limited to the field of education. Our goal is to make the generative AI ecosystem more accessible overall. One example of this is the Wikidata Embedding Project.[31] Wikimedia Deutschland has teamed up with the following partners for this project: DataStax, an IBM company and provider of AI and data solutions, and Jina AI, a Berlin-based AI-powered search specialist. Together, we have created the technical conditions necessary to make the high-quality data from Wikidata usable for AI applications—freely accessible to everyone. Not only tech giants but also open-source initiatives can use this as a foundation for developing AI solutions that are verifiable, fair and for the common good.

Wikidata is an impressively large collection of knowledge. It consists of over 119 million structured data sets (as of December 2025)[32], which are both understandable to humans and machine-readable. Wikipedia also accesses this data to automatically update information such as population figures and dates of birth. Over 12,000 volunteers around the world work to verify, update, and expand this database. As a source of knowledge, Wikidata has long had an enormous influence on our everyday lives. Digital assistants like Siri and Alexa, for example, would not exist in their current form without the immense amount of information provided by Wikidata.

Wikidata as a source of facts
How the Embedding Project works

The Embedding Project converts the open knowledge from Wikidata into a vector database—a step that open-source developers usually cannot manage on their own. DataStax provides a powerful vector database for this purpose, while Jina AI supplies an open-source model for vectorizing the text data.

The conversion of data into vectors allows developers to perform semantic searches more efficiently and integrate Wikidata data into AI models. This enables carrying out faster and more precise searches including queries in natural language. It also facilitates integration of Wikidata into so-called RAG (Retrieval-Augmented Generation) applications. These applications minimize AI errors by not only accessing training data for search queries, but also incorporating current and verified facts into their results. Generative AI can use RAG to directly access verified data from Wikidata, thereby reducing incorrect answers and hallucinations.

What will now be possible

There are many practical applications for the Embedding Project. Developers can use it as a basis, for example, to develop a chatbot that works similarly to ChatGPT, but that uses Wikidata's huge data pool for truly fact-based answers. It would also be conceivable to develop a data visualization tool that uses vectorized data and is controlled by natural language commands to filter out all relevant information from Wikidata without errors. Another goal of the project consists of better detecting vandalism on Wikidata. Vectorizing the data makes it possible to identify and correct potentially harmful changes to entries more quickly.

Wikidata as a model for openness
A fairer digital future

With the vector database, developers can identify Wikidata as a source, which is transparent for everyone. The source code is also available under an open license. Wikidata is maintained and expanded by an active community on a daily basis. This means that the results of generative AI queries can be more up-to-date than those from systems that only draw on their trained „knowledge." And last but not least, thanks to its international community, Wikidata also covers underrepresented topics and perspectives, thus creating a diverse, multilingual database for the development of generative AI.
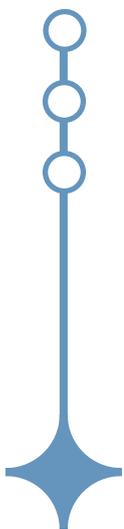
The Embedding Project strengthens open-source projects against the dominant tech giants. Simplified access to reliable data for everyone means democratizing the digital world based on fundamental values such as openness and transparency. This not only advances the developer community, but society as a whole.

Outlook
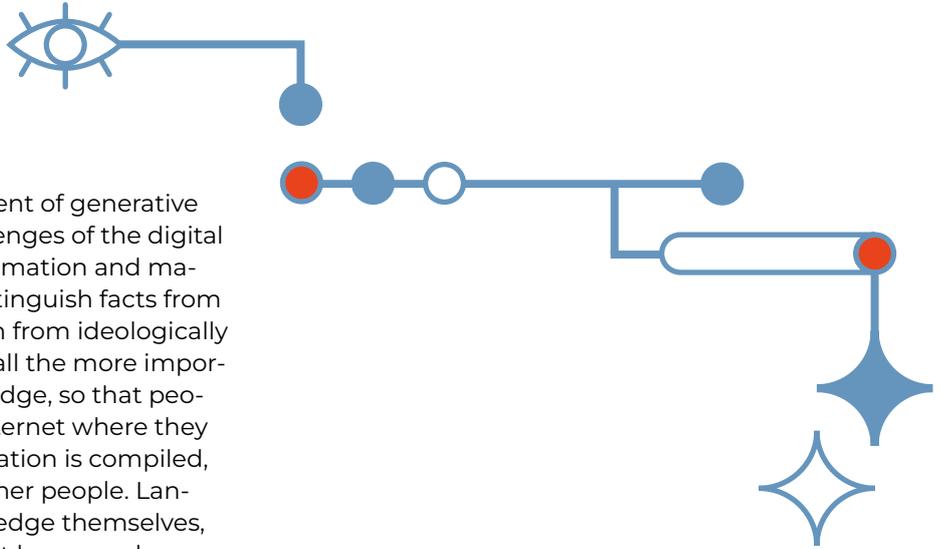
It is not yet clear where the development of generative AI will lead in the future. But the challenges of the digital age are likely to increase. With disinformation and manipulation on the rise, how can we distinguish facts from hallucinations and reliable information from ideologically biased claims? This makes Wikipedia all the more important as a free source of reliable knowledge, so that people continue to have a place on the internet where they can find valid information. This information is compiled, discussed, reviewed and verified by other people. Language models do not generate knowledge themselves, but rather draw on the knowledge that humans share.

That is why, now more than ever, we need as many people as possible to get involved in Wikimedia projects. This explicitly includes the younger generation, who must be invited to help shape the digital future. We will continue to discuss the unresolved issues surrounding the development of generative AI in close collaboration with the community. Who is involved in the development of generative AI? What biases and privileges does it reproduce? How can data and knowledge equity be achieved? No chatbot can provide the answers. Knowledge is and will remain human.

1   https://meta.wikimedia.org/wiki/Statistics
2   https://diff.wikimedia.org/2025/04/01/how-crawlers-impact-the-operations-of-the-wikimedia-projects/
3   https://enterprise.wikimedia.com
4   https://t3n.de/news/ki-quellen-analyse-google-reddit-1703438/
5   https://www.tryprofound.com/blog/ai-platform-citation-patterns
6   https://netzpolitik.org/2025/google-menschen-klicken-halb-so-oft-auf-links-wenn-es-eine-ki-zusammenfassung-gibt/
7   https://t3n.de/news/referral-traffic-chatgpt-websites-1703629/
8   https://www.heise.de/hintergrund/Wie-KI-Zusammenfassungen-zivilgesellschaftliche-Vielfalt-einschraenken-10668166.html
9   https://meta.wikimedia.org/wiki/Research:Knowledge_Gaps_Index/Measurement/Readers_Survey_2024#Age_2
10 https://diff.wikimedia.org/2025/10/17/new-user-trends-on-wikipedia/
11 https://www.404media.co/wikipedia-says-ai-is-causing-a-dangerous-decline-in-human-visitors/
12 https://meta.wikimedia.org/wiki/Wikimedia_Futures_Lab
13 https://blog.wikimedia.de/2026/02/19/meinungsbild-zum-einsatz-von-ki-in-der-wikipedia/
14 https://de.wikipedia.org/wiki/Wikipedia:Belege#Was_sind_zuverlässige_Informationsquellen?
15 https://en.wikipedia.org/wiki/Wikipedia:WikiProject_AI_Cleanup
16 https://wikimediafoundation.org/news/2025/04/30/our-new-ai-strategy-puts-wikipedias-humans-first/
17 https://de.wikipedia.org/wiki/Wikipedia:Schnelll%C3%B6schantrag
18 https://de.wikipedia.org/wiki/Spezial:Seiten_mit_ungesichteten_Versionen
19 https://de.wikipedia.org/wiki/Wikipedia:Beobachtungskandidaten
20 https://de.wikipedia.org/wiki/Spezial:Missbrauchsfilter/453
21 https://de.wikipedia.org/wiki/Wikipedia:WikiProjekt_KI_und_Wikipedia
22 https://de.wikipedia.org/wiki/Spezial:Missbrauchsfilter/453
23 https://en.wikipedia.org/wiki/Wikipedia:WikiProject_Korea/WikiVault
24 https://diff.wikimedia.org/2025/07/15/introducing-wikivault-a-new-chapter-in-wikipedia-contributions-with-ai/sd
25 https://meta.wikimedia.org/wiki/Proposals_for_closing_projects/Closure_of_Greenlandic_Wikipedia
26 https://diff.wikimedia.org/2025/08/14/creating-kannada-wikipedia-article-using-generative-ai/
27 https://www.wikimedia.de/forum-offene-ki-bildung/
28 https://www.wikimedia.de/publikationen/offene-ki-fuer-alle-10-handlungsempfehlungen-fuer-offene-ki-technologien-im-bildungsbereich/
29 https://meta.wikimedia.org/wiki/Offene_KI
30 https://www.wikimedia.de/offene-ki-in-der-schule/
31 https://www.wikidata.org/wiki/Wikidata:Embedding_Project
32 https://www.wikidata.org/wiki/Wikidata:Statistics/de

About Wikimedia Deutschland

Wikimedia Deutschland is a non-profit association with over 111,000 members and 180 employees that is committed to promoting freely available knowledge in the digital space. As the largest national representative of the international Wikimedia Movement, the association promotes the volunteer Wikipedia communities and other Wikimedia projects in Germany. Wikimedia Deutschland develops and maintains free software and the free knowledge graph Wikidata. The association is committed to promoting digital and education policy frameworks that enable free access to knowledge and data. We also cooperate with cultural institutions to make more cultural heritage freely accessible.

All links in this publication were last accessed on 3. March 2026.