Wikipedia reader motivations and behavior

QR code

https://goo.gl/IJcm2q

WIKIMEDIA
FOUNDATION

To follow along or to checkout the slides later

“Pop quiz, hotshot”

1. What % of wikipedia readers come for:
   a. Quick fact
   b. Topic overview
   c. In-depth learning

If it’s 90% quick fact, that would really change what we build

“Pop quiz, hotshot”

2. What % of wikipedia readers come because of:
   a. work/school
   b. saw something in the media
   c. bored

Going in, we assumed people had very directed questions.

Agenda

1. Who we are
2. What we knew
3. The project
4. Method
5. The results
6. What next

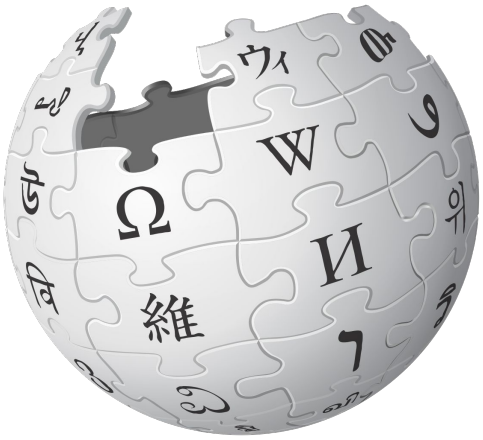

WIKIMEDIA
FOUNDATION

Who we are

Who I am

Jon Katz
Product lead, Reading team
Wikimedia Foundation
jkatz@wikimedia.org

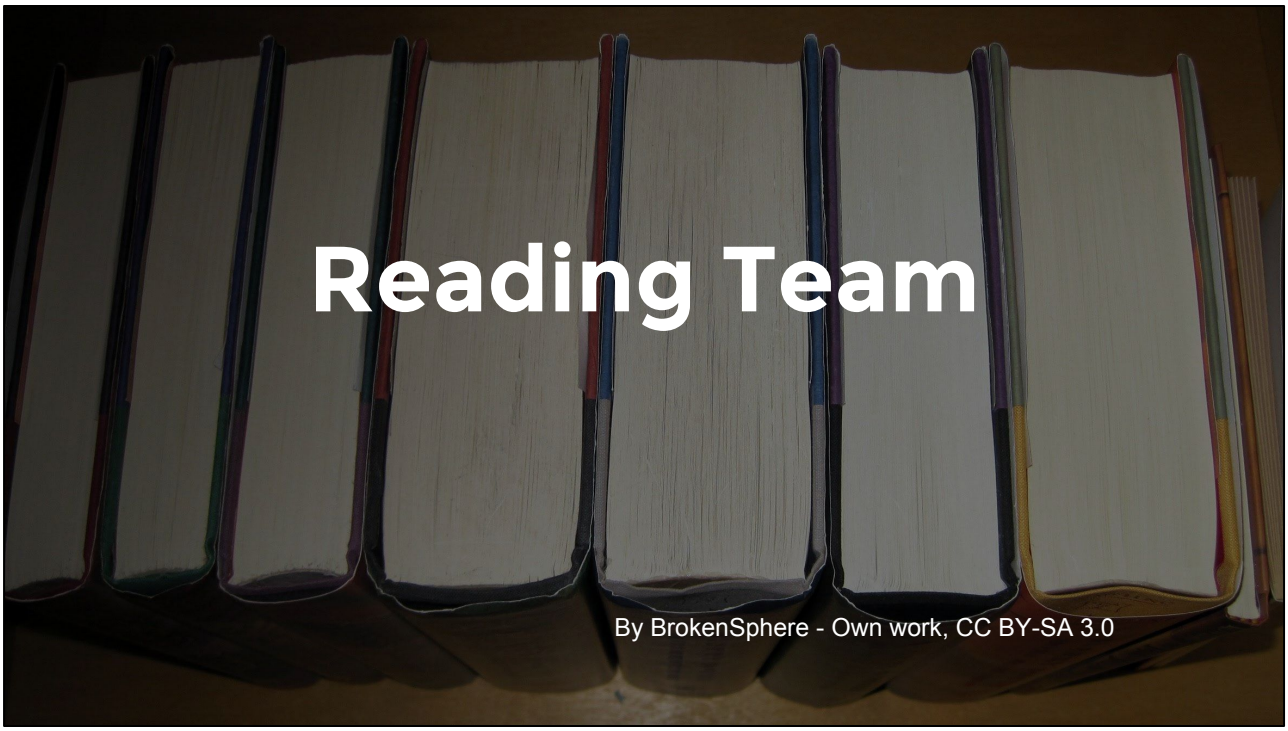

Image: https://commons.wikimedia.org/w/index.php?curid=9970866

Reading team impact

Help people learn more using Wikimedia projects

So far focus has been on Wikipedia and on Foundation hosted destinations, but that is not set in stone and as we create services we think about 3rd party developers.

Global View

Engage and retain readers

Enable new readers

Very basic strategy: in many places we are well known and widely used. The goal there is retention. In other places, we are not well known or widely used at all. In those places the goal is driving awareness and removing barriers to learning with wikipedia.

This research was focused 100% on existing users.

Shameless plug

New Reader's project overview
https://meta.wikimedia.org/wiki/New_Readers

Research results deck + Video of results presentation

For more details on what we're doing to reach new readers.

We just know the tip of the iceberg.

Image: https://commons.wikimedia.org/w/index.php?curid=757764

Image:
By
http://wellcomeimages.org/indexplus/obf_images/3c/ec/b46cf3ebc9a71fa6569fcd56f5c3.jpgGallery: http://wellcomeimages.org/indexplus/image/L0012069.html
Location, https://commons.wikimedia.org/w/index.php?curid=35958411

Research team

Design and test new technologies, produce empirical insights to support product and engineering decisions, and publish research informing the organization's and the Movement's strategy.

https://www.mediawiki.org/wiki/Wikimedia_Research

They did the bulk of the work here. Paper pending.

What we knew

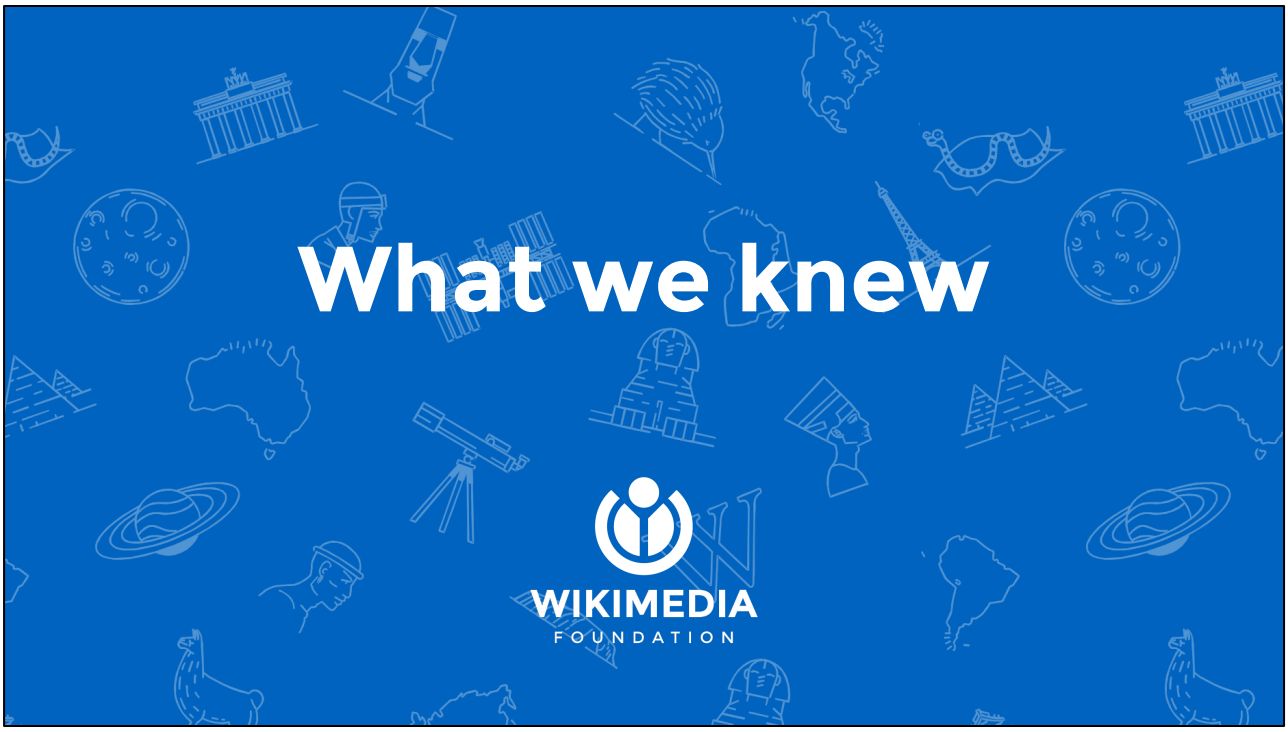

Image: https://commons.wikimedia.org/wiki/File:TheMap(bySylwiaBartyzel).jpg

Summary

- We are huuuge:
  - 15B pvs a month, ~1B unique devices/month*
- We are half mobile, mobile is growing, desktop shrinking
- Mobile is shallow
  - <½ the pageviews per device of desktop
- Overall usage is flat, even in locations where things grow

* using the device numbers for only the top wiki in each country

Long-term traffic trends

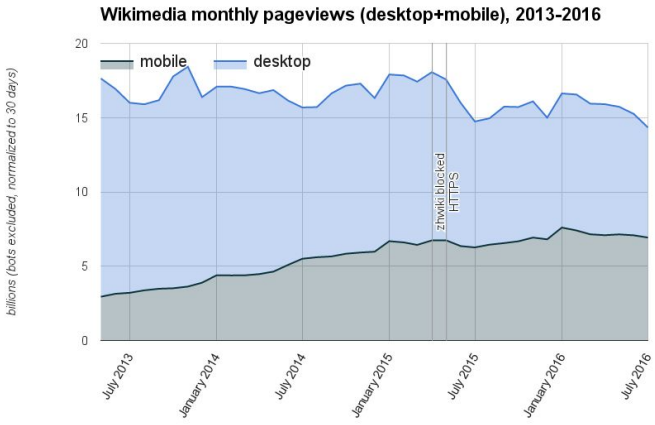

	Annual change (2013-16, linear)
Total	-3%
Desktop	-18%
Mobile (web+apps)	+25%

Reading is shallow

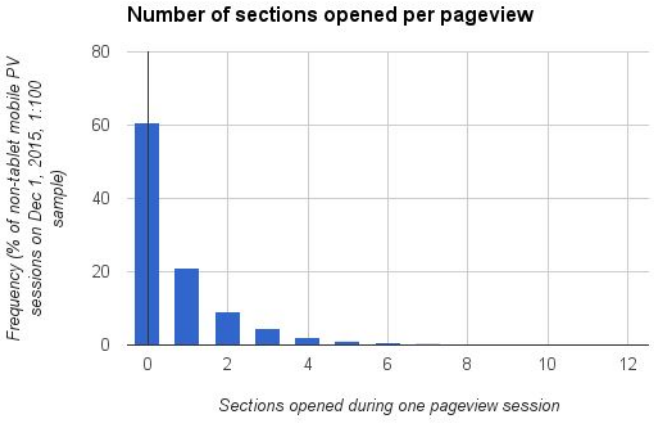

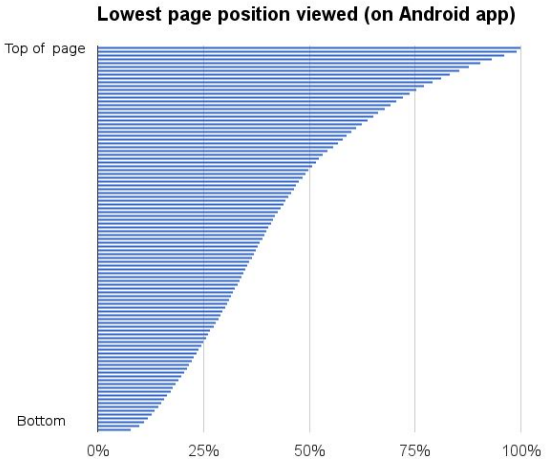

More detail and other findings:
[[File:New readership data (Wikimedia Foundation Tech Talk).pdf]], [[m:Research:Which parts of an article do readers read]]

What’s new
WIKIMEDIA
FOUNDATION

Why do people read Wikipedia?

Are there different kinds of readers?

Do different articles appeal to different kinds of readers?

Some external research suggested the answer is yes

Reader Preferences and Behavior on Wikipedia[1] found distinct browsing patterns and associated topics among readers. They also found that the most popular topics did not always align with the activity of editors.

[1] Janette Lehmann , Claudia Müller-Birn , David Laniado , Mounia Lalmas , Andreas Kaltenbrunner, Reader preferences and behavior on Wikipedia, Proceedings of the 25th ACM conference on Hypertext and social media, September 01-04, 2014, Santiago, Chile

Why it matters

- How articles are written
- What features we prioritize for readers
  - You'll see some of the feature decisions below….
- Avoid the danger of "averages"

Averages are dangerous

"7% improvement, let's push it live!"

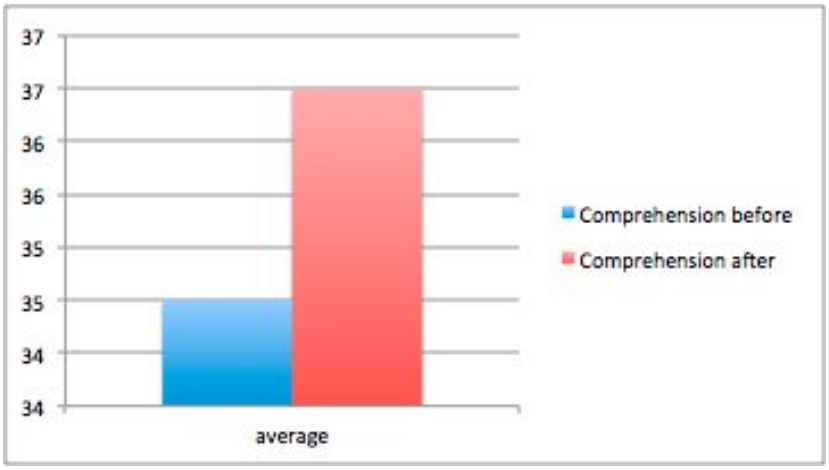

If you look at the average impact for this hypothetical change, it's an improvement

Averages are dangerous

"Woah...serious issues."

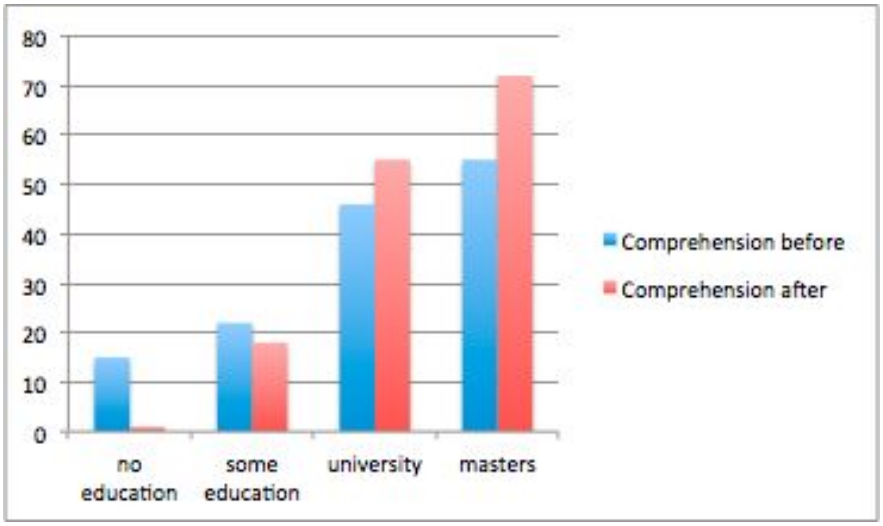

However, if you look at the segments, the segment with the least comprehension to begin with is actually hurt by the change...definitely needs some additional thought.

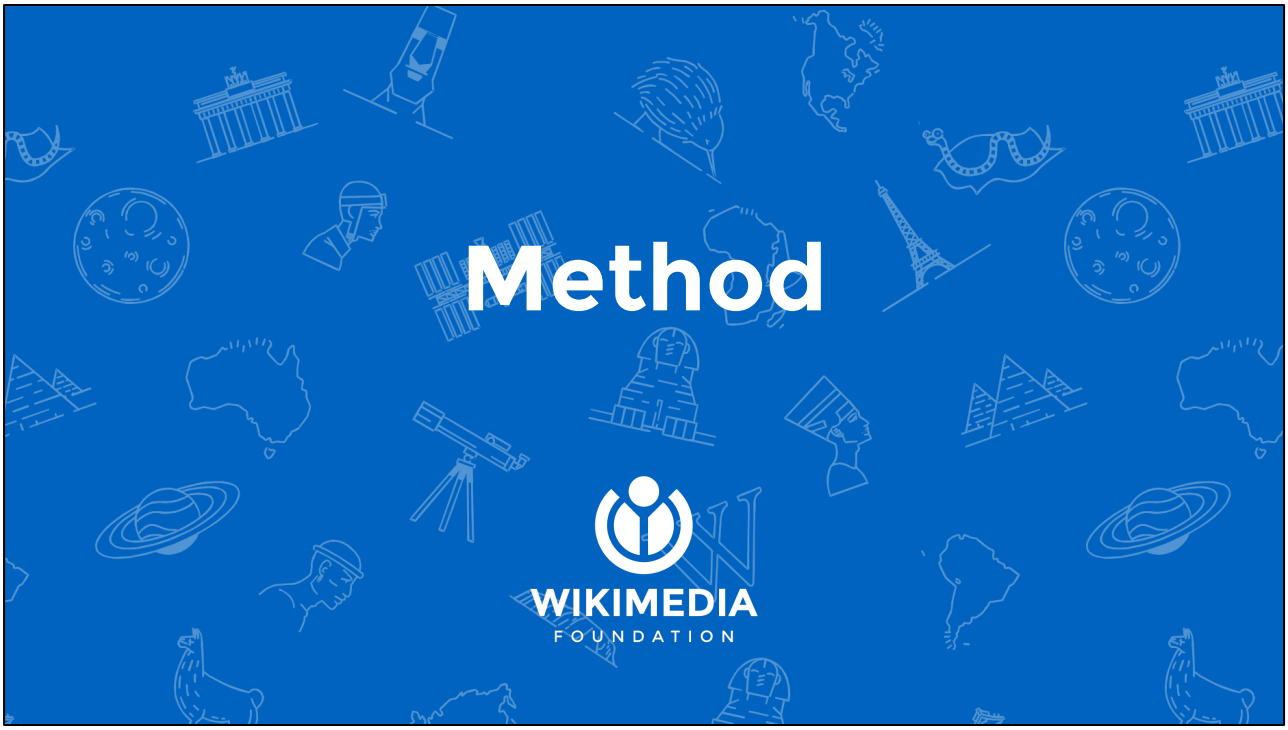
Method
WIKIMEDIA
FOUNDATION

Research Process

- Ask reader questions
- Record browsing
- Analyze

Details

- 1 week
- enwiki, eswiki, fawiki
- 5:1000 sampling rate for enwiki

WIKIMEDIA FOUNDATION

Ran on 3 different languages to code open text answers - the answers fell in the same categories in each language.

Bias

Coverage: we only look at En Wiki for final analysis

This would have to be replicated on more languages and ideally broken out more by country to consider this full coverage

Bias

Response bias: people who responded had much longer sessions. Used propensity scoring to compensate, but the impact was not significant. Only the data mining results below use the corrected data.

We used propensity score matching (PSM) and showed that the bias in the survey responses is not statistically significant.

- This is a blurb on how propensity score weighting works:
  - we have a sample of users that are not in the survey, and those that are in the survey. we run a classification, that probabilistically predicts if someone is in the survey or not. The reverse predicted probability of a response being in the survey is then used to reweight the survey responses. If this reweighted response set shows the same results as the original response set, we say that there is no bias in the survey data. https://en.wikipedia.org/wiki/Propensity_score_matching

Initial round

Image: https://commons.wikimedia.org/w/index.php?curid=7743547

For work!

Studying for my med school test

To extend my research, comparing the April 2015 Nepal Earthquake and the September 11 attacks.

I want to build a deck “wooden” and support a hot tub at one end of it.

Because I’m currently watching this show and wondered when the episodes originally aired.

Curiosity.

Bored.

cuz yolo.

Answers fit into 2 dimensions

External Motivation

By germanny - Flickr: bg_twitter, CC BY-SA 2.0

Image:
https://commons.wikimedia.org/w/index.php?curid=31565341

I am reading this article because...

work/school	I have a work or school-related assignment.
personal-decision	I need to make a personal decision based on this topic (e.g., to buy a book or game, to choose a travel destination).
current event	I want to know more about a current event (e.g. Black Friday, a soccer game, a recent earthquake, somebody's death).
media	the topic was referenced in a piece of media (e.g. TV, radio, article, film, book).
conversation	the topic came up in a conversation.
bored/random	I am bored or randomly exploring Wikipedia for fun.
intrinsic_learning	this topic is important to me and I want to learn more about it. (e.g., to learn about a culture).

Replicated open text in Persian and Spanish and found that the categories did not change. We set these as the multiple choice options and then redid the survey in english.

Image: https://commons.wikimedia.org/w/index.php?curid=5999158

I am reading this article to...

fact	look up a specific fact or to get a quick answer.
overview	get an overview of the topic.
in-depth	get an in-depth understanding of the topic.

Replicated in Persian and Spanish

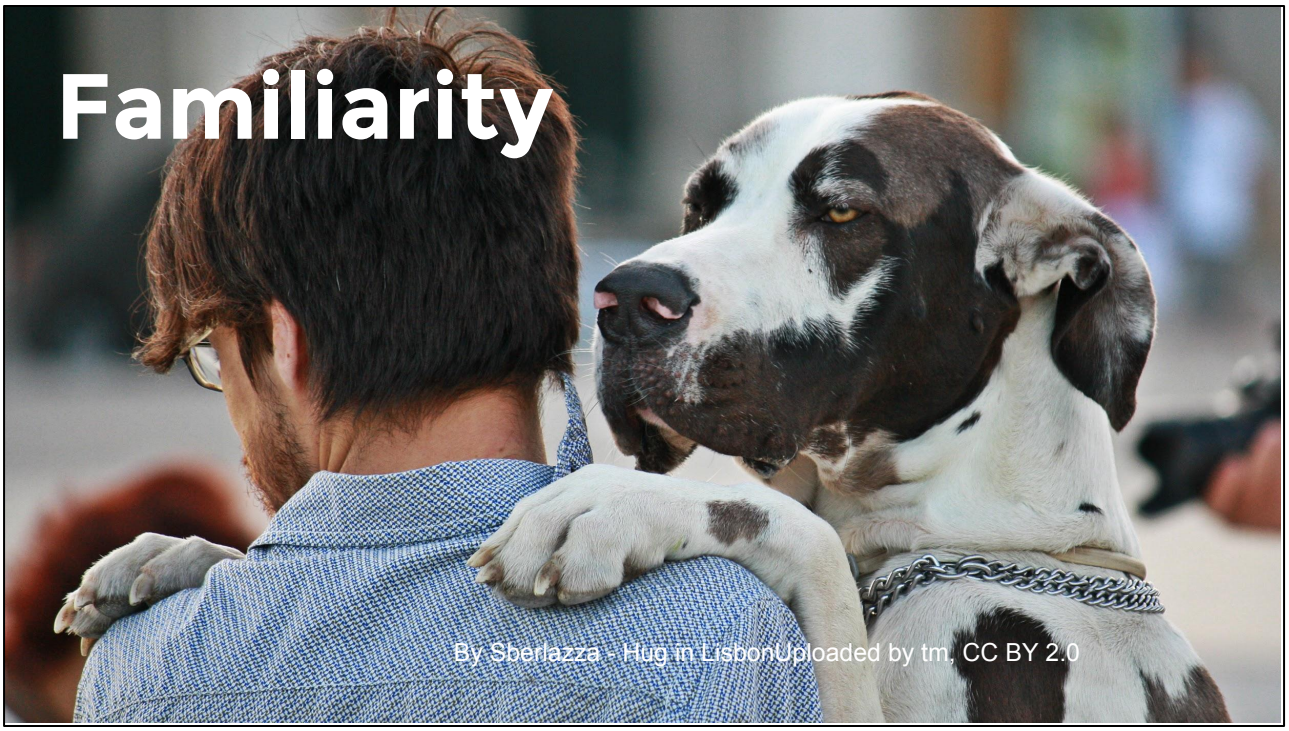

Image:
https://commons.wikimedia.org/w/index.php?curid=29891775

Prior to visiting this article ...

familiar	I was already familiar with the topic.
unfamiliar	I was not familiar with the topic and I am learning about it for the first time.

We then looked at other characteristics

User:

- Mobile/Desktop
- Time of day, day of week
- Country, Continent

Session:

- Pages looked at
  - Topic, centrality, page rank
- Session time, pageview #
- Branching within session

Details, caveats and more here:
https://meta.wikimedia.org/wiki/Research:Characterizing_Wikipedia_Reader_Behaviour/S3-English_Large_Scale
More results: https://github.com/ewulczyn/readers/blob/master/src/analysis/Metrics.ipynb

Topics created by LDA method (not so great or comprehensive, actually)
Centrality is as it relates to network theory.
Pagerank, what links in. Like google.

The results

Depth

Image: https://commons.wikimedia.org/w/index.php?curid=5999158

Only ¼ of users come for in-depth reading

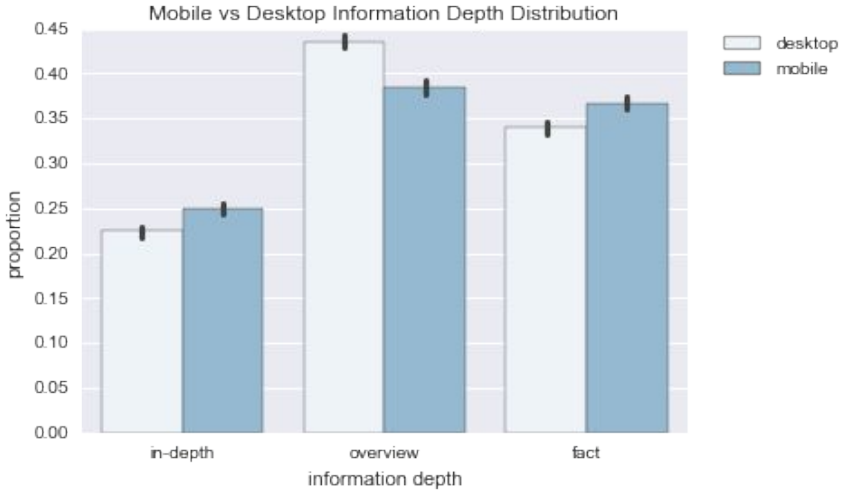

People spend more time based on purpose

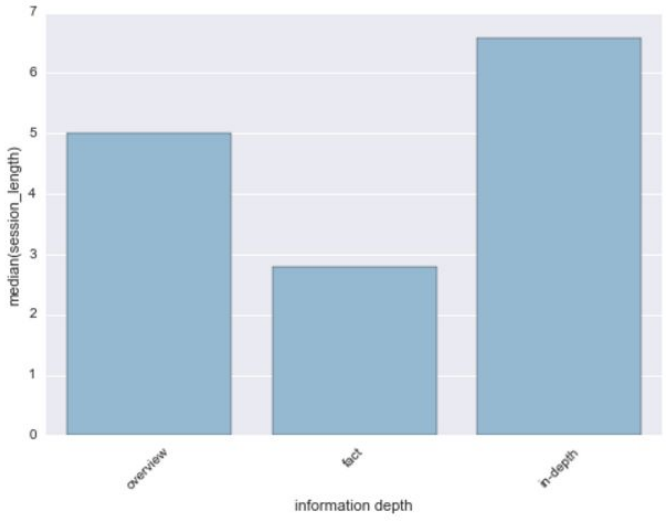

Motivation

Image:
https://commons.wikimedia.org/w/index.php?curid=31565341

Media is the leading motivator

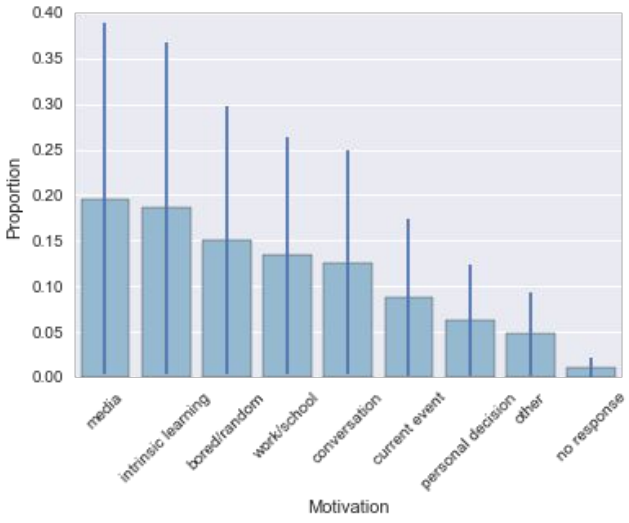

*results limited to respondents who listed one external motivation (70%)

The error bars are huge here, but when we look at the same data corrected for potential bias, they shrink significantly and the results are not significantly different.

Intrinsic learning is the biggest driver of *in-depth*

Media is the leading motivator for everything else...followed by *work/school*

Weekday v. Weekend is predictable

On weekends, work/school goes down and media goes up.

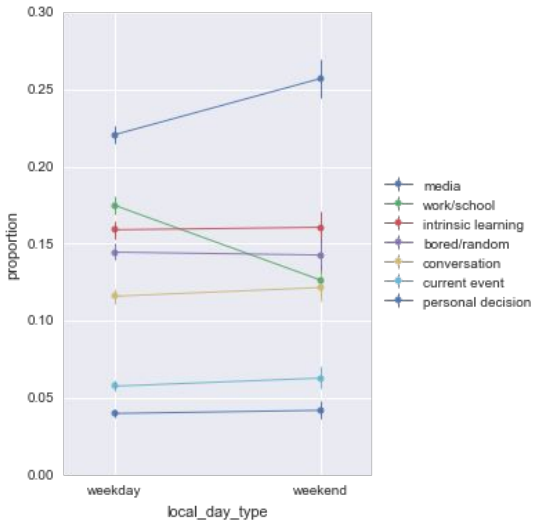

This just confirms our experimental setup has some merit :)

The bar bet hypothesis...

Users motivated by a conversation are more likely to be looking up a fact.

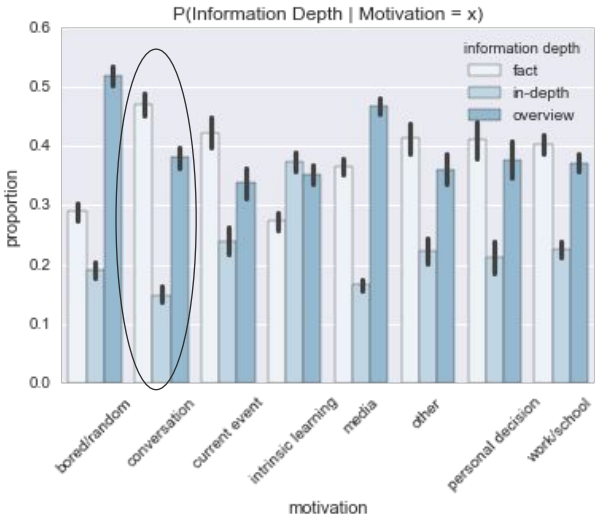

Bored people surf more

If you look at averages, then work/school is also hefty

Work/school spend more time per page.

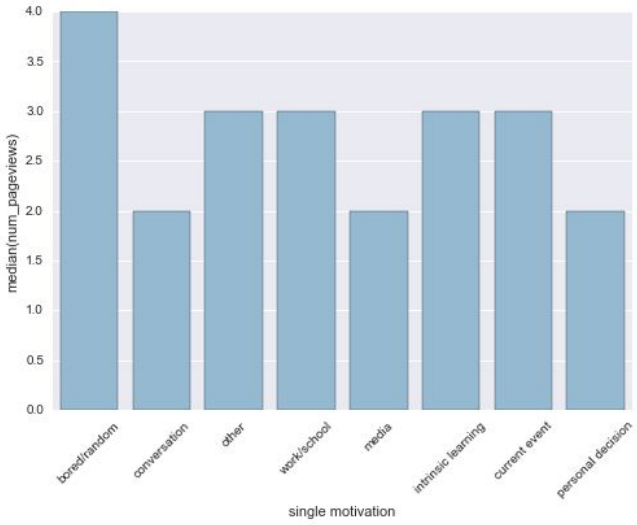

Work/school means more time spent

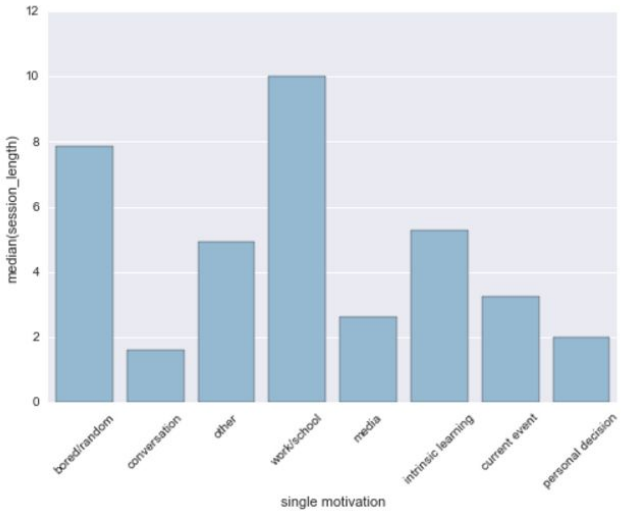

Work/school means more external searches

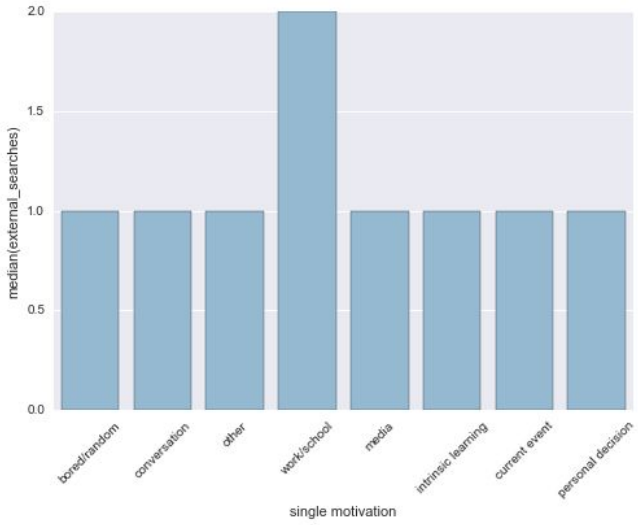

These are re-entries during a session from an external search engine!!

Co-related motivations*

- Bored & Media
- Work/School & Intrinsic learning

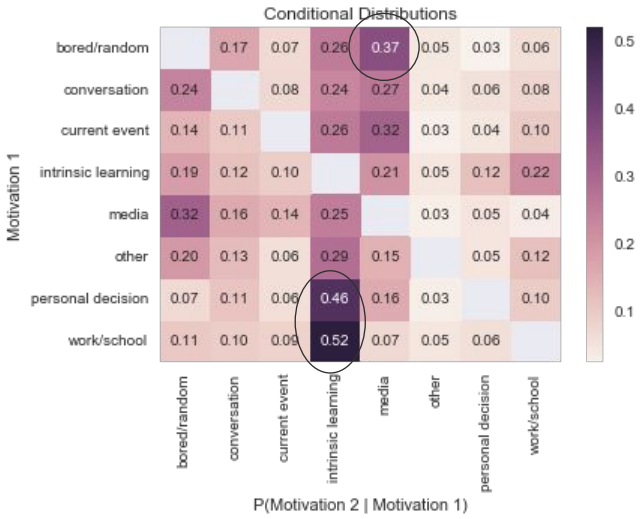

*For the 30% of readers who listed >1 external motivation:

People intrinsically interested in what they study/work on!

We also asked the software to identify the characteristics for which various motivations overrepresented.

Images:
mine
https://commons.wikimedia.org/w/index.php?curid=10485596
Dog
https://commons.wikimedia.org/w/index.php?curid=20741549

Media: __ times more likely...

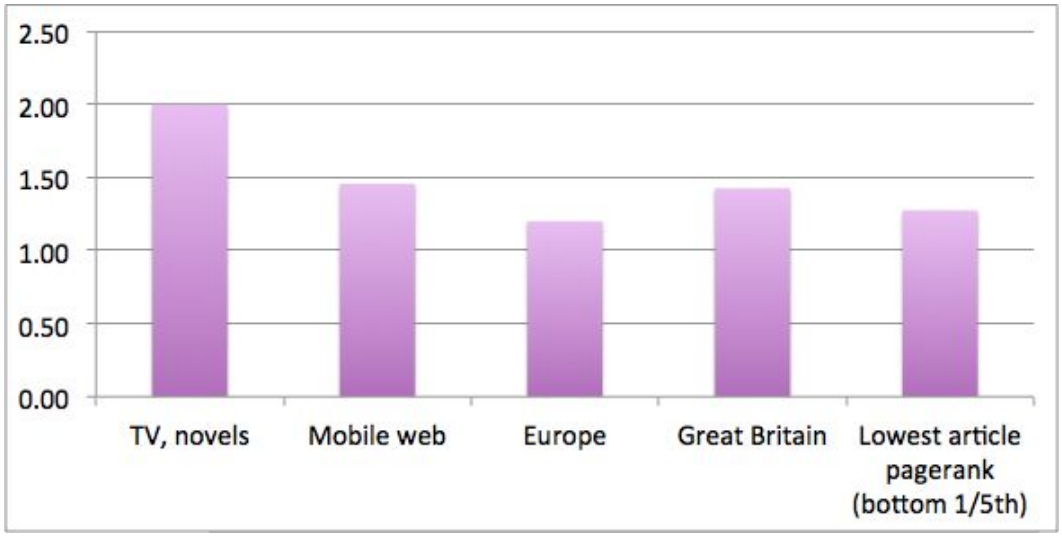

For instance, media was ~20% of users, but 40% of all visitors on tv/novels. Nothing crazy suprising here, but validation that we picked proper things from the graphs.

Work/school: __ times more likely

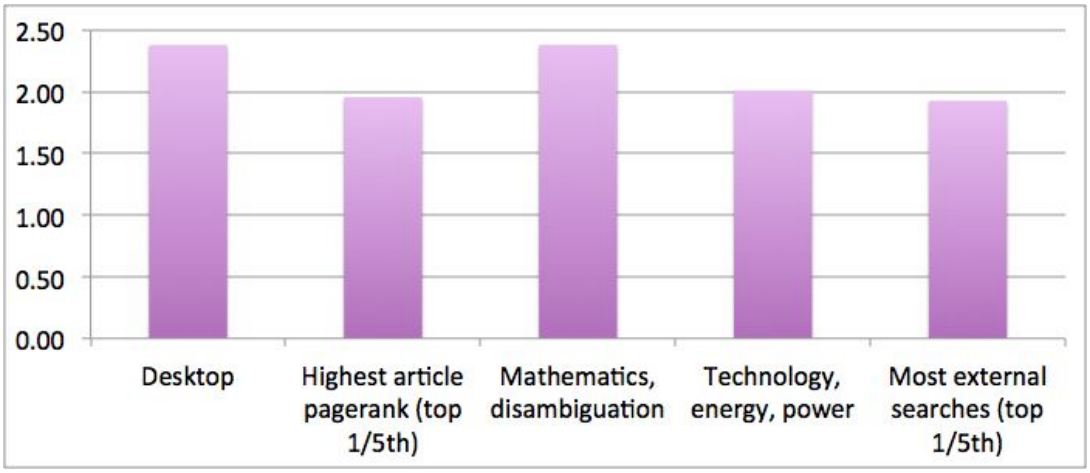

Bored: __ times more likely...

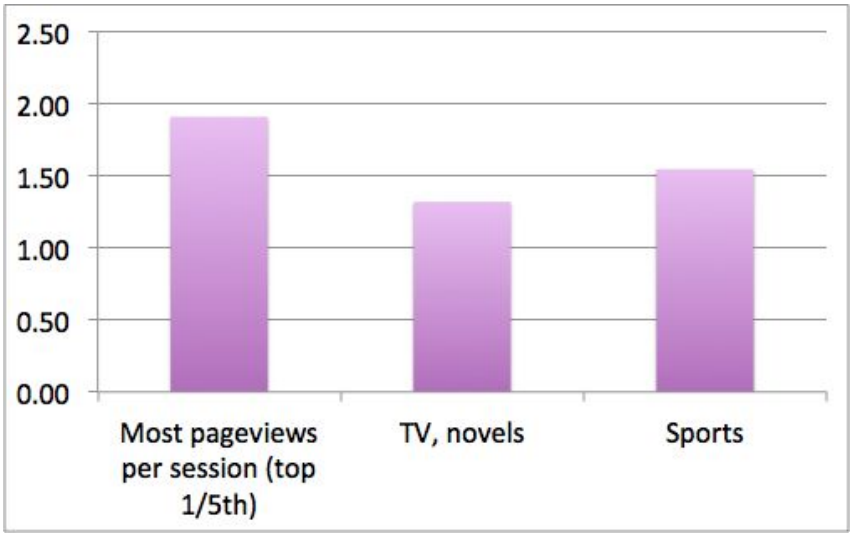

Image: https://commons.wikimedia.org/w/index.php?curid=5356352

Platform matters

We see predictable differences between desktop and mobile.*

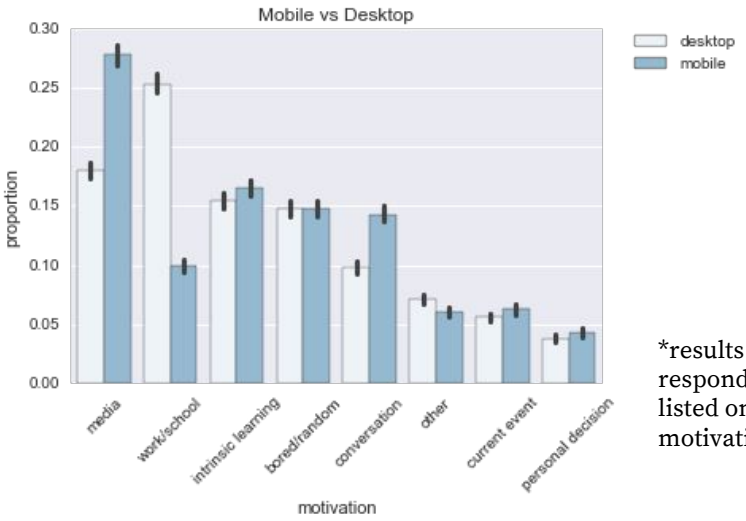

*results limited to respondents who listed one external motivation (70%)

Desktop has deeper sessions

Also more searches, more branching

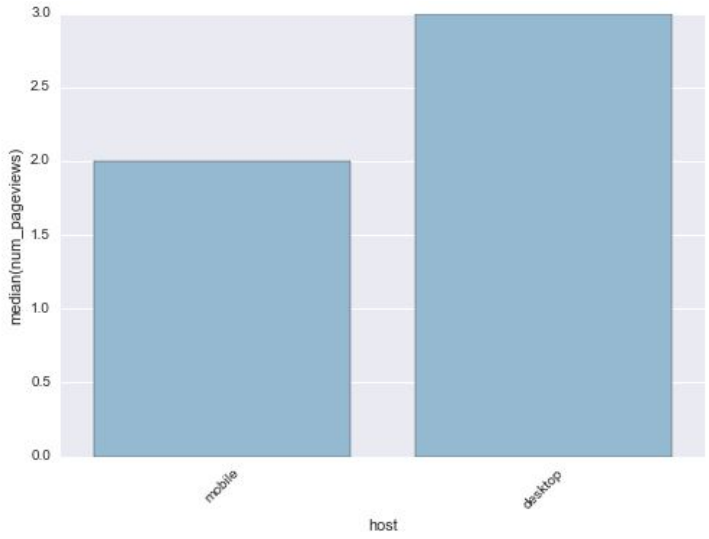

Desktop has longer sessions

More dramatic when you look at time spent

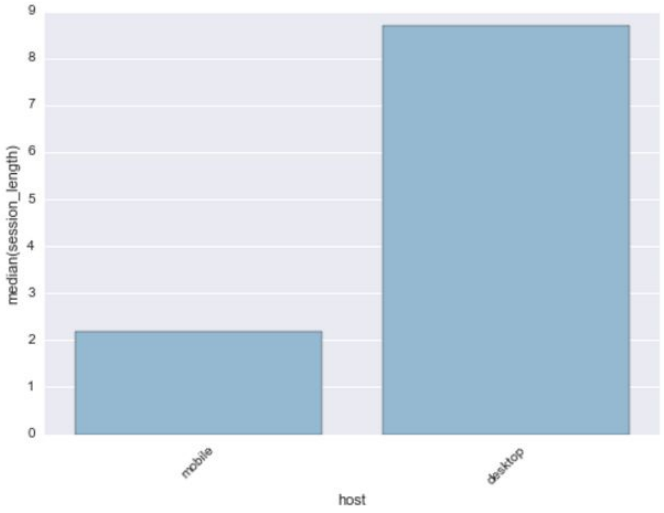

And more time per page

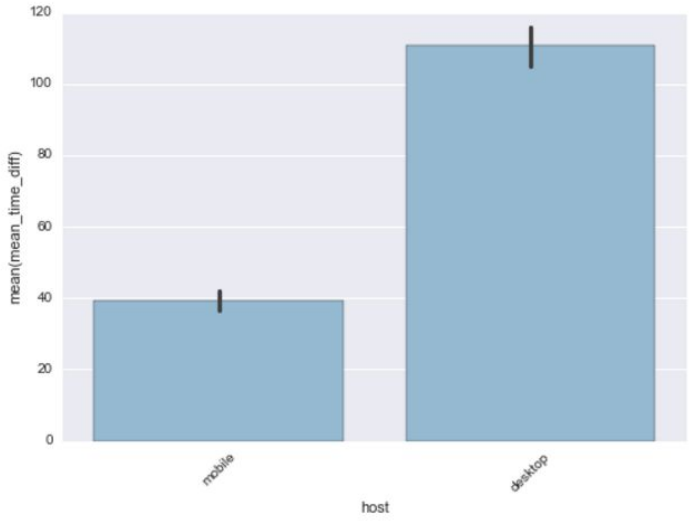

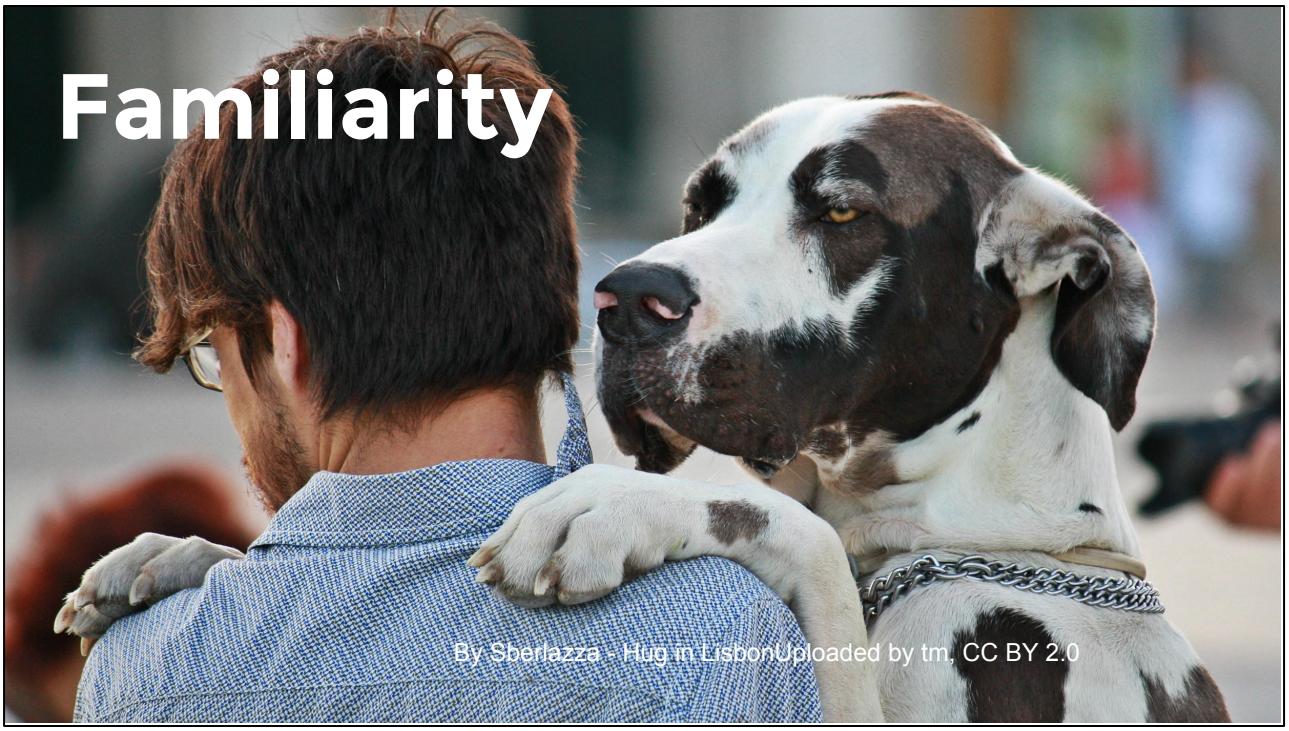

Image:
https://commons.wikimedia.org/w/index.php?curid=29891775

Roughly 50:50

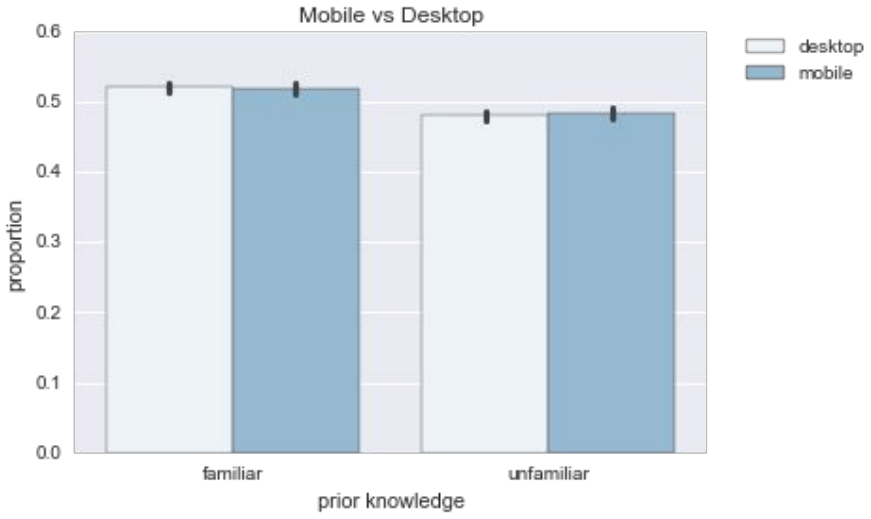

Users unfamiliar with a topic are likely to be driven by media, looking for an overview and spend less time on smaller, less central pages.

Intrinsic learners focus on familiar

Strong connection between media and unfamiliar

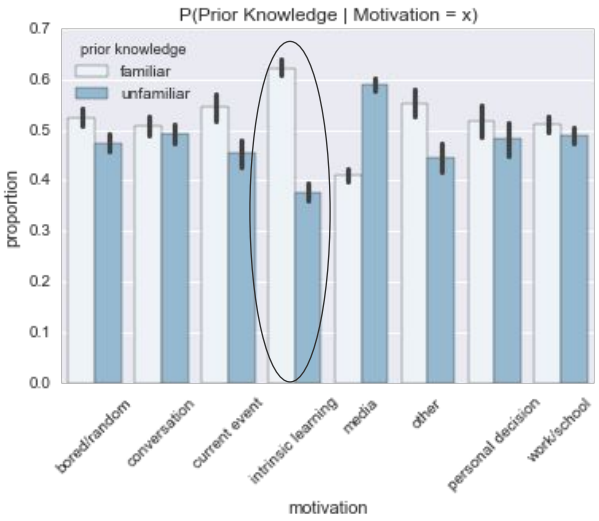

Moar:

Details, caveats and more here:
https://meta.wikimedia.org/wiki/Research:Characterizing_Wikipedia_Reader_Behaviour/S3-English_Large_Scale

More results:
https://github.com/ewulczyn/wiki-readers/tree/master/src/analysis

Paper pending.

What next

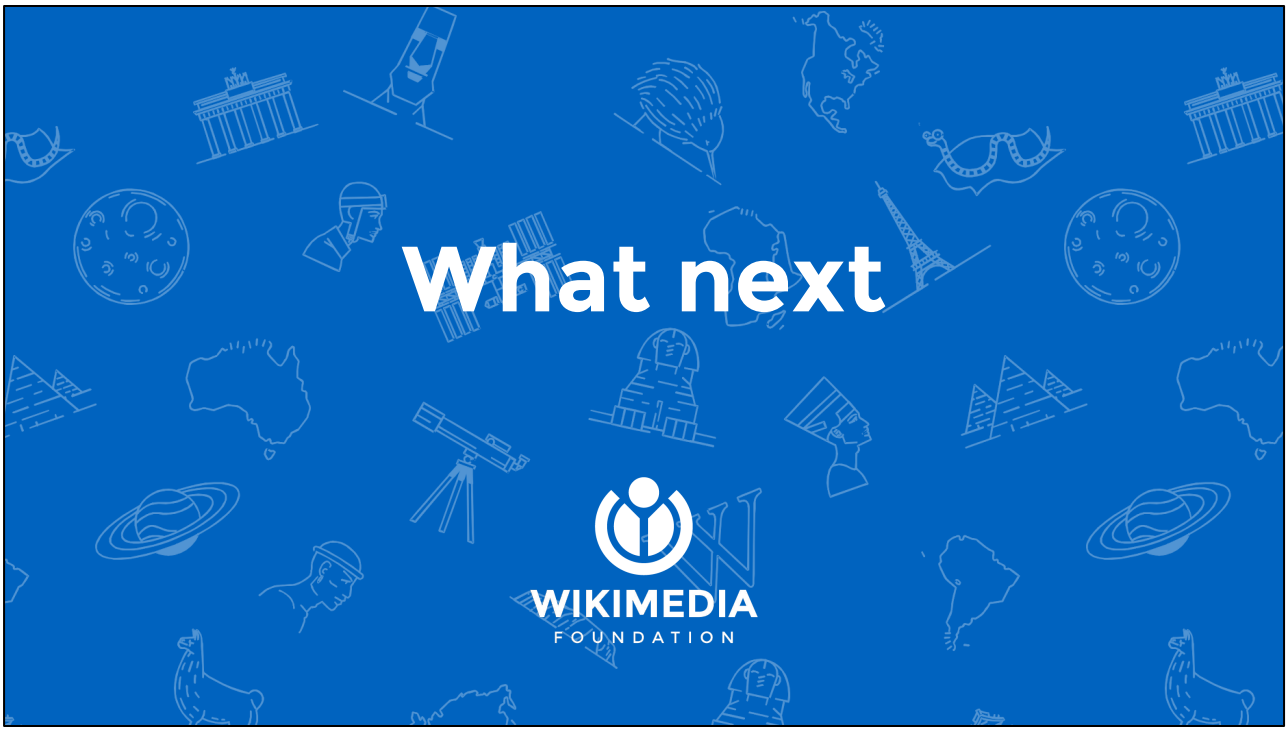

This informs every decision we make.

Some examples

Top of article

What is a Barack Obama?

Before

After

Research Result: Many people want overviews and don't get that far in the page.

Issue: This is slightly facetious, but if you lived on Mars and didn't know who Barack Obama was, you might not even know a Barack Obama was a person...maybe it is a suit, a posture, or a teeth whitener.

Solution: wikidata description, limited image at top, lead paragraph before infobox

Page preview

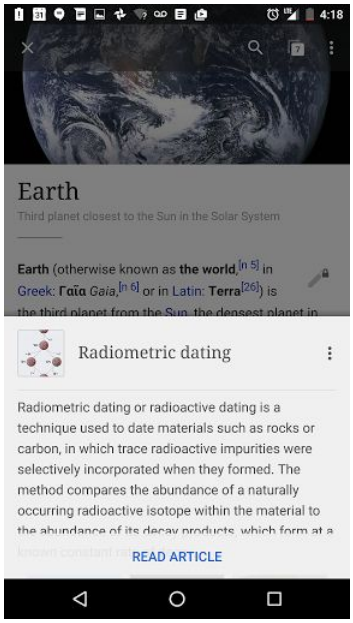

Research Result: users want overviews, mobile users don't open many tabs or click on many links, they don't rabbit hole.

Issue: Presumably, the cost of opening a new page on mobile is greater: context switching, tabs suck, wait for page to load, incur data costs.

Solution: surface the first paragraph when someone clicks on a link, users don't have to switch contexts. The result: links clicked per page jump ~15%.

The feed

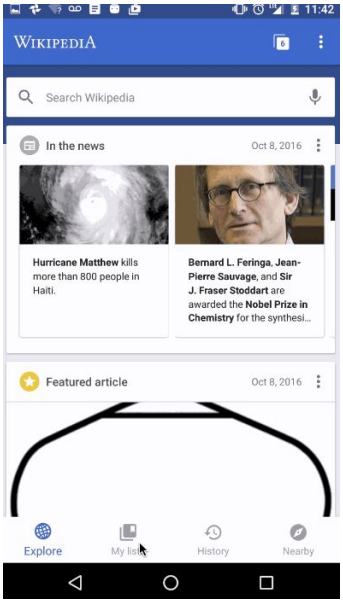

Research Result: users come because they are bored and they browse around, not because they have something specific in mind

Issue: The burden is on the user to ask a question or search a specific topic...if you're bored you want someone to suggest something to you.

Solution: a feed of suggested topics based on the news, article of the day, trending articles, things related to articles you recently read. Already contributing ~5% of pageviews on the app

Related pages

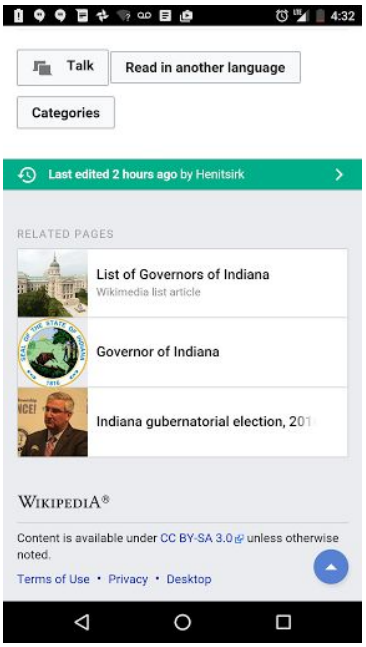

Research Result: users come because they are bored and they browse around, not because they have something specific in mind. They tend not to see multiple pages on mobile.

Issue: The burden is on the user to ask a question or search a specific topic...if you're bored you want someone to suggest something to you.

Solution: on mobile suggest an article based on the one they are currently reading. Only show it to readers who reach the bottom of the article without leaving so as to not compete with "See also sections". For the article above, US Vice Presidential candidate Mike Pence, there was no see also section as of 10/10/2016. See a click through rate of 20% of people who reach the bottom of the page.

Reading lists

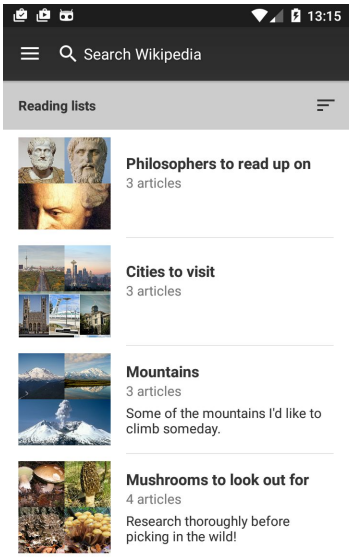

Research Result: 25% of users come for in-depth learning. >10% for school or work. They tend not to visit multiple pages on mobile.

Issue: Hard to save something for later on mobile, because bookmarking isn't robust.

Solution: on mobile, create saved pages feature. Can save groups of saved pages so that users can collect pages by topic for revisiting later. Additional benefit: pages are saved for offline use.

Opportunities to tailor content are limited

How do you think we should use this information?

Image: https://commons.wikimedia.org/w/index.php?curid=43476900

Image: https://commons.wikimedia.org/w/index.php?curid=14

THANK YOU
WIKIMEDIA
FOUNDATION

Skipped slides

Regions

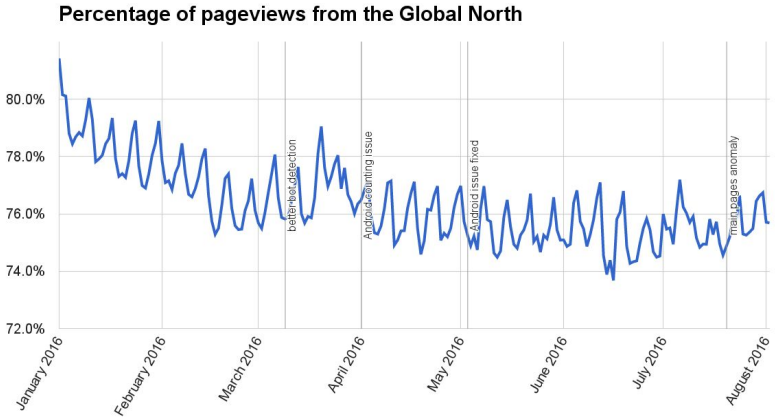

After a notable decrease earlier this year (or, increase of the Global South ratio), this has been relatively steady in recent months.

NB: We are currently rethinking this metric and might replace it with a different country selection constructed as part of the work on the New Readers project.

For more detailed updates on our core metrics, see [[mw:Reading/Readership metrics reports]] and [[mw:Wikimedia Product]].

Size and breakdown

- 478m pageviews/day, across all projects
  - Desktop 51.7%
  - Mobile web 46.9%
  - Mobile apps 1.4%

For detailed updates on core Reading metrics, see [[mw:Product]] and [[mw:Reading/Readership metrics reports]]

Week until May 29, 2016

Unique devices

Project	Monthly unique devices (desktop + mobile web) July 2016	% mobile
English Wikipedia	559 million	65%
Spanish Wikipedia	112 million	62%
German Wikipedia	78 million	56%
Japanese Wikipedia	54 million	66%
Russian Wikipedia	61 million	55%

Monthly pageviews/device English Wikipedia July 2016 (not corrected for main pages anomaly)	
Desktop	22
Mobile web	9

Top 5 projects by overall unique devices - no global number available yet)
NB: devices <> users

Recently introduced metric, no long-term trends yet. But interesting that desktop pageviews/device have gradually increased on enwiki over the last 7 months.

Initiatives

Strategic initiatives

Here are our more specific initiatives. The bulk of our efforts are in the first initiative. 2nd initiative is in research mode. The 3rd is smaller and still kicking off.
Icons belong to the noun project

Strategy

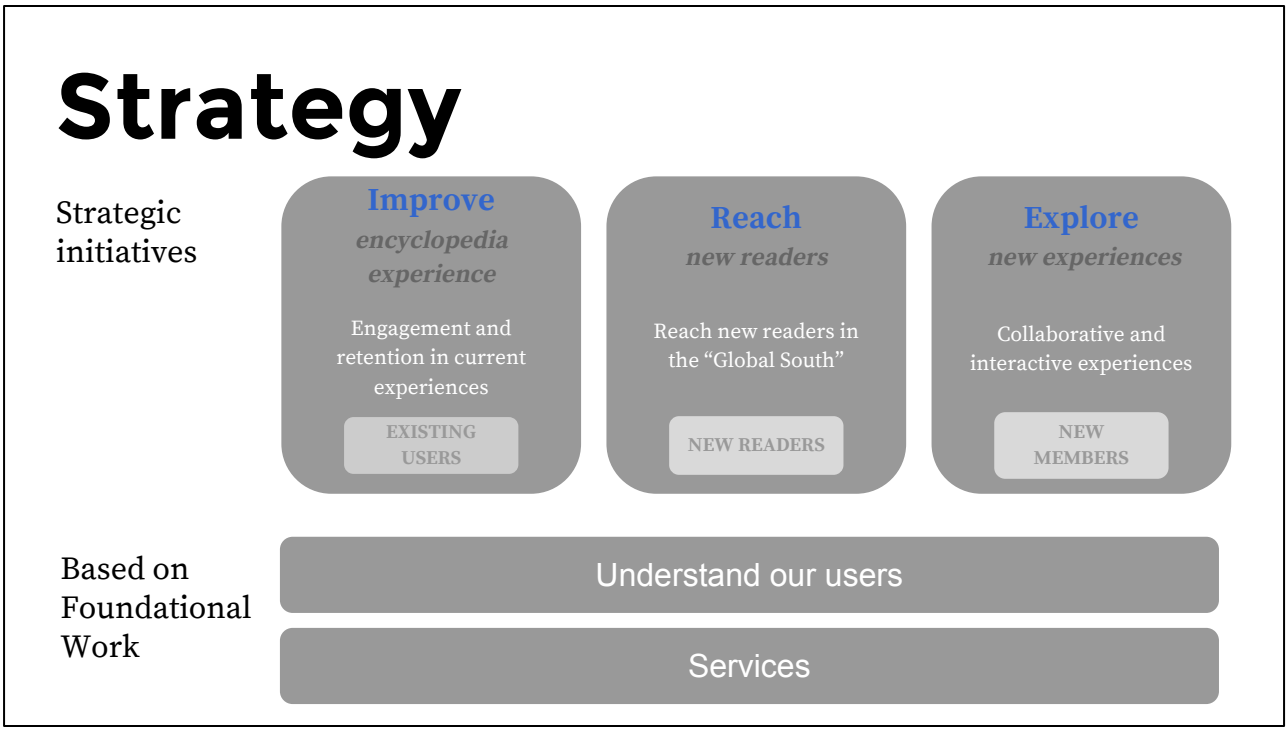

Here are our more specific initiatives. The bulk of our efforts are in the first initiative. 2nd initiative is in research mode. The 3rd is smaller and still kicking off.
Icons belong to the noun project

Intrinsic motivation: __ times more likely…

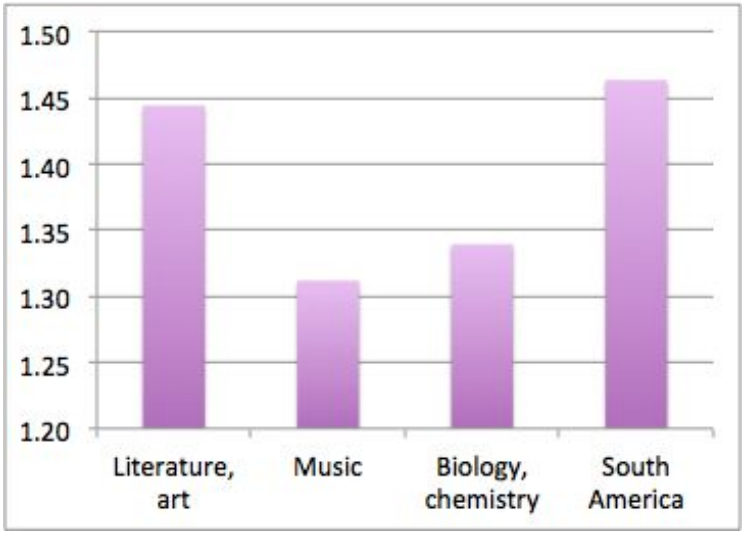

Desktop has deeper sessions

More dramatic when you look at time spent

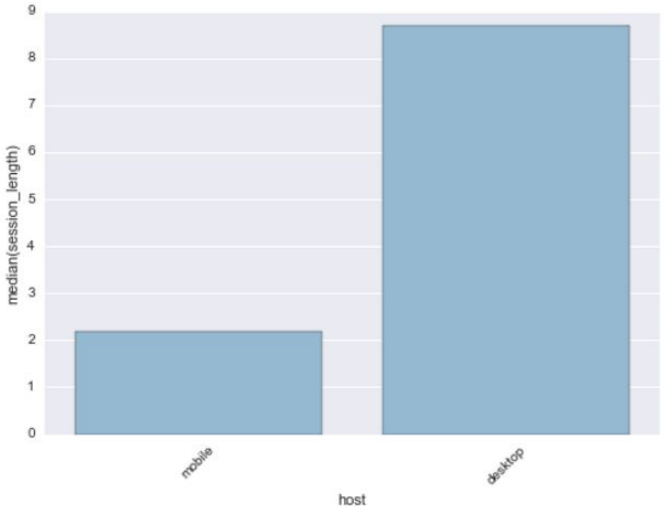

Unfamiliar leads to overviews

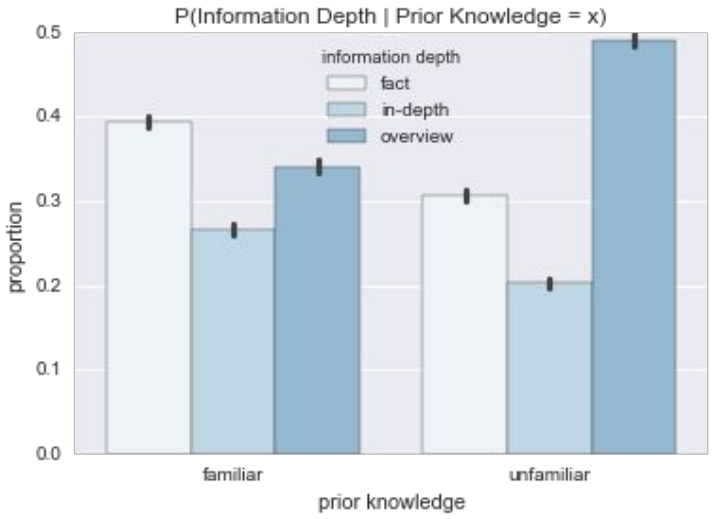

Unfamiliar leads to shorter sessions

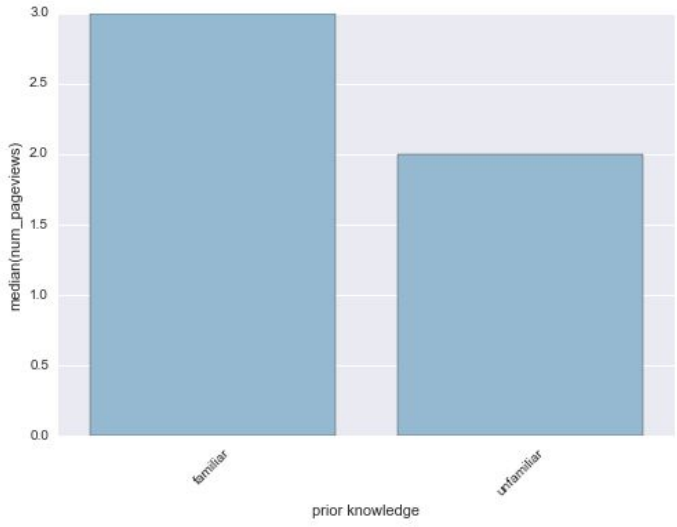

Traffic is highly seasonal

Wikimedia pageviews year-over-year comparison, May 2013-July 2016

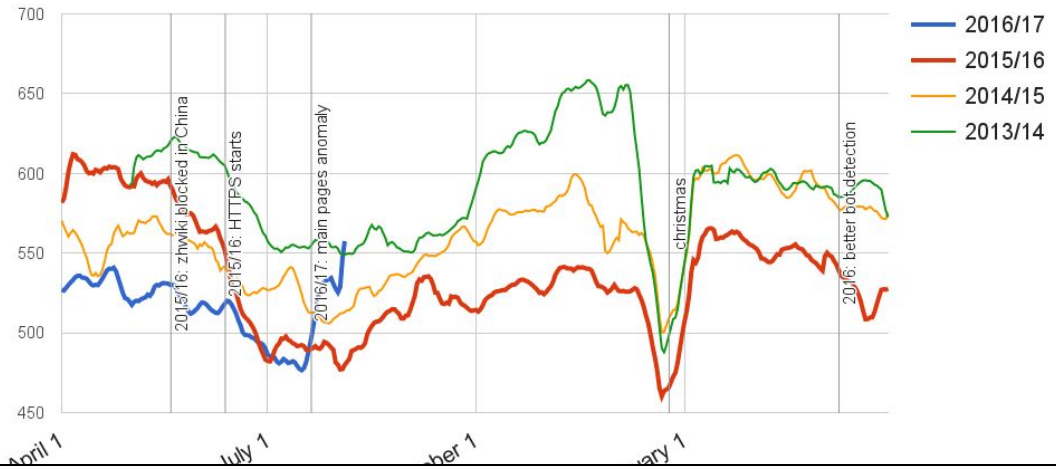

Why are you reading this article today?

I am reading this article to

- ○ look up a specific fact or to get a quick answer.
- ○ get an overview of the topic.
- ○ get an in-depth understanding of the topic.

Prior to visiting this article

- ○ I was not familiar with the topic and I am learning about it for the first time.
- ○ I was already familiar with the topic.

I am reading this article because

Please select all answers that apply

- ☐ I want to know more about a current event (e.g. a soccer game, a recent earthquake, somebody's death).
- ☐ I am bored or randomly exploring Wikipedia for fun.
- ☐ the topic was referenced in a piece of media (e.g. TV, radio, article, film, book).
- ☐ this topic is important to me and I want to learn more about it. (e.g., to learn about a culture).
- ☐ I have a work or school-related assignment.
- ☐ I need to make a personal decision based on this topic (e.g. to buy a book, choose a travel destination).
- ☐ the topic came up in a conversation.
- ☐ Other: