

5-10 years from now, what is your vision? What's different about Wikimedia?

Personally, I would like to see Wikimedia in the top five for reach in every country. I want to see a broad, rich, deep encyclopedia that's demonstrably meeting people's needs, and is relevant and useful for people everywhere around the world.

In order for that to happen, a lot of things would need to change.

The community of editors would need to be healthier, more vibrant, more fun. Today, people get burned out. They get tired of hostility and endless debates. Working on Wikipedia is hard, and it does not offer many rewards. Editors have intrinsic motivation not extrinsic, but even so, not much is done to affirm or thank or recognize them. We need to find ways to foster a community that is rich and diverse and friendly and fun to be a part of. That world would include more women, more newly-retired people, more teachers – all different kinds of people. There would be more ways to participate, more affirmation, more opportunities to be social and friendly.

We need a lot of tools and features to help those people be more effective. Currently, there are tons of hacks and workarounds that experienced editors have developed over time, but which new editors don't know about, and would probably find difficult to use. We need to make those tools visible and easier to use, and we need to invent new ones where they are lacking. Imagine being able to tell at a glance where another editor is from and how experienced they are. Imagine being able to get help from friendly editors in real-time, whenever you are stuck. Imagine having watchlists with way more functionality than we have today – which point you towards new work you've self-identified as wanting to do. And so on. We need to make it easier for people to be productive and powerful.

We need a robust platform. We need to be available and easily usable on mobile devices. We need to figure out how to enable people using mobile phones to participate with us.

We need to enable flexibility around look-and-feel, so that people can have a Wikipedia with a skin that makes sense to them – their culture, their language, their age and gender and sense of personal style. Currently the Korean Wikipedia looks very, very different from other Korean websites: I believe that could be damaging to usage and participation in Korea.

What's the most controversial part of the vision you just described?

Diversity of participation.

Wikipedia is successful because it is open to participation by anyone. That is what makes it strong. And it is theoretically true that anyone can participate in Wikipedia.

But in practice, the reality is that anyone can *not* participate in Wikipedia. In order to be a constructive useful Wikipedia editor, there are a bunch of prerequisites, basically, essential skills and experiences. You need to be moderately tech-centric, and you need to have overcome various usability barriers and learned wiki syntax. You need to have read and absorbed a bunch of rules. You need to really care, because you wouldn't bother if you didn't. And you *couldn't* do it—you wouldn't be capable—if you weren't smart and thoughtful and capable of intense focus. So there's a self-selection bias, in that the only people who successfully overcome the hurdles to learn to edit Wikipedia are by necessity really smart and really caring.

However, there are millions of smart, caring people in the world who don't edit Wikipedia, just because they don't have time to learn all the rules, or they are afraid of wiki syntax, or they don't meet one of the other criteria I talked about. We don't want arbitrary, non-deliberate prerequisites to keep those people from editing. We want them to edit Wikipedia. If they do, Wikipedia will be better, stronger and smarter, in large part because its editor base will be broader and more diverse. I don't know all the groups or types of people who are systemically underrepresented inside Wikipedia by, for example, the requirement to be a little tech-centric, but I am sure, for example, that it includes women and older people.

The problem is that anything that makes it easier to edit Wikipedia will make it easier for *anyone* to edit Wikipedia. And I know that worries the core editing community. The core editing community experienced a massive influx of new contributors over the past five years, and many of those contributors were horrible. Pranksters, little kids, self-promoters. I think that experience really profoundly shaped the core community of editors, and one effect was to leave people afraid of opening the floodgates up to new people, for fear that good editors will drown in bad edits. Nonetheless, my belief is that we need to encourage new participants, to ensure a vibrant, healthy community of contributors, and to mitigate against systemic bias.

Are there any communities or examples we could learn from?

There are things we could learn from the Linux community, for example, about making it *fun* to participate. Linux has done a lot of meet-ups, beer-drinking, social events, fostering of team—there's fun.

But Wikimedians are an earnest, serious-minded group, and while they do meet, it's not very visible unless you're actively looking for it.

We also have a lot to learn from successful social networking sites. I know this is controversial. Our job is not to create a space for people to have fun. Our goal is not participation for its own sake. It's not to amuse people and to make our site "sticky." We do not have the normal incentives for doing that; we are not monetizing eyeballs. Having said that, we can borrow technical tools, features, and so on to make our site more engaging and to make our site easier to use. If you're interested in X, you might like Y. If you like copy-editing, here are 10 articles you can work on. If you live in Mainz, there is a meet-up there this weekend.

Wikimedia could also learn a lot from traditional journalistic enterprises. A lot of its structures have developed organically to resemble traditional newsrooms, and it could be that more study of this could be useful to find out what hasn't developed naturally, but could still be useful.

We can learn a lot from other organizations with an international focus. For example I talk with people from eBay's international operations. There is a lot we can learn from them. We can learn from internationally-active non-profits, about how they fundraise internationally. That kind of thing.

What do you think is holding the projects back today?

I believe the policy proliferation is a really serious problem, and I don't know what the answer is. Part of the culture is you add, but you don't remove. You add, edit, refine, but you don't do a lot of removing, synthesizing, simplifying. I hunger for a structure of volunteers with authority to synthesize and simplify policies. It's tough because every policy is there for a reason, and most have value. But they really deter newcomers: the deck is stacked against you. Something has to be done to reduce and simplify policies.

Another piece is the general feeling of welcoming. The kind of people who want to build an encyclopedia in their spare time are fact-based, writing-focused people. They tend not to be particularly warm or extroverted. Obviously the answer is not to drive out the introverts, but it is rather to additionally pull in new people. In a task-oriented culture like ours, to just say "Thank you" or "I agree" on a mailing list is a controversial act because warmth is not valued and "getting on with the job" *is* highly valued. The lack of welcoming behavior, the lack of kindness and friendliness, is a vicious circle that seriously deters some types of people. Again, for example, women. Girls are socialized to be highly sensitive to social cues and to care what other people think of them. Not all women ultimately do end up caring, of course, but

many do. For example, I consider myself reasonably thick-skinned, but if I had come to Wikipedia as a volunteer editor, I would probably have left due to the sometimes hostile atmosphere.

What are your priorities for quality content?

Two things: One is related to individuals. We need to broaden our contributor demographic because that will increase quality. The end goal is not participation; the end goal is better quality. Here, I'm interpreting quality as relevance, breadth, utility. I know this from my experience as a journalist. If the group of people who decide what the news is, is a completely homogeneous group—if, for example, it is all 25-year-old western European men—then the output of the newsroom will be most relevant and interesting to that group. Therefore, if we want everyone to be reading Wikipedia, we need a wide diversity of people creating the material. Not just gender and geography and age, but an array of different skills, too. We need great editors to restructure and smooth out and tighten the writing. We need copy editors. We need people who can make great illustrations, take great photos, add citations. And so on.

There's also the institutional piece. That is about partnering with other organizations to get content we don't have yet. The German chapter has made good progress here, and so have some others. In my view—and this is just my view—the Wikipedia projects will always be fundamentally text-based. That's because text is manipulable. It's immediately updatable, it's easy to collaborate and remix and reuse and share. But I believe the projects will always be supplemented with other media: video, graphics, and so on. In my perfect world, a really rich article is supplemented by multiple related videos, by fantastic photographs, by a historic map, by a really great diagram. There is a rich palette of options. That transformative change will come from partnerships with galleries, libraries, museums, archives and media organizations. There is a very good fit between them and us. Their goal is to achieve distribution of their material, to reach people. In many cases, particularly for educational institutions, they have a hard time reaching young people in particular, and they have a particular mandate or desire to reach the young. Walled gardens don't work for them, locking up the content does not work. Meanwhile, we have the audience: 350 million people visit the Wikimedia projects every month. They should release their material under a free license, and let Wikimedia editors incorporate it into our projects.

Obviously we're in a massive period of disruption, and the traditional business models, especially for media, are crumbling. They don't know yet what will replace them; nobody does. Their best bet right now is to be doing massive experimentation and learning, and I believe sharing content with us should be one of their experiments.

We've heard that, in the short-term, the desire to be entirely open (in terms of copyright) might actually restrict the type of content that comes in. But in the long-term, Wikimedia may have a role to play in bringing more content into public use if it stands firm and pushes licenses like Creative Commons. Where do you stand on this issue?

I was talking to my board earlier this year, and we were discussing using Flash-based video playback. Right now, there's hardly any video anywhere on our projects. At the time, I made the argument to the board that we should consider non-free video players and formats, at least as a secondary option, because they are much more widely used than the fully open source alternatives. My view was, our job is to distribute information, and if we can't distribute, then we are not doing our job. I believed that successfully getting useful, educational, instructional video to end users of the projects, was more important, and more core to our mission, than advancing the goals of the free software movement.

The Board said back to me, we can change the world. We can change this dynamic. So we decided we would focus our energies on strengthening support for the free Ogg Theora format. And Firefox 3.5 shipped with built-in support for Ogg Theora. That changed my mind. We have enormous clout, and we can use it. Every week, thoughtful Wikipedians have long conversations with museum administrators and persuade them they should adopt more open licenses. It's water on a stone, and it's a long game.

There's some risk on that. If the game is too long to play, and we don't get video for 20 years, we may become irrelevant and clunky. But at this point, we're in a strong enough position that we can afford to be principled. We feel an obligation to the free content movement, as well. And there's no question that we have more leverage than anybody else.

What's your perspective on offline distribution?

The offline challenge will go away in time, although nobody knows how long that will take. It's a mid-term challenge, not a permanent problem. It's not our job to get people online, but it's an important job that paves the way to transformative knowledge-sharing and communication among people who didn't previously have it. There are five billion people who have no Internet access. We can't help some of those people. If you are unable to get your basic needs met due to poverty or famine or war, the Wikimedia projects cannot help you. But there's a layer of people who are not online yet that we could help. So we have to focus on the medium-term in getting materials to people who are not on the Internet. The big challenge is that they can't participate in the Wikimedia projects, or it's very difficult. That is problematic for us. The most active Wikimedians cluster in richer parts of the world, where they have

good Internet access, good educational systems, lots of leisure time. We do not want to be a set of projects in which rich people create educational material that is read by poor people.

Then there's the question around allocation of resources. If it takes \$1,000, or \$10,000, to reach one person in sub-Saharan Africa, and it costs three cents to reach a person in Denmark, what should we do with that? If we were a for-profit company, the answer would be easy. But we're not. We need to figure out how much energy to focus on reaching very or moderately poor people who don't have Internet access, versus how much energy to focus on the people who do have Internet access. We know that providing educational materials for people who don't have Internet access is not our core strength, and it should not be our core strength. Fortunately, there are lots of NGOs and other organizations who are focused on exactly that work. So our goal needs to be to partner with organizations that are closer to the ground, and encourage them to utilize Wikipedia in their work.

What's your perspective on how important expanding the number of language Wikipedias is?

First, I will say that there are many people who know more about this than I do, and frankly who have thought about it more deeply than I have. So I do not pretend to have the right answer. But here's what I think.

The idea of having an encyclopedia available for every single person in the world in their own language is enormously powerful. But at this moment, some language versions of Wikipedia are flourishing, and others seem stalled. And I don't know if the stalled versions are doing anyone any good, or if there is any reason to think they will flourish in the future.

There are currently about 270 language versions of Wikipedia available. Some are broad and deep and high-quality, for example, the German Wikipedia. Some are very quickly growing, for example, the Russian Wikipedia. But some are very small, and faltering, and it is not clear if they will ever grow enough to offer sufficient breadth and depth to achieve a reasonable standard of quality. That might be okay for languages that are almost no-one's first language – for example, something like Esperanto. But in the case of a language that is widely used as a first language, like Swahili or Finnish or Tamil, I don't know how helpful it is. If people experience their language-version of Wikipedia as tiny and broken and poor-quality, year after year after year, well, then we are not helping those people, and we are possibly damaging the reputation of the projects and hurting perception of the other bigger Wikipedia language versions.

To the extent that this is a problem, I think it is obviously going to get worse. Gravity will continue to pull attention and participation to the globally-used language versions of Wikipedia. India is a terrific example of this. My understanding is that Indians who are online generally speak English. English is the language in India of the courts, of higher education, the police, business. It makes sense to me that English-speaking Indian people would be generally drawn to edit the English-language Wikipedia rather than for example the Tamil or Kannada or Telugu ones. Because more people read the English one, and one of the big motivators for participation is that people will read your work. I would expect also that if your culture is less-known to the rest of the world, that would be a particular motive – for example, as a Canadian I am always happy when I can help non-Canadians better understand Canada.

So ultimately, I wonder about the value of small, imperfect language versions that are not growing and developing. Perhaps ultimately their purpose and usefulness is slightly different from the global language versions.

What is the danger in keeping these projects “alive”?

It's a brand management issue. We want to be quality. We want to be credible. We want schools to trust us. We want to be a good-quality product and used and understood as such. Things that are new and experimental, I applaud that, obviously. And different projects are meant to be different sizes, like Wiktionary. Wiktionary has a significantly narrower scope than Wikipedia: it shouldn't surprise or disappoint anyone that its readership is significantly lower. That is not a sign of failure. But to me, the risk of having projects that have atrophied or failed is that readers experience them as such, and that will necessarily infect and damage their perception of all the Wikimedia projects.

Currently, we have no mechanism of assessing what makes a project a “failed” one, or what to do if a project has failed. I worry about that. Partly for the brand reasons I already talked about, but also because I believe that if a diseased, failed community develops, it can have ramifications for other communities. I do not think it's the role of the Wikimedia Foundation to make these assessments, but I am uncomfortable with the fact that no one is making these assessments.

Also, there's no formal mechanism to create new projects. I think the last project was created about four years ago, when Erik Moeller created Wikinews. As far as I can tell, he did that through sheer force of will, by continually talking about it and drafting and refining proposals and explaining the advantages and potential, and so on. I don't know if it would be possible to create a new project today. I believe the Wikimedia community has a strong bias towards the status quo, and I do not believe it is possible to innovate and invent in our current context.

As a result, I personally strongly believe that the Wikimedia community needs to create a council of community members to think about these issues. To develop mechanisms for innovation. Mechanisms for assessment and improvement. That kind of thing.

What are the key problems in the delineation and execution of different roles throughout the community?

Well, the key problem is that there *is* no delineation of roles.

Basically, there are currently three players: the Wikimedia Foundation, the 27 international, mostly-volunteer-run chapters, and “the community,” which is an amorphous contested term that essentially includes “everybody else” (e.g., editors, volunteer system administrators, volunteers who speak with the media, volunteers who answer reader e-mails, the people who moderate our mailing lists, etc.) Depending who you talk to, and depending on the specific context, “the community” may or may not also include paid staff and board members.

There is currently zero consensus on roles and responsibilities among those three players.

For example, in December 2007, the Wikimedia Foundation surveyed the chapters, asking them about their goal and where they were planning to focus their energy. Legend has it that one chapter wrote back, saying it couldn't possibly achieve internal consensus, and needed to fill in multiple versions of the survey, one for each board member. That's a big problem!

Here's what I think: I believe the Wikimedia Foundation should handle core operations of keeping the website up-and-running. So it should pay for the servers and bandwidth, maintain MediaWiki, that kind of thing. That's our most core and basic responsibility. The Wikimedia Foundation has a responsibility to ensure the projects stay up-and-running – which requires legal support, administrative support, fundraising support, and so forth.

Further, I would say the Wikimedia Foundation has a responsibility to do whatever important, priority work the volunteers are not doing. Someone from the Red Cross once told me that the difference between a Red Cross staff person and a volunteer is simply that, when the call comes at 2AM, the staff person *has* to get out of bed, while the volunteer has a choice. That sounds really obvious, but it's actually a good and useful guideline. At a macro level, it requires us to take a look at our shared mission, figure out where

we are failing to make sufficient progress towards it, and decide what intervention is required to fix the problems, the gaps where we are failing. So for example, one big gap is usability. Open source developers have a natural bias towards making improvements to MediaWiki that support super-users, because super-users are visible and valued and vociferous. Similarly, editors have a bias towards the super-user as well, for the same reasons. I believe that one of the important functions of the Wikimedia Foundation is to counteract that bias, to deliberately and explicitly represent the interests of readers, non-editors, new editors. Another gap, a different gap, is furthering awareness and understanding of the Wikimedia projects in geographic regions where there are no chapters yet. If readership is low in India, and there is no Indian chapter, then arguably that is a problem the Wikimedia Foundation should help to fix. It should help to fill the awareness gap in the short term, and it should help to nurture a chapter into existence in the longer term. Then it can step back.

We need to focus energy on active chapter development in developing countries. Chapters are not spontaneously springing up in developing countries, for many obvious and non-obvious reasons, and we need to start experimenting and learning about what structures will work in those countries. This is really important, because without it, the voices of the established chapters are ever-increasingly amplified inside the community and its internal dialogue. The Western European voices get louder and louder. The rich get richer and richer. For example, some of the European chapters paid to send their members to Wikimania, which resulted in a strong overall European presence. That is perfectly natural and makes sense for the European chapters to do; they have a responsibility to their members, and sending members to Wikimania will help the chapter evolve and flourish. But someone has to be counterbalancing that by deliberately actively supporting Wikimedians in developing countries. There are people who are not at the table that I think WMF has to bring to the table.

I believe the role of the chapters is different from that of the Wikimedia Foundation, and fairly easy to define. Chapters, by definition, are geographic entities. So I believe they should handle work linked to geography and physical location, such as partnerships with galleries, libraries, museums, archives and media organizations. If a German archive wants to release images to Commons, it should talk with Wikimedia Germany. If the French public TV station wants its documentaries in Wikipedia, it should talk with Wikimedia France. I also believe the chapters have a responsibility to the Wikimedians who live in their region – to provide support of various kinds, coaching, training, outreach events, quality-improvement activities. They have a responsibility to help recruit new Wikimedians. And they have a responsibility to fundraise, to become experts in fundraising in their geographic locations. I imagine a world—we don't live in it today, I am talking about 10 years from now—in which the established chapters are terrific at fundraising, and a substantial portion of their revenues goes to the Wikimedia Foundation, to

be spent on activities to develop and support younger chapters, and countries in which there are no chapters.

Beyond chapters, are there other types of structures that you think are needed within the Wikimedia community?

Yes, I also believe there is a role for a “volunteer council” or some such structure. This is a fairly controversial notion that has been proposed many times and so far, has always been eventually abandoned. I find that sad, and I also see it as an indicator of Wikimedia's inability, at this point, to take action, to do anything brave or experimental. I believe that some delegation of responsibility has to happen, and there are conversations at a meta level that need to be happening, that are not.

I have said many times that I feel sorry for the community-selected board members. They are supposed to channel the will of “the community,” which is flat-out impossible; there are no mechanisms or systems to support them in doing that. So I believe a “volunteer council” could play a very useful role in terms of supporting the board in knowing what the community is thinking, and what it wants. I do not believe a “volunteer council” would solve every problem, but I think it makes good sense to have a structure, however imperfect, that is charged with thinking at a more meta level about all the projects and all the language versions. It would be great to have a body responsible for thinking about how to foster quality. It would be great to have a body developing and prioritizing feature requests for the MediaWiki developers. It would be great to have a body making recommendations about policy to different projects.

I believe we also need to enable very loose voluntary organizational structures, that might take the form of, for example, “friends of Wikimedia” social clubs, groups of people who drink beer together and talk about Wikipedia, or high school or college Wikipedia clubs, that kind of thing. For example, Mozilla has outreach to schools, where each school has a Mozilla rep. We have a huge missed opportunity in secondary education. This is fertile terrain to us: societies, clubs geared to a younger age level seems like an obvious thing for us to do. Basically, we need to support organizational structures that are more formal than a group spontaneously gathering for a meet-up, but less formal than a chapter. This would enable people to self-identify as Wikimedians, and to have support for that identification. We need this because people tell us they want it. Wikimedia Brazil, for example, didn't want to become an NGO or to sign a legal agreement with the Wikimedia Foundation, but it did want some form of acknowledgement that it exists. I believe the movement should support and nurture all kinds of different associations.

How will the strategy process help with these kinds of problems?

The purpose of the strategy process is to have a high-level conversation. When you try to have a high-level conversation with no structure supporting it, you run the risk of getting bogged down in detail, or diverted by edge cases, or of the conversation just petering out into nothing. My hope here is that the strategy project will provide the structure to enable a successful high-level conversation.

First, it will provide a solid bedrock of high-level data and analysis. Simple high-level data. For example: very few people in China read Wikipedia. Participation is stagnant across the majority of projects. The Wikimedia projects do not offer the kinds of features and functionality Internet users increasingly have come to expect.

Second, it will provide some structure for the conversation itself. A little facilitation help, some help finding or carrying out research, other process support of various kinds.

And third, it will help us drive towards tough decisions. It's hard to set priorities, it's hard to say we are going to do X instead of Y. We will never do Y, or we will do it in five years but not now.

What is top of mind for technology during this process?

I would like some big-picture thinking about technology.

The general user experience on Wikipedia in the last five years hasn't changed much. During that same period, the Internet has changed a lot. And we have not kept pace.

I want us to do some thinking about that. How can we radically improve MediaWiki, and Wikipedia functionality generally. Who is best placed to do that, and how can it best be done. Would it help MediaWiki if there were—entirely unrelated to Wikimedia—a for-profit firm focused on evolving MediaWiki for commercial uses, and supporting its many commercial users? Would that drive adoption, and thereby also usability and innovation? How can we radically improve our mobile presence, and how can we enable participation from mobile devices? How can we make the Wikimedia projects more permeable, more porous, so that they are easily interoperable with other projects and tools such as the Encyclopedia of Life, identi.ca, Flickr, OpenID, and not walled off from everything else?

What do you see as the role of WMF in advocacy? What are the top issues? And what are some challenges?

The Wikimedia Foundation is too small to do much advocacy. And it doesn't need to do much. There are a lot of organizations whose primary purpose is advocacy, or who are large enough to have departments devoted to government relations and so forth, that have positions that overlap with ours. I am thinking of the Electronic Frontier Foundation, Google, Creative Commons, Yahoo, Facebook. It is not our core competency or purpose, nor are we large enough to develop it as an ancillary function.

The biggest and most obvious core issue for us is censorship. Governments or companies such as ISPs standing between people and the Wikimedia projects are one of a very few obstacles that can utterly prevent us from fulfilling our mission. This is our main issue when it comes to advocacy: defending people's right to go online and get the information that they want, free of interference.

Beyond opposing censorship, there is very little clarity and very little consensus about Wikimedia's stance on various issues. My experience thus far is that individual volunteers tend to represent a wide spectrum of opinion. The chapters seem to have more consensus and are comparatively conservative, probably because they are used to being the "face" of their projects to the general public, the media, and cultural organizations.

Where should WMF advocacy efforts be focused?

Honestly, just based on sheer capacity, Google, the EFF, and Creative Commons will take care of 90% of U.S.-based issues. It would be a waste of Wikimedia's energy to lobby Washington. I get invitations to go to Washington, and I don't see the value for us. Ultimately, it's something that an American chapter would probably lead.

It's also not our core competency to focus on really difficult issues like censorship in China.

I believe the Wikimedia Foundation is best off to continue doing what it does today. We do occasionally, rarely, take a public stance, when an issue is particularly core to the Wikimedia projects, or when a response is required from us due to a lot of media interest. The volunteers are in many ways better suited to advocacy than the Wikimedia Foundation.