

121 countries (at least) represented		Languages used in fully completed surveys	
		Language	n
Question:	n=	English	1,600
<u>gender...</u>	<u>2,831</u>	German	486
<u>countries...</u>	<u>2,284</u>	French	258
<u>are you a...</u>	<u>2,958</u>	Spanish	126
<u>S3</u>	<u>3,011</u>	Russian	116
<u>all responses*</u>	<u>3,286</u>	Italian	110
		Polish	63
		Portuguese	52
		Ukrainian	39
		Japanese	35
		Chinese	27
		Other Asian	46
		n=	2,958

* This number includes partial responses.

Before we start, I will give you the URL for downloading the data and summary statistics on the last slide of this presentation, so don't start googling for the data just yet.

The OpenStreetMap Foundation's 2021 community survey ran from January 16 to February 14, 2021. It was advertised and announced via the OSM talk list, via social media, direct emails to working groups, local chapters and communities, OSMWeekly, and the OSM wiki front page. Initially the survey was offered in 14 languages, but within the first week and a half we added four more, for a total of 18 languages. This is the broadest array of languages offered for any OSM survey to date. Native speaker volunteers either translated from scratch or edited machine translations of all languages except Russian and Farsi. We never found a Farsi speaker, and I translated the survey into Russian.

I had hoped for 1,100 responses over the course of a month. We had that many token requests in about the first four days. We had well over 4,000 token requests by the end of the survey, but about 1,200 of them were either not used or the respondent abandoned the survey on the first page. This may indicate that some potential respondents didn't consider the questions relevant, decided they didn't know enough or care enough about the issues described in the questions, or that they were put off by the prospect of answering 18 questions (even though 9 of them were optional and demographic and thus easy to answer). A few emails and social media comments came in that indicated some respondents simply didn't expect questions of a policy nature and didn't feel like wading into OSMF politics.

By the end of the survey, we had received over 3,000 responses. Some were incomplete and so were not completely useful in terms of doing cross-sectional analysis and comparisons, but we still ended up with over 2,200 responses with complete demographic data and full responses to all mandatory questions.

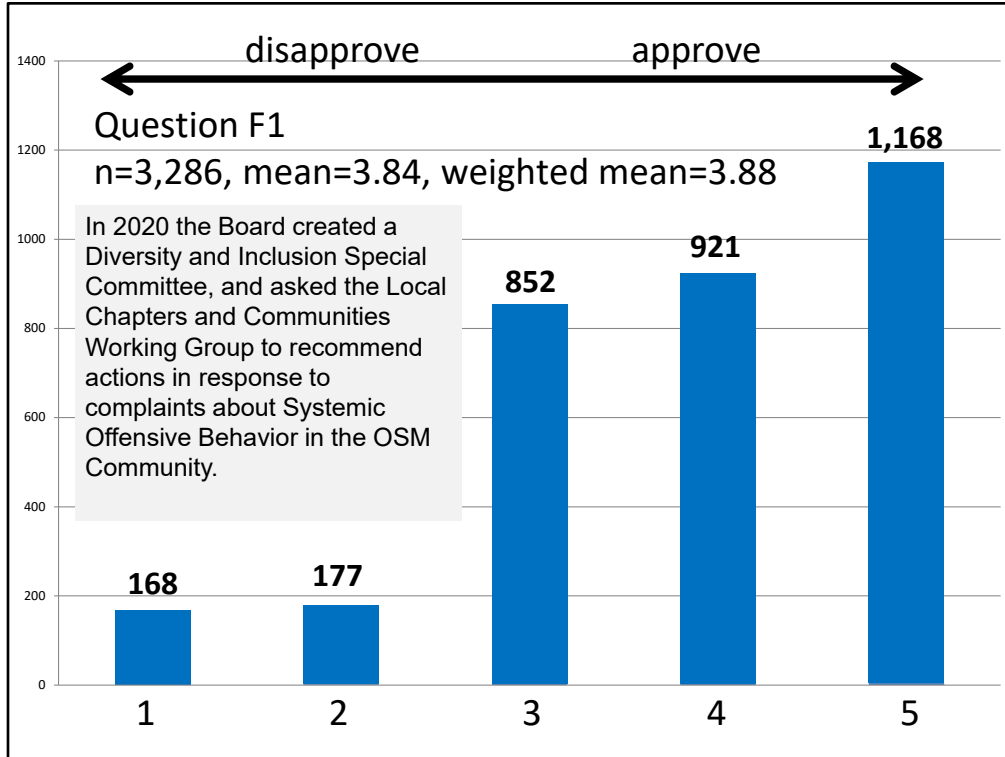
Caveats

- 121 countries are identified in survey vs. 249 presumed OSM-active countries based on Jennings Anderson's data.
- OSMF is overrepresented: 615/2958 or ~21% of sample, roughly 1/3 of OSMF membership.
- Paul Norman's analysis indicates Russia, Japan, Italy, and Indonesia are underrepresented in the survey, while the U.S. and UK are slightly overrepresented.
- Normalization using Jennings Anderson's geodata indicates, however, that the results are not heavily biased.
- Normalization was not attempted using any other variables than declared country of residence.

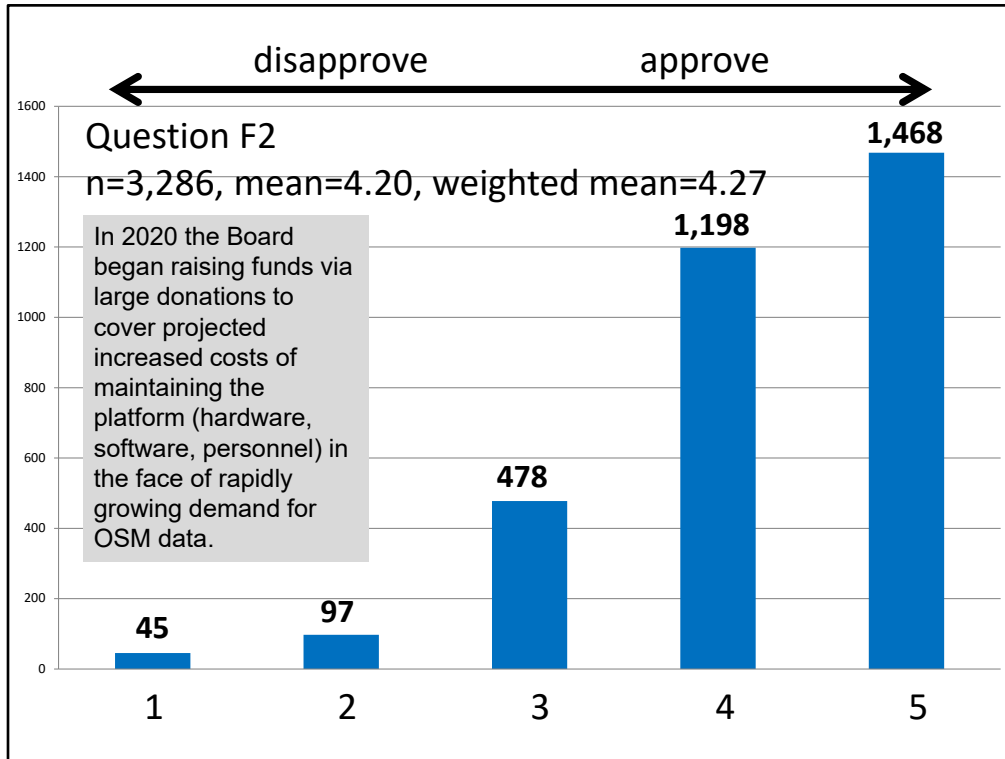
We received strong indications in the summary statistics that although there is evidence of selection bias, notably the heavy overrepresentation of Foundation members, overall the survey does not betray a high degree of statistical bias. We infer that from a comparison of the means of responses to questions F1 through F5 by the entire sample, the subset of the sample that includes demographic data, and the subset of Foundation members.

An obvious selection bias issue is the fact that, as a web-based survey, this survey was accessible to individuals with relatively easy Internet access, speakers of one of the 18 languages offered, and those interested in OSMF matters. This in fact, however, was the target audience. This survey did not target the general population, and since both contribution to and using data from OSM require a computer, Internet access, and knowledge of a major language, we did not consider that an undue source of sample bias.


Tobias Knerr used Jennings Anderson's data on changesets to normalize the data in terms of geographic coverage. This resulted in changes of a few hundredths of a point in the means of questions F1 to F5, as we will show you. These differences, between raw means and weighted means, are small.



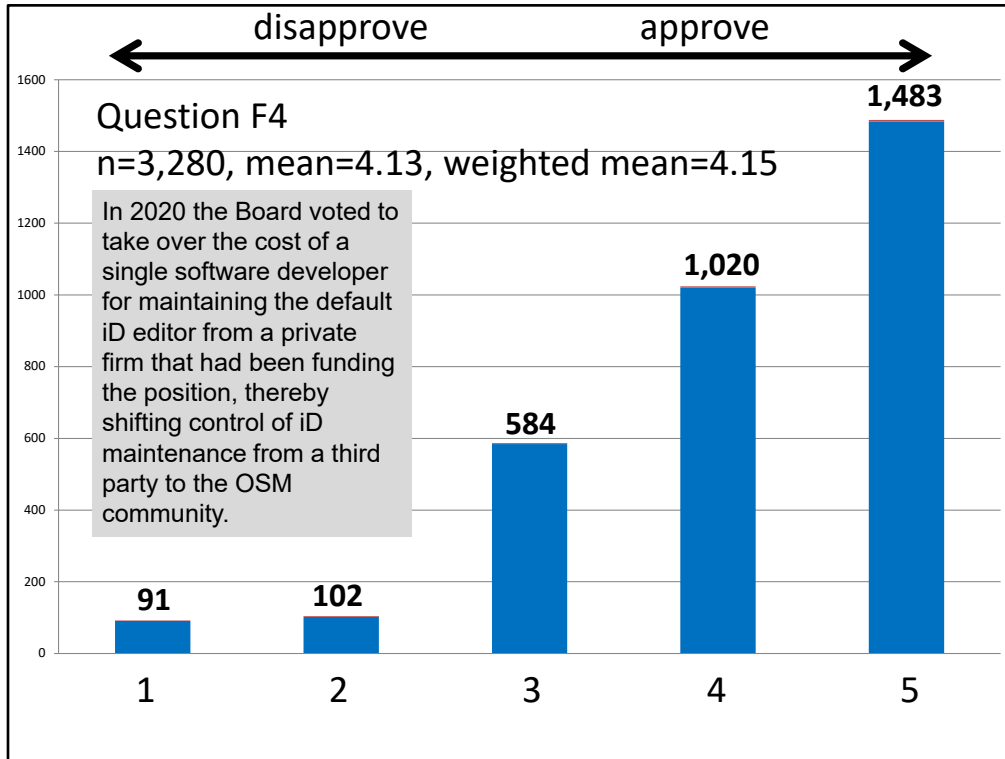
Question F1, feedback on the Board’s decision in 2020 to address diversity and the allegations of systemic offensive behavior, was the most controversial, if it can be called that. 63.6% of respondents approved or strongly approved of the Board’s decision, versus 10.5% who disapproved or strongly disapproved. The rest, roughly a quarter of respondents, were neutral on the subject. You can see that the weighted mean differs from the raw mean by 4/100ths of a point.




Just over 81 percent of respondents approved or strongly approved of the Board’s decision to begin raising funds via large donations. 4.3% disapproved or strongly disapproved. The raw mean differed from the weighted mean by 7/100ths of a point.



81.4% of respondents approved or strongly approved of the Board’s funding of short-term grants to upgrade Nominatim, osm2pgsql, and Potlatch. Slightly under 4 percent disapproved or strongly disapproved. This was the least controversial feedback question. The raw mean and weighted mean differed by 3/100ths of a point.



Just over three quarters of respondents approved of the contract to engage Quincy Morgan as a full-time iD maintainer. About 5.9% of respondents disapproved or strongly disapproved. The raw mean and weighted mean differed by 2/100ths of a point.



72.4 percent approved or strongly approved of the microgrants program. Five percent disapproved or strongly disapproved. The raw mean and weighted mean differed by 2/100ths of a point.


**Score S1 - priority order of tasks of
the Board of Directors**

11,249	Stability of the core infrastructure
4,004	Takeover protection
2,093	Outreach
1,631	Fund-raising
1,092	Attribution guidelines
728	Recruitment for Working Groups
287	Brexit

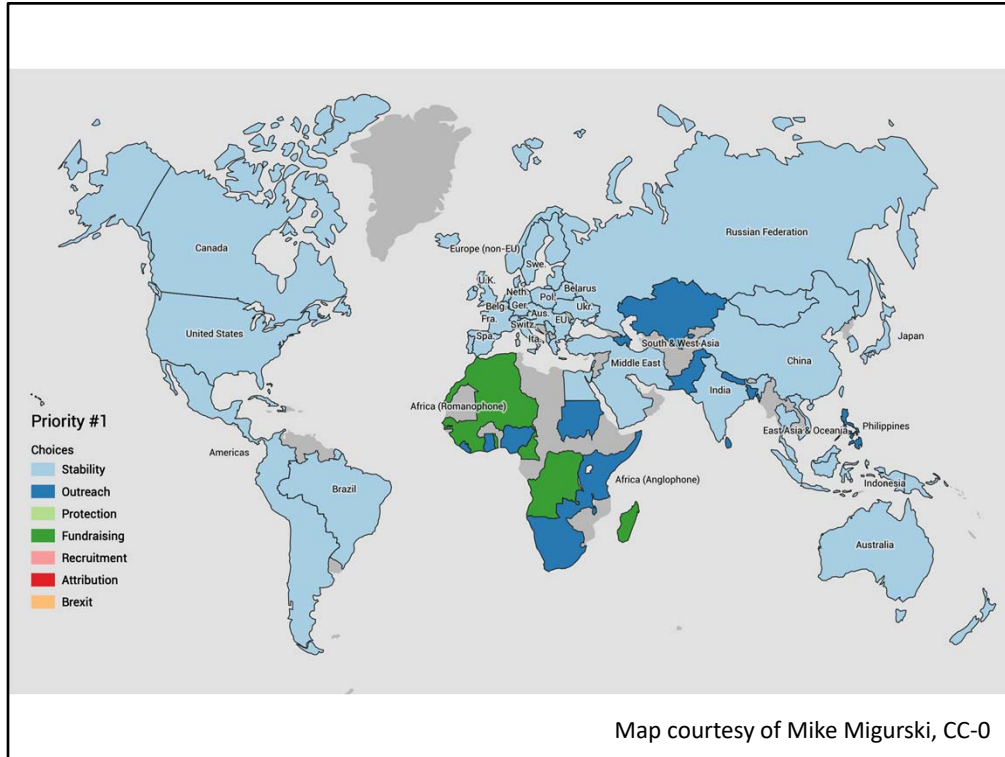
n=3,011

Shifting to the community sentiment questions, the first asked for a sense of what priorities the Board should set for 2021. Stability of the core infrastructure was a clear winner across the three demographics we have checked so far, which are OSMF members, respondents with more than 15 years in the project, and mappers. No other issue comes close.

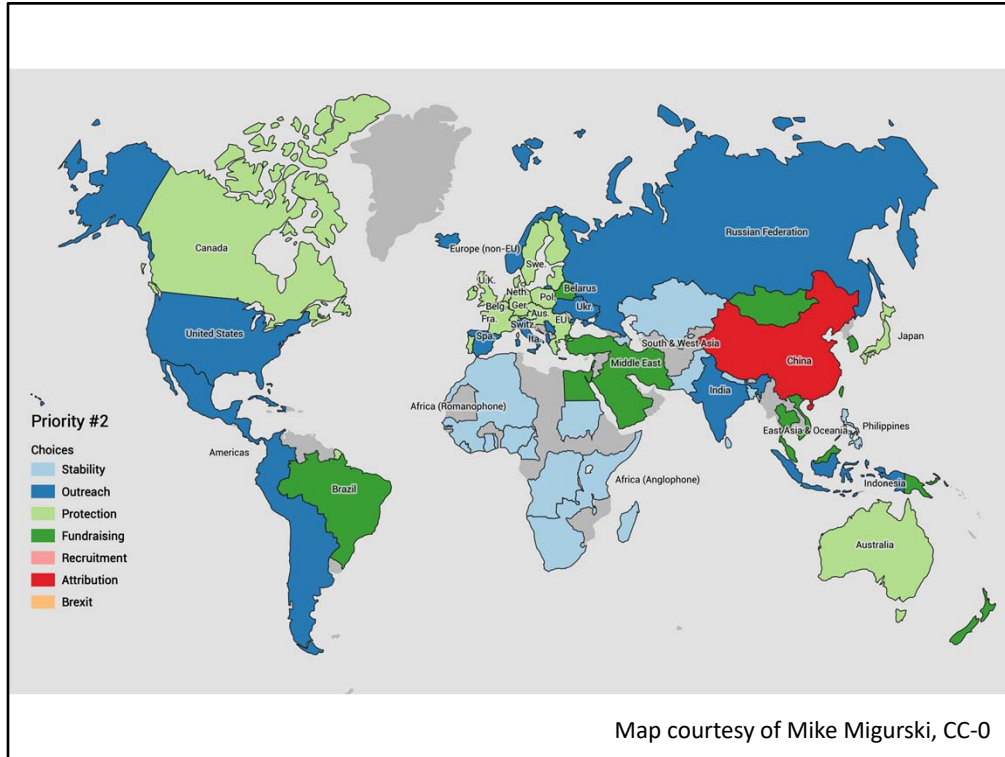
These rankings were calculated by ranking each response with a score of 1 to 7, with 7 assigned to first choices, 6 to second choices, and so on. These were summed into a score and then ranked.



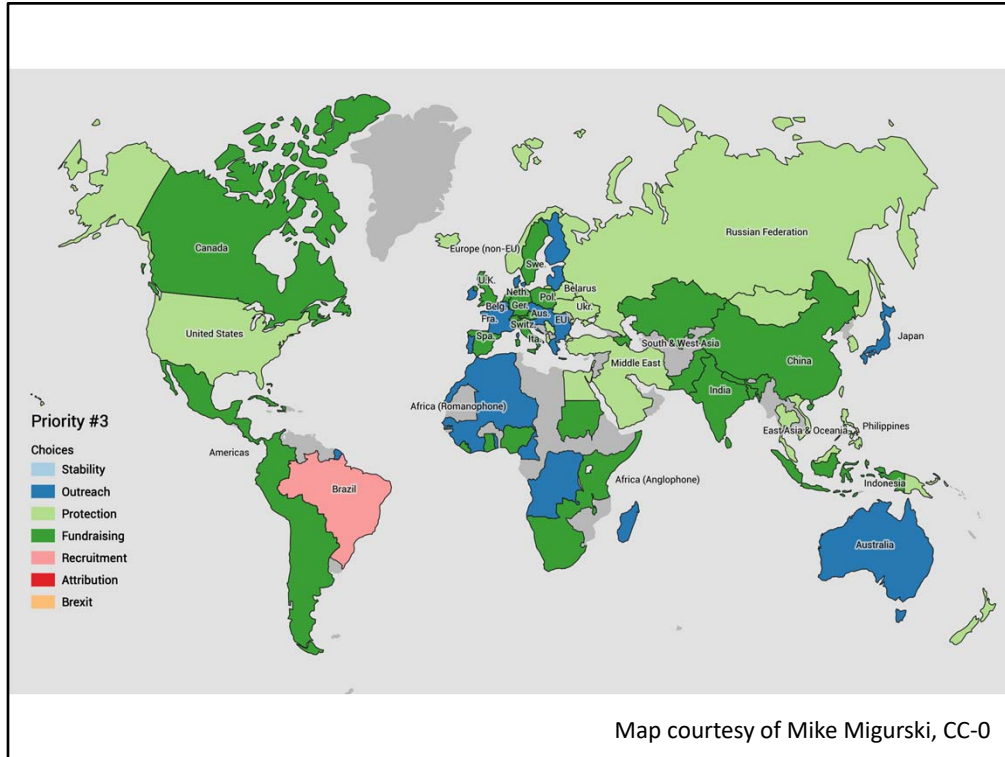
These are the same data in graphical form. You can clearly see how stability of infrastructure outweighs everything else.




Mike Migurski ran a Scottish STV ranking, the same system we use to elect board members in the OSM Foundation, and then used the results to plot these maps. As you can see, in most of the world, stability of the platform is the top priority, but in Africa we find fundraising and outreach to be higher priorities.



This is a similar map for the second priority based on the STV model of counting ballots. Stability of the platform is number two in Africa, South and West Asia, and outreach is important to most of the Americas, plus Russia, Belarus and Ukraine, as well as India and Indonesia.




In third place the big winners are protection from hostile takeover and fundraising. Outreach to local chapters and communities gets a fair amount of attention as well.




The second community sentiment question tested attitudes toward artificial intelligence and machine learning. Almost 38 percent want the Board to take no position on artificial intelligence and ML, which de facto means leaving the matter to local communities. Another 31 percent favor a formal policy of approving artificial intelligence and ML, but with the proviso that local community approval is needed before artificial intelligence and ML can be deployed for a locale. In short, 69 percent, or over 2/3 of respondents, want the matter left to local communities, whether the Board is involved or not. The only clear losers are a ban and blanket approval.

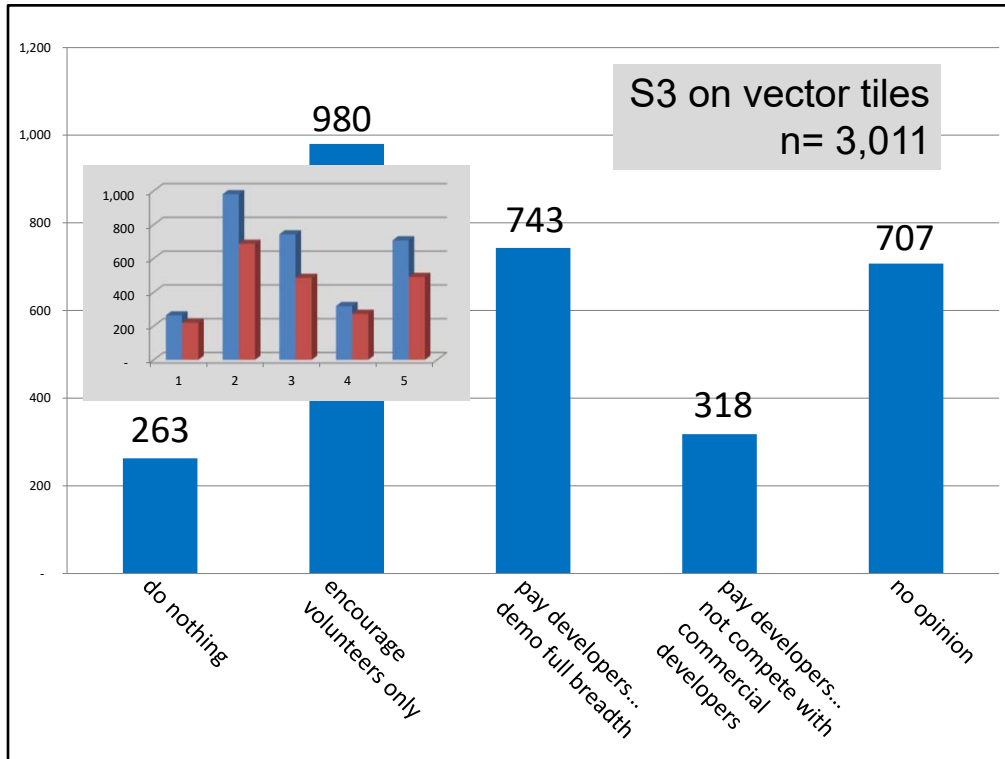
I have my own interpretation of this, which is that the community emphasizes local knowledge as the cornerstone of data quality. If my perception is correct, then the community appears to believe that the tool is less important than the on-the-ground human mapper who can verify data by putting eyes on the location.




If we plot the geographically weighted numbers alongside the raw numbers, we can see that the proportions remain pretty much the same even if the absolute figures change. This gives us a clue that the data are pretty robust across geographic lines. In the small inset chart, the blue columns are raw data and the red columns are weighted or normalized data.



The vector tiles sentiment question showed a plurality of just under a third in favor of leaving matters strictly to volunteers, but not anywhere close to a majority. A slightly larger number of about 35% favored paying developers to produce vector tiles, but split roughly two to one between those in favor of demonstrating the full breadth of OSM data and those who don't want to see us competing with commercial providers of vector tile services. If you add the "do nothings" to volunteers only, you have just over 41% favoring no substantive role for the Foundation in putting vector tiles on the website. In short, there is no clear consensus on the issue.

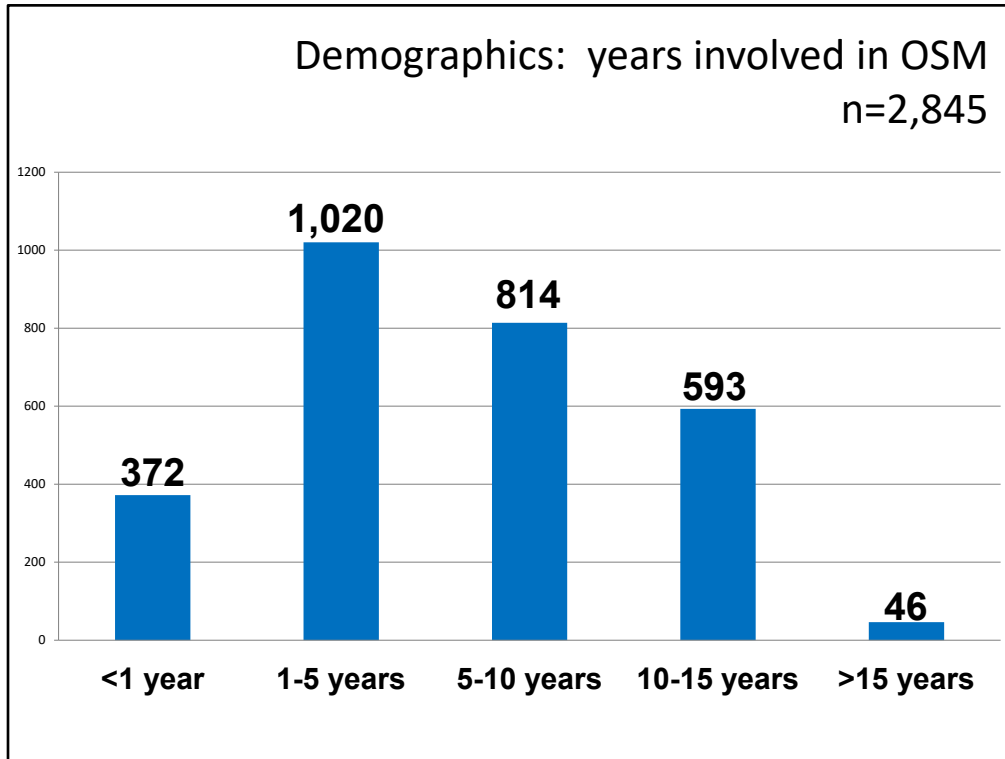


Same as for question S2, we see that although the absolute numbers change a bit, the proportions or percentages remain pretty consistent when the data are normalized against known activity by country. In the small inset chart, the blue columns are raw data and the red columns are weighted or normalized data.



The demographic data are in certain respects stunning. We actively solicited female respondents by sending emails to women known to be active in the OSM community and asking them to spread the word. We used social media to reach out to local communities, particularly those on Twitter and Slack. The results infer that 8.2% of the OSM community consists of women, 89% consists of men, 1.6% is non-binary, and 1.2% is “other”. In other words, given that roughly half the global adult population is female, women are grossly underrepresented in the OSM community.

The “other” category deserves a comment. The survey allowed individuals to enter a text response describing what “other” means. Some respondents indicated that they are attack helicopters. This is apparently intended as a disparaging term for people who consider themselves other than cis-male or cis-female.



This is perhaps the most uplifting of the data points. It infers that the project is growing.

One other comment: I watched this chart change over the four weeks of the survey. Initially, the three middle bars were roughly even, but starting around the end of the second week they began to diverge, and the 1-5 year and 5-10 year bars began to outstrip the 10-15 year bar. This told me that in the latter half of the survey we were reaching deeper into the community, getting to those members who are not as well plugged in as OSMers of longer standing. The old-timers were the first to take the survey, and others followed suit in the fullness of time.

<i>Are you a member of...</i>	
221	[a working group]
615	[the OpenStreetMap Foundation]
485	[a local chapter]
832	[a local community without chapter status]
70	[a corporate sponsor of OSM]
289	[a commercial company using OSM data]
443	[a non-profit organization using OSM data]
<i>Are you a...</i>	
2,489	[mapper]
306	[communicator]
1,617	[data user]
407	[developer / maintainer]
319	[event organizer]
115	[hardware / systems operator]
n=2,959	

This list simply shows you the raw numbers of respondents who self-identified in one or more categories. The largest single category is “mapper”, with 87.5% of respondents. Next are data users, with 54.6% of respondents. I ran crosscuts on these data, and to give one example, found that 1,371 mappers also identified themselves as data users. The summary statistics spreadsheet gives additional detail, and more can be derived if desired from the anonymized data spreadsheet as well.

I should also point out another important data point. A total of 615 respondents self-identified as members of the Foundation. That’s only a little less than a third of the Foundation membership, which gives us a very good sense of what the Foundation membership thinks. That is not a large sample, it is a huge sample. As one former board member admonished me recently, “The Board doesn’t report to the community, it reports to the Foundation members.” This survey gives us excellent insights into the Foundation membership.

Examples of possible subset- and cross-sectional analysis

Structure of the community

How certain groups answered questions

Comparisons of responses between groups

The data lend themselves to some deep dives into the structure of the community. Statisticians consider a sample of 1,067 to be a “large sample”, and our sample containing full demographic data is slightly more than twice that size. It is a robust sample, and based on Jennings Anderson’s geodata, appears to be relatively unbiased. Genderwise, the 8% figure for women is in line with previous estimates, so it is also likely reasonably accurate.

We can thus use these data to answer some questions. Here are a few examples of things I was curious about.

Are you a...?	
mapper	2,489
mapper and...	
commercial company employee	249
non-profit employee	346
corporate sponsor employee	57
local community w/o chapter status	768
local chapter	455
OSM Foundation	582
working group	189
mapper and...	
working group + commercial company	32
working group + corporate sponsor	13
mapper and...	
male and German	330

Here is an example of how we can crosscut the data. We have 2,489 mappers, and we can query the dataset to determine how many mappers are also employees of companies or NGOs, or members of the Foundation, and so on. We can ask how many are simultaneously mappers, members of working groups, and employees of a commercial company using OSM data. We can do this for all the demographic variables.

Because of the need to anonymize the data, not all demographic variables can be crosscut in the anonymized data set. However, if you have a desire or need for a specific crosscut, I'll be happy to run those numbers for you using the raw original data.

Percent of Community Members with 1-5 Years in Project, by Country							
Africa-angl	83	53	63.9%	Indonesia	20	10	50.0%
Africa-rom	34	18	52.9%	Italy	125	42	33.6%
Americas, other*	80	29	36.3%	Japan	27	7	25.9%
Australia	38	12	31.6%	Middle East	40	14	35.0%
Austria	47	13	27.7%	Netherlands	39	14	35.9%
Belarus	26	10	38.5%	Philippines	27	9	33.3%
Belgium	28	11	39.3%	Poland	56	25	44.6%
Brazil	52	24	46.2%	Russia	61	19	31.1%
Canada	41	15	36.6%	S&W Asia, oth*	25	12	48.0%
China	24	15	62.5%	Spain	53	17	32.1%
EAsia/Oceania*	35	16	45.7%	Sweden	24	7	29.2%
EU, other*	132	50	37.9%	Switzerland	50	10	20.0%
Europe	24	8	33.3%	Ukraine	36	13	36.1%
France	165	50	30.3%	United Kingdom	146	49	33.6%
Germany	411	106	25.8%	United States	294	113	38.4%
India	41	17	41.5%	n=	2,284	808	35.4%

This chart shows in the blue columns how many respondents there were from each country or region, then in the gray columns how many of them have been in the project for 1 to 5 years, and then what percent column 2 is of column 1. This gives us a sense of where the growth is. Green numbers are above the average for the whole sample.

I was curious about where the new mappers are coming from, since there is worry about new mappers being paid mappers who are on the payrolls of corporations. We can see that the greatest growth is in anglophone Africa, China, romanophone Africa, Indonesia, and in South and West Asia outside India. We can examine these data further to see how many of these mappers are affiliated with an NGO, how many are on the payroll of a company using OSM data, and so on. Some of this growth may be due to NGO activity, or corporate activity, or it may be something else.

**Responses to Question S1 (Board
Priorities) by Respondents in the
Project > 15 years**

	score
Stability of the core infrastructure	274
Takeover protection	211
Fund-raising	200
Outreach	190
Recruitment for Working Groups	156
Attribution guidelines	147
Brexit	82

n=45

I was curious about the responses of the old-timers to the project, so ran those numbers. This is the old-timers' response to question S1, on Board priorities for 2021. They closely mirror the priorities of the whole sample, with stability of core infrastructure on top and Brexit at the bottom. There are some shifts in between. You can see these numbers in detail in the summary statistics spreadsheet.

Responses to Question S1 (Board Priorities) by OSMF Members

	score
Stability of the core infrastructure	3,704
Outreach to Local Chapters/Communities	2,670
Fund-raising	2,589
Takeover protection	2,561
Recruitment for Working Groups	2,328
Attribution guidelines	2,212
Brexit	1,121

n=614

Foundation members had similar reactions to this question. There is a slight preference to outreach and fund-raising over takeover protection, but not by very much, compared to the old-timers. The conclusions we can draw from examining these crosscuts is that stability of core infrastructure is one priority everybody seems to share, Brexit is an issue few people understand or care about, and the three other issues of broad concern are takeover protection, fund-raising, and outreach to local chapters and communities, not necessarily in that order.

Given the amount of noise about attribution guidelines, I was frankly stunned that it consistently fell close to the bottom across the board. I did not expect that the Foundation membership would rate outreach to local chapters and communities as their second-highest priority. So, yes, there were some surprises in the data.

French and German females F1-F5							
mean	4.25	4.35	4.2	4.05	4.3	F1 mean difs	
std dev	1.07	0.93	0.83	1.05	0.98	Difference	-0.42
n	20	20	20	20	20	Standard error	0.244
French and German males						95% CI	-0.8989 to 0.0589
mean	3.83	4.10	4.23	4.25	3.98	t-statistic	-1.723
std dev	1.07	0.88	0.84	0.91	0.88	DF	541
n	523	523	523	523	523	Significance level	P = 0.0855
Females	F1	F2	F3	F4	F5	F1 mean difs	
mean	4.17	4.22	4.21	4.14	4.12	Difference	0.32
std dev	1.12	0.98	0.89	1.01	1.07	Standard error	0.076
n	232	232	232	232	232	95% CI	0.1718 to 0.4682
Males						t-statistic	4.234
mean	3.85	4.24	4.25	4.16	4.05	DF	2752
std dev	1.10	0.86	0.86	0.98	0.90	Significance level	P < 0.0001
n	2522	2522	2522	2522	2522		

One criticism of the data even before they were published was that they would not be comparable across geographic regions because of cultural and linguistic differences. I decided to test that hypothesis by comparing responses of German men and German women, and chose German because it is the largest geographic cohort in the sample. Unfortunately, there are so few German women in the sample that I could not assure anonymity, so decided to combine German and French men and women in order to get a sample size of 20 for the women. The results are in the upper table. The t-statistic does not reach the critical value of 1.9644, so we cannot reject the null hypothesis, even though the probability that the means are the same is only 8.55 percent. I then tested the hypothesis across the entire sample, all 232 women and 2,522 men. The t-statistic was robust, well over the critical value, with a probability that the null hypothesis is valid of 1 in 10,000. In layman's terms, this means that women are on average more approving of the Board's decision described in feedback question 1 than are the men.

I determined that to reach significance at the 95% confidence level, I would need a minimum sample of between 40 and 45 women. To achieve this, I aggregated all EU member states and ran the same statistics for question F1.

**EU Men vs. EU Women
Regarding Question F1**

<i>Response</i>	<i>No.</i>	<i>Score</i>	
EU women			
1	2	2	mean=
2	1	2	4.28
3	4	12	
4	12	48	std dev =
5	24	120	1.05
n= / total=	43	184	
EU men			
1	43	43	mean=
2	57	114	3.79
3	281	843	std dev =
4	291	1,164	1.09
5	318	1,590	
n= / total=	990	3,754	

F1 mean difs	
Difference	-0.490
Standard error	0.170
95% CI	-0.8227 to -0.1573
t-statistic	-2.890
DF	1031
Significance level	P = 0.0039

In the EU member states, 43 women responded with demographic data. These are the comparisons of the mean or average responses to question F1 by women and men from EU member states. The average EU woman approved of that Board decision with a score of 4.28. The average EU man approved with a score of 3.79. The individual counts of the responses also reveal that while 28% of EU men were ambivalent, or neutral, to the question, only 9% of EU women were ambivalent or neutral. Another telling statistic is that while 61.5% of EU men approved or strongly approved of the Board action, nearly 84% of EU women approved or strongly approved. These are significant differences, as shown by the t-statistic which is quite robust, and the significance level indicating that the probability of the null hypothesis is roughly 4 in one thousand.

This is an example of a data examination the Board can do because of its access to raw data that have not been anonymized. If there is interest in summary statistics of this nature, we can run them for you. We cannot, because of our privacy policy, release the raw data outside the board unless the data have been anonymized.

Mean responses		Mean responses by group to question F1	
Brazil	4.27	Mean responses	
French and German females	4.25	all respondents	3.84
Indonesia	4.25	East Asia/Oceania*	3.83
South & West Asia*	4.24	French and German males	3.83
Females	4.17	Ukraine	3.83
Non-binary	4.16	OSMF members	3.83
United States	4.12	Belgium	3.82
Latin America and Caribbean	4.10	non-mappers	3.80
India	4.05	European Union	3.80
Netherlands	4.00	Germany	3.77
Philippines	4.00	Middle East	3.73
Americas, other*	3.99	China	3.71
Africa, anglophone	3.98	Belarus	3.69
France	3.98	Japan	3.67
Africa, romanophone	3.97	Sweden	3.63
Italy	3.96	Switzerland	3.58
Spain	3.96	Other European Union*	3.52
Other Europe (Non-EU)*	3.92	Poland	3.38
Australia	3.89	Russia	3.31
Asia minus India	3.89	other gender	2.85
Asia, aggregated	3.88		
males	3.85		
Canada	3.85		
mappers	3.85		
demographic respondents	3.84		
all respondents	3.84		
		* Not elsewhere specified	

I was curious about how various groups reacted to question F1, since it was the most controversial of all the feedback questions. This list is in order of the means of the responses by group or country. This table is also on the summary statistics spreadsheet.

As you can see, the Brazilians were most approving of the Board's decision, followed by French and German women, Indonesians, and South and West Asia outside India.

Search the OSMF wiki for “2021 Survey Results”

page URL is

https://wiki.osmfoundation.org/wiki/2021_Survey_Results