



Unpacking Stitching between Wikipedia and Wikimedia Commons: Barriers to Cross-Platform Collaboration

YIHAN YU, Human Centered Design & Engineering University of Washington, USA

DAVID W. MCDONALD, Human Centered Design & Engineering University of Washington, USA

346

The study of work practices across two or more collaborative platforms is relatively rare. Participants and researchers have to be competent, or even experts, in both just to begin to make sense of what is happening. With the growing popularity of peer-production systems, the integration of resources across various platforms is more and more common. The framework of stitching is one analytical stance that has been used to describe the cross-platform work to build and highlight informational and social networks. Through a qualitative study with 32 participants who have different foci on Wikipedia and Wikimedia Commons, we reveal their practices in three essential stitching processes, production, curation, and dynamic integration. We highlight how their practices enact stitching and extend the conceptual framework to explain barriers that inhibit effective stitching across platforms. We further discuss implications for research in cross-platform work and design to facilitate Wikipedia-Commons collaboration.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**.

Additional Key Words and Phrases: Collaboration, work study, qualitative, cross-platform, multi-platform

ACM Reference Format:

Yihan Yu and David W. McDonald. 2022. Unpacking Stitching between Wikipedia and Wikimedia Commons: Barriers to Cross-Platform Collaboration. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 346 (November 2022), 35 pages. <https://doi.org/10.1145/3555766>

1 INTRODUCTION

Spencer is a Wikipedia editor who loves airplanes. He had been working on an English Wikipedia article for the Boeing 777. Besides writing the article text, he wanted a picture showing the airplane. He went to Wikimedia Commons and searched for “Boeing 777.” The search result returned 22,572 images for Boeing 777 with 5,686 categories and multiple pages created by different curators who work to sort images. Spencer got lost in a sea of potential images. Spencer wondered which image would be best to illustrate his developing article and why Commons would have 22,572 picture of a Boeing 777?

This simple scenario shows one problem in working and collaborating across platforms. Often a contributor has to be an expert in more than one set of platform content, rules, and norms, and collaborate with contributors who hold diverse goals, perspectives, values, and expectations. Prior literature in CSCW has considered cross-platform work but with a focus on comparing rules [11, 12, 22, 29], technology affordances [10, 31, 39], and collaboration patterns [3] within the *individual* social computing platforms. Studies that consider how individuals conduct their work that crosses platforms and the extra efforts necessary for cross-platform collaboration are still rather rare.

Authors' addresses: Yihan Yu Human Centered Design & Engineering, University of Washington, 3960 Benton Ln NE, Seattle, Washington, USA, 98195, yyu2016@uw.edu; David W. McDonald Human Centered Design & Engineering, University of Washington, 3960 Benton Ln NE, Seattle, Washington, USA, 98195, dwmc@uw.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).

2573-0142/2022/11-ART346

<https://doi.org/10.1145/3555766>

Stitching is a framework that has been used to help describe and characterize cross-platform work to build organizations and also build awareness of topical content. There are three processes of stitching, *production*, *curation* and *dynamic integration* that enable resources to be distributed and utilized across different technical platforms and social networks. A few studies in CSCW have adopted this framework to investigate organizing activity as it moves between platforms [8, 9, 19]. But these studies have not dug deeply into the three underlying processes of stitching, and have instead focused on stitching to feed individuals' information needs.

In this study, we consider cross-platform work in the Wikimedia ecosystem with a focus on how editors work between Wikimedia Commons and the English edition of Wikipedia. Wikimedia Commons (aka, Commons) is the world's largest online repository of free multimedia files for anyone to contribute and use. To date, there are more than 10.5 million volunteers and over 77 million media files on Commons¹. Commons also acts as a shared repository for over 300 Wikimedia platforms including different language editions of Wikipedia and their sister platforms such as Wikidata, Wikiversity, Wikivoyage, and Wikisource. Volunteers from these platforms, along with Commons contributors work collaboratively on Commons to share, curate, and reuse multimedia content for their individual projects. This study uses the framework of stitching to analyze the cross-platform work necessary to bring together Wikimedia Commons and the English Wikipedia. We ask the following research questions:

RQ1: How do practices of editors with different platform focuses and goals enact three essential stitching processes across Wikipedia and Wikimedia Commons?

RQ2: What are the barriers that inhibit effective stitching between Wikipedia and Wikimedia Commons?

Our findings demonstrate how editors collaborate across platforms to produce, curate and integrate resources. We also identified five barriers that inhibit effective stitching, lack of communication across networks, differing perspectives, multilingual resources, cross-platform vandalism, and differing policies.

In the remainder of this paper, we first situate the context of "platform" and cover prior work in cross platform research and use prior literature to define stitching. We then describe the method we used in this study. The majority of this paper focuses on reporting our findings in stitching Commons and Wikipedia, and barriers to stitching. We follow up by discussing implications for research in cross-platform work and design to facilitate cross-platform collaboration.

2 LITERATURE REVIEW

Our study is focused on the activities of users as they work "across" or "between" distinct platforms. The notion of what constitutes a platform is therefore important to our framing. We review one specific framing of platform that helps us distinguish Commons from Wikipedia. We follow that with literature that has performed some cross-platform research. Research that studies "cross platform" collaboration and work activity can be roughly grouped into two broad categories; research that compares somewhat similar platforms and work that considers specific focused activity as it moves between platforms. Based on our review, true multi-site or cross-platform studies are still somewhat rare in the CHI and CSCW literature. As such, we attempt to consider the prior work somewhat broadly to frame the space of what is already known. Lastly, we cover the prior work that defines stitching. The term and its definition is aligned with work that considers multiple platforms and how they are used to achieve specific information seeking and provisioning goals.

¹<https://commons.wikimedia.org/wiki/Special:Statistics>

2.1 Platform

The term “platform” was traditionally used to describe a computational service that allows code to be written or run like operating systems and mobile devices. In 2010, Gillespie [15] reconsidered the way the term platform was used. He noted that the term platform had detached from its traditional computational meaning to describe online services of content intermediaries such as Facebook, Twitter, and Youtube. These service providers use the term as a “discursive resting point” to “frame their services and technologies” rather than simply indicating a technically functional shape. Defining “platform” for this new era, Gillespie [15] highlighted three additional meanings that complement and extend its traditional computational meaning: architectural, figurative, and political. One might read “architecture” to imply “software architecture,” but this is not quite the same as Gillespie’s definition. According to Gillespie [15], *Architectural* concerns the specific use for which the technical structures of a platform are designed. Platforms with the same technical architectures could be “architecturally” different because of their designers and builders conceptualized uses. This could be made manifest by the same technical infrastructures having distinctly different interfaces and interactions that enable or prohibit some uses. *Figurative* indicates the positionality or situation brought about by the technical structures of a platform. This meaning deals with the figure ground or framing of a platform – how people conceptualize what it is. And, lastly, *political* refers to the political valence carried by a platform. While the term platform is often regarded as politically “neutral,” users of different platforms articulate different beliefs and they behave differently where positions are taken regarding content and contexts external to the platform.

Adopting Gillespie’s extended definitions of platform, we argue that Wikipedia and Commons are two distinct platforms. Although these two platforms are based on nearly the same technical structure, they differ across the three extended meanings defined by Gillespie [15]. Although Wikipedia and Commons have similar technical structure, their architectural meanings are different. The two platforms are built upon the same MediaWiki software infrastructure, but have different usages, have distinctive interfaces, and distinct user experiences. The interface of Wikipedia was built to assist encyclopedic text editing and is built around the assumptions that text editing and text contributions are of primary importance. It features text editing functions such as creating sections, inserting citations, previewing changes, and linking to other pages. In contrast, Commons interface focuses on facilitating image uploading, image annotating, metadata tagging and categorizing. It has specific functions that allow users to upload and download files, add license information and link images to other structured data sources. Wikipedia and Wikimedia Commons differ along Gillespie’s *architectural* meaning. As well, the two platforms differ in the *figurative* meaning. The figurative meaning of Wikipedia is to create a “reference,” a “widely accessible and free encyclopedia written in 300 languages by volunteers around the world ².” Wikimedia Commons aims at building a “collection” that “makes available public domain and freely-licensed educational media content (images, sound and video clips) to everyone ³.” The figurative meaning of Wikipedia as “reference” and Wikimedia Commons as “collection” shape key understanding of the participants and what they do in each platform. While one might think that the *political* meaning of Wikipedia and Commons is driven by the Wikimedia Foundation (WMF), that stance undercuts the role that the editors and contributors play in the individual platforms. Wikimedia Commons and Wikipedia have very different politics as manifest in divergent and sometimes contradicting norms, policies and guidelines ⁴⁵. For example, Wikimedia Commons does not censor around the viewpoint an

²<https://wikimediafoundation.org/our-work/wikimedia-projects>

³<https://wikimediafoundation.org/our-work/wikimedia-projects>

⁴https://en.wikipedia.org/wiki/Wikipedia:Policies_and_guidelines

⁵https://commons.wikimedia.org/wiki/Commons:Policies_and_guidelines

image holds, as its project scope policy ⁶ states “Commons is not Wikipedia, and files uploaded here do not necessarily need to comply with the Neutral Point of View ⁷ and No Original Research ⁸ requirements imposed by many of the Wikipedia sites.” In contrast, Wikipedia asks all contributors write from a Neutral Point of View and cite content from reliable, published sources. In short, the political (normative and policy regimes) are distinctly different between these two different platforms.

2.2 Cross Platform Comparisons

Prior work in CSCW has performed some cross-platform research. One major line of this research focuses on comparing somewhat similar social computing platforms in terms of policies and guidelines, technology affordances, and collaboration patterns.

Jiang et al. [22], Pater et al. [29], Fiesler et al. [11], and Fiesler et al. [12] have compared policies and guidelines adopted by different social media platforms. They found policy is a collaborative social construct of individual platforms that reflects unique values of the platform. For example, social media platforms such as Facebook, YouTube, Twitter and Instagram chose to focus on different misbehavior in policy making and content moderation in their community guidelines, possibly prioritize policies and guidelines that are most reflective of the platform’s values [22]. Educating individual users about the copyright policies is a consistent usability issue across social media platforms featuring different media types. The rules about copyright are opaque but users’ discussion around copyright and problems caused by copyright is prevalent [12].

Researchers also cross-compared “technology affordances [14]” of social media platforms for users’ self-presentation. They found that most social media users combine different social media platforms to conduct different tasks (collaboration or conversation), share different types of content, and reach out to different audiences, though social media platforms were commonly treated as individual tools or services [31, 39]. They argued that future social media research should consider system affordances within a communication ecosystem, in which different affordances (e.g., policy and practice) of different platforms need to be coordinated to maintain the stability of the ecosystem [10, 31, 39].

Another approach to cross-platform study relies on a framework or theoretical stance to understand how similar platforms vary. Prior studies in CSCW applied the same collaboration framework [25] to attempt to understand collaboration on different Wikipedia language editions. For example, Bipat et al. [3] compared the use of power plays across English, French, and Spanish Wikipedia. They suggest that editors of different Wikipedia versions valued different policies while collaborating and trying to get control over edits.

The prior work in understanding cross-platform collaboration has largely been comparative. Though researchers in CSCW have made efforts in conducting multi-site research that probes more than one social computing platform, few prior work in CSCW focused on how collaboration happens across multiple platforms that are designed to be used in tandem. Our study begins to fill the gap in understanding the way collaborators and their collaborative activity move between and across platforms that must be used in tandem to satisfy collaborative goals [2].

2.3 Understanding Cross Platform Work through Stitching

A small number of studies have considered how work activity crosses different technical platforms. The framework *stitching* has been coined to descriptively convey the way work and content are

⁶https://commons.wikimedia.org/wiki/Commons:Project_scope/Neutral_point_of_view

⁷https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view

⁸https://en.wikipedia.org/wiki/Wikipedia:No_original_research

interleaved to archive a goal that cannot be simply met by a single person in a single platform. Stitching is “the production, curation and dynamic integration of various types of information content and other resources that become distributed and utilized across the crowd” [1]. The framework considers the organization of the individuals involved by saying that stitching includes “... the organizational processes of crowd-enabled connective action” [1] in the absence of individuals as recognized leaders. There are three elemental processes of stitching that enable resources to be distributed and used across platforms and crowd, production, curation and dynamic integration. *Production* refers to the creation and sharing of resources within one social computing community/platform. *Curation* entails preserving, maintaining and sorting resources generated in the production process. It also involves constructing norms and boundaries for individual communities/platforms. *Dynamic Integration* is the process that allows resources to be transmitted and linked between platforms, creating social and informational networks that better integrate individual participants as well as content across different technical infrastructures [1].

Not all cross platform work could be described as stitching. There are several types of work that people do which requires that they move or curate information between different platforms. But some of them are not stitching. For example, a social media user might copy and paste a post from Twitter to Facebook to reach diverse audiences. Though the user moved across two platforms to conduct this activity and they curated content by selecting and moving content to another platform, the activity is not stitching. Because it does not include two other essential processes of stitching, production and integration. The user did not generate any new information nor did they integrate content of the two platforms. More importantly, the activity fails as stitching because the framework is interested in the organizational aspects of the work (i.e., the “... organizational processes of crowd-enabled connective action” [1]). Without the context of crowdwork, or a shared goal, or some other shared production process, the individual activity would not be appropriately characterized as stitching. Thus, not every kind of work that includes two platforms could properly be analyzed as stitching behaviors.

Most of the research that attempted to define and describe stitching as a theoretical framework has focused on crisis situations or transmission of misinformation [8, 9, 19]. However, these studies use stitching as a metaphor and focus their analysis on “seams [6, 34]” through which multiple social computing platforms were strategically aligned by individual actors to meet their information needs. That is, the focus of the analysis is more commonly on information needs. These studies have not illustrated how three essential processes of stitching play out. Being fair to that work, each of the studies do illustrate at least one of the three stitching processes. For example, Dailey et al. [9] adopted the stitching framework to investigate how social media platforms afford actors in crisis to assemble heterogeneous platforms for their highly dynamic information work. In their study of the 2014 Oso Landslide, Dailey et al found that government workers “stitch” multiple social media platforms to reach different audiences. They used Twitter to monitor and communicate with media and journalists, and Facebook to engage local communities. Though this study explained how individual government workers use multiple social media platforms, it did not investigate the process of organizing a crowd across platforms. As well, this study focused on two processes of stitching and did not bring all three essential processes of stitching together. Government workers produced, made posts on both platforms and curated, fact-checked and promoted tweets on Twitter. But their activities did not demonstrate the integration process of stitching. That is, there were no dynamic links created that allow resources to be used and integrated across two platforms. This illustrates a gap in understanding how three essential processes of stitching work together to create organization of a crowd in cross-platform collaboration.

In the following, we use stitching as a lens to study collaboration across Wikipedia and Wikimedia Commons. Through interviews with participants who have different perspectives on the work and

the platforms, we are able to show the organizational processes of two peer production platforms that enact stitching. We demonstrate all three elements of stitching across both platforms and their participants. We extend the conceptual framework in two ways. First, we illustrate the distinct platform shaped perspectives of the participants who mostly focus on one of the platforms and how their work enables stitching across platforms. Second, we extend the conceptual framework of stitching to explain how specific barriers may inhibit effective collaboration across platforms.

3 METHODS

Wikimedia Commons (aka "Commons") and the many language editions of Wikipedia are related through the activities of the editors work "across" or "between" two or more wiki platforms. Our research focused on individuals who work in Commons and the English language edition of Wikipedia ("Wikipedia"). Our choice of this pair over others was simply a function of our familiarity. We recognize that there are many language editions and that there is prior work showing that collaborative activity in the English Wikipedia may not generalize to other language editions [3]. However, we were interested in the editors experiences of the work they conducted on or across both platforms, not just their collaborative activities. We conducted qualitative interviews oriented around their general experiences contributing, a situated exploration of some specific contributions, and their views of the relation between Commons and Wikipedia. We explain our methods in more detail below.

This study was reviewed and approved as "exempt" by the Institutional Review Board (IRB) at our University. We also followed Wikimedia Foundation's (WMF) guidelines for conducting research on WMF related websites. We created a research project page on Meta-Wiki that describes the project and facilitates community review on the project talk page. The project page states our research questions and scope, participants, recruitment and data collection methods, and potential impacts of the study. Both authors monitored the talk page and replied to a small number of posts from community members. During our research we updated our Meta-Wiki project page with our research progress noting the number of individuals we had contacted and our interviewing progress.

3.1 Recruitment

Our recruitment targeted two groups of editors: Commons focused editors who work primarily on Commons to upload, edit and organize multimedia contents. And Wikipedia focused editors who work primarily on Wikipedia to edit articles and use multimedia content from Commons to illustrate Wikipedia articles. Our initial recruitment effort focused on Commons focused editors. We extended our recruitment in subsequent rounds to include Wikipedia focused editors.

Since we are interested in participants who have different perspectives on the work and the platforms, we tried to recruit participants through different strategies. We identified Commons focused editors and Wikipedia focused editors for possible recruitment using six different strategies (Appendix A): (a) Commons active users list, (b) Commons autopatrolled users list, (c) Commons Help Desk, (d) Wikiscan tool for English Wikipedia, (e) English Wikipedia Featured Article list and (f) snowball sampling based on participant suggestions. Our diverse recruitment strategies allowed us to recruit across a number of important dimensions: highly active editors, editors who contributed high quality content, less experienced editors and newcomers, and random selected editors who made at least three contributions.

We contacted editors selected for an interview through the built in "Email this User" function of Wikimedia Commons and English Wikipedia. We emailed each editor a personalized message that includes their username, purpose of the study, link to our Meta-Wiki study page, how we identified them as a potential participant, and an invitation to the interview. Editors could reply to us either

Table 1. Recruitment strategies and the number of editors contacted, replied and interviewed from each channel

Recruitment Strategy	#Contacted	#Replied	#Interviewed
Commons Active User List	81	6	5
Commons Help Desk	43	13	7
Commons Snowball Sampling	12	2	2
Commons Autopatrolled User List	4	1	1
English Wikipedia Wiki Scan	20	7	6
English Wikipedia Featured Article	75	18	11
Total	235	47	32

through the email address we provided or the email function of the underlying MediaWiki software. We contacted a total of 235 editors, of which 47 replied to us with 32 agreeing to participate in an interview. Table 1 shows the number of editors contacted, replied and interviewed from each recruitment strategy. We conducted our semi-structured interviews with 32 participants from April 2020 to January 2021.

3.2 Participants

Table 2 summarizes the basic demographics of our interview participants. We assigned a pseudonym to each participant for confidentiality. Although we strove to recruit participants with diverse backgrounds, of our 32 participants only five self-identified as a woman. This representation roughly corresponds to the gender distribution of Wikimedia projects' global editor population, in which roughly 13% of the contributors were women.⁹ Besides gender, our six different recruitment strategies allowed us to interview participants with diverse platform participation, activity levels, and experiences with different Wikipedia language editions.

Below we describe several demographic and participative dimensions of our participants. We rely on cohorts because providing individualized demographics would make many of our participants individually identifiable. Our approved IRB process included specific wording that we would do our best to protect the confidentiality of participants' participation in the research. Even so, we recognize that some of our participants contribute in very specialized ways that may be more obvious to well established Wikipedia contributors and researchers. We have done our best to reduce potential risks through demographic and participative cohorts, and by abstracting away some of the specifics of contributions in the quotes that we use to illustrate our analysis.

3.2.1 Platform Focus. During the interview, we asked participants to describe their experience in contributing to Wikipedia and Wikimedia Commons. The resulting discussion often resulted in the participant telling us where they believe they focused their contributions. Of all the participants, 14 described themselves as focused on both platforms. Six participants self-identified as Commons focused, and 12 said they were focused on Wikipedia. Two participants we interviewed specifically stated that they felt like newcomers to Wikipedia and/or Commons. We wanted to understand, to what degree we could see their self-identified platform focus through their specific editing activities. We considered this by collecting the participants' edit counts on both English Wikipedia and Wikimedia Commons. We created edit count cohorts to simplify comparisons, and for many of our participants the self-identified focus roughly corresponds to the difference in edit counts between Commons and Wikipedia.

⁹https://meta.wikimedia.org/wiki/Community_Insights/Community_Insights_2020_Report/

3.2.2 Activity Level. Participants' edit counts reflect their activity levels on both platforms. Our participants' editing activity ranges from highly active editors who have made tens-of-thousands of contributions to newcomers who have less than 1k edits on both platforms.

3.2.3 Experience with Different Wikipedia Language Editions. Though we targeted English-speaking editors, eight participants mentioned that they contributed to other language editions of Wikipedia besides English, such as Czech, Spanish, French, and Italian. Those participants sometimes compared what they did in the English Wikipedia to the other language edition that they edited. These participants provided some valuable insight to the challenges of using Commons when the primary language of the curator or editor was not English. We note some of these insights as we analyze our data below.

3.2.4 Location. Our participants are located in nine different countries around the world. Most of them are from English-speaking countries including US (9 participants), UK (10 participants), Australia (4 participants) and Canada (2 participants). We also have participants from five non-English speaking countries, Germany (3 participants), Italy (1 participant), Israel (1 participant), Czech (1 participant), and China (1 participant).

3.2.5 Edit History. We checked the edit history of participants on both platforms to see when the first edit was made. In English Wikipedia, the longest time a participant was engaged is 19 years. The shortest is less than a year. And the average is 12 years. In Wikimedia Commons, the longest time of engagement is 16 years. One participant does not have a Commons account and five participants edited Commons for less than three years. The average time of engagement is 10 years.

3.3 Data Collection and Analysis

Through several iterations and pilot interviews we developed an interview protocol that consists of three phases (see Appendix B). In phase 1, we asked about the participant's experience in contributing to Wikipedia and/or Wikimedia Commons. In phase 2, we share 4-6 examples in which the participant used an image(s) from Commons to illustrate a Wikipedia article or made a contribution directly to Commons and asked questions about the examples. In phase 3, we guided the participant to discuss the connections between Wikimedia Commons and other Wikimedia projects such as Wikipedia. We tested the interview protocol with a Wikimedia Foundation research staff in a pilot interview, gathered their feedback and revised the protocol in order to inspire more insightful responses from the Wikipedia and Wikimedia Commons communities.

The first author conducted 32 semi-structured interviews in English through participants' preferred tele-conference applications (e.g., Zoom, Skype and Google Hangouts) or phone calls. We recruited participants through six strategies, but we did not use the six strategies all at the same time. Instead, we had multiple recruitment attempts with each strategy. For each strategy, we recruited and interviewed participants. The research team debriefed after each interview and discussed participant insights and discussed any emergent concepts. We ended a recruitment strategy when we found participants repeating insights and concepts. We then discussed participative and demographic characteristics of the participants we had interviewed. We would then consider a possibly new recruitment strategy to help reach participants with different experiences. We believed we reached data saturation when we could not figure out a new strategy to include different types of participants than we had already interviewed. This yields a data collection of 32 interviews. We recognize that our approach allowed us to reach what we concluded was data saturation, but we realize there is potential diversity in both platforms that may not have responded to our recruitment efforts.

Table 2. Participants' Demographic Information

Pseudonym	Recruitment Channel	Gender	WP Edit Count Cohort	Wikipedia Language Edition	Commons Edit Count Cohort	Participation Focus
Cliff	Help Desk	Man	1K-5K	EN	1K-5K	Focused on Both
Clara	Help Desk	Woman	5K-10K	CS, EN, ES	10K-50K	Focused on Both
Jasper	Active Users	Man	100K-500K	EN	500K-1,000K	Focused on Both
Max	Commons Snowball	Man	10K-50K	EN	10K-50K	Focused on Both
Stefan	Active Users	Man	50K-100K	EN	500K-1,000K	Focused on Both
Amos	Wikiscan	Man	5K-10K	EN, IT, FR	10K-50K	Focused on Both
Arthur	Wikiscan	Man	100K-500K	EN	100K-500K	Focused on Both
George	Featured Article	Man	10K-50K	EN	10K-50K	Focused on Both
James	Wikiscan	Man	100K-500K	EN	100K-500K	Focused on Both
Mike	Wikiscan	Man	100K-500K	EN	<1K	Focused on Both
Nick	Wikiscan	Man	50K-100K	EN	1K-5K	Focused on Both
Grace	Help Desk	Woman	<1K	EN	<1K	Focused on Both
Andy	Active Users	Man	100K-500K	EN	100K-500K	Focused on Both
Gavin	Featured Article	Man	100K-500K	EN	10K-50K	Focused on Both
Frank	Active Users	Man	50K-100K	EN	5,000K+	Commons Focused
Alana	Active Users	Woman	50K-100K	EN	500K-1,000K	Commons Focused
Christopher	Help Desk	Man	5K-10K	EN	50K-100K	Commons Focused
Jiminy	Commons Snowball	Man	50K-100K	EN; ES	100K-500K	Commons Focused
Jeff	Wikiscan	Man	50K-100K	EN; DE	1,000K-5,000K	Commons Focused
Winston	Autopatrolled	Man	5K-10K	EN	10K-50K	Commons Focused
Duke	Help Desk	Man	50K-100K	EN; ES	1K-5K	Wikipedia Focused
Arlo	Help Desk	Man	10K-50K	EN; HE	1K-5K	Wikipedia Focused
Dale	Wikiscan	Man	100K-500K	ES	1K-5K	Wikipedia Focused
Eric	Featured Article	Man	10K-50K	EN	<1K	Wikipedia Focused
Cody	Wikiscan	Man	100K-500K	IT; SV; DE	5K-10K	Wikipedia Focused
Philip	Featured Article	Man	50K-100K	EN	<1K	Wikipedia Focused
Jack	Wikiscan	Man	100K-500K	EN	1K-5K	Wikipedia Focused
Kim	Wikiscan	Woman	100K-500K	EN	50K-100K	Wikipedia Focused
Louie	Featured Article	Man	10K-50K	EN	<1K	Wikipedia Focused
Marlon	Featured Article	Man	50K-100K	EN	<1K	Wikipedia Focused
Camilia	Help Desk	Woman	<1K	EN	<1K	Wikipedia Focused
George	Featured Article	Man	1k-5k	EN; DE	<1K	Wikipedia Focused

In 30 interviews conducted through tele-conference applications, we shared screens with the participants when we needed to open and discuss a webpage and video recorded the interviews. For two interviews conducted via phone calls, we sent a list of useful links and asked participants to open them on their own devices prior to the interview. We audio recorded the phone calls. The interviews took between 35 - 100 minutes. The first author transcribed the interviews and documented all pages (e.g., user page, user contributions page, Wikipedia article page, and Commons project page) discussed in the interviews for data triangulation. Both authors open coded the transcripts using a thematic analysis approach [4] and made memos. We iterated on the themes and memos and decided to report only on the emergent themes related to participants' practices, perspectives and challenges they encountered in stitching.

In the next two sections, we report our findings by first presenting evidence of stitching Commons and Wikipedia, focusing on participants' practices and perspectives. We then discuss about barriers that inhibit effective stitching across the platforms.

4 STITCHING COMMONS AND WIKIPEDIA

In this section, we illustrate several ways that stitching occurs between Wikimedia Commons and the English edition of Wikipedia. We talk about practices and perspectives of participants in three essential stitching processes, *production*, *curation* and *dynamically integration*. As discussed in Section 3.2, we group participants by their different platform focus, Commons focused editors, Wikipedia focused editors and editors focused on both platforms. In this section we highlight how they talk about their work and enact stitching to achieve their goals.

4.1 Production

The key production processes on each platform are different. Commons supports a wide range of different media types, but most of the effort is on the creation and uploading of pictures, or images. Commons has associated work practices that enhance the availability of those images through the addition and editing of important metadata. On Wikipedia much of the effort is on the creation and editing of article text. But as well, there are important associated work practices that surround editing activities [24].

4.1.1 Commons Focused Editors. In the production process, Commons focused editors work mainly to create and upload multimedia content to Commons. There are three prevalent ways in which they produce content. Some of them upload and share images they create, either for documenting a subject, or for use in Wikipedia articles. Some transfer files from other freely usable sources for similar documentation purpose or to preserve free image collections and for potential use in articles. A few Commons focused editors also serve as an OTRS (Open-source Ticket Request System) agent who helps upload images donated to Commons via email.

Create & Upload Images. Commons focused editors who create and upload images hold two different objectives. One objective is *documenting* a "significant" subject, without consideration for how the images might be used. In their view the subject has its own intrinsic value that is captured by the image.

Often if I find anything like a public lecture or something like that, and there's people of any significance really at all, I'm taking their photos and getting them up there, unless they actively object to having pictures taken in that context. -Jiminy

Another objective is to *illustrate* Wikipedia articles. Christopher is a significant contributor to Commons, with many of his pictures currently being used in Wikipedia articles. As Christopher points out, he does not bother to upload an image if he does not think it would be used,

... [many] of my images are on English Wikipedia. ... And so my objective, I normally don't upload an image unless it is good enough to go into a Wikipedia article. So if I take a picture of a butterfly, that is not as good as one that is on Commons already, and is not more useful, I don't bother to upload it. -Christopher

Commons focused editors sometimes noticed that a Wikipedia article was in need of illustrations. They would create and upload some "decent" images of the subject to Commons for Wikipedia editors to find and use,

I was reading the article on Wikipedia of [someone] and it didn't have a decent photograph of him at all. I think there was either no photograph or not a very good photograph. I've got quite a few photographic books at home. So I looked up some. And I think I might have scanned this picture from one of my own books..from National Portrait Gallery, people in camera. It's an exhibition catalog that I had. And I found this photograph in it. And I thought it was quite a nice one. I scanned it and uploaded it. And that's been used. -Mike

Mike noticed that an article he was reading needed a good illustration. He scanned a picture from his exhibition catalog and uploaded the picture to Commons for Wikipedia editors to use. The picture was later found by Wikipedia editors and is now used in the article.

Transfer Files. Another way in which Commons focused editors produce content is transferring free licensed multimedia files from other sites and databases to Commons. As with Commons focused editors who upload their own images, some transfer files with an intent to document the subject.

So I'm an employee of the [country] government, [government jobs website] is the site that you use to apply for [government jobs]. I know a lot of the stuff for [government] agencies. I just have an interest in documenting how agency owners oriented... I'll go to the webpage, or find a PDF that they uploaded that's in a vector format. And I used Inkscape to extract it from a PDF. And I might have actually had to ask for help getting some of this fixed... When I found that the [government jobs website] logo was publicly hosted on [website], I downloaded a copy of it and uploaded it. -Winston

Winston said that many of the resources he was interested in saving and documenting, such as logos of various government agencies, are in the public domain. As such, he felt his contributions to Commons were adding resources that others might not otherwise easily access.

Commons focused editors who perform large scale file uploads, called batch uploaders, find free licensed databases or archives and "harvest" them, transferring all files to Commons regardless of content. They do this to "preserve" free images and "get a lot of good quality content available very quickly."

... there was some discussion about [stock photography site] changing their terms and conditions. So I looked at that, and it was a blog post about it from [stock photography site]. And I think it was 2017, they moved from cc zero to having their own statement about free use, ... So actually, I wrote a program just to harvest all of these images from [stock photography site], ... actually, the harvesting was just a month after they changed the licensing, ... there were 32,000 photographs, all very good quality, where the photographer had released them as cc zero. And it was important to harvest them because one of the things that happens with [stock photography site], as it happens on other sites, is that photographers disappear along with their photographs. ... if I hadn't uploaded them on mass, if somebody had later uploaded it, there would have been a debate about whether it was really free or not, whether it was really cc zero or came under the new license, and that debate would keep on happening. So this was a way of dealing with a copyright problem. And a way of getting a lot of good quality content available very quickly. -Frank

Frank "harvested" all 32,000 photographs from a stock photography site before the website changed its licensing. Frank believed that he contributed by making a large number of high-quality images available on Commons and eliminating licensing issues that would inhibit editors from reusing the images in the future.

Upload Donated Images. A small number of Commons focused editors serve as an OTRS agent who processes emails sent into Wikimedia Commons and helps upload free licensed images donated to Commons.

I'm also an OTRS agent. So I've got some uploads in there. That is stuff that people sent the emails into Wikimedia and go like I'd like to donate this. It's public domain because of this, but I don't want to upload it myself. So we upload stuff on their behalf, ... I've got things like [stamp]... well, not a topic I would ever be interested in myself, obviously... I just uploaded on their behalf. -Nick

In the production process, through creating, transferring and uploading multimedia files, Commons focused editors produce and share (some hard to access) resources on Commons so they can be freely used by other platforms without licensing issues.

4.1.2 Wikipedia Focused Editors. Wikipedia focused editors focus on producing text for Wikipedia articles. Their activities of editing articles are diverse and collaborative [23, 24, 26]. However, few studies of Wikipedia have investigated multimedia resources that they use to illustrate articles. Work by Kriplean et. al [24] is one exception. They identified some of the multimedia efforts, but without explicitly considering Commons. Wikipedia focused editors sometimes produce images on Commons with the singular intent of using the images in Wikipedia articles,

I only upload photos because I intend to use them in Wikipedia articles. I don't upload them just because it's handy. [whereas] some people just upload photos because they've taken photos and they upload them. I don't do that so much. -Kim

They produce images on Commons in three ways, hunting images, taking pictures, and uploading interesting illustrations.

Hunt Images. When a Wikipedia focused editor cannot find an image from Commons to illustrate a Wikipedia article, they go hunting for images from other free licensed websites such as GLAM (Galleries, Libraries, Archives, and Museums") websites, government-owned websites, and online multimedia sharing communities (e.g., Flickr). If an appropriate illustration is obtained, they upload it to Commons and link it to the article, as Louie explained,

Mostly, what I do is in very specific cases, for instance, if I write articles about recent sporting events, there are never or very seldomly photos on commons already. So what I do is I go in flickr, I search for images that might work. And either they already have a license where I can upload them to Commons, or I have to write the author of the images and ask them for permission. And thankfully, people are very often quite kind and give permission. Yeah, you can find quite a lot of images that way and then I have to upload them myself. -Louie

Take Pictures. When there were no great illustrations available either on Commons or other free licensed websites, some Wikipedia focused editors go take pictures themselves and upload them to Commons. Kim and her collaborator call it a "photo drive,"

I write content about all of the places in [State] I've put every town and locality in [State] onto Wikipedia... [my partner] took the photo, but I was driving the car. So when I was working with someone else, what we did is we would look at Wikipedia articles in an area, see which one needed photos. And then we would take a day where I would do the driving, and he would do the photography to get some more photos. So what we call a photo drive. -Kim

This is clearly a focused effort by Kim to take and contribute photos that they believe are needed for specific articles. We asked about whether they were certain that these images were essential, the participant elaborated their answer,

I'm not going to drive 100 kilometers to take a photo without checking first. No, we will, typically we were specifically looking for articles that didn't have photos. And obviously, if there had been a photo on Commons, we would have already used that. -Kim

Upload Illustrations. Wikipedia focused editors sometimes create interesting illustrations for their articles. If the illustrations are not on Commons already, they would upload them,

I think what I've done is I have a sense, or an idea of what illustrations are interesting. These would be some articles that I was writing. Like I wrote an article about a college that existed from 1850 to 1860. There was actually an old photograph of this college. And I made an illustration out of that. And then I put that in the article. And that's mostly what I'm doing, is things that I noticed that are interesting, that are not there. -Duke

While Wikipedia focused editors produce and share multimedia content on Commons with a singular intent to illustrate articles, their activities help make hard-to-find resources accessible on Commons. So editors who work on related articles or editors from other language editions could easily reuse them.

4.1.3 Editors Focused on Both Platforms. The practice of editors who focused on both Wikipedia and Wikimedia Commons have some overlaps with Commons focused editors and Wikipedia focused editors. We will not include the overlaps here. However, editors who focused on both platforms have a distinctly unique view of both projects. This allows them to produce multimedia resources that align closer with Wikipedia's needs and interests. Similar to Commons focused editors, editors focused on both platforms create and upload images to commons. But they have a unique focus on filling the content gap between Wikipedia and Commons. Instead of taking generic pictures of well-illustrated subjects, editors focused on both platforms know what illustrations are "missing" from Wikipedia articles. They would produce the "missing" content,

... since I've seen a lot of articles and a lot of images, I kind of have a feeling of what is missing. So every time I feel something very specific that can be missing, I try to take pictures. -Amos

For example, Amos knows that a lot of good quality images on Commons are artistic, but sometimes these images do not have details that are necessary to illustrate Wikipedia articles. Therefore, he prefers to produce high quality, "boring" pictures that are not artistic enough to be featured, but capture details that help Wikipedia editors develop their ideas,

As a photographer, I know what makes a good image. Technically, what looks right. What describes the subject that I'm writing about... I picked that one out, because it showed the flower and the leaf structure, which are two key ID features. It may be different for another plant. Or for a bird, it may be certain coloring or location of the coloring... -Amos

In addition, editors focused on both platforms have a better understanding of the notability guideline¹⁰ of Wikipedia. They know what subjects, though do not have a Wikipedia article yet, would be notable and warrant their own articles in the future. They work to make illustrations of these subjects available on Commons,

...so do you know Douglas? ... Pamela Harris' husband? I found images of them ... And when they created an article on him, it was useful that I already had found images of them. Because otherwise, you wouldn't have had photos of them. Someone would have had to dug deep, I had already found a photo of them and categorized it ahead of time. -Stefan

In summary, participants with different platform focuses produce resources for different purposes. But collaboratively, they increased the topical coverage of Commons and the amount of resources available for reuse.

¹⁰<https://en.wikipedia.org/wiki/Wikipedia:Notability>

4.2 Curation

The curation process of stitching entails preserving, maintaining and sorting resources, and constructing norms for individual platforms [1]. In Commons-Wikipedia stitching, curation involves maintaining and sorting multimedia resources and establishing platform policies.

4.2.1 Commons Focused Editors. The curation process on Commons is conducted mostly by Commons focused editors. They verify and edit metadata, categorize images and construct policies and guidelines of Commons.

Edit Metadata. A group of Commons focused editors go through images uploaded in the production process, verify and correct metadata such as data, location, and description of the files to improve the accuracy of Commons content. For example, Jiminy checked resources uploaded by a university library and had corrected hundreds of metadata pieces curated by the initial uploader,

I've been through several thousand photos from [a university library collection], which it turns out, were not very well curated by the library. And I have literally fed back hundreds of corrections to the library about the metadata on their photos, including things as extreme as something being misdated by 20 plus years, a location being entirely wrong, where something was described as being a particular neighborhood of Seattle that wasn't even in Seattle, misidentified people, things like that. -Jiminy

Besides, Jiminy would zoom into an image to identify "prominent things" depicted in small regions of the image. He would add his discoveries to the metadata because someday someone might be interested in and search for the specific things and find the image useful,

One of the things I've been doing as I go through this library stuff, because at most they might have picked up half a dozen prominent things in these and in some of them I've made a lot of really interesting discoveries by these short lived businesses. Found one just the other day on the corner of [a street], the corner where [a building] is now, some sign on an awning saying the evil eye who knows what it was, but I suspect some historian at some point is going to be really interested in that and probably would not have ever found that if I hadn't made it something searchable. -Jiminy

Categorize Images. We found a large group of Commons focused editors who categorize images. Categories is "the primary way to organize and find files on Commons ¹¹." Commons focused editors create categories and categorize images with a goal to improve searching. They believe that every image on Commons should be found by browsing the category structure. They also consolidate and clean up categories to eliminate miscategorization and redundancies,

I really enjoy organizing things. So to me working with the categories is organizing. And I think that's important because people who are using Commons as opposed to working on it and administering it and that kind of thing. [And they] need to be able to find the things they want to find. And to do that you need to have them in the right categories, because that's the main way people find things I think. -Alana

Alana is one of the Commons focused editors who specialized in Categorizing. She argues that categorizing is important because it helps people to find the right resource they need.

Construct Norms. We also found a small group of Commons focused editors that work to create and maintain norms of the platform. They help define the scope of Commons, write policies and guidelines, and make them clear to other contributors. Max is one of the Commons focused editors

¹¹<https://commons.wikimedia.org/wiki/Commons:Categories>

who collaborated on writing Commons policies and guidelines. He said the policies and guidelines make Commons resources "*more reliable*",

I think initially, we found that some people on Wikipedia got a bit cross with Commons, because they found that the images were just being deleted. They didn't know why they were being deleted. So it needed a bit more structure. And some definite rules, it shouldn't just be deleted, because an administrator doesn't think this is a good picture. So it has become a bit more formal. But I think that helps because it makes it more reliable. -Max

In the curation process, through activities such as editing metadata, categorizing images and constructing norms, Commons focused editors improved the accuracy of commons resources and made them more accessible, searchable and reliable. These largely facilitate the users to find and reuse Commons resources.

4.2.2 Wikipedia Focused Editors. Wikipedia focused editors focus mainly on curating Wikipedia articles and constructing policies and guidelines for Wikipedia. These activities have been well studied by prior literature [5, 35]. However, in our study, we found Wikipedia focused editors sometimes curate Commons content by editing images and categories.

Edit Images. Wikipedia focused editors edit Commons images to make better illustrations for Wikipedia articles. There are two ways in which an image could be edited on Commons, overwriting an existing file and uploading it as a new file. If an edit is considered as minor, for example, removing a watermark, minor cropping, and increasing resolution of the same file without using artificial methods, the existing file could be overwritten¹²,

I really wanted to illustrate a little bit of the history of [a place], so I found this picture, which I can't remember if I adjusted it. I may have cropped it or adjusted the contrast and focus a little bit. I quite often do that. -Cody

However, if the change is major (e.g., a radical crop, updating a map), editors need to upload the edited file under a new name (e.g., "<image name> (crop).png") without changing the original file,

[That picture] came from this not very promising image. And I've taken a little piece of this image. And then I selected the black background... And then I cropped the image. And I saved it as a gif with a transparent background. And then I pasted that into my SVG. And now if you go back to our SVG... I took this extract of the rather untidy photograph, and I needed the tail a bit higher. So I moved the tail into sections. So the tail is a little bit higher... And then I drew some extra tails to show where the thing would spray on. And I just put dotted lines on there. So I think it's worked quite nicely. It gives quite a good impression of this rather strange piece of biology actually. -Cody

Cody edited a "not very promising image" he found on Commons – he took part of the image and made an illustration from it. He uploaded the illustration as a new version to Commons and used it in his article to give a "*good impression*" of the subject he was writing about.

Edit Categories. Besides editing images, we also found Wikipedia focused editors editing categories on Commons. Dale had a hard time browsing a category of commercial buildings in a city. He had to check all images in the category to identify pictures of a specific building. Therefore, he created a new category that more specifically identify the building for "*anybody else to find it*",

So for this city, I'll browse in categories and maybe I'll find a category of commercial buildings in the city. And I'll look through all of them and say, Oh, is this the one that I'm

¹²https://commons.wikimedia.org/wiki/Commons:Overwriting_existing_files

looking for? And if I do find it, I probably will identify it, I'll probably create a category, more specifically identify it ... to make it easier for anybody else to find it. -Dale

While Wikipedia focused editors curate Commons resources to assist their own editing activities, their work benefit others. The improvements they made to images and the categories they created help other editors to find and reuse resources.

4.2.3 Editors Committed to Both Platforms. Editors who committed to both platforms curate content on both platforms. Their activities overlap with Commons focused editors and Wikipedia focused editors. But we found they did additional work to align Commons categories with Wikipedia article names,

I've managed by various means to narrow something down to the point where you're more likely to find it as a result of that...the category names themselves feed into the search process. -Jasper

Jasper narrowed Commons categories down so there would be a corresponding category on Commons for each Wikipedia article. This makes it easier for Wikipedia editors to find all images related to an article at one place without having to navigating through a broader category.

4.3 Integration

In Commons-Wikipedia collaboration, editors integrate resources of two platforms by inserting Commons images into Wikipedia articles.

4.3.1 Wikipedia Focused Editors. The integration process is mainly conducted by Wikipedia focused editors. Wikipedia focused editors go to Commons, search for images that would illustrate an article, and insert the images into the article. Their search often starts with a known-item search, typing "*some obvious keywords (Cody)*" of an article in the search bar "*like searching Google (Cody)*." However, the known-item search always returns pages of results,

If you search on Wikimedia Commons, you'll find the pages, the categories and lots of images. So it can be kind of bombarded with information. -Jack

Occasionally, Wikipedia focused editors could find "*a really good image*" from one of the first few search results. But this did not happen very often. Most of the time, Wikipedia focused editors had to browse categories to find a good illustration,

The same way I find all of my commons images. I go on Commons, I search for the particular user. And then I see if they've got a category associated with them. Sometimes they already have a really good image, that's the first one you find. But generally, I just look for the category. -Eric

Wikipedia focused editors also tried to reuse images that had already been used in related Wikipedia articles. They did this because they believed that editors of the other articles had already done the hard work to find the best image,

But another (way) is looking at Wikipedia articles. Because then what I can do is I can, you know, other people have already done the work of finding the best images. And so, for example, you show me the picture of [a person] before, I'm going to guess that I didn't search Wikimedia Commons for that. But instead, I went to the Wikipedia article about [the person], and just used the image that they use in that article. Because normally, not always, but normally, the editors of that article will have already found the best image of that person. -Jack

Some Wikipedia focused editors also linked a Commons category to a Wikipedia article,

Somebody had uploaded 30 pictures of some historic building out in Nebraska. And he only used a couple to illustrate the article. But ... he had created a category and he'd put in one of these links to the category. So I learned that you could do that. And in general, I wanted to provide as much information as possible to the Wikipedia readers, but if there were too many photographs to include directly in the article..., you can just link to a category over there. -Dale

Dale explained that this would help "provide as much information as possible to the Wikipedia readers."

4.3.2 Commons Focused Editors. Compared with Wikipedia focused editors, we found Commons focused editors integrated resources of the two platforms in a more indirect way. For example, if Winston thought images he uploaded to Commons are better than the existing images used in a Wikipedia article, he would leave a message on the article's talk page saying "*here's some photos that I've identified on Commons, maybe you should go to this category and take a look and pick one,*"

On occasion ... when I think this should be the best photo, I will oftentimes either leave a talk page said, here's some photos that I've identified on Commons, maybe you should go to this category and take a look and pick one. Some of them do get trapped in. -Winston

Winston further explained that he would not make an image change himself because he might not fully understand the topic of the Wikipedia article that needs to be illustrated.

4.3.3 Committed to Both. Similar to Wikipedia focused editors, editors who committed to both platforms search and reuse Commons images to illustrate Wikipedia articles. But compared with editors who focus on only one platform, editors who committed to both platforms integrate resources of the two platforms in a more efficient way, because they have "*a much more multi-platform view,*"

... that usually start with a much more multi platform view. They know where to find all the pictures. They have everything precisely. They know how to find all sources. They know all the copyright tags. They're just a small environment on their own. -Amos

As Amos said, people who committed to both platforms are experts of the content and rules of both platforms, so they know how and where to find the best illustrations for articles "*precisely.*"

5 BARRIERS TO STITCHING

In the last section, we demonstrated stitching across Commons and Wikipedia. Next, we are going to discuss five barriers we identified that inhibit effective stitching across the platforms.

5.1 Barrier 1: Lack of Communication Across Networks

One challenge to successful stitching is the disparate means and diverse locations of participant communication. Stitching links resources across different technical platforms, but also links individuals and groups as a means of organizing [1]. Stitching forms dynamic connections between the micro-networks that enable content and other resources to be circulated through a network of networks. We found various micro-networks that handle divergent tasks in the Wikipedia ecosystem. For example, networks of photographers focused on producing images of different subjects, a network of Commons admins that handle copyright issues, a network of categorizers that work to sort pictures on Commons, and networks of Wikipedia editors who write articles in specific subject areas. These micro-networks establish their own ways of communicating and organizing their activities. However, participants argued that there was an absence of communication between these distributed micro-networks. For example, there was no formal way for Wikipedia editors and Commons curators to discuss the imagery needs of Wikipedia articles,

I think, obviously, what could be done is if there would be more central oversight in a way... there could be attempts to specifically say, on Wikipedia, a number of articles is working on these and these topics. So editors, on commons, could learn about that, and then help them specifically by trying to find media for those purposes. -Louie

Louie explained that Commons curators could not really know what topics Wikipedia editors were working on. But that was also a problem for Wikipedia. Louie continued to say that Wikipedia editors did not have a way to inform Commons curators about what articles were missing illustrations,

And Wikipedia, you have projects about certain topics, where there are to do lists. And those to do lists often include articles that do not have any photos or images yet. So by going through that, you can sort of lead to articles that would still need media that would then be uploaded to Commons. ... So there could be attempts to strengthen these sorts of pathways to connect the two projects more. -Louie

While Louie highlights how cross-platform awareness is problematic, other participants also pointed out that communication within platform micro-networks is opaque. Participants found it hard to engage in discussions held in other networks to understand their goals and practices,

I am a member of wikimedia foundation Australia, and I have the occasional discussion there. And I also have had worthwhile discussions on particular topics in Wikipedia, but not on Wikimedia Commons. Although currently, I'm trying to engage in some discussion about Australian copyright and copyright of government documents that I've found that extremely difficult, in fact impossible so far, to attract anyone who's interested in that topic. ... But it would be good to have an area where you can go and say it was say, hey, we've got some discussions going on soon. So would you like to join? -Cliff

The lack of mutual awareness of project needs inhibits the formation of broader networks from existing micro-networks. Though Commons curators produce images with an intent to support Wikimedia projects and Wikipedia editors rely on the images to illustrate articles, the communication channels between micro-networks and across the platforms are hard to find.

5.2 Barrier 2: Differing Perspectives

There are many different motivations for Commons curators and Wikipedia editors. That participants have different perspectives on the goals of each of these different platforms is not that surprising. However, the issue is that differing perspectives has an influence on how stitching proceeds.

Editors who focused on developing Wikipedia articles said that Commons has no goals in and of itself. Their view is that the only goal of Commons was to support other Wikimedia platforms. In contrast, editors who focused their work on Commons, stated that the goal of Commons was “*broader than simply being a media library for Wikimedia*” (Jasper). Commons focused editors view Commons primary goal as an online repository of freely licensed multimedia files for open knowledge. Arthur lists at least four different ways that Commons supports open knowledge: (1) support of Wikimedia projects, (2) individual use, (3) documenting states of things, and (4) educational uses,

Well, there are two things really. The first one is to support the Wikipedia and other Wikimedia projects by being the image repository for those projects. But the other thing and to some people this is more important. It is a repository of free media for anybody to use. If you want a picture to put on a T shirt or your own blog, or your academic paper or whatever it might be, just to print out and put on the wall to look nice. You can use Wikimedia Commons. And it documents things. So if you want to know what style of

works artists produced, you can go and look at their images on Commons. And that has an educational value in its own right. So it's part of the free culture movement that makes knowledge available freely to people, which is what Wikimedia's mission is. And it exists. You know, if Wikipedia went away, Wikimedia Commons would still have that purpose. And sometimes there is a tension between the two purposes. Most of the time, they get along nicely, but sometimes they get pulled in different directions, and people have different priorities. -Arthur

Further, Arthur points out that these differing perspectives can create conflicts around producing and curating resources on Commons.

5.2.1 Conflict of Production. The utility or value of any given piece of media differs as a function of whether a participant views Commons as just infrastructure in support of other projects or as an individually useful collection of media. This, in turn, results in a conflict in what should be included or deleted from Commons. People who view Commons as just infrastructure maintain that the value of an image is directly related to how it would be used in a Wikimedia project,

There's an awful lot of rubbish on there, in topics that, frankly, have vastly over represented ... I think you could say that any image that's been put on Commons that is used somewhere else, on any language Wikipedia, ... exemplifies the benefit of Commons that says it can be used. -Nick

These participants suggested that Commons should improve its guidelines to encourage uploading images that would be useful for other Wikimedia platforms. And if an image was not used "at some stage," it should be removed. Because excessive images are an extra cost of time and space, and it makes the category system overly complicated and may obscure the existence of the more useful content,

I think it can improve the guidelines for uploading photographs... And I think It should encourage people to upload photographs that may have some use for a Wikipedia project. ... And maybe if a photograph is not used, then it will be deleted at some stage... I mean, if we have 800 pictures of the lion, then we probably, ... we probably only need about 50. Yeah. So we need lions, eating and sleeping and mating and eating things and climbing trees and having babies and we need them in different countries. And then we probably need some famous ones, like Elsa. And then anyone coming to find a picture of a lion can then find that he's got 50 pictures to choose from. ... Otherwise, if you never find the best picture of a lion, because the categories are too complicated. -Christopher

On the other hand, people who focus on contributing to open knowledge believed evaluating the value of Commons media resources by reuse was problematic and "haphazard." Because the scope of Commons "isn't notability, isn't the Wikipedia concept notability. It's a lot more expansive" (Winston). Not all media in Commons has to be immediately and obviously important or notable. In fact, it is often hard to imagine all of the ways that something could be used or reused making it difficult to determine what should be produced or kept,

I think as a metric reuse is haphazard. ... There are about three occasions I can recall, where suddenly, those amateur photographs of that particular aircraft with this airframe number became important. And that was, of course, when those aircraft crashed. And it became very very newsworthy. And suddenly, that photograph of that aeroplane we were hosting on Commons appeared in newspapers and all over the place, because it was the only freely available photograph they could find of that exact aircraft. And people would say, why do you need 400 photographs of 400 different Boeing triple sevens? Well, each aircraft is unique. And any one of those aircraft might be the one where a wing

drops off one day, you know, so actually, that turns out to be quite important. But from the perspective of use on Wikipedia, yeah, we only need one really good photograph of a Boeing triple seven. But, you know, different usage. -Frank

In this case, a picture of a Boeing triple seven aircraft with a specific airframe number was not used by anyone until that specific aircraft became newsworthy. The picture preserved by Commons became extremely important for the mass media. This illustrates the difficulties of judging the utility of Commons resources as a function of their use in other WMF platforms.

5.2.2 Conflict of Curation. In stitching the act of curation goes beyond the simple addition of metadata. Curation is an act of contextualization, where a statement, an image, a person, is put into a context that allows the viewer to understand something that the curator is attempting to explain or communicate. For editors who view Commons as an infrastructure, media from Commons was seen as key depictions that assists the words of an article. From this perspective Commons media provides important contextualization for words,

I think it's what I would call an honest photo. So as I said before, not the best of a place, not the worst of the place, but the ordinary typical place. ... some people upload photos to Commons that are very artistic. Yeah, but like the sunset photo I mentioned before. They're a pretty photo, but they don't help the person really understand the content of the article. They don't assist with that. So I see the pictures as assisting the words of the article. -Kim

That a piece of media exists, but that it is not quite right to assist with the article creates a challenge to curation. Editors sometimes edit an image (i.e. curate), make “minor improvements” to better illustrate an article. This modification could result in conflicts as a function of what the original image depicts,

Sometimes somebody will crop an image because they want to use it on Wikipedia, to illustrate part of the subject. But from the Commons point of view, the entire image was a historical artifact. And really, the cropped version should be a separate image rather than overwriting the original. Or people will edit images to improve or enhance the contrast or the colors or whatever. And that causes tensions because it's not what Commons wanted. -Arthur

There is a clear policy about what constitutes a “minor edit¹³” in Commons. And, perhaps, in some of these cases the individual who edited the image followed the policy correctly. However, depending on the participants’ perspective of Commons, the act of curation was a type of violation that resulted in conflict. As Arthur is explaining in the quote above, there are special classes of images that have more value than as simply an illustration. The image itself has a unique contextual frame, for example, a historical context, that should be maintained. For these images what constitutes a minor edit is not the same as for other images and curatorial modifications should be limited.

5.3 Barrier 3: Multilingual Resources

Wikimedia Commons enables participants from 300+ Wikimedia projects, including 285 language editions of Wikipedia to collaborate on sharing and curating multimedia content. Participants can contribute in any language they prefer,

And remember, there was no way prior to Commons, for the Wikis to share the same image. The Commons is the means by which the various Wikis can share an image. So that's also extremely valuable to be able to upload it once. Put together metadata in whatever

¹³https://commons.wikimedia.org/wiki/Commons:Overwriting_existing_files

number of languages people may wish, and be able then to share that across all of these different Wikipedia without having to do work redundantly. -Jiminy

Media resources on Commons are technically multilingual, but practically, this is a barrier. Our participants mentioned several ways that the multilingual flexibility of Commons resulted in barriers to stitching. One problem is that the search engine of Commons is key-word based and is not capable of searching “*in the middle of all the languages.*” Some participants had to try different translations of a keyword to enlarge their search result,

So almost anything, sort of specific and detail that you look for, if you do sort of searches to see what the category is missing? You’ll usually find it ... quite a lot over a number of searches. And of course, that’s just in English. Sometimes I do them in French, or even in German or Italian. But I’m not doing in the middle all the languages. If you did them across all the languages, you’d find even more I’m sure, depending on what it is. -James

This issue severely impacted participants from other language editions of Wikipedia who have limited or no English proficiency. They would find “*so little of Wikimedia Commons*” was available for them to search and use in their own language. Because Commons was written in English and the majority of its content was produced and curated by English speakers,

And then one of the other challenges with Commons, of course, is that it’s multilingual. So it’s quite easy for me if I’m searching in English, because a lot of things on Wikimedia Commons are in English. But if someone from ... the Bulgarian Wikipedia wants to find content, they’re going to be really challenged by the fact that so little of Wikimedia Commons is going to be available in Bulgarian they’re going to have to try searching in English. -Jack

This barrier is not just a one-way street, it impacts English-speaking contributors as well. It is difficult for English speakers to find materials about non-English speaking countries because most of the related content was produced and curated in the language spoken in the respective country,

Well, again, the linguistic issue, I mean, if people are not English speakers, it’s much harder. Commons is in theory a multilingual project, probably over 50% of the work done there is done in one language. There’s a few other languages that are pretty serious contenders. And I would say conversely, if you’re looking for material about China, you do a lot better to have Chinese to do your search than English, because an awful lot of the material about China is documented only in Chinese. And that’s less so for some other countries. That’s probably the extreme of it in my experience, is you know, you go looking for things about take a country that I have some from connecting to Romania. You know, it’s probably in Romanian English. -Jiminy

Wikidata¹⁴ is a structured data project, backed by WMF, that was integrated with Commons in 2017. The integration efforts encourage contributors to add “multilingual information about a media file that can be understood by humans, with enough consistency that it can also be uniformly processed by machines¹⁵” to support multilingual search,

The fact that Commons was born as a multilingual platform in a moment where everyone language was just on its own. So people were considering multilingualism. Like, I write in English and Spanish. And that’s fair. ... And so now we have wikidata. We know how really multilingual flexible environment actually works. But the time in 2005, there was no culture about that. And it takes time to replace that old culture with a new one. And there is some resistance. -Amos

¹⁴https://www.wikidata.org/wiki/Wikidata:Main_Page

¹⁵https://commons.wikimedia.org/wiki/Commons:Structured_data

While many recognize the potential value of clear structured data, Amos suggests that on Commons it is largely ignored, if not opposed. Amos explained that the project did not align with the “old culture” of Commons. There is still a long way to go before editors can search and use multilingual resources in any language they speak.

5.4 Barrier 4: Cross-Platform Vandalism

Commons offers the convenience for multiple WMF platforms to share and reuse the same media resources without having to upload it to their local servers redundantly. The resource is uploaded to and held by Commons and editors of other projects can reuse it by simply inserting the appropriate Wikitext. When the original image on Commons is improved, all Wikipedia platforms that use that same image automatically benefit from the new improved version,

And it means if somebody improves an image, if somebody takes a photograph and restores it by removing dust and scratches from an old print, they don't just improve it for the Wikipedia in their language, but they improve it for all of Wikipedia, and all the other Wikimedia projects. So that shared effort makes the work on all the other projects a little bit easier. -Arthur

However, this also creates an opportunity for cross-platform vandalism. Because anyone can overwrite an image on Commons to change what is displayed on other platforms for malicious purposes,

... there was an incident of a celebrity ... porn outing on their Wikipedia page. So there were nude photographs of this actress. I can't I don't remember who it was. ... And for some period of time for like five minutes. Maybe it was longer. This unfortunate woman had her nude photographs on her Wikipedia article. I don't think the press picked up. The concern for Wikipedia, English Wikipedia, was that it was a type of vandalism. It was a form of inter project vandalism that was undetectable on Wikipedia. So recent change editors on Wikipedia wouldn't see anything changing if the image on Commons was overwritten. And there's nothing to stop somebody even even a new editor from overwriting images for the vast majority of images, including portraits of celebrities. -Frank

Stitching, bringing resources and communities together through resource and communicative linking is just as open to misuse, vandalism, and misinformation as any other web or internet related activity. And as Frank explains, there's no way for Wikipedia editors to track when an image is overwritten on commons. This makes this stitching activity vulnerable to vandalisms that are difficult to detect technically.

5.5 Barrier 5: Differing Policies

Social media platforms develop policies and guidelines that reflect the platform's values [11, 12, 22]. While Commons and the different language editions of Wikipedia share the same underlying philosophical stance toward open and free knowledge, policies and guidelines still reflect some of the individual character of the differing online communities. The, sometimes subtle, differences in policies and guidelines can be a barrier to effective stitching across the platforms.

One key misalignment between Commons and Wikipedia is how copyright is treated. Commons implements a “Precautionary principle¹⁶” which states that “where there is significant doubt about the freedom of a particular file, it should be deleted.” This delete first and discuss it later approach is in contrast to Wikipedias’ “Assume Good Faith” policy¹⁷ that encourages discussion first. An image

¹⁶https://commons.wikimedia.org/wiki/Commons:Project_scope/Precautionary_principle

¹⁷https://en.wikipedia.org/wiki/Wikipedia:Assume_good_faith

uploaded by a Wikipedia editor to illustrate an article could be deleted because of uncertainty about copyright, which could cause a problem,

The way they treat copyright and the so called precautionary principle on commons, where if there is a doubt about copyright, rather than a certainty that something is in copyright, it still gets deleted. On Wikipedia, sometimes that causes a problem because the needed images disappeared. -Arthur

Another misalignment is in the differences between Wikipedia and Commons reliance on data sources. The practice on Wikipedia is to “citing sources¹⁸”, and in particular, “reliable sources¹⁹” all in the service of making statements “verifiable²⁰”. Media resources on Commons do not need to satisfy all of these standards and there is no judgement as to the validity or correctness of the media artifact. From one perspective similar versions of something like a map or a deep fake image, might have high utility when contrasted with alternate versions. Frank helps make this point,

If you're aware Commons has a policy where don't judge the correctness of content. So if somebody makes a map, and said, and says China rules, all these territories, and this is a map, it is a political map. And on Commons, Commons does not judge whether it's correct or not. So we don't make the decision. But what we do say is actually it is fine to have several maps of the same territory, as presented by different political parties, or political groups, even if that's offensive, and on commons, we therefore include offensive content. And that kind of comes under a form of free speech. The problem with it is then how people use it on other projects. If someone says this is the truth and puts in an article about here is a map of this region and their political challenge, you know, regions, which happens all the time. -Frank

Given this generally inclusionary standard, Commons contributors sometimes produce images without including information about the sources of the data as part of the content metadata. Without this key metadata, media is then suspect under Wikipedia's stricter policies,

Well, largely because one of the issues with maps on Wikipedia, or on Wikimedia Commons, for that matter, is that a lot of them have been made by people. And they might have been based on some reliable source. But often the source is not actually provided in the file description. So in order for a file to be acceptable at the particular map file to be acceptable at a featured article class, what's required is that the map has to be properly sourced, you have to be able to say, Well, I copied this from a map on this page of this book, or whatever... That's a common problem with maps. And people often, yeah, upload maps that they've made themselves within Skype or whatever. But they're not. But they're not able to be used, particularly in featured articles, because they're not, it's not clear where the information that's on the map has come from. -Philip

While Wikipedia and Commons share similar ideals, the subtle policy differences and the slightly different goals of each platform are a barrier to stitching resources and community members across these platforms.

6 DISCUSSIONS AND IMPLICATIONS

Through an empirical analysis of 32 interviews with Wikipedia and Wikimedia Commons editors, this study demonstrates how actors with different perspectives on the work and the platforms enacted stitching. We extended the conceptual framework of stitching by articulating five barriers

¹⁸https://en.wikipedia.org/wiki/Wikipedia:Citing_sources

¹⁹https://en.wikipedia.org/wiki/Wikipedia:Reliable_sources

²⁰<https://en.wikipedia.org/wiki/Wikipedia:Verifiability>

that inhibit effective stitching across the platforms. In the following section, we first reflect on Gillespie's definition of platform and our findings. We then discuss stitching as a lens for studying cross-platform collaboration and illustrate how our contributions extend the prior work on stitching. This, in turn, links our study to the existing body of CSCW work. Lastly, We provide design implications for mitigating barriers of stitching across Wikipedia and Wikimedia Commons.

6.1 Definition of Platform and Barriers in Stitching

One of the significant problems with studying cross-platform collaborative work is defining what constitutes a platform. Prior work that compares platforms and the work done on those platforms has been able to rely on the fact that the technical infrastructures they studied were owned by different organizational entities (e.g., Facebook, YouTube, Twitter, Snap). In our case, the Wikimedia Foundation (WMF) is the entity that manages and hosts both Wikipedia and Wikimedia Commons. Therefore, we could not easily rely on the same "obviousness" of a claim to platform uniqueness. We worked to resolve this by adopting Gillespie's [15] definition of platform and explained the differences between Wikipedia and Commons in terms of political, figurative and architectural meanings of platform. We believe that Gillespie's extension and our explanation (c.f. Section 2.1) provide one clear example of how the definition distinguishes different platforms, even when they are owned or managed by the same entity.

Our findings with regard to the barriers to stitching allow us to reflect on Gillespie's [15] definition and what it means for cross-platform work in an interesting way. Thinking about our data and Gillespie's definition we hypothesize two different things. First, we believe that different barriers to collaboration can be found in each of the different definitions of platform (i.e., *architectural*, *figurative*, *political*). That is, the different meanings articulate spaces where barriers arise because of the differences in the platform. Second, the scale of the differences in meaning would seem to contribute to the severity of the barrier. Our work cannot prove these hypotheses, but we would support the potential of them as follows.

Politically, Wikipedia and Commons have very different politics, which led to differing policies and communication patterns. We found differing copyright and data source policies of Commons and Wikipedia inhibit effective content reuse across the two platforms. Besides, networks of these two platforms defined means and locations of participant communication — what to discuss, "openness" of the discussion, where should the discussion happen, and who should be included in the discussion differently. This resulted in a lack of communication across networks, which inhibits stitching coordination.

Architecturally, the user interface of Commons is in English. Participants noted that this creates a specific type of barrier. This enables English speakers to produce, curate and reuse resources on Commons quite easily. However, Commons supports more than 280 language editions of Wikipedia, each of which has interface written in its own language to support local editors. The English Wikipedia is just one edition. Language is a significant obstacle for non-English speakers from non-English Wikipedia versions to contribute and reuse resources on Commons. The language architectural difference is not the only one. Commons carefully versions images and shows the edit history of them. However, where Wikipedia relies on an image, it does not allow editors to track when an image is overwritten/edited on Commons. Therefore, the stitching activity across two platforms is vulnerable to a type of vandalism that is very difficult to detect.

Figuratively, Wikipedia and Commons have different goals. The goal of Wikipedia is to create an encyclopedia. But Commons aims to be a collection that provides freely-licensed educational media content to everyone including Wikipedia editors. Different goals shape editors' perspectives and practices differently. As a result, Commons editors and Wikipedia editors have conflicts around what constitutes a valuable resource and how edits on images should be made.

Considering our use of Gillespie's definition and what we believe our findings illustrate about the potential for differences to create barriers, our findings also suggest a way to think through design problems that could mitigate barriers. Our findings suggest that in addition to seeking technical solutions that bring multiple platforms in collaboration, researchers and designers should also carefully inspect and design to either mitigate the differences in the definitions across platforms or bring the platforms into better alignment with regard to the political, architectural and figurative meanings of the platforms. That is, by considering Gillespie's definition and analyzing how platforms differ across the additional meanings, designers can identify potential cross-platform collaboration challenges.

6.2 Stitching as a Lens to Study Cross-Platform Collaboration

Prior literature in CSCW has investigated the use of multiple social media platforms by individuals in situations like crisis response [8, 9] and sharing personal content [31, 39]. This thread of work has grown to consider how a crowd collaborates across platforms to achieve goals such as crisis communication [20], fact checking [17], misinformation correction [30, 38], information operations [32, 37], and disinformation campaigns [33]. All these studies show that researching and designing to support (or inhibit) such activities is challenging due to their cross-platform nature and the potentially large scale of the collaboration. Our study demonstrates that stitching is a promising framework for understanding the cross-platform collaborations of the crowd. While a few prior studies in CSCW have used the framework of stitching to uncover some aspects of cross-platform work, none of them looked at all three processes of stitching. In this study, through an analysis of the three stitching processes, we were able to understand the collaborative work that happens between Wikipedia and Commons. We identified how resources were produced and shared, the ways to maintain, preserve and sort resources and how resources were transmitted and integrated between the two platforms. We also explained how editors with different goals and platform focuses contributed to each of the stitching process. These allowed us to identify five barriers for effective stitching, which could not be fully revealed by studying only one process of stitching or only one type of actor involved. The barriers we identify extend our understanding of stitching and the utility of the framework for considering both the effectiveness of stitching and what may prevent stitching.

Though our study focuses on stitching between Wikipedia and Wikimedia Commons, the barriers we identified are generalizable to other cross-platform collaborations. For example, the barrier of differing policies has also been observed in the effort to combat cross-platform malicious information operations. Studies [33, 37] found that while larger, more established social media platforms have policies to censor and remove misinformation, non-mainstream platforms that do not have such policies enable the strategic narratives to build resilience and be sustained in the information ecosystem. Besides, the barriers of cross-platform vandalism and lack of communication between networks exist in cross-platform disinformation identification. The detection and determination of vandalism behaviors often stay within an individual platform, but combating cross-platform disinformation campaigns needs data sharing, communication and collaboration between platforms [32, 33]. The current strategies seem to mostly rely on humans to recognize and link the efforts in their own minds. This effectively relies on a type of cross-platform moderation that is unlikely to happen if platforms are differentiated by the owner and/or host. Moreover, understanding multilingual resources is a barrier for cross-platform fact-checking work as well. And this calls for multilingual systems that are capable of understanding misinformation that spans both geographic and linguistic boundaries [17]. A closer look at the stitching processes would help understand the social and technological causes of these barriers and design to facilitate collaboration.

6.3 Designing to Mitigate Barriers to Stitching across Wikipedia and Commons

In the following, we discuss a few design implications that could help mitigate some of the barriers to stitching across Wikipedia and Wikimedia Commons.

6.3.1 Contextualizing Use and Reuse. The current design of Commons and Wikipedia allows an image to be contextualized differently by the contributor, curators and users. Our study shows that without knowing something of the "other" platform context, editors with perspectives shaped by their own platform could disagree with editors from the "other" platform. An example is arguments between Commons and Wikipedia editors around what constitutes useful resource and where a minor edit could be made. Prior work has also identified the contextualization problem, for specific topics. For example, to address contextualization challenges of illustrating articles about women's health and anatomy, Menking and McDonald [28] suggest that medical images should be hosted directly on Wikipedia instead of on Commons. This would help avoid Commons' censorship and give Wikipedia editors more control over where such images could be reused. While this approach might work for the specific case they studied, it does not generalize well to Wikipedia-Commons stitching. An alternative approach could be a process similar to Wikipedia's notability voting. The process would enable editors from both platforms to figure out whether an image warrants significance in any contexts collaboratively, rather than relying on judgement of editors from one platform or the other.

6.3.2 Improving Multilingual Support. As we mentioned in Section 5.3, Wikimedia Foundation has recognized the multilingual issue across its platforms and introduced the Structured Data project²¹ to tackle it. However, the project made little progress on Commons because many contributors simply did not know about it or did not care. Many participants explained that they ignore the project because they felt that structured data is a layer of work imposed on them by the foundation. Further, our participants felt that structured data is redundant to Commons existing category system, which has been developed, maintained and used by volunteers for years. Editors preferred their "own" system over a new structure designed by the foundation. In many ways this illustrates a classic CSCW challenge: who does the work versus who gets the benefit [16]. One potential solution is for the foundation to investigate ways to incorporate Commons existing categories into the Structured Data Project. This case demonstrates that while designing stitching technologies to facilitate cross-platform collaboration, one should carefully investigate the existing processes and practices adopted by the crowd to avoid design solutions that might be ignored.

6.3.3 Supporting Technical and Social Collaboration. In this paper, we identified a set of barriers to stitching in Wikipedia ecosystem. But any given barrier has different implications for stitching as a function of the different platform from which it is considered, barriers are multifaceted. Considered from Wikipedia's point of view, cross-platform vandalism could be an integration problem, in which an image used to illustrate an article is changed, or vandalized. But the same barrier considered from Wikimedia Commons point of view could be a curation problem; an image is overwritten. The problem of cross-platform vandalism illustrates another key challenge in design to support stitching. Essentially, one platform alone is rarely able to address a stitching barrier. Though Wikimedia Commons can track and detect when an image is overwritten, it is hard to evaluate the legitimacy of the overwrite because the context of reuse is unknown. That is, all of the locations where the image is being used may be technically known, but how those specific reusages are integrated with text is unknown. From the other side, Wikipedia editors can easily for a suspicion that an image modification is vandalism when they see a suspect image in the context of an article. However, a

²¹https://commons.wikimedia.org/wiki/Commons:Structured_data

Wikipedia editor cannot easily track when images are modified and overwritten. The 'event' of the overwrite on Commons does not provide an event notification to the other Wikipedia platforms where it is used. Even if these notifications existed, these platforms would need to collaborate on addressing the problem. In the general case, resolving barriers will require technical and social collaboration across or between platforms.

6.3.4 Design to Foster Communication. Wikimedia Commons and many language editions of Wikipedia have dedicated places like the village pump, the help desk, and many policy talk pages, for participants to hold discussions. However, participants do not feel like there are channels for discussions that effectively cross platforms. WMF presents MetaWiki as a potential platform for discussions that impact all platforms, but what our participants wanted may not clearly fit that purpose, as MetaWiki deals mainly with the admin side of the projects rather than the content side. While creating talk pages seems to be an easy solution for holding discussions within individual Wiki platforms, there's no truly obvious solution to address this cross-platform communication barrier. Talk pages need to be held on a specific platform, but there are multiple platforms involved in stitching. Finding a preferred platform for each type of the discussions requires a joint effort of the communities involved. As we've discussed earlier, resolving a barrier of stitching often needs collaboration across platforms.

Prior studies in CSCW/HCI have investigated similar situations in which stakeholders of a design problem were distributed. For example, Walsh et al. [36] explored ways to include a large group of children that were distributed globally in the design of new technologies. Heintz et al. [18] created Pdot, an online participatory design tool for widely distributed project stakeholders. Macleod et al. [21, 27] promotes the Asynchronous Remote Communities (ARC) method that allows researchers to run online participatory design sessions with distributed participants. These studies use tools and methods that apply participatory design principles in a distributed, asynchronous design context rather than traditional, synchronous, in-person participatory design sessions [7, 13]. WMF could explore these approaches to engage editors distributed across platforms in a participatory design process to address their communication needs.

6.4 Limitations

As noted above, although we aimed to recruit a diversity of participants, only 5 out of 32 participants we interviewed self-identified as a woman. This reflects a similar gender gap in the Wikimedia ecosystem. Our study focuses on stitching across Wikimedia Commons and English Wikipedia creating a clear language diversity limitation. Wikimedia Commons is supposed to be a media repository for all Wikipedia language editions. While six participants mentioned that they contributed to language editions of Wikipedia other than English, our findings are unlikely to generalize to stitching across Commons and other language editions of Wikipedia. Future work could extend our findings by considering different aspects of these limitations.

7 CONCLUSION

In this paper, we present an analysis of 32 participants' experience of working across or between distinct platforms. Our work contributes an empirical understanding of stitching across Wikipedia and Wikimedia Commons and extends the conceptual framework to explain barriers that inhibit effective stitching across platforms. This study shows an exciting potential for future work to study stitching technologies that support cross-platform work and collaboration in peer-production systems.

REFERENCES

- [1] W Lance Bennett, Alexandra Segerberg, and Shawn Walker. 2014. Organization in the crowd: peer production in large-scale networked protests. *Information, Communication & Society* 17, 2 (2014), 232–260.
- [2] Avijit Bhattacharjee, Sristy Sumana Nath, Shurui Zhou, Debasish Chakroborti, Banani Roy, Chanchal K Roy, and Kevin Schneider. 2020. An Exploratory Study to Find Motives Behind Cross-platform Forks from Software Heritage Dataset. In *Proceedings of the 17th International Conference on Mining Software Repositories*. 11–15.
- [3] Taryn Bipat. 2020. A Comparative Language Study: Characterizing Collaboration Models in the EN, FR, and ES Language Editions of Wikipedia. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 81–86.
- [4] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [5] Brian Butler, Elisabeth Joyce, and Jacqueline Pike. 2008. Don't look now, but we've created a bureaucracy: the nature and roles of policies and rules in wikipedia. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1101–1110.
- [6] Matthew Chalmers and Ian MacColl. 2003. Seamless and seamful design in ubiquitous computing. In *Workshop at the crossroads: The interaction of HCI and systems issues in UbiComp*, Vol. 8.
- [7] Luigina Ciolfi, Gabriela Avram, Laura Maye, Nick Dulake, Mark T Marshall, Dick van Dijk, and Fiona McDermott. 2016. Articulating co-design in museums: Reflections on two participatory processes. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 13–25.
- [8] Dharma Dailey and Kate Starbird. 2014. Journalists as crowdsourcers: Responding to crisis by reporting with a crowd. *Computer Supported Cooperative Work (CSCW)* 23, 4-6 (2014), 445–481.
- [9] Dharma Dailey and Kate Starbird. 2017. Social media seamsters: Stitching platforms & audiences into local crisis infrastructure. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 1277–1289.
- [10] Michael A DeVito, Jeremy Birnholtz, and Jeffery T Hancock. 2017. Platforms, people, and perception: Using affordances to understand self-presentation on social media. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*. 740–754.
- [11] Casey Fiesler. 2007. Everything I need to know I learned from fandom: How existing social norms can help shape the next generation of user-generated content. *Vand. J. Ent. & Tech. L.* 10 (2007), 729.
- [12] Casey Fiesler, Jessica L Feuston, and Amy S Bruckman. 2015. Understanding copyright law in online creative communities. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 116–129.
- [13] Hugo Fuks, Heloisa Moura, Debora Cardador, Katia Vega, Wallace Ugulino, and Marcos Barbato. 2012. Collaborative museums: an approach to co-design. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. 681–684.
- [14] William W Gaver. 1991. Technology affordances. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 79–84.
- [15] Tarleton Gillespie. 2010. The politics of 'platforms'. *New media & society* 12, 3 (2010), 347–364.
- [16] Jonathan Grudin. 1994. Groupware and social dynamics: Eight challenges for developers. *Commun. ACM* 37, 1 (1994), 92–105.
- [17] Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2021. A survey on automated fact-checking. *arXiv preprint arXiv:2108.11896* (2021).
- [18] Matthias Heintz, Effie L-C Law, Sten Govaerts, Adrian Holzer, and Denis Gillet. 2014. Pdot: participatory design online tool. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. 2581–2586.
- [19] Jeff Hemsley, Jenna Jacobson, Anatoliy Gruzd, and Philip Mai. 2018. Social media for social good or evil: An introduction. *Social Media+ Society* 4, 3 (2018), 2056305118786719.
- [20] Amanda Lee Hughes, Kate Starbird, Alex Leavitt, Brian Keegan, and Bryan Semaan. 2016. Information movement across social media platforms during crisis events. *CHI'16 Extended Abstracts 1* (2016).
- [21] Haley MacLeod Ben Jelen, Prabhakar Annu, Oehlberg Lora, Siek Katie, and Connelly Kay. 2016. Asynchronous remote communities (arc) for researching distributed populations. In *10th EAI International Conference on Pervasive Computing Technologies for Healthcare*. [Google Scholar].
- [22] Jialun'Aaron' Jiang, Skyler Middler, Jed R Brubaker, and Casey Fiesler. 2020. Characterizing Community Guidelines on Social Media Platforms. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 287–291.
- [23] Aniket Kittur, Bryan Pendleton, and Robert E Kraut. 2009. Herding the cats: the influence of groups in coordinating peer production. In *Proceedings of the 5th international Symposium on Wikis and Open Collaboration*. 1–9.

- [24] Travis Kriplean, Ivan Beschastnikh, and David W McDonald. 2008. Articulations of wikiwork: uncovering valued work in wikipedia through barnstars. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. 47–56.
- [25] Travis Kriplean, Ivan Beschastnikh, David W. McDonald, and Scott A. Golder. 2007. Community, Consensus, Coercion, Control: Cs*w or How Policy Mediates Mass Participation. In *Proceedings of the 2007 International ACM Conference on Supporting Group Work (Sanibel Island, Florida, USA) (GROUP '07)*. Association for Computing Machinery, New York, NY, USA, 167–176. <https://doi.org/10.1145/1316624.1316648>
- [26] Shyong K Lam, Jawed Karim, and John Riedl. 2010. The effects of group composition on decision quality in a social production community. In *Proceedings of the 16th ACM international conference on Supporting group work*. 55–64.
- [27] Juan F Maestre, Haley MacLeod, Ciabhan L Connelly, Julia C Dunbar, Jordan Beck, Katie A Siek, and Patrick C Shih. 2018. Defining through expansion: conducting asynchronous remote communities (arc) research with stigmatized groups. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [28] Amanda Menking and David W McDonald. 2020. Image Wishlist: Context and Images in Commons-Based Peer Production Communities. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–21.
- [29] Jessica A Pater, Moon K Kim, Elizabeth D Mynatt, and Casey Fiesler. 2016. Characterizations of online harassment: Comparing policies across social media platforms. In *Proceedings of the 19th international conference on supporting group work*. 369–374.
- [30] Emily Saltz, Claire Leibowicz, and Claire Wardle. 2020. Encounters with Visual Misinformation and Labels Across Platforms: An Interview and Diary Study to Inform Ecosystem Approaches to Misinformation Interventions. *arXiv preprint arXiv:2011.12758* (2020).
- [31] Manya Sleeper, William Melicher, Hana Habib, Lujio Bauer, Lorrie Faith Cranor, and Michelle L Mazurek. 2016. Sharing personal content online: Exploring channel choice and multi-channel behaviors. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 101–112.
- [32] Kate Starbird, Ahmer Arif, and Tom Wilson. 2019. Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–26.
- [33] K Starbird and T Wilson. 2020. Cross-platform disinformation campaigns: lessons learned and next steps. *Harvard Kennedy School Misinformation Review* 1, 1 (2020).
- [34] Janet Vertesi. 2014. Seamful spaces: Heterogeneous infrastructures in interaction. *Science, Technology, & Human Values* 39, 2 (2014), 264–284.
- [35] Fernanda B Viégas, Martin Wattenberg, Jesse Kriss, and Frank Van Ham. 2007. Talk before you type: Coordination in Wikipedia. In *2007 40th Annual Hawaii International Conference on System Sciences (HICSS'07)*. IEEE, 78–78.
- [36] Greg Walsh, Craig Donahue, and Zachary Pease. 2016. Inclusive co-design within a three-dimensional game environment. In *Proceedings of the The 15th International Conference on Interaction Design and Children*. 1–10.
- [37] Tom Wilson and Kate Starbird. 2021. Cross-platform Information Operations: Mobilizing Narratives & Building Resilience through both 'Big' & 'Alt' Tech. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–32.
- [38] Wenqing Zhao. 2019. Misinformation Correction Across Social Media Platforms. In *2019 International Conference on Computational Science and Computational Intelligence (CSCI)*. IEEE, 1371–1376.
- [39] Xuan Zhao, Cliff Lampe, and Nicole B Ellison. 2016. The social media ecology: User perceptions, strategies and challenges. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 89–100.

A RECRUITMENT STRATEGIES

A.1 Commons Active Users List

The Active User list²² is a Special Page that shows curators on Wikimedia Commons who have made more than 10,000 edits in the last 30 days. The first author manually checked the user page of the top 200 most active curators from the list and selected 81 of them who self-identified as a native English speaker or who indicated that they were able to contribute with a near-native/advanced/intermediate level of English.

²²<https://commons.wikimedia.org/wiki/User:F%C3%A6/Userlist#top>

A.2 Commons Autopatrolled Users List

The Autopatrolled Users List²³ shows possible candidates for autopatrolled status on Wikimedia Commons. “Autopatrolled is a user right given to prolific creators of clean articles and pages in order to reduce the workload of New Page Patrol. Articles and pages created by autopatrolled editors are automatically marked as “reviewed” or “patrolled” in the system.” We use this list to recruit prolific commons curators who were nominated as contributing high-quality content by the community.

A.3 Commons Help Desk

The Help Desk²⁴ is a community page where individuals can go to ask questions and attempt to have problems resolved. We recruited less experienced curators by considering the individuals who were using the Help Desk to ask what appeared to be basic questions. The first author read through recently archived threads posted on Commons Help Desk and identified curators who asked questions about contributing to Commons in English. The first author checked their Commons user page and contribution page and selected 43 of them who made at least 5 contributions in the last two years.

A.4 Wikiscan Tool for English Wikipedia

Wikiscan²⁵ is a tool that shows edit statistics of Wikimedia projects and their contributors. To identify English Wikipedia editors who have used content from Wikimedia Commons to illustrate Wikipedia articles, we use Wikiscan on English Wikipedia to find editors who made “File:” changes most frequently. The first author manually checked the contribution list of the top 300 editors and invited those who used at least five images from Wikimedia Commons to illustrate articles in their most recent 2,500 contributions.

A.5 English Wikipedia Featured Article List

We manually checked the edit history of articles in English Wikipedia’s Featured Article archive for 2019²⁶. We identified editors who made the most recent imagery change to the featured articles and checked their contribution list. We invited editors who used at least three images from Commons in the last two years. In this way we could reach out to editors who have used content from Commons to illustrate Wikipedia articles but not as frequently as users we identified from Wikiscan.

A.6 Snowball Sampling

At the end of each interview, we asked participants to recommend one or two Wikimedia Commons curators and/or English Wikipedia Editors who we should contact and potentially interview. We pointed out that we were interested in a diverse range of participants (e.g., non-native English speakers, female editors, controversial figures...)

B INTERVIEW PROTOCOL

B.1 Participants Recruited through Wikimedia Commons

Opening Script

²³https://commons.wikimedia.org/wiki/Commons:Requests_for_rights/possible_autopatrolled_candidates

²⁴https://commons.wikimedia.org/wiki/Commons:Help_desk

²⁵<https://wikiscan.org/>

²⁶https://en.wikipedia.org/wiki/Wikipedia:Today%27s_featured_article/2019

Hi! I'm <Researcher 1 self-introduction, Introduce Researcher 2>. Today I'd like to talk with you about your experiences working with Wikimedia Commons. I have a set of questions that outline our conversation, but I'll also ask follow-up and clarification questions.

I want you to know that your participation in this study is completely voluntary. If at any time you would like to stop, just tell me and we will stop the interview. Also, if there is a specific question you would prefer not to answer - let me know and I will move to the next question.

With your permission I'd like to audio record the interview. The only people who will have access to the audio recording are me and my advisor. We will transcribe the interviews and erase the recording.

In case of poor network connection during the interview: If my connection drops, please stay on line and wait for at least five minutes for me to rejoin the interview; If your connection drops, please try to rejoin the interview using the same URL/Meeting ID. I'll wait for you for up to five minutes for you to rejoin. If we cannot resume our interview in five minutes, I will contact you to reschedule the interview. I'll email the details to you.

It is our practice to make a good faith effort to maintain the confidentiality of this interview. We will not tell outsiders that you were part of this research. We'll do our best to anonymize the data. However, we feel it is important to tell you that we know that Wikipedians sometimes figure out who and what is being described, even though we have worked to anonymize the data.

Do you have any questions for me before we start?

<Start recording . . . Turn on the recorder> Just for the record, is it ok for us to audio record your interview?

Phase 1: Becoming a Curator on Commons

Q1: Tell me, how long have you been editing Wikimedia commons?

Q2: How has your work in Wikimedia Commons changed over time?

Follow up:

What did you contribute to Commons when you started?

Why did you choose to contribute to Wikimedia Commons?

What types of work do you do now?

How would you describe the difference between what you worked on when you started and what you are working on right now?

Why did things change?

Q3: What contribution have you made to Wikimedia Commons that makes you most proud or happy?

Follow up:

Could you show us your favorite contribution to Commons?

About when did you make this contribution?

Could you show us something that you are working on now - or something that you worked on recently?

Q4: How do you collaborate with other Commons editors?

Follow up:

Do you communicate or interact with others?

How do you interact?

How do these interactions help you decide what work to do in Commons?

Phase 2: Tasks and Activities

As part of our effort to find people to interview we reviewed lots of contributions to Wikimedia Commons. I would like to show you some of the contributions you made and ask you some questions

about them. I'm going to paste a URL into the chat that will link to the first example. I'm also going to share my screen so that you can see what I see.

<Link to Example 1>

<Link to Example 2>

<Link to Example 3>

<Link to Example 4>

<Link to Example 5>

<For each example, ask:>

Q5: What is the story behind this contribution?

Q6: Why did you choose to contribute this? Why is this important to have in Commons?

This is very interesting ... <a little transition to get to Q7>

Q7: If you were to characterize the majority of your work in Commons, what would it be?

Follow up:

What type of work is the majority of your work in Commons?

Q8: Do you think there is a type of work that is most valued in Commons?

Follow up:

Is there a hierarchy of contributions - some more valued than others?

Phase 3: Connection with Other Wikimedia Projects

Q9: In your opinion, how does Commons support other Wikimedia projects, like Wikipedia?

Q10: Do you know of an example of Commons content that is currently being used in another Wikimedia project such as Wikipedia?

Follow up:

Do you know if something you contributed is being used?

Q11: For the next question, I'm going to show you a list of the most linked-to files. I'll paste the URL into the chat so you can click through to the page <copy and paste this url into the chat:

<https://commons.wikimedia.org/wiki/Special:MostLinkedFiles>>

I'm also going to share my screen so that you can see what I see.

Could you briefly describe for me what you see on this page?

How does this list show Commons support for Wikimedia projects?

Do you think this is a good example of Commons supporting Wikimedia projects?

If not, can you think of a better list?

Q12: Do you know how Wikipedia editors find content from Commons to support their work?

Q13: How does your work in Commons support use (or reuse) by Wikipedia editors?

Follow up:

What makes supporting use (or reuse) difficult?

Q14: How would you describe the goals of Commons in relation to Wikipedia?

Follow up:

What if we flip the question - How would you describe the goals of Wikipedia in relation to Wikimedia Commons?

Closing Script

Do you have any questions for me?

Do you know of 1 or 2 Wikimedia Commons contributors who we should contact and potentially interview? <snowball sampling>

Can I say that you were the person who referred us to <user>?

Note if we use your name we are breaking confidentiality.

We are acknowledging the time and effort of our participants by making a contribution to one of three ‘like-minded’ organizations. Would you like us to acknowledge your participation with a contribution to either:

- Wikimedia Foundation
- Creative Commons
- Internet Archive

Thank you for taking the time to talk with me about Wikimedia Commons. Do you have anything else you would like to add?

B.2 Participants Recruited through Wikipedia

Opening Script

Hi! I’m <Researcher 1 self-introduction, Introduce Researcher 2>. I would like to talk with you about your experiences using Wikimedia Commons. I have a set of questions that outline our conversation, but I will also ask follow-up and clarification questions.

I want you to know that your participation in this study is completely voluntary. If at any time you would like to stop, just tell me and we will stop the interview. Also, if there is a specific question you would prefer not to answer - let me know and I will move to the next question.

With your permission, I would like to audio record the interview. The only people who will have access to the audio recording are me and my advisor. We will transcribe the interviews and erase the recording.

In case of poor internet connection during the interview: If my connection drops, please stay online and wait for at least five minutes for me to rejoin the interview; If your connection drops, please try to rejoin the interview using the same URL/Meeting ID. I’ll wait for you for up to five minutes for you to rejoin. If we cannot resume our interview in five minutes, I will contact you to reschedule the interview. I’ll email the details to you.

It is our practice to make a good faith effort to maintain the confidentiality of this interview. We will not tell outsiders that you were part of this research. We will do our best to anonymize the data. However, we feel it is important to tell you that we know that Wikipedians sometimes figure out who and what is being described, even though we have worked to anonymize the data.

Do you have any questions for me before we start?

<Start recording ... Turn on the recorder> Just for the record, is it ok for us to audio record your interview?

Phase 1: Wikipedia Participation

Q1: Tell me, how long have you been editing English Wikipedia?

Q2: How has your work in English Wikipedia changed over time?

Follow up:

What did you contribute to Wikipedia when you started?

Why did you choose to contribute to Wikipedia?

What types of work do you do now?

How would you describe the difference between what you worked on when you started and what you are working on right now?

Why did things change?

Phase 2: Tasks and Activities

As part of our effort to find people to interview we reviewed lots of contributions to Wikipedia. I would like to show you some of the images you used to illustrate Wikipedia articles and ask you some questions. I'm going to paste a URL into the chat that will link to the first example. I'm also going to share my screen so that you can see what I see.

<Link to Example 1>

<Link to Example 2>

<Link to Example 3>

<Link to Example 4>

<Link to Example 5>

<For each example, ask:>

Q3: What is the story behind this illustration?

Q4: How did you find this image?

Q5: Why is this the right image for the article?

This is very interesting ... <a little transition to get to Q6>

Q6: How would you describe the way you use Wikimedia Commons?

Follow up:

How have you contributed to Wikimedia Commons?

Q7: What are the types of media in Wikimedia Commons?

Follow up:

Which of the media types have you used in Wikipedia articles?

In your opinion, what is the most valued media type?

Q8: Is there an image that you believe exemplifies the value of Commons to Wikipedia?

Phase 3: Connection with Commons

Q9: How does Commons support other Wikimedia projects, like Wikipedia?

Q10: For the next question, I'm going to show you a list of the most linked-to files. I'll paste the URL into the chat so you can click through to the page.

<copy and paste this url into the chat: <https://commons.wikimedia.org/wiki/Special:MostLinkedFiles>>

I'm also going to share my screen so that you can see what I see.

Could you briefly describe for me what you see on this page?

How does this list show Commons support for Wikipedia?

Do you think this is a good example of Commons supporting Wikipedia?

Could you think of a better list?

Q11: How do you find content from Commons to support your work?

What makes finding images or media difficult?

How do other Wikipedians find media to support their work?

Q12: How would you describe the goals of Commons in relation to Wikipedia?

Follow up:

What if we flip the question.

How would you describe the goals of Wikipedia in relation to Wiki Commons?

Closing Script

Do you have any questions for me?

Do you know of 1 or 2 Wikipedians who we should contact and potentially interview? <snowball sampling>

Can I say that you were the person who referred us to <user>?

Note if we use your name we are breaking confidentiality.

We are acknowledging the time and effort of our participants by making a contribution to one of three ‘like-minded’ organizations. Would you like us to acknowledge your participation with a contribution to either:

- Wikimedia Foundation
- Creative Commons
- Internet Archive

Thank you for taking the time to talk with me about your work on Wikipedia. Do you have anything else you would like to add?

Received April 2021; revised November 2021; accepted March 2022