Automatic approaches to bridge knowledge gaps in Wikimedia projects

WMF Research Team

Wikimedia Research Showcase October 2021

The Research Team

...and many formal collaborators

Leila Zia
Director, Head of Research

Pablo Aragón
Research Scientist

Martin Gerlach
Research Scientist

Isaac Johnson
Research Scientist

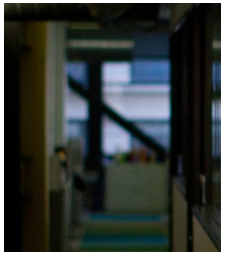

Fabian Kaelin
Senior Research Engineer

Emily Lescak
Senior Research Community Officer

Miriam Redi
Research Manager

Diego Sáez-Trumper
Senior Research Scientist

Learn more: https://research.wikimedia.org

The strategic direction: Knowledge Equity
[from Movement Strategy]

Knowledge equity: As a social movement, we will focus our efforts on the **knowledge and communities that have been left out by structures of power and privilege**. We will welcome people from every background to build strong and diverse communities. We will **break down the social, political, and technical barriers** preventing people from accessing and contributing to free knowledge.

The research program: Addressing Knowledge Gaps

Identify gaps

September 2020:

https://w.wiki/4DWU

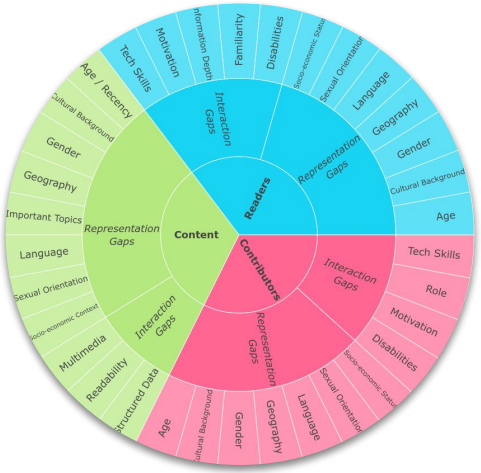

Measure gaps

Ongoing:

https://w.wiki/4Exh

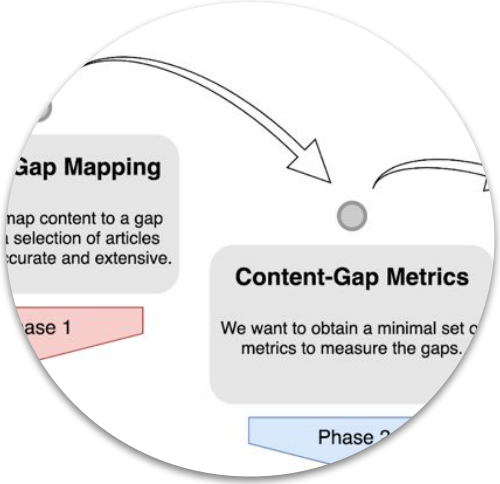

Bridging gaps

Today:

https://w.wiki/4DWe

Bridging knowledge gaps: Automated tools for content discovery and recommendation.

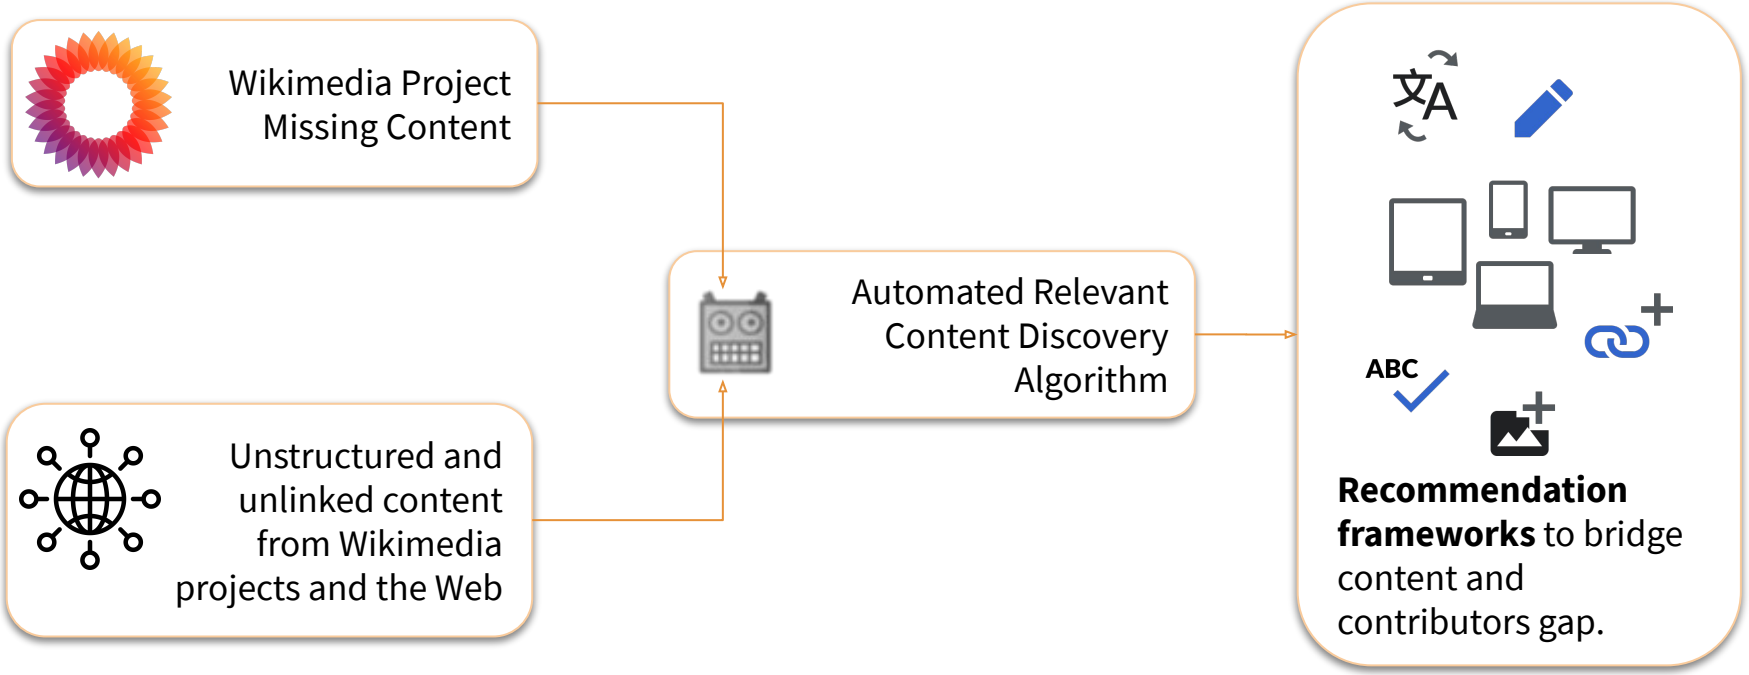

Bridging knowledge gaps: Previous research work from our team

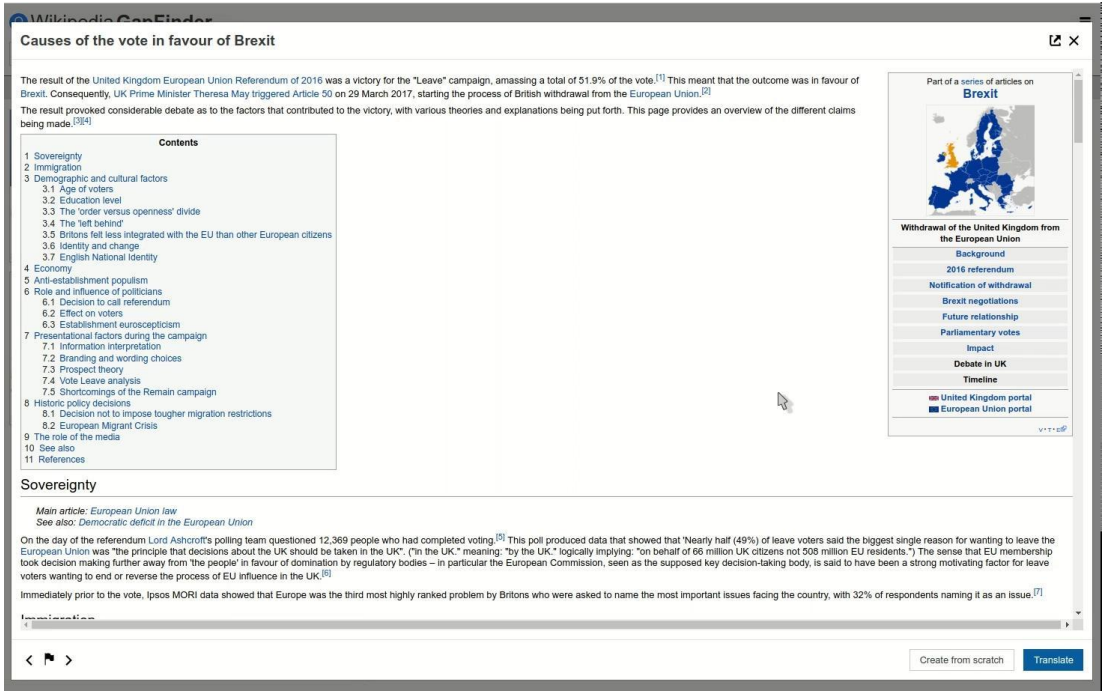

Bridging gaps: Developing automatic tools

- Section alignment (Diego Sáez-Trumper) https://w.wiki/Bb6
  - In Collaboration with the Language team
- Link recommendation (Martin Gerlach) https://w.wiki/4BCP
  - In Collaboration with the Growth team
- Image recommendation (Miriam Redi) https://w.wiki/4DW8
  - In Collaboration with the Growth and Structured Data teams
- Equity in recommendations (Isaac Johnson) https://w.wiki/$4P

Section alignment

Diego

WIKIMEDIA
FOUNDATION

Problem: Translation engines can not be directly used on Wikipedia

- We want to compare sections and share content across languages.
- We can’t compare sections across languages.

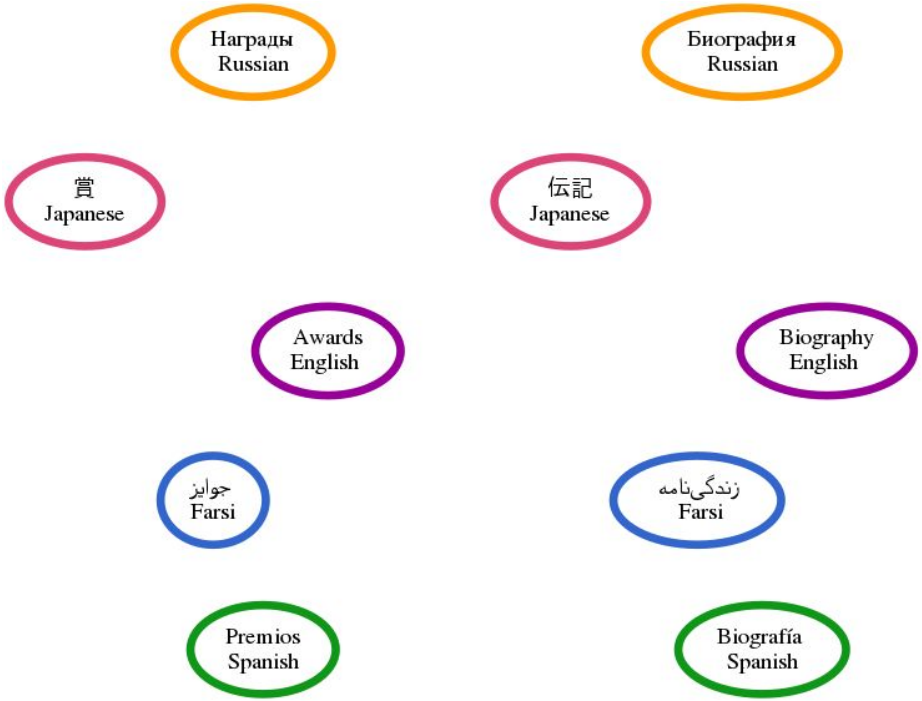

Solution: Cross-lingual section alignment

- Create an ML system to align sections across languages.

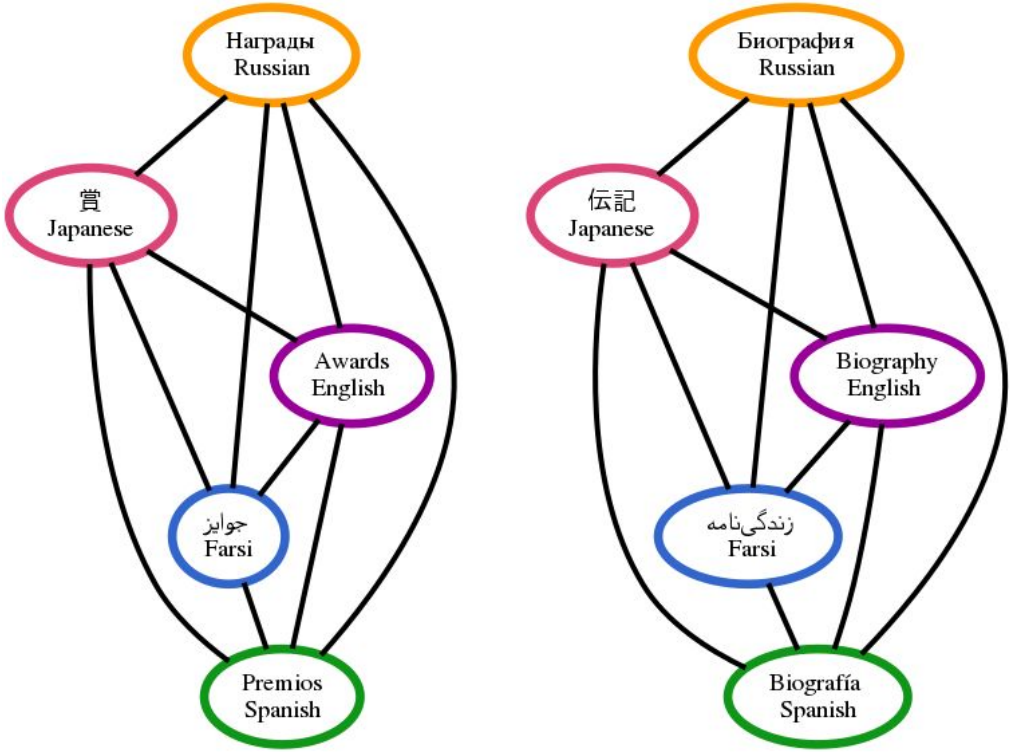

Step 1: Let’s try to do better than Automatic Machine Translation

- Ask the community to help annotate data:
  - ar,es,fr,ja,ru,en (all pairs)
- Cross-lingual word-embeddings + Wikipedia specific features.

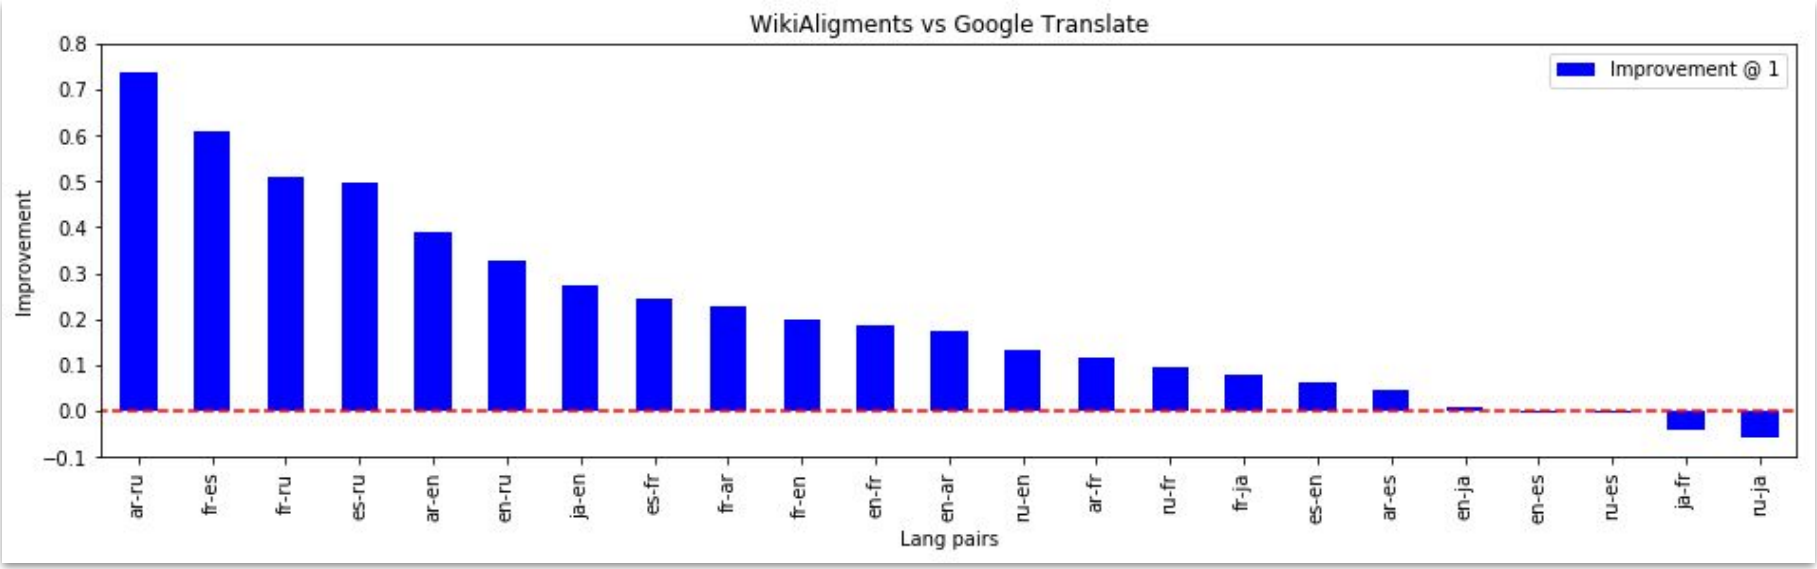

Step 1: Let’s try to do better than Automatic Machine Translation

- Ask the community to help annotate data:
  - ar,es,fr,ja,ru,en (all pairs)
- Cross-lingual word-embeddings + Wikipedia specific features.

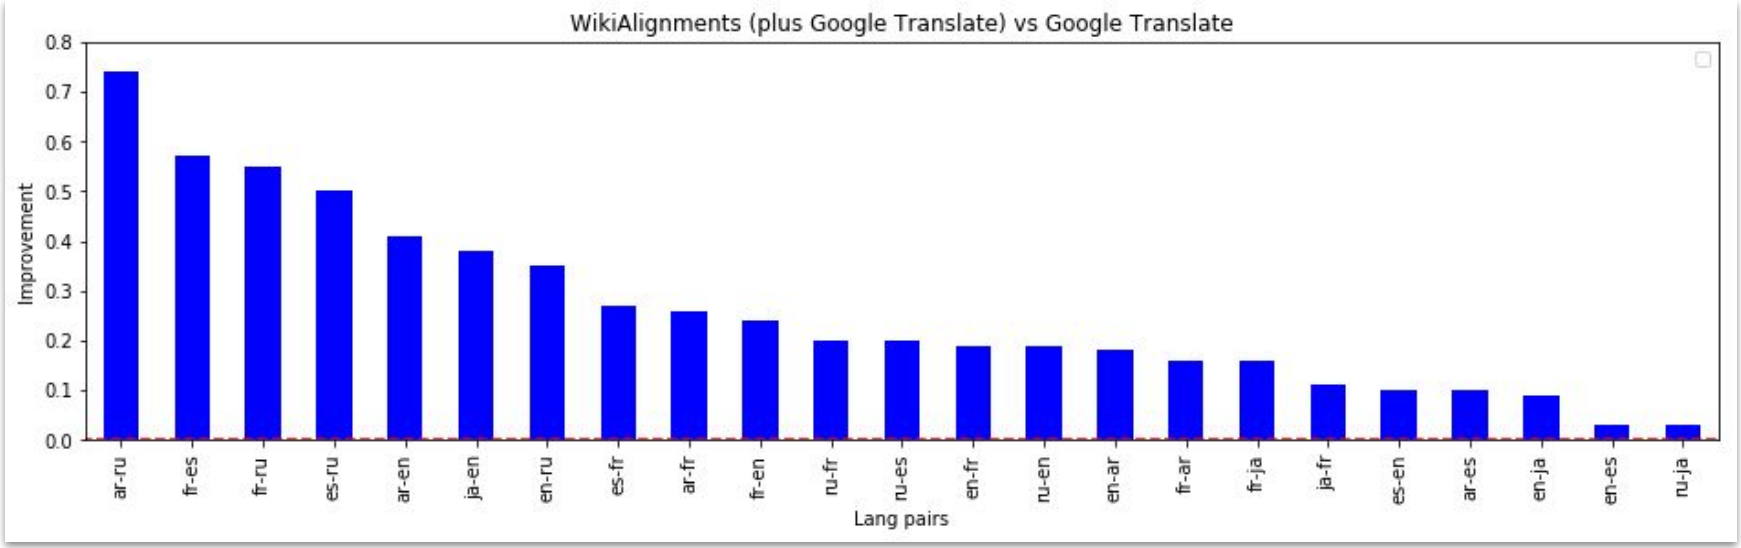

Learn more and test the system!

[[m:Research:Expanding_Wikipedia_articles_across_languages/Inter_language_approach]]

https://secrec.wmflabs.org

Section Recommendation Demo

Wiki: en.wikipedia Article: Quilombo Try Another Article

Quilombo - en.wikipedia.org

Article Summary:

A ***quilombo*** is a Brazilian hinterland settlement founded by people of African origin, and others sometimes called Carabali. Most of the inhabitants of quilombos, called quilombolas, were maroons, a term for escaped slaves.

Go to Wikipedia and edit this article

Recommended sections:

- **Etymology** ○ Good ○ Redundant ○ Not Related
- **Argentina** ○ Good ○ Redundant ○ Not Related
- **Infrastructure** ○ Good ○ Redundant ○ Not Related
- **Mexico** ○ Good ○ Redundant ○ Not Related
- **Organization** ○ Good ○ Redundant ○ Not Related
- **Economy** ○ Good ○ Redundant ○ Not Related

Evaluate

Number of sections in other languages:

- **es:** 20 Quilombo
- **fr:** 6 Quilombo (esclavage)
- **ru:** 1 Киломбу

Current sections and subsections in this article are:

- ['Slavery in Brazil', 'History', 'Quilombos', 'Palmares', 'Mola', 'Curiaú', 'Cunani', 'Film', 'Constitution of Brazil', 'In Spanish', 'See also', 'References', 'Further reading', 'External links']

Step 2: Scale up to 100+ languages

- Add new languages and deploy.
- Test new embeddings and language models.
- Improve the recommender system.

Link recommendation

Martin Gerlach (Research Team)
Growth Team (Marshall Miller, Rita Ho, Kosta Harlan, many more)
Djellel Difallah (NYU Abu Dhabi)

WIKIMEDIA
FOUNDATION

Problem

Editing is hard

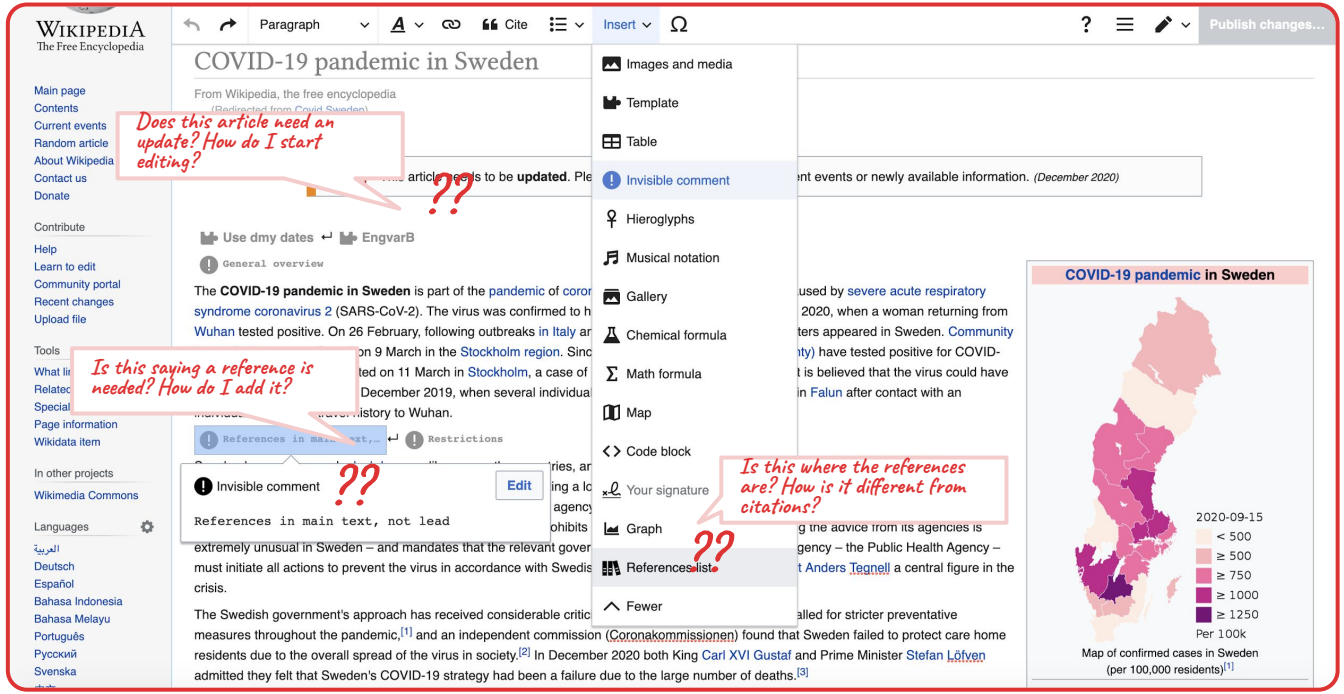

Technical
What is an infobox?

Conceptual
What is notability?

Cultural
Why are people so mean?

Solution

Structured task editing

- Break down editing into simpler tasks
  - Easier to: understand, do on mobile, get positive experience
- We picked the task of: "adding a link"
  - Well-defined, frequent, and attractive task type
- Machine-in-the-loop:
  - Generate recommendations using ML approaches.
  - Editors verify the output and validate the insertion.

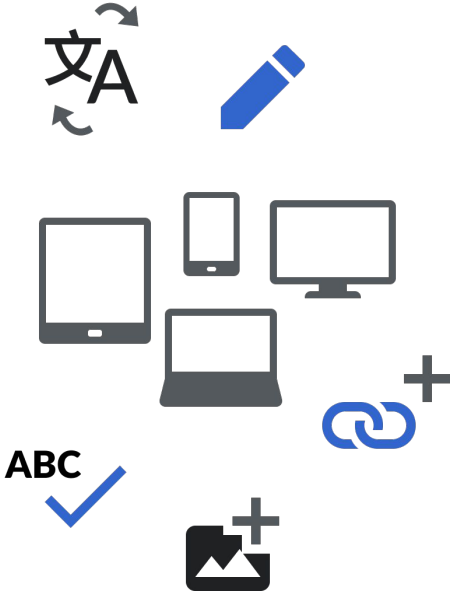

Link recommendation

Hypatia (born c. 350–370; died 415 AD) was a Hellenistic Neoplatonist philosopher, astronomer, and mathematician, who lived in Alexandria, Egypt, then part of the Eastern Roman Empire.

astronomer → Astronomy
astronomer → Astronomer
astronomer → --no link--

?

Entity-linking task

- Mention detection
- Link generation
- Link Disambiguation

The Add-a-link Task in Wikipedia

- Language support
  - 300+ language version of Wikipedia; highest impact for smaller communities.
- Other considerations
  - Manual of style constraints
  - Prefer simpler over complex models (scalability, transparency, etc)
  - Utility: find a balance between precision over recall

Step 1: Mention detection

- Build mention dictionary from all existing links
  - e.g. English Wikipedia: 7M anchors, 170M links
- String-matching sweeping window for all possible n-grams (N=1...10)
- Give preference to larger N

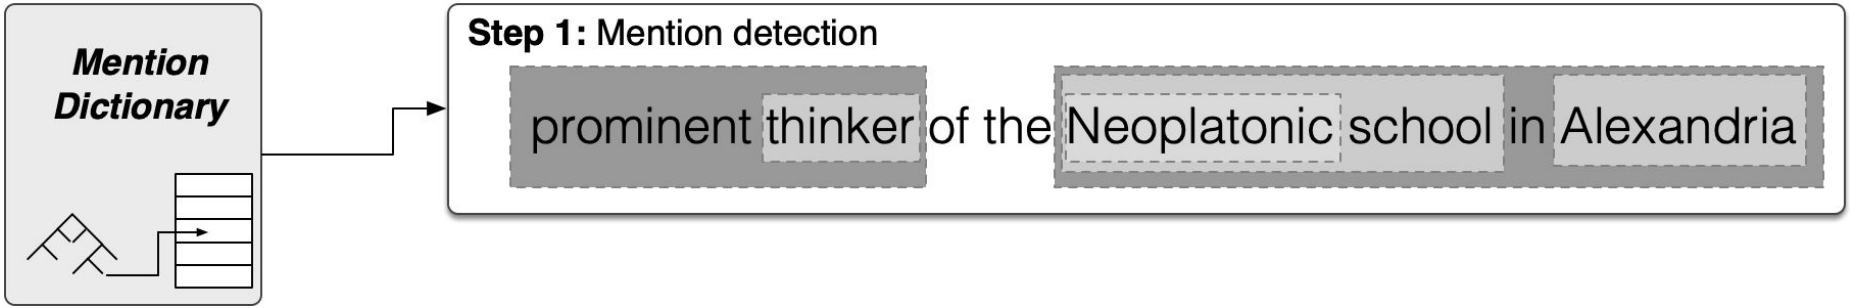

Step 2: Link generation

- From the anchor-dictionary extract all used links
- Drops links based on constraints (type-based)
- Drop links based on link probability heuristic:
  - Text to Link ratio < 6.5% (picked empirically, and supported by previous work)

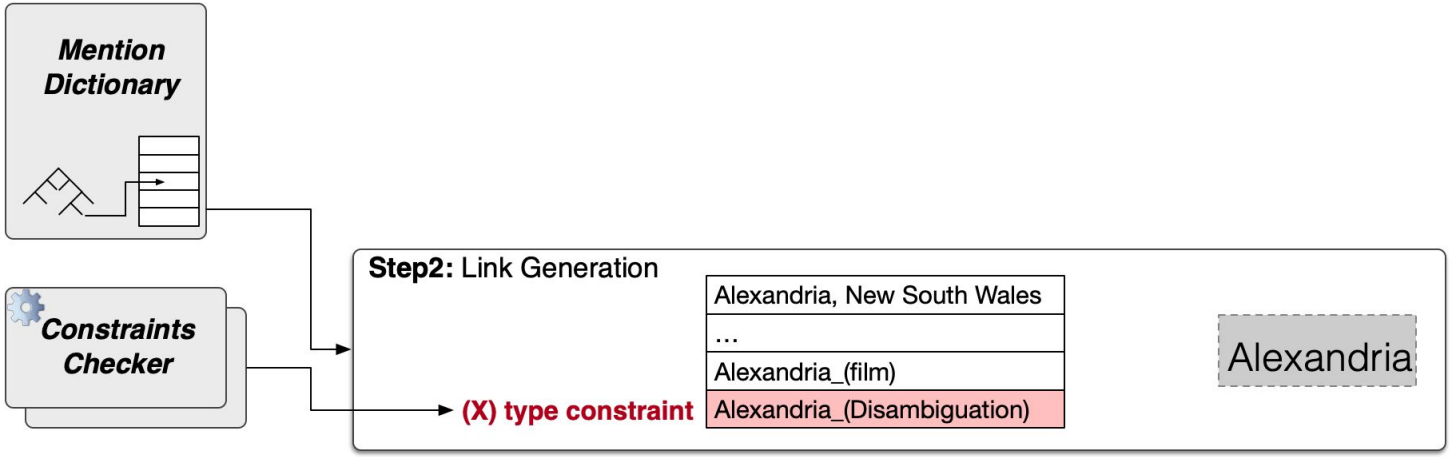

Step 3-a: Link disambiguation- features

- **N-gram size:** the number of tokens in the anchor (based on simple tokenization).
- **Frequency:** count of the anchor-link pair in the anchor-dictionary.
- **Ambiguity:** how many different candidate links exist for an anchor in the anchor-dictionary.
- **Kurtosis:** the kurtosis of the shape of the distribution of candidate-links for a given anchor in the anchor-dictionary
- **Levenshtein-distance:** a string similarity measure between the anchor and the link, e.g., the Levensthein-distance between “kitten” and “sitting” is 3.
- **Wiki2Vec Distance (entity embedding):** similarity between the article (source-page) and the link (target-page) based on the content of the pages.

Step 3-b: Link disambiguation- classifier

- Extract fully linked sentences from the lead sections
  - Positive example: linked mention with correct link
  - Negative example: linked mention with incorrect link, unlinked mention
- Train a binary classifier (XGBoost)

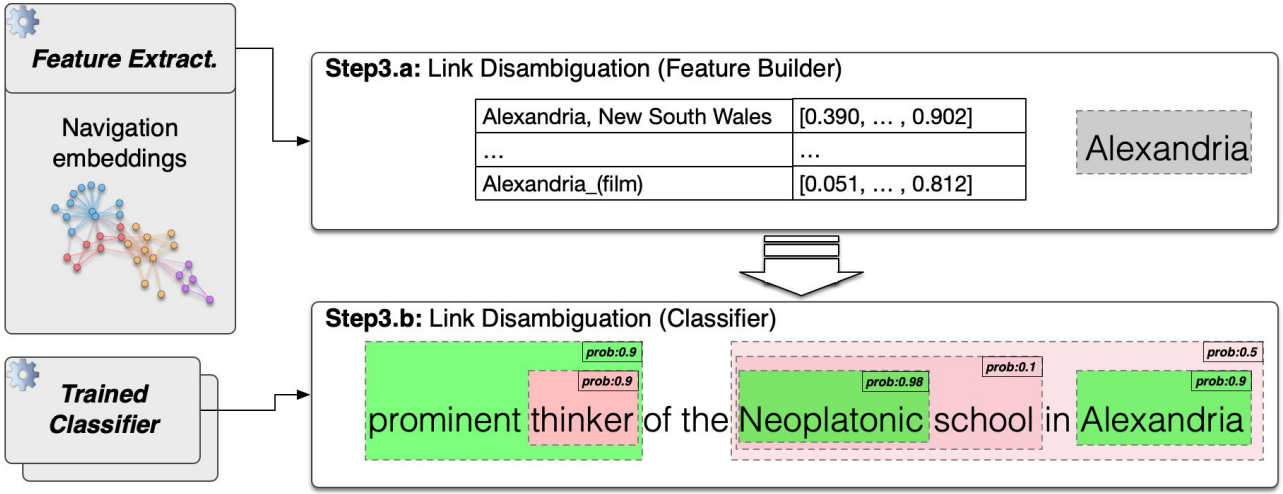

Evaluation

KaiO.Reid, CC BY-SA 4.0

Held-out testset + **Manual evaluation**
(thanks: Bennoit Evellin, Habib Mhenni, Martin Urbanec, Bluetpp, -revi)

Tested Wikis: Arabic, Bengali, Czech, English, French, Vietnamese

Precision: 70% - 92%
How many suggestions are correct?

Recall: 30% - 66%
How many of the possible links captured?

Link recommendation model

- Training pipeline for each language.
- Models/datasets published publicly
- Link recommendation API on kubernetes

Get link recommendations

Discussion Updated 24 June 2021

GET `/service/linkrecommendation/v1/linkrecommendations/wikipedia/{language}/{title}`

Get a set of possible links that can be added to improve an article on Wikipedia.

https://api.wikimedia.org/wiki/API_reference/Service/Link_recommendation

User interface

Evaluate the suggestion

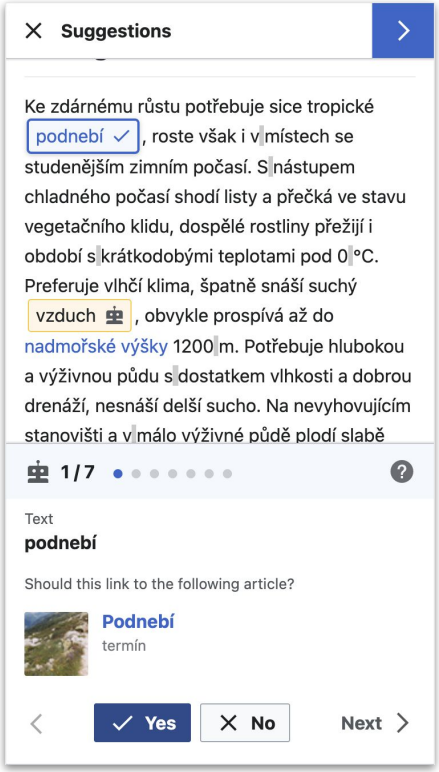

Feedback on algorithm

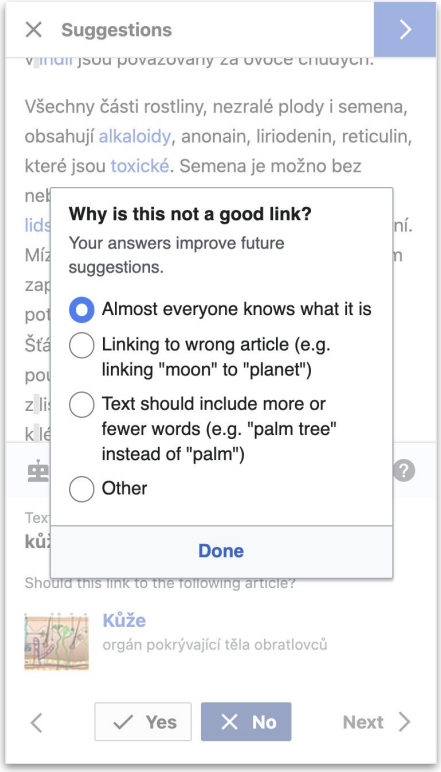

Edit summary

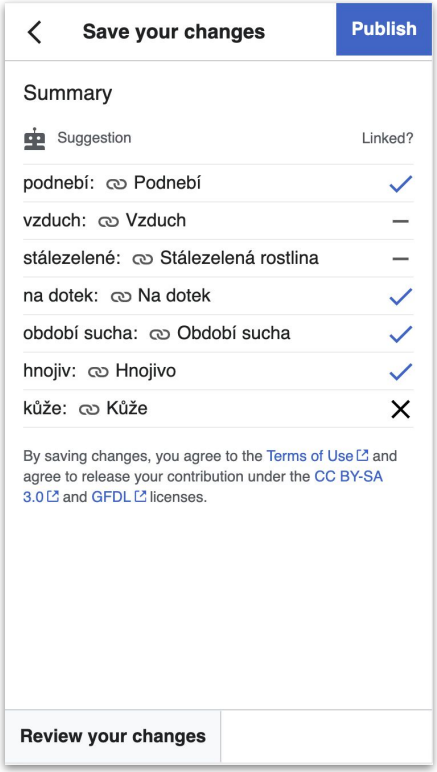

Next suggestion

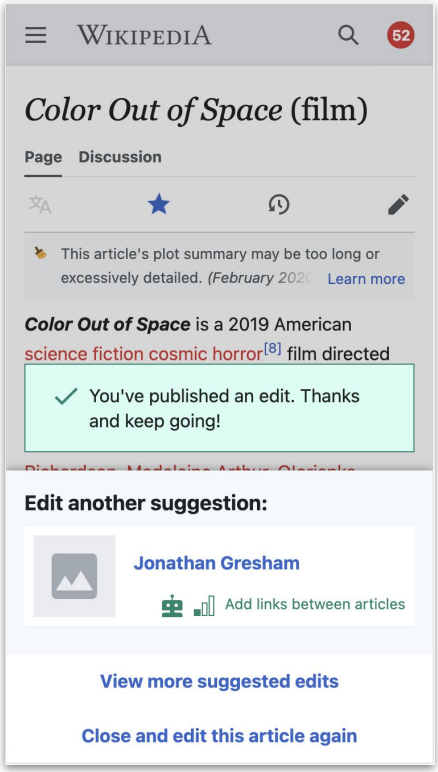

In practice

- Results from pilot-wikis (Arabic, Bengali, Czech, Vietnamese)
  - Newcomers prefer structured editing
  - Revert rate of structured edits is much lower (7.9% vs 25.5% for unstructured tasks)
  - Careless editing is rare
  - Reactions from community mostly positive
- Deployment
  - Currently: Arabic, Bengali, Czech, French, Hungarian, Persian, Polish, Romanian, Russian, Vietnamese
  - Planned: Catalan, Hebrew, Hindi, Korean, Norwegian, Portuguese, Simple English, Swedish, Ukrainian

Image recommendation

Miriam

WIKIMEDIA
FOUNDATION

Problem: **Wikipedia is missing images**

Wikipedia Pages Wanting Photos 2021

Home | Participating Communities | Organizing Team | Participate | Results | Resources | FAQ & Campaign Rules

Contents

Please do not add photos in a language you do not speak. Do not add images without a good caption. In your edit summary, add a good edit summary, and the hashtag #WPWP.

Wikipedia Pages Wanting Photos (WPWP) is an annual campaign where Wikipedia editors across the world, Wikipedia language projects and communities add photos to Wikipedia articles lacking photos. This is to promote the use of digital media files collected from various Wikimedia photography contests, photowalks organized by the Wikimedia community, on Wikipedia article pages. Photos help to grasp the reader's attention better than a wall of text, illustrate content, and make the article more instructive and engaging for readers.

Thousands of images have been donated and contributed to Wikimedia Commons via various advocacy programs, photowalks, and contests including international photography contests such as Wiki Loves Monuments, Wiki Loves Africa, Wiki Loves Earth, Wiki Loves Folklore, etc. Yet relatively few of these photos have been used on Wikipedia articles. Today, the Wikimedia Commons hosts millions of photo images but only a tiny portion of these have been used on Wikipedia article pages. This is a huge gap that this project aims at bridging.

Many Images!

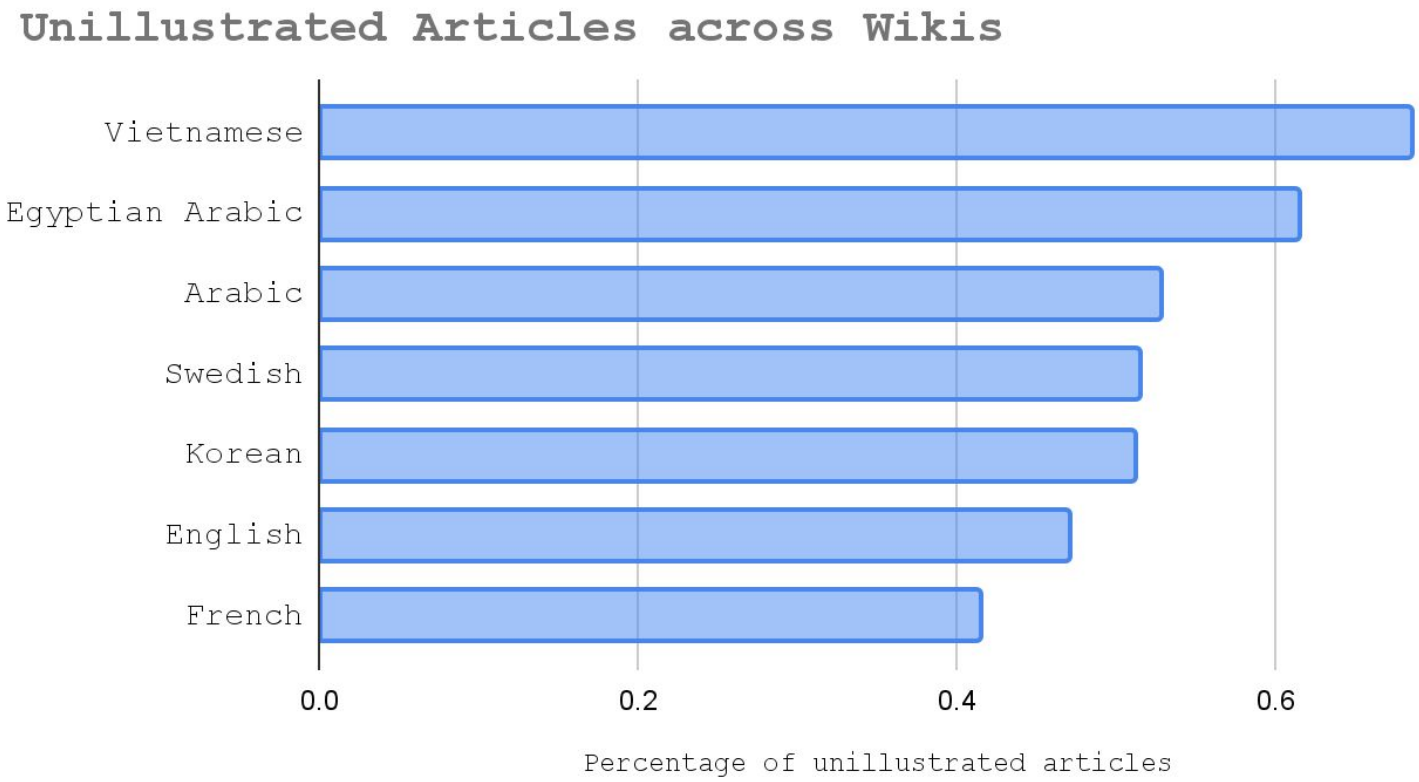
Unillustrated Articles across Wikis
Vietnamese
Egyptian Arabic
Arabic
Swedish
Korean
English
French
0.0
0.2
0.4
0.6
Percentage of unillustrated articles

“Add an Image” Structured Task

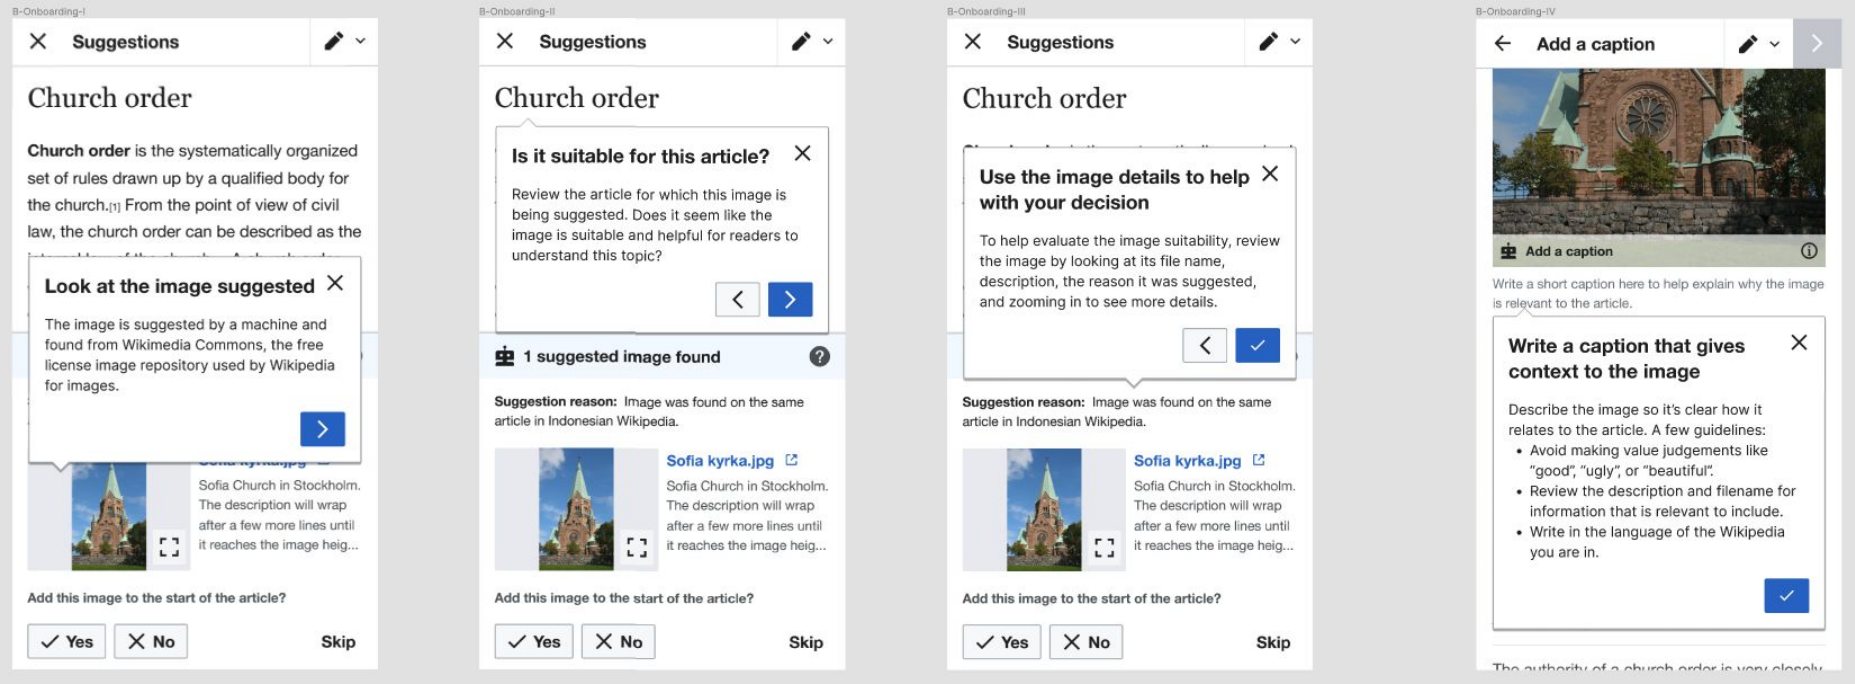

By RHo (WMF) - Mockups from WMF designer, CC BY-SA 4.0,
https://commons.wikimedia.org/w/index.php?curid=108208100

“Add an Image” Algorithm Logic

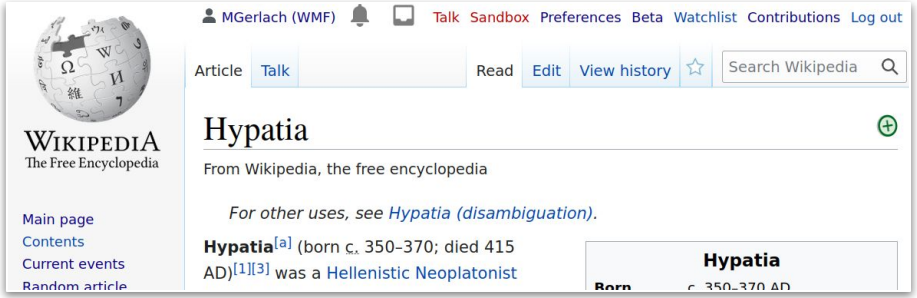

Unillustrated Article

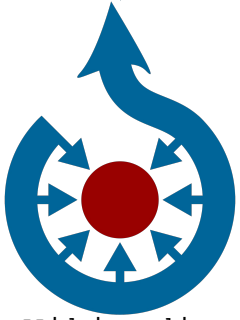

Wikimedia Commons

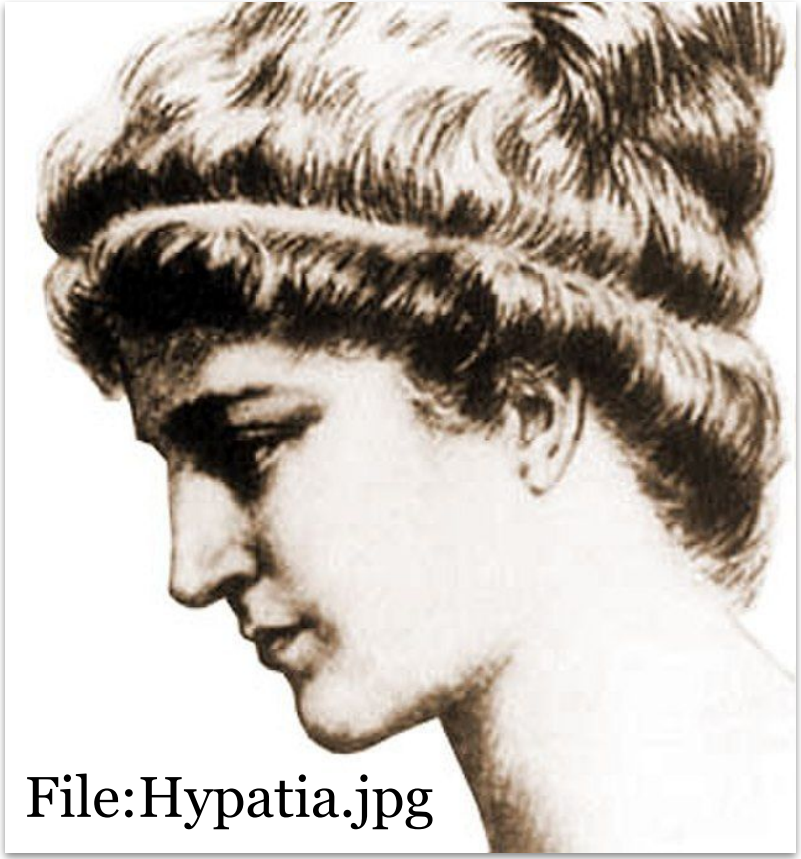

The “Add an Image” Task in Wikipedia

- Manually finding good images for articles is hard:
  - The space of search for good free-licensed image matches can be huge, especially for newcomers
  - Sometimes finding matches is difficult: image search is not very effective in absence of metadata
- Automatically finding good images for articles is also hard:
  - We need to have “explainable” image-article matches
  - Due to the impact and visibility of images, accuracy/quality of matches is crucial
  - Computer vision technologies are not able to identify fine-grained objects, and simpler models are preferred
- Language support
  - Similar to links 300+ language version of Wikipedia

A Simple Solution: **Leveraging Existing Matching Signals**

Discovering unillustrated articles as articles without images or with *icons* only.

Generate list of icons for each Wikipedia

Unillustrated Article

Discovering related images from different sources

Filtering

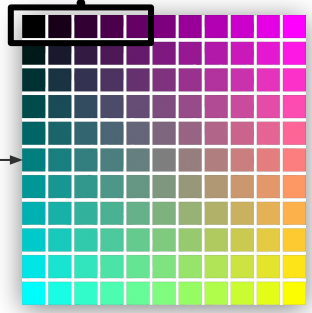

Ranking images according to Relevance and Quality

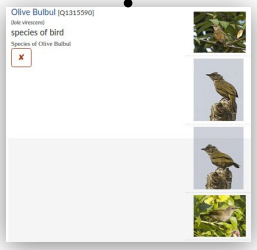

Task recommendation

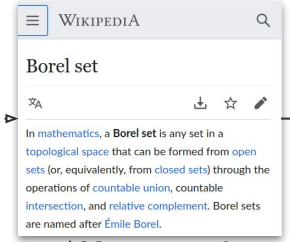

Unillustrated article

Has a **Wikidata** item?

NO

No recommendation

Dynamic unillustrated article detection based on Wiki size

ويكيپيديا
الموسوعة الحرة

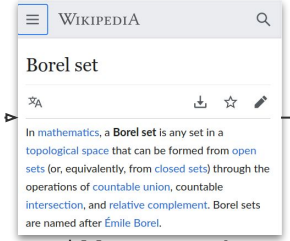

Unillustrated article

Has a **Wikidata** item?

YES

Discover image candidates:

- **P18** (Image) of Wikidata item
- **P373** (Commons Category) of Wikidata Item
- **Lead Images** of same article in other languages that are on Commons (i.e. can be used on all Wikis)

Filter out bad images: placeholders, graphics, flags, maps **Remove articles about years, disambiguations, and lists**

More than 0 candidates left?

NO

No recommendation

NO

Dynamic unillustrated article detection based on Wiki size

ويكيپيديا
الموسوعة الحرة

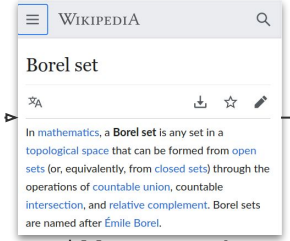

Unillustrated article

Dynamic unillustrated article detection based on Wiki size

Has a **Wikidata** item?

YES

NO

Discover image candidates:

P18 (Image) of Wikidata item
P373 (Commons Category) of Wikidata Item
Lead Images of same article in other languages that are on Commons (i.e. can be used on all Wikis)

Filter out bad images: placeholders, graphics, flags, maps **Remove articles about years, disambiguations, and lists**

More than 0 candidates left?

NO

YES

No recommendation

Rank images according to Reliability:

- Prioritize sources according to their reliability
- Number of languages referring to the same image

Assign image recommendations and confidence scores

Up to 3 recommendations

ويكيپيديا
الموسوعة الحرة

Evaluation

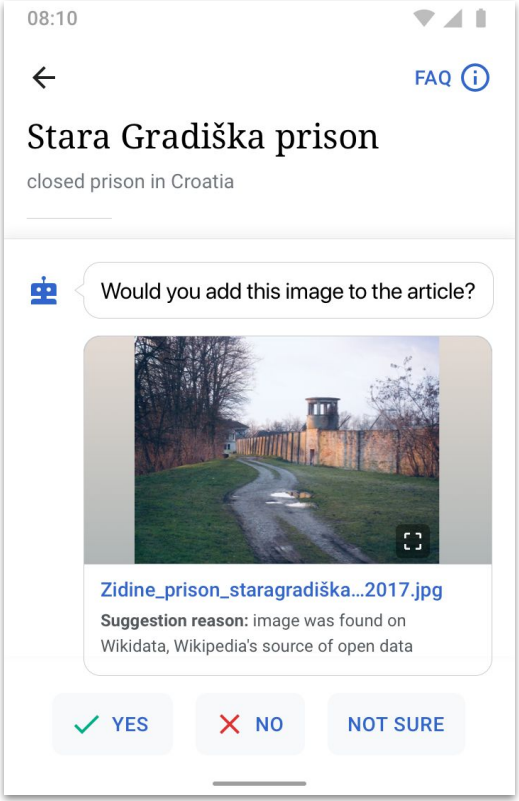

Android MVP "Train Image Algorithm"

Image Rec Task deployed on plwiki, eswiki, ruwiki, hewiki, ptwiki, fawiki, viwiki, frwiki, trwiki, itwiki, arwiki, enwiki, dewiki (many thanks to the Growth, Android and Platform Engineering Teams!!)

Engagement is high: average 9 edits/day

Precision is high but depends on source:

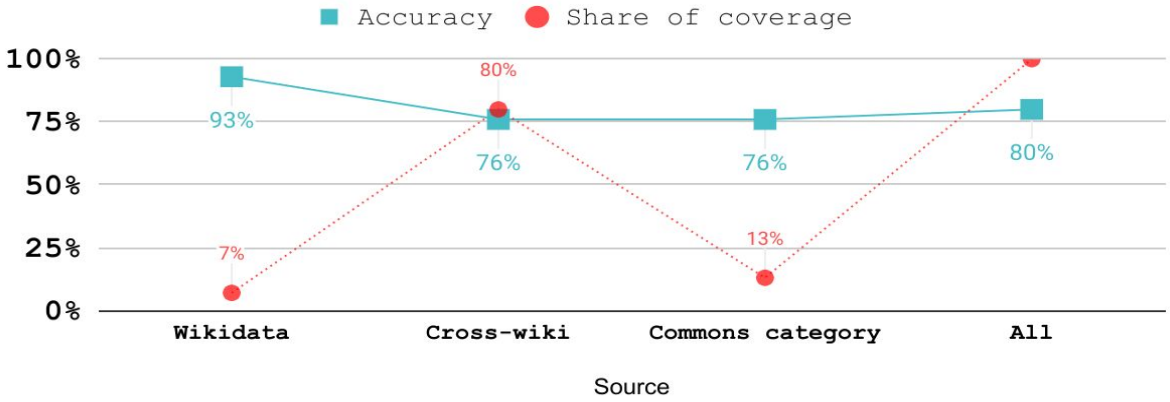

Coverage: 4% - 26% of unillustrated articles

Gathering our Scientific Communities around the problem of image/text matching

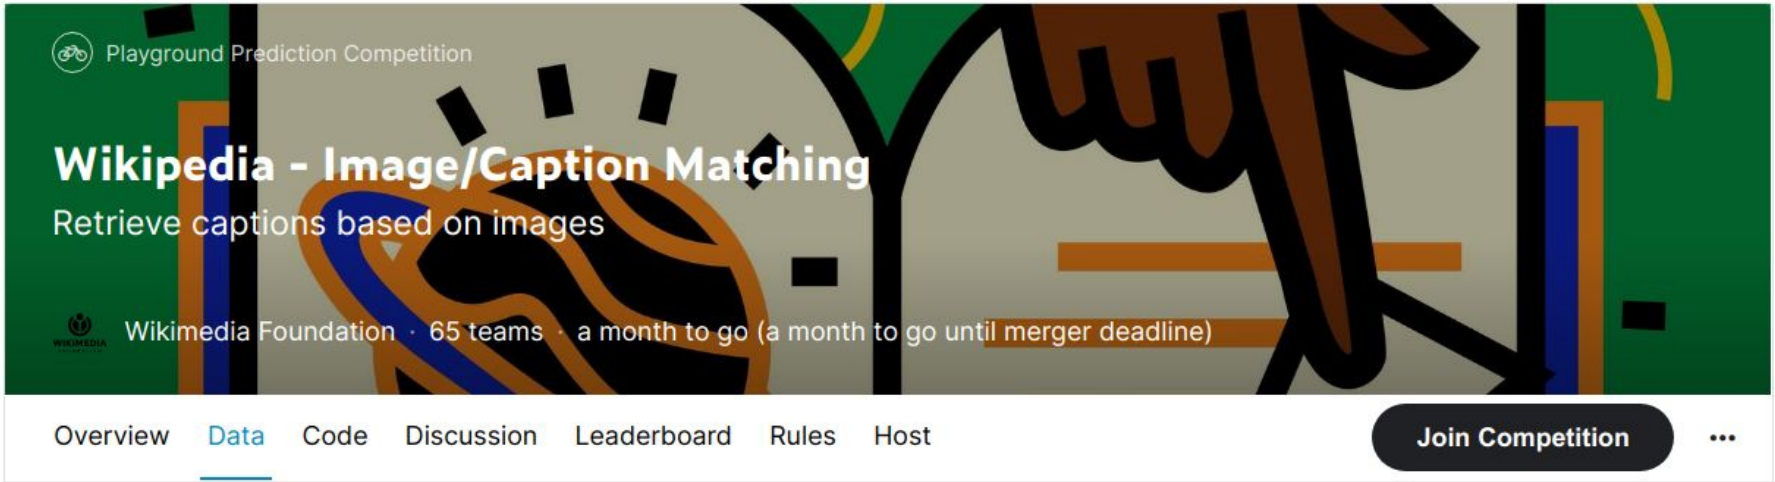

Overview Data Code Discussion Leaderboard Rules Host **Join Competition** •••

Data Description

Edit

The objective of this competition is to predict the target `caption_title_and_reference_description` given information about an images. The targets for this competition are in multiple languages.

Equity in recommendations

Isaac

WIKIMEDIA
FOUNDATION

Equity in Recommendations

- Recommender systems facilitate a growing number of Wikimedia edits:
  - Suggested Edits (300K edits Commons; 825K edits Wikidata)
  - Newcomer Tasks (190K edits across 70 Wikipedias)
  - Content Translation (1M articles created across all Wikipedias)
- How do they align with knowledge equity goals?
  - What languages are these tools available in?
    - Read more: https://w.wiki/4Hhg
  - What types of content receives the greatest benefits? should be prioritized?
    - Read more: https://w.wiki/$4P

Newcomer Tasks -- Russian Wikipedia -- Topic Filters

Flow of geographic "bias" induced by algorithmic filters (Baseline -> Impressions) and editor behavior (Impressions -> Clicks -> Edits)

Baseline
54% of Russian Wikipedia articles are geographic

Impressions
Proportion of geographic impressions by country

Clicks
Proportion of geographic clicks by country
(5.7% of all impressions)

Edits
Proportion of geographic impressions by country
(56.5% of all clicks)

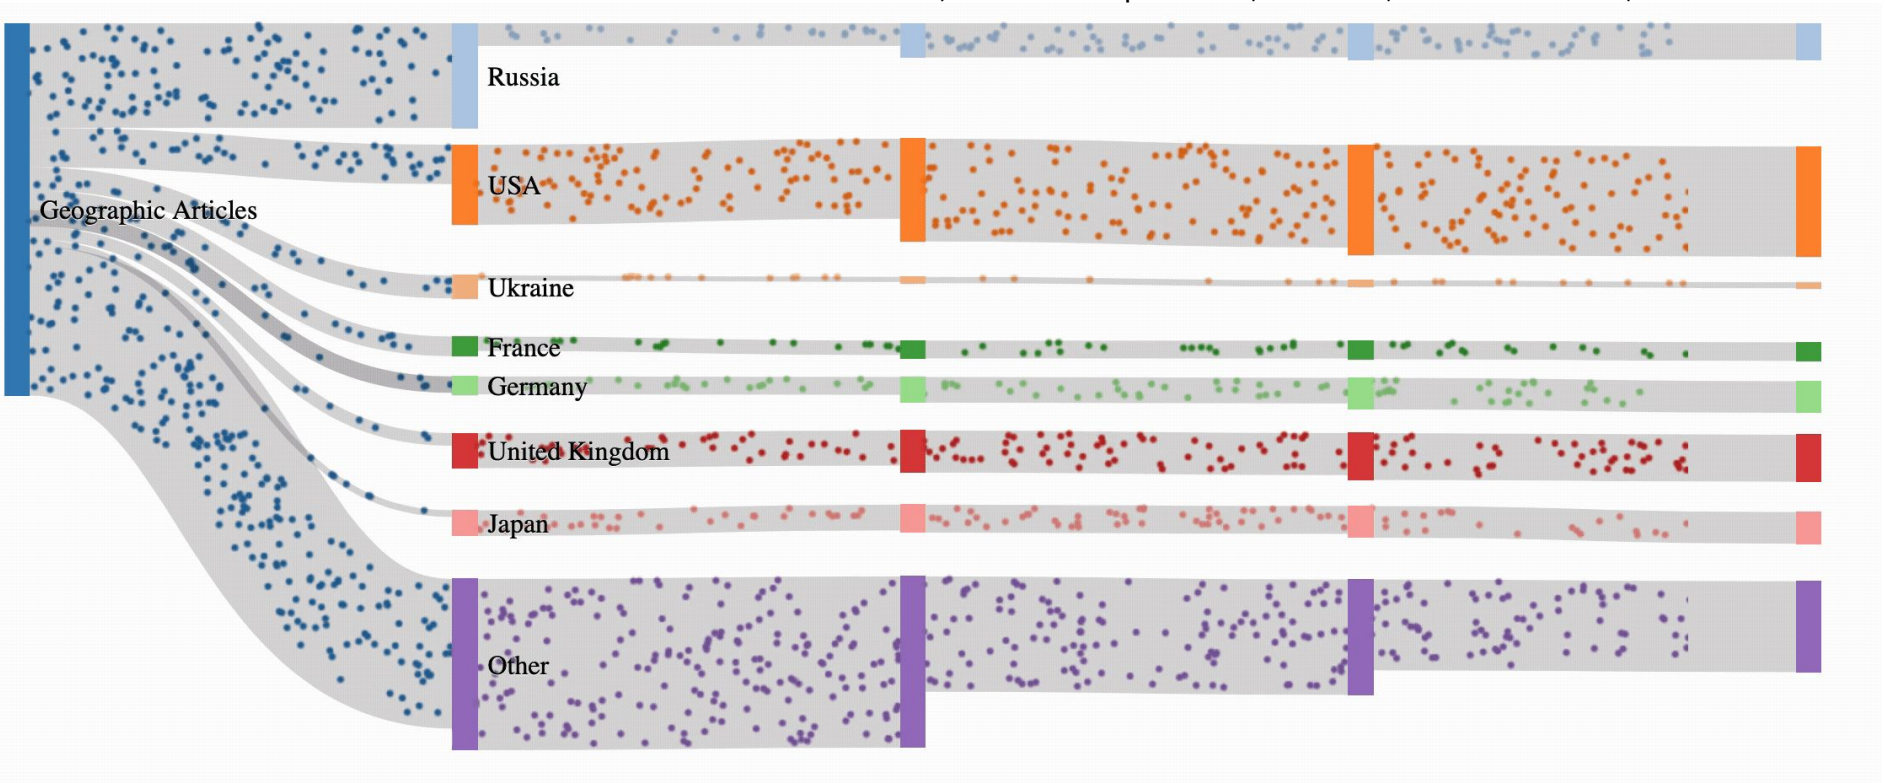

What challenges do you face in expanding section alignment to more Wikipedia languages?

Images of English and Vietnamese article for Modern Art showing how content can be aligned across Wikipedia articles through mapping wikilinks to their respective Wikidata IDs. More information: https://arxiv.org/abs/2103.00068

How might we prioritize links to be added to Wikipedia articles?

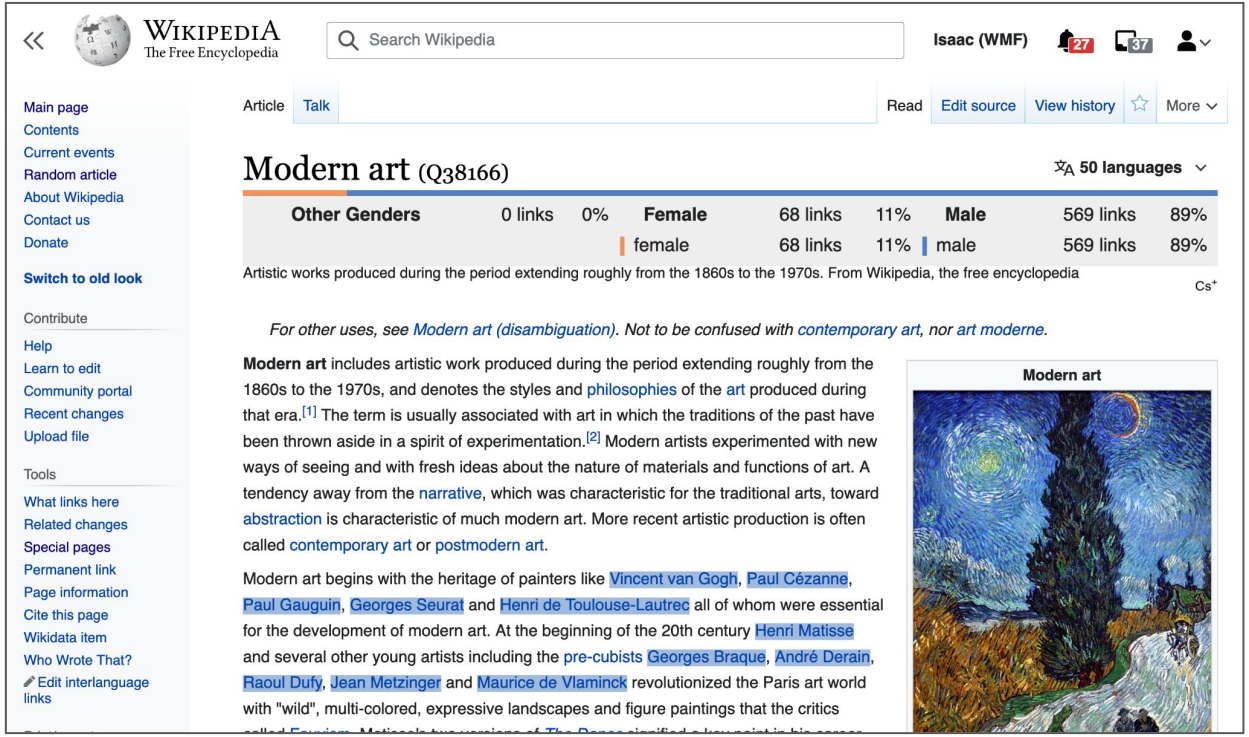

Snapshot of English Wikipedia Modern Art article with links highlighted based on gender of the person they are about. Links to articles about men are often more prominent than links to articles about women or non-binary individuals. More information: [Youtube:WikiConference North America 2021]

What special considerations do we make when considering the usage of recommendation or AI models for images?

Logo for VisibleWikiWomen, an annual campaign organized by Whose Knowledge? that focuses on adding images to articles about women.
CC-BY-SA 4.0. Whose Knowledge?

Next steps ...
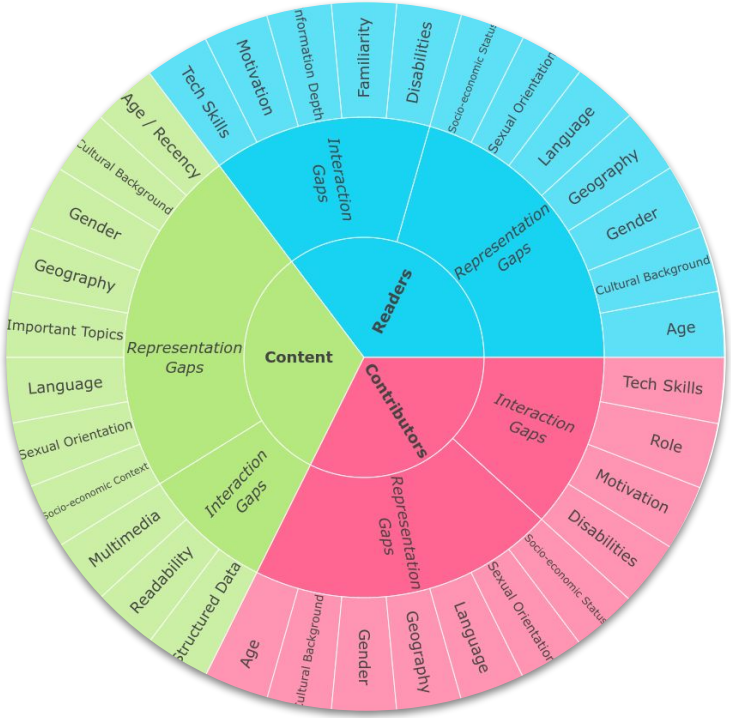
We need more and more research to Address Knowledge Gaps!