

Visual Narratives and Collective Memory across Peer-Produced Accounts of Contested Sociopolitical Events

EMILY PORTER, University of Washington
P. M. KRAFFT*, University of Oxford, UK
BRIAN KEEGAN, University of Colorado

Studying cultural variation in recollections of sociopolitical events is crucial for achieving diverse understandings of such events. To date, most studies in this area have focused on analyzing variation in texts describing events. Here, we analyze variation in image usage across Wikipedia language editions to understand if, like text, visual narratives reflect distinct perspectives in articles about culturally-tethered events. We focus on articles about coup d'états as an example of highly contextual sociopolitical events likely to display such variation. The key challenge to examining variation in images is that there is no existing framework to use as a basis for comparison. To address this challenge, we use an iterative inductive coding process to arrive at a 46-item typology for categorizing the content of images relating to contested sociopolitical events, and a typology of network motifs that characterizes structural patterns of image use. We apply these typologies in a large-scale quantitative analysis that establishes clusters of image themes, two detailed qualitative case studies comparing Wikipedia articles on coup d'états in Soviet Russia and Egypt, and four quantitative analyses clustering image themes by language usage at the article level. These analyses document variation in imagery around particular events and variation in tendencies across cultures. We find substantial cultural variation in both content and network structure. This study presents a novel methodological framework for uncovering culturally divergent perspective of political crises through imagery on Wikipedia.

CCS Concepts: • **Human-centered computing** → **Wikis**; *Collaborative content creation*; Social network analysis; • **Social and professional topics** → *Geographic characteristics*; *Cultural characteristics*;

Additional Key Words and Phrases: Wikipedia, politics, visual content analysis, cultural differences

ACM Reference format:

Emily Porter, P. M. Krafft, and Brian Keegan. 2020. Visual Narratives and Collective Memory across Peer-Produced Accounts of Contested Sociopolitical Events. *ACM Trans. Soc. Comput.* 3, 1, Article 4 (February 2020), 20 pages.

<https://doi.org/10.1145/3373147>

*Author did not participate in data collection or analysis. Writing was conducted at the University of Washington.

This work was supported in part by the Washington Research Foundation, by a Data Science Environments project award from the Gordon and Betty Moore Foundation (Award 2013-10-29) and the Alfred P. Sloan Foundation (Award 3835) to the University of Washington eScience Institute. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

Authors' addresses: E. Porter, University of Washington, Box 352840, Mary Gates Hall 370, Seattle, WA 98195-2840; email: Eporter@uw.edu; P. M. Krafft, Oxford Internet Institute, University of Oxford, 1 St Giles', Oxford OX1 3JS, United Kingdom; email: p.krafft@oii.ox.ac.uk; B. Keegan, University of Colorado, INFO 129 1045, 18th St., Boulder CO, 80309; email: Brian.keegan@colorado.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

2469-7818/2020/02-ART4 \$15.00

<https://doi.org/10.1145/3373147>

1 INTRODUCTION

Culturally shared recollections or “collective memories” play a significant role in shaping the perception and interpretation of current events [17, 71], and images are powerful tools in this process [1]. Images document important events, communicate complex messages, and symbolize shared values [3, 63, 68, 69]. Culture is an identity structure defined by varying beliefs, social institutions, values systems, and communication styles [6] that influence the way we perceive and assign meaning to imagery [63, 68, 69]. Imagery, in turn, also impacts culture by shaping political outcomes [49], consumerism [56], and social movements [70]. Collective memory is also deeply embedded in shared cultural experience and identity [12, 23, 47], forming a collaborative depiction of history through the lens of group identity typically rooted in shared social experiences [16, 19, 34, 46, 58, 71].

Because images can influence public opinion, traditional visual media have historically been closely controlled by governments, businesses, and other gatekeepers to support specific messages [65, 75]. However, documenting historical events is no longer a privileged role for elite actors; increasingly ubiquitous mobile phone cameras enable amateurs to document historical moments and then share the images over social media platforms [1, 11]. These images can support dominant viewpoints or can surface overlooked perspectives, challenging popular narratives [3, 63, 68, 69].

Images of political crises elicit strong and divergent reactions about how well they represent the grievances, actors, and consequences of major turning points in a country’s political history [1, 11, 14, 51]. Culture is inextricably linked to the formation of political groups and ideologies [6], as well as their historical role through a collective memory lens [23, 47]. Peer-produced encyclopedias like Wikipedia provide a unique environment to examine differences in rhetoric across controversial historic and culturally ingrained political events like coups d’état. Wikipedia’s encyclopedic style requires the consolidation and distillation of these events into a single, authoritative narrative [33, 62]. However, the English Wikipedia is not the only Wikipedia: more than 200 language editions exist, each with their own policies and substantive differences in content [7]. Wikipedia’s various language editions exhibit a regular pattern of “self-focus bias” by unintentionally prioritizing knowledge encoded with locally relevant information rather than distant information [27–29]. This self-focus bias is unaffected by efforts of multilingual editors who work to standardize content across language editions [24] and produces discrepancies in topical coverage [2]. To date, these multilingual comparisons have privileged textual content and hyperlinks as the primary constructs for measuring content biases [19, 26]. Much less is known about the variance in the use of imagery across Wikipedia language editions, despite the importance of visual rhetoric as a form of cultural knowledge production and collective memory [26]. We hypothesize groups experiencing a specific tumultuous socio-political event in the same way or who share the same cultural identity and, therefore, exposure to collective memory, to perceive the event and encode its visual narrative differently than those who do not take part in that cultural experience. [12, 17, 23, 69, 75]. By examining the diversity in image usage, rather than textual content, on articles about the same events, we may shed light on alternative and potentially biased processes of cultural knowledge production on Wikipedia.

This article compares peer-produced accounts of political crises across language editions on Wikipedia to examine cultural variation in imagery. We analyze 761 images across 1,164 articles about coups d’état in 97 language editions. Our contributions are:

- (1) Inductively developing a “visual vocabulary” in the form of an image typology that we can use to facilitate image content analysis and comparison.
- (2) Identifying three categories of structural patterns within the network that characterize some of the typical ways the same or different images are used across languages.

- (3) Empirical evidence suggesting that imagery of politically contested events used on Wikipedia exhibits cultural collective memory bias.
- (4) Development of a novel method for detecting divergent perspectives of culturally tethered events through imagery.

A mixed methodological analysis is used to cluster image usage and features across languages. Two case studies about the 1991 Soviet coup d'état attempt and 2013 Egyptian coup d'état highlight differences in image adoption among “insider” language versions where these events have strong cultural significance and “outsider” languages that do not share these cultural associations. In spite of the availability of common visual records, variation in imagery foregrounds opportunities to use peer production systems to surface differences in the formation of collective memory, community maintained visual media collections, and enrich current understandings of cultural salience in visual media.

2 BACKGROUND

We situate our study within the literature on visual imagery, collective memory, and cultural variation on Wikipedia. Wikipedia is a peer-produced online encyclopedia with active editions in 284 languages, each of which has its own administrators and regulations [33, 62]. Wikipedia promises to be a global collaborative platform for generating encyclopedic knowledge grounded in norms and policies of notability, neutrality, and verifiability [7, 33, 35]. Although the interface encourages cross-language translation efforts, Wikipedia’s volunteer editors still have substantial latitude to edit the breadth and depth of topics independently of other languages’ coverage. This results in low levels of standardization in how ostensibly identical topics are written and a strong “self-focus bias” toward covering topics relevant to the language’s parent culture [2, 24, 28, 53, 54]. No longer just a bug, this variability in topical coverage across Wikipedia’s language editions has become a powerful feature for comparing the cultural contexts of online knowledge production [8, 25, 28, 33, 50, 54].

In addition to the variability of its topical coverage across languages, Wikipedia unexpectedly became a platform for collaboratively documenting current events [38–40, 53]. Wikipedia editors’ self-organized responses to events like the 9/11 terrorist attacks, 2004 Indian Ocean earthquake, and 2007 Virginia Tech shooting deeply influenced its rules and practices to balance between its commitments to notability and neutrality, gratifying editors’ humanitarian impulses, and general users’ information-seeking needs [35, 40]. Collective memories also profoundly shape the production and consumption of Wikipedia articles about current as well as historical events; readers re-visit these articles on anniversaries of the events and editors adjust content to prioritize some representations of the event over alternative narratives [19, 47, 58].

The role of imagery in the production and consumption of Wikipedia articles has not been deeply analyzed or theorized despite the significant efforts the community invests in collecting and integrating photos, visualizations, and other imagery into its articles [26, 72]. Projects such as the Wikimedia Commons play a critical role in archiving and documenting imagery for use across Wikipedia language editions. Started in 2004, Wikimedia Commons provides a multilingual repository for public domain and freely-licensed images, video, and audio clips for all Wikimedia projects to be “used by anyone, anywhere, for any purpose [9].” However, not all files are acceptable for Wikimedia Commons; files not in the public domain or lacking open content licences are not permitted. As of 2019, Wikimedia Commons houses over 55 million media files that can be used in any Wikimedia project across all languages [9]. The unique archiving process of Wikimedia Commons allows files to be embedded into Wikipedia articles without having to upload each image separately to a particular language page [10, 72]. This method of file sharing allows Wikimedia

Commons to track how the file is being used across languages and is provided under the “Global Usage” information for each file [9, 10].

2.1 Visual Rhetoric and Online Imagery

Images exhibit tremendous power in their ability to transcend oral and written language, delivering several different levels of information simultaneously [61, 65, 67]. Photography has historically provided the opportunity for individuals to bear witness to an event with one’s own eyes even as it takes place on the other side of the globe [30, 67]. Technological developments have also enhanced the capacity to digitize and organize image archives enabling new forms of sharing, curation, and re-use on platforms like Flickr or Instagram [21, 31, 52, 66].

A photograph does not only document, it also encodes narratives, opinions, behaviors, and perceptions. The meaning of these narratives is not static but rather dynamic in the sense that both the framer of the image and the viewer of the image play a part in how it is interpreted [14, 55, 65, 67]. A single image can narrate divergent stories when employed in different circumstances; an image of a crowd waving flags could depict a protest or a celebration. The capability of images to serve as rhetoric is not limited to a textual framing of an image but extends to the pre-existing cultural and historical perceptions of the framer and viewer as well [30, 57, 61, 67]. For example, American photojournalism exhibited bias in the visual narration of the war in Kosovo; photos in Western media outlets depicted Serbs in military uniforms causing destruction or suffering and far fewer images of Serbian civilians than Albanian civilians. These differences helped to legitimize and sustain support for the official American policy at the time [51]. Today, amateurs are using commoditized cameras in mobile phones as well as digital platforms lacking strong gatekeepers to share imagery [1, 11].

Previous work has examined imagery during crises, but this has been traditionally limited to a single event or language [4, 51]. Viégas [72] surveyed image contributors on the English Wikipedia to understand their motivations and practices, and found image-related collaboration is largely isolated from the text-focused collaboration. This isolation between image and text-editing groups enables a stronger sense of community and promotes independent negotiations about topical relevance of new information within each group [72]. An analysis of the 25 largest and most-active Wikipedia language editions finds significant cross-language image diversity exceeding that found in text [26], but engaged in little in-depth analysis of this variation. Given the evocative nature of imagery for contextualizing and narrating events, it is important to understand the usage of imagery as semiotic or rhetorical links connecting languages and articles.

2.2 Collective Memory in Peer-produced Accounts of Crisis

Wikipedia’s unique affordances as the “encyclopedia that anyone can edit” makes it a compelling setting to study how contentious events are documented as they unfold. This capacity is vividly illustrated in collaborations following natural disasters where articles are created within the hour [38], unique coordination practices support high-tempo editing of a shared artifact [17, 39], organizational structures are rapidly regenerated from previous crises [40], specific expertise is adapted to fill to emergent social roles [36], handling regular but unpredictable events like the deaths of celebrities is routinized [37], and single events can be integrated into larger narratives that require supporting the migration of editors and readers between articles [71].

Sociopolitical contextual factors play an enormously important role in understanding disaster response, but even more so, coups d’état and revolutions are emblematic in the degree to which they are steeped in the context of social, historical, and other human-mediated processes, each of whom have a stake in legitimizing or contesting the outcome of an event [6, 60]. As fundamental turning points in a country’s history, images play an important role in documenting contested

political events to support particular narratives: a coup is remembered as violent, justified, popular, and so on [1, 51, 61]. The ways in which these narratives are framed—including the images they use—contributes to a shared understanding about how individuals’ experiences should be remembered [63, 68, 69]. The collaborative reconstruction of the past or collective memory emerges from social interactions in the present to frame people, places, or events to fit with prevailing understandings and identities [16, 23, 34, 46, 58].

The meaning of an image is influenced by culture [63, 68, 69], and imagery serves as a vehicle for memory [30, 68, 69]. When an event is reconstructed in an environment like Wikipedia, different groups may produce narratives of the event reflecting their memories—including imagery with different thematic content. Because different Wikipedia language editions maintain their own norms and communities, and because translation between language editions is the exception rather than the norm [20, 28, 54], collective memory processes may diverge and generate significantly different narratives about the same events—a possibility we examine empirically in the present work.

2.3 Cultural Biases in Online Peer Production

Wikipedia’s different language editions operate relatively autonomously, and translations of identical parent articles have been rare for most of Wikipedia’s existence to date [20, 28, 54]. This autonomy across language edition communities means that editors in each language can independently collaborate on topics, resulting in different information [20, 28, 50, 54, 74]. This variation across language in the online community reflects the larger cultural traits of the languages as well [8, 25, 42, 59]. Examples of variation in coverage can range from the absence of an article that exists in language i but not language j or investing in developing high-quality “Featured Articles” in one language but not another. The intersection of topics across language editions of Wikipedia is incredibly small and even where overlaps exist; the information each contains tends to be unique [2, 7, 27, 28].

While language and culture are not interchangeable and neither language nor culture is explicitly tied to a particular geographical territory in the world, there is still substantial evidence that supports their overlap. Research on “cultural contextualization” suggests that differences between Wikipedia language editions are brought about by the cultural identities of the editors [5, 24, 29, 50]. These editors are biased in contributions relative to nearby locations or toward territories inhabited by those who share their cultural identity [44]. Contributions reflect their experiences, culture, geography, and history. Researchers argue this variation in topical coverage and content is reflective of inherent cultural bias within Wikipedia language communities [8, 50, 54]. Even with the Wikipedia movement’s commitment to a “Neutral Point Of View” in article coverage, it has been unable to address unintentional information bias resulting from a diverse range of cultural and linguistic backgrounds [8, 54].

3 DATA

Wikipedia organizes its articles by categories, which we leverage for data collection. We use seed articles in the English Wikipedia categories associated with each decade from “1950s coups d’état and coup attempts” through “2010s coups d’état and coup attempts”. We retrieved all articles in these categories and filtered for only those pages containing words such as “coup”, “putsch”, “incident”, “crisis”, or “revolution” to generate a list of 169 English Wikipedia articles about coups over the time range. Since all researchers were native English speakers, this strategy prioritizes topical relevance of the sample over sample size at the cost of excluding events that do not have an article in the English Wikipedia. We used custom Python scripts to retrieve data about inter-language

Table 1. Top 10 Events in the Corpus, Sorted by Coverage Across Language Versions (L)

Topic	L	R	E	S	Y
2016 Turkish coup d'état attempt	60	25	7	14,397	2016
Carnation Revolution	43	41	24	7,376	2007
1991 Soviet coup d'état attempt	42	51	30	14,353	2009
Romanian Revolution	38	67	32	12,880	2007
1973 Chilean coup d'état	31	37	20	9,632	2011
2013 Egyptian coup d'état	31	32	13	13,820	2013
1993 Russian constitutional crisis	28	35	22	11,049	2010
Tulip Revolution	27	34	21	4,775	2007
1953 Iranian coup d'état	27	59	26	8,404	2009
Saur Revolution	25	21	14	7,131	2011

We report the following statistics using the median values for each article across all language editions: number of unique revisions (R), editors (E), largest page size in bytes (S), and article creation year (Y).

links,¹ page content,² image usage,³ external links,⁴ and revision histories⁵ from the MediaWiki API. We also used custom Python scripts to download the images themselves for content analysis.⁶ The data we analyze consisted of the state of each of these articles as of late March 2017. From the seed set of 169 English Wikipedia articles, we identified 1,164 articles in 97 language editions via inter-language links—we start with each seed article and visit the other languages in which the seed article is represented. Descriptive statistics about the most widely-covered coups are reported in Table 1.

We then downloaded the images used across all the 1,164 articles we had retrieved. Not all articles have images; of the original 169 seed coup topics, 71 (42%) have no images in any language. Similarly, out of the 1,164 coup articles across languages, 537 articles (46%) contained no images. The absence of images in any given language version's article about a coup does not mean that no images are available. Of the 169 seed coup topics, 79 have at least one language version missing an image when a differing language version has at least one image, but the number of images available in these other languages is still relatively small (1.71 images on average). In total, there were 761 images used in the articles from our cross-language dataset. The 761 coup-related images were adopted on different language versions to different extents. The most widely-used images in the corpus appeared in 25 different languages about coup events.

4 TYPOLOGIES

Image data is challenging to analyze compared to text data because, unlike text data, images do not have a set existing vocabulary since fewer tools exist for automated semiotic analysis of imagery. Our strategy to analyze cultural variation in image use on Wikipedia is therefore first to develop two typologies that allow us to transform our image data into representations that are more amenable to analytical processing. We develop one typology for analyzing image content and one typology for analyzing the network structure of image usage across articles.

¹API:Langlinks.

²API:Parsing wikitext.

³API:Imageusage.

⁴API:Extlinks.

⁵API:Revisions.

⁶A Jupyter Notebook with the code for replicating this data collection is available on [GitHub](#).

4.1 Content Typology

To develop our content typology, we draw upon several existing typologies for other types of sociopolitical contexts. We constructed initial candidate codes using Nikolaev's methods for analyzing bias in war photography as well as Zahedi and Bansal's cultural signifiers [51, 75]. We selected a random subset of 50 images from our full dataset of images to iteratively code for emergent themes. The coding identified themes including violence, political leaders, locations, and demographic differences. We then mapped themes arising from the 50 image subset back to the frameworks from existing literature, and added new codes where no appropriate parallels existed. We further refined our framework by drawing on additional literature in the fields of cultural and political visual communication in order to capture the most salient generalizable features of the images [32, 63, 67]. Our final typology consisted of seven categories of codes and 46 sub-codes. One author and one additional coder coded the remaining images outside the 50-image subset with the final typology. The inter-rater reliability on the full set of images was calculated using Cohen's Kappa and the scores for each sub-code are reported in Table 2. The combined Kappa for all codes was 0.839, indicating significant agreement [41].

The 46-item framework in Table 2 was applied to all 761 images in our sample. To better understand the structure of coup-related imagery, we performed a dimensionality reduction task using the t-Distributed Stochastic Neighbor Embedding (t-SNE) algorithm [48]. By reducing the dimensionality of the image-by-code matrix from 46 dimensions (all undiscarded codes from the framework) to 2 dimensions, latent clusters of images with similarly coded features can be identified. K-Means was used to identify different candidate clusters in the two-dimensional embedding. The $K = 12$ clusters parameter was determined by computing silhouette scores [64] for different candidates K ; 12 had a maximal score across multiple t-SNE and K-Means runs. The t-SNE algorithm is stochastic and sensitive to different hyperparameter specifications [73]; these parameters were iteratively tuned, estimated, and interpreted. Figure 1 visualizes a representative embedding for the 761 images in two dimensions and colored by their K-Means cluster membership.

The x and y dimensions as well as the distances between clusters should not be over-interpreted [73], but nevertheless capture relative similarities of image content as coded by the crisis visuals framework in Table 2. Representative images from each of the six clusters were selected based on their proximity to the centroids of each of the 12 clusters and are visualized in Figure 2. Cluster 3 is cleanly separated from the rest of the clusters and contains 42 pictures of various types of maps. Clusters 0, 6, and 8 containing 98, 73, and 80 images respectively sit off to the left. These clusters clearly depict leaders either as portraits without a clear expression (cluster 0), engaged in diplomatic activity (cluster 6), or smiling (cluster 8). Clusters 10 and 7 with 61 and 18 images are largely symbolic in nature, emphasizing symbols like flags, patches, and commemorative artifacts like stamps. Focusing on different types of imagery of the events themselves, clusters 1 and 9 show death and destruction as well as active military engagement. Clusters 5 and 2 provide different perspectives of demonstrations, with cluster 2 (84 images) featuring images of the crowds themselves and cluster 5 (91 images) zooming in on individuals within the movements. Cluster 4 contained 39 images of buildings while cluster 11 (47 images) highlights the appearance of written documents in the articles. These results demonstrate the existence of stable genres of images used to document coups and political crises across languages and events.

4.2 Network Typology

In addition to our content typology, we also developed a typology to categorize patterns in the structure of the relationships between images, the articles they appear in, and the languages having articles about the event. To develop this typology, we first computed the bipartite relationships between Wikipedia articles and the images they use. We then visualized these

Table 2. Framework for Coding Crisis Imagery by Sub-Code, Count of Usage Across Images (Multiple Codes Can be Applied Per Image), Inter-Rater Reliability (Cohen's Kappa), and Description of Sub-Codes for $N = 761$ Images from Coup and Crisis Articles

Code	Sub-codes	Count	IRR	Descriptions
<i>Activity</i>	Active	206	0.719	Protesting or fighting
	Passive	549	0.679	Portraits, meetings, or landscapes
<i>Demographics</i>	Adult male	447	0.897	Primary subject(s) of image is an adult man or men
	Adult female	81	0.846	Primary subject(s) of image is an adult woman or women
	Child male	7	0.492	Primary subject(s) of image is a young boy or boys
	Child female	15	0.697	Primary subject(s) of image is a young girl or girls
	Cannot discern*	68	0.777	There is no primary subject in the image
	Non-human	224	0.896	There are no humans in the image
<i>Facial expression</i>	Frown	33	0.488	Primary subject(s) of image are frowning or distressed
	Smile	114	0.842	Primary subject(s) of image are smiling or pleased
	Unclear*	299	0.667	Primary subject(s) of image have ambiguous expressions
<i>Foreign influence</i>	Foreign leaders	31	0.905	Images of leaders from different countries meeting
	Foreign flag	15	0.931	Images of symbols from different countries
<i>General topics</i>	Alleged atrocities	14	0.819	Images of dead bodies
	Buildings	84	0.828	Images of government, religious, etc. structures
	Celebrations	14	0.853	Images of celebration
	Demonstrations	125	0.865	Images of protesters or crowds
	Destruction	22	0.741	Images showing explosions, fire, or ruins
	Diplomacy	98	0.820	Images of negotiations, press conferences, etc.
	Documents	21	0.592	Images of announcements, laws, etc.
	Humanitarian aid	6	0.798	Images showing distribution of food, etc.
	Infographics/maps	45	0.973	Images of locations, statistics, etc.
	Memorials	44	0.806	Images of statues, plaques, etc.
	Military actions	101	0.838	Images of uniformed soldiers
	Military technology	114	0.905	Images of tanks, guns, etc.
	Police/Prisoners	5	0.798	Images of uniformed police, people being detained
	Portraits	286	0.859	Formal images of political leaders
	Signs/symbols	135	0.817	Images of graffiti, protest signs, flags, etc.
Suffering	12	0.913	Images of injuries, distress, etc.	
Other*	55	0.748	Images of recreation, relaxation, etc.	
<i>Image focus</i>	Individual	250	0.956	Clear facial features with few/no other people
	Group	212	0.837	Not focused on any individual and many people
	Individual in group	37	0.691	Individual facial features as part of a larger group
	Unclear*	16	0.434	Focus is unclear
<i>Type of people</i>	Civilians	118	0.858	Images of people in civilian clothing
	Journalists	14	0.943	Images of people with professional A/V equipment
	Political leaders	337	0.892	Images of portraits, official actions, etc.
	Military personnel	86	0.880	Images of people in fatigues, with weapons, etc.
	Police	8	0.727	Images of people in uniform, with badges, etc.
	Unclear*	222	0.562	Images of people with unclear roles
<i>Violence</i>	Violent	48	0.795	Depicting acts of physical or intended aggression
	Non-violent	725	0.778	Not depicting acts of physical or intended aggression

*denote miscellaneous categories subsequently discarded from the analysis.

networks, incorporating language as color in these plots. We examined these networks derived from all our images and identified common network structural patterns, highlighted in Figure 3. These structures are not exhaustive of all possible configurations, but they illustrate emergent patterns of multilingual image usage that will be discussed in subsequent sections.

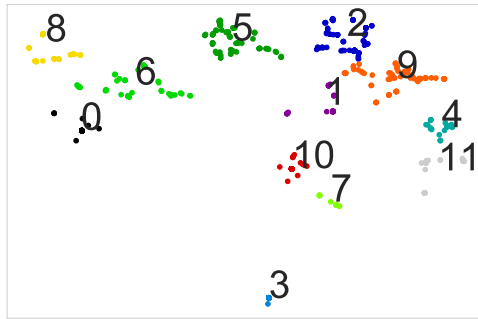


Fig. 1. t-SNE of images and their codes.



Fig. 2. Examples of images from each cluster in Figure 1.




Fig. 3. Examples of image usage patterns. Images are black nodes and the articles they appear in (as well as the edges connecting them) are colored by their language edition. The “bloom” on the left is an example of multiple articles across languages using the same image. The “cluster” in the middle is an example of a subset of languages invoking the same images across multiple articles. The “monoculture” on the right is an example of a single language (in brown) using multiple images that appear in no other languages.

4.2.1 Blooms. The first pattern we identify is the “bloom” characterized by a single central image used across multiple articles in multiple languages. The bloom is illustrated on the left in Figure 3 and has the image “Day after Saur revolution in Kabul.jpg” at the center; it is connected to articles like “War in Afghanistan (1978–present)” (English), “Surrevolution” (German), and other articles about the 1978 coup precipitating the Soviet invasion of Afghanistan. The image at the center was taken by an amateur at the street level and depicts civilian cars entering a courtyard outside the presidential palace in Kabul guarded by soldiers and military vehicles. Among images that appear on coup pages, this one distinctively captures an event to the exclusion of any other images.

4.2.2 Clusters. The second pattern is the “cluster” characterized by several images being used across multiple articles in the same subset of languages. The cluster is illustrated in the middle in Figure 3 and has the images “Guinea Bissaus interim President Raimonda Pereira (cropped).jpg”, “Carlos Gomes Junior.jpg”, and “Guinea bissau sm03.png”. All three images are used by the English, Catalan, and Ukrainian language versions of the “2012 Guinea-Bissau coup d’état” article (at center, connected to all three images) and the images are surrounded by articles about lists of political leaders, elections, and United Nations resolutions. The use of the same set images among these languages—to the exclusion of other languages and images—suggests a high level of standardization, potentially resulting from multi-lingual editors or article translations.

4.2.3 Monocultures. The third pattern is the “monoculture” characterized by images used exclusively by a single language across multiple articles. The monoculture is illustrated on the right in Figure 3 and has more than two dozen coup-related images used by no other language but Arabic articles about the 1958 and 1963 coups in Iraq. These Arabic Wikipedia articles use images about political leaders that no other coup page uses, despite the historical importance of these coups for the emergence of the Ba’athist party that dominated Iraqi politics for 40 years until the 2003 U.S.-led invasion. This pattern is not wholly isolated from the rest of the coup image usage network; it interfaces with English, Farsi, and Turkish articles at the top through the images “Arif with Qasim.png” (leaders of the 1958 coup) and “Abd_al-Karim_death.jpg” (showing Qasim’s executed body following the 1963 coup). The extensive use of images that are adopted by no other language versions about the same events demonstrates their salience within this language edition as well as blind spots in other languages’ coverage.

5 CASES

Example images from two events, their codes from this framework, the languages in which they appear, and other articles using these images are given in Tables 3 and 4. Despite the surprise and energy associated with coups, 83% of the images in our sample depicted passive activities like portraits or landscapes, rather than acute action such as protests or violence. The gender composition of the photographs where humans were present and gender could be discerned was also heavily skewed toward men; women or children only appeared in 9% of images. Coups and crises can be occasions for violence as well as celebration, but ambiguous facial expressions were the most common expression, although images depicting smiling subjects out-numbered frowning subjects by a factor of 4.5. Coups are often consequences of outside influence and support [60], but depictions of foreign influence with leaders or symbols of multiple countries combined make up less than 8% of the sample. The most common topics found in the images included portraits, military actions and technology, signs and symbols of groups, and demonstrations. There was wide variability in the focus of imagery, but individuals and individuals within groups were the most common codes typically reflecting old and new leadership. The types of people depicted

Table 3. The Photo on the Left Containing Boris Yeltsin on a Tank is Representative of an “Outsider” Image

Topic	1991 Soviet coup d’état attempt	
		
File name	“Boris Yeltsin 19 August 1991-1.jpg”	“1991 CPA 6369.jpg”
Codes	VI-Non-violent Image; IT-Passive Image; IF-Group; FE-Unclear; D-Adult Male; TOP-Military or Political Leaders; TOP-Civilians; TOP-Journalists; GT-Military Technology; GT-Diplomacy; GT-Buildings	VI-Non-violent Image; IT-Passive Image; D-Adult Male; IF-Individual; FE-Smile; D-Adult Male; TOP-Military or Political leaders; GT-Portraits; GT-Signs and Symbols
Languages	Arabic; Bavarian; Bulgarian; Catalan; Czech; German; English; Esperanto; Spanish; Estonian; Basque; Persian; Finnish; French; Galician; Hebrew; Croatian	Bulgarian; Czech; Chuvash; Norwegian; Polish; Romanian; Russian; Ukrainian
Other articles	1991; Boris Yeltsin; History of the Soviet Union; Moscow; Stalin (1992 film); Russia; White House (Moscow); Alexander Korzhakov; Civil resistance; Nikolai Vorontsov (et, ru); Mikhail Gorbachev; Communist Party of Great Britain; Cold War	None

in these images unsurprisingly tended toward political leaders and military personnel, with less representation from civilians, journalists, or police.

5.1 Case Study 1: 1991 Soviet Coup d’état Attempt

The 1991 Soviet coup d’état attempt ranks third in terms of language coverage for articles within the data, appearing in 42 languages and housing 32 unique images across these languages. The Russian version of this article used the most images (11) of any language edition to discuss the event, followed by Norwegian (10), Polish and Arabic (7), Spanish (6), and Bulgarian, Korean, Latvian, Portuguese, Romanian, Turkish, and Chuvash (5). Less than half of all language versions of the article contain five or more images. With the exception of Korean and Arabic, article versions with the most imagery are in European or Central Asian languages with historical ties to Russia and the Soviet Union.

Although many languages did not use large quantities of imagery within articles, most used at least one image. Among the languages that used imagery (31), most displayed a “lead image” in the introduction within a box, which also contained very basic information regarding the event. The position of imagery conveys its importance in relation to the rest of the narrative to both introduce the topic and prioritize it over alternatives [15]. Throughout the “1991 Soviet Coup d’état attempt” pages, two photos appeared the most frequently as the lead images across most languages depicting similar themes of military tanks next to government buildings. However, the Russian article used neither of these images but instead a photo of a coin commemorating the victory of democracy in Russia. This image of a coin was used only in the Russian version of the article even though it is housed on Wikimedia Commons and available for other languages.

Table 4. The Left Shows a Popular “Outsider” Image, a Portrait of General Al Sisi, while the Right is an Example of What Might Be Considered an “Insider” Image for this Particular Event

<i>Topic</i>		2013 Egyptian coup d'état
		
<i>File name</i>	“General Al Sisi.jpg”	“R4bia sign used in solidarity with victims of Rabaa crackdown 23-Aug-2013.jpg”
<i>Codes</i>	IT-Passive Image;D-Adult Male; FE-Smile; GT-Portraits; GT-Signs and Symbols; VI-Non-violent Image; TOP-Military or Political Leaders; IF-Individual	IT-Active Image; D-Adult Male; D-Adult Female; GT-Demonstration; GT-Signs and Symbols; IF-Group; TOP-Civilians; VI-Non-violent Image
<i>Languages</i>	Arabic; Belarusian; German; English; Farsi; French; Hebrew; Armenian; Italian; Dutch; Norwegian; Polish; Russian; Uzbek	Arabic; German; English; Spanish; Farsi; French; Hebrew; Indonesian; Romanian
<i>Other articles</i>	None	Post-coup unrest in Egypt (2013–2014); Rabia sign; 2017 Mohammad Bagher Ghalibaf presidential campaign; Qatar–Saudi Arabia diplomatic conflict

Examining every image on each language page of the “1991 Soviet Coup d'état attempt” topic provides a much more vivid story than the lead image alone. For example, there was a stark contrast in image narratives for languages with historical ties to the Soviet Union such as Russian, Polish, and Bulgarian. These language editions’ articles emphasized commemorative imagery and used the same images of stamps, monuments, and celebratory events. Latvian, a language for a culture with deep historical ties with the Soviet Union, used similar commemorative imagery but also included official documentation like proclamations that did not appear in other languages.

Departing from the slavophone sphere, there was a general shift in imagery toward discrete events and elite actors. Articles in Norwegian, Spanish, and English depicted a mix of imagery focused on Soviet military activity and Boris Yeltsin’s leadership. Images of Boris Yeltsin and his supporters on a tank (Table 3 on left) were found in 17 of the 42 languages about the event. Almost all outsider languages with more than one image on their respective article pages used multiple photographs of military weaponry (tanks in particular) in front of Soviet government buildings like the Kremlin. These themes were minimal or non-existent on insider pages.

The higher usage rates of imagery, differences in lead image usage, and common imagery themes for language editions closely related to the country where the coup took place provides additional support for Hecht and Gergle’s self-focus bias [27, 28]; editors emphasize details in their peer produced account of a historical event that are most relevant to their own cultural experiences and compatible with their collective memory. The contrasting imagery patterns between slavophone “insiders” versus “outsider” languages explains how each experienced and, thus, reproduced memories of the failed coup differently—heroes’ portraits versus men atop tanks.

5.2 Case Study 2: 2013 Egyptian Coup d'état

The 2013 Egyptian coup d'état article is currently the sixth most widely-covered coup event in the dataset; it appears in 31 languages with 35 unique images. German Wikipedia used the most images (14) of any language edition to discuss the event, followed by Arabic (10), English (7), Farsi (6), and Spanish (5). Of these languages, only two, Arabic and Farsi, have historical ties to the region; other related languages like Hebrew, Turkish, or Amharic do not appear. The 2013 Egyptian coup does have the same correspondence between the quantity of image usage and cultural proximity as was the case with the higher usage of Soviet coup imagery among slavophone languages. The lack of a self-focus bias in image intensity may be due to a number of factors, including the area affected by the event was felt in a much smaller region. Although Egyptian Arabic Wikipedia language edition only used four images in total for the article compared to German's 14 images, this is still greater than the average (1.7) number of images across all pages [27]. While the 2013 Egyptian coup d'état may have fewer language pages than its Soviet counterpart, its imagery proved to be much more diverse. Visual rhetoric surrounding the Egyptian event was generally thematically distinct from pages covering the 1991 Soviet coup d'état, with the exception of English featuring political leadership in both accounts. Interestingly, the two photos used most frequently as the lead image across languages were thematically very different from the Soviet case above. A smiling portrait of ousted president Mohamed Morsi and anti-Morsi protest march appeared most commonly for the 2013 Egyptian coup d'état. The three languages most closely associated with the geographical region did not use either photo as a lead image; the Egyptian Arabic language edition used a photo of U.S. Secretary of Defense Chuck Hagel meeting with President Morsi from Commons while the Arabic language edition used an image clipped from a YouTube video of Al-Sisi announcing Morsi's removal. Farsi used no imagery in this position.

Unlike the majority of images used on article pages, the Arabic Al-Sisi image was locally uploaded and not hosted on Wikimedia Commons. As discussed earlier, this means the image and its usage are visible only on the Arabic Wiki. Numerous photos of Al-Sisi are available for global use through Wikimedia Commons but were not chosen as the lead image in Arabic. These photos highlight an interesting dichotomy in lead imagery; the images themselves appear very different (one a still frame taken from a YouTube video and the other a photo published by the U.S. Secretary of Defense), yet, may actually be interpreted as thematically similar. For example, both images are calling attention to political leadership instead of the protests or military actions. The context of the full visual narrative provides a richer insight into thematic similarities and differences across languages than a standalone image.

Although there are fewer languages with close regional ties to the 2013 Egyptian coup d'état, overarching themes between these languages appeared very similar. Arabic, Farsi, and Egyptian Arabic highlight the overall event, using imagery of protests and destruction while few, if any, images of political figureheads appear throughout the narratives. Languages more prevalent outside of the region reflected a greater variety of content within narratives. German, English, and other western languages emphasized the role of foreign as well as domestic leadership throughout their respective articles but fewer images of protests during the event. Many Western languages chose to crop head-shots of important political figures from larger photographs containing additional people, situations, or general themes that did not appear in the rest of the narrative (Table 4 on left). While German did not include violent imagery, the English article included graphic scenes of destruction and death. Few East Asian languages covered the topic and included imagery. Japanese (three images) and Vietnamese (four images) appear to follow a narrative similar to Western languages, using images of both political leadership and civilian protests; however, neither language used any violent imagery.




Fig. 4. t-SNE of the language editions of “1991 Soviet Coup d’état attempt” clustered by thematic similarity of imagery used.

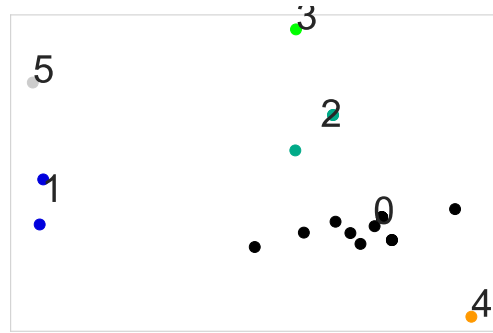


Fig. 5. t-SNE of the language editions of “2013 Egyptian coup d’état” clustered by thematic similarity of imagery used.

5.3 Thematic Variation at the Article Level

Given the appearance of “insider” and “outsider” languages in the case studies, we use a micro-level quantitative analysis to provide empirical backing for these findings. This section examines language editions’ image usage on the same article to quantify variance in topical coverage. Following the same method used for Figure 1, we generate t-SNE plots of languages’ imagery for the two case studies discussed in Sections 5.1 and 5.2. Four language clusters were identified for the “1991 Soviet coup d’état attempt” in Figure 4. Of the 44 languages with an article about this event, 35 occur in cluster 0. Their co-occurrence in this cluster reflects “blooms” from Section 4.2 where identical images or images with similar themes are used. Large and active language editions like Japanese and Spanish are in cluster 2. Turkish, Norwegian, Vietnamese, French, Finnish, and Portuguese occur in cluster 3. The single member of cluster 1 is Russian, a major outlier in the kinds of imagery used to discuss this event compared to the other language editions. The Russian language outlier is an example of a “monoculture” using unique types of imagery to document Russian history distinctively from other languages’ historical accounts, matching the qualitative findings from the case study in Section 5.1.

Six language clusters were identified for the “2013 Egyptian coup d’état” in Figure 5. Of the 32 languages with an article about this event, 24 language editions are in cluster 0. These language editions cluster together because they employ similar imagery like the “blooms” and “clusters” discussed in Section 4.2. Large language editions like Hebrew, Japanese, and French are in cluster 2 and English is in cluster 3. The outsider languages show relatively high levels of similarity compared to smaller and more local language editions. Two outliers are Tajik in cluster 4 and Korean (no images) in cluster 5. Cluster 1 contains Arabic and Farsi, two languages with strong cultural ties and political interest ties to Egypt. Again, these results reflect distinctive structural patterns discussed in Section 4.2. Arabic and Farsi are an example of an insider group displaying “cluster” behavior, selecting imagery similar to one another, and isolating themselves from alternative narratives. “Outsider” languages’ are characterized by “blooms” with many using the same themes to talk about the same event. This analysis method was then applied to imagery from two additional articles—the “Carnation Revolution” and the “Saur Revolution.” Findings mirrored the same patterns found in the case studies. Visual narratives of languages deeply connected to the event historically and geographically appeared together in a single cluster while outsiders were characterized by the bloom pattern, using imagery at odds with the insider clusters.

6 DISCUSSION

Wikipedia's role as a multilingual mediator of contested sociopolitical events like coups and political crises makes it an ideal setting for examining the impact of culture and, in turn, collective memory on the visual rhetoric of these important events. Culture impacts the way we perceive imagery and assign meaning [63, 68, 69]. Coups are typically domestically situated political events embedded in cultural identity [6, 60]. These events seek to displace the current executive power from within the institution of the state [60]. With the understanding that culture and collective memory are tightly intertwined [17, 18], we might expect those that experienced a catastrophic socio-political event such as a coup d'état to assign meaning to imagery about that event differently than those who do not take part in that cultural experience [63, 68, 69]. Our findings suggest that, like textual segments of articles [28, 50, 54, 74], languages do not paint a thematically consistent visual narrative across events—language editions used different imagery and image themes to document the same events.

Using previous multilingual research on self-focus bias in Wikipedia as a reference point, it would stand to reason that for every coup event, there would be many different narratives with each language edition providing its own visual history of the event [7, 27, 28, 53]. However, to the other extreme, a "global consensus" mechanism [28] would predict the prevalence of many "blooms" characterized by languages sharing similar imagery. We instead find something in between. Languages tend to cluster themselves into camps. "Insider" languages, those more closely tied to the history, cultures, and states where these events occurred tended to display visual narratives closely aligned with one another. These narratives appeared at odds with, and distinct from, "outsider" languages, those which had few or weak historical, cultural, or geographical ties to an event. "Outsider" languages also exhibited intriguing thematic clustering patterns among themselves, though less distinct than differences illustrated between "insider" and "outsider" languages.

The existence of multiple "monocultures" highlights how an image's isolation, despite the availability of these images to the global community through Wikimedia Commons, demonstrates a high level of cultural signification. This is reflective of Peesapati et al.'s finding that cultural familiarity affects image tagging practices [57]. Coups remain sensitive events for individuals culturally tied to those experiences [17], and this sensitivity appears to be reproduced through image choice in documenting the events creating discernible structural differences in usage patterns. As with text, the lack of a "global consensus" about the images to be included across language versions about historical events like coups encodes important information about cultural salience and collective memory in ways that are collaboratively encoded into formal knowledge about the events themselves. The exclusion of widely distributed "bloom" imagery from the lead positions across "insider" languages contradicts the adoption behavior of "outsiders" in both case studies. This suggests an important but overlooked mechanism in how languages introduce and characterize the essence of an article.

Our findings raise questions about how these structural patterns develop and their connection to collective memory. Are these image sharing patterns observable across all types of articles or reserved specifically for event-based narratives with easily identifiable cultural connections? What role does time play in the structural development of these patterns? For example, does being "first past the post" when assigning visual rhetoric to current events influence adoption patterns as other language version are subsequently developed? Our findings suggest that these structural patterns may be a useful indication of divergent narratives and collective memories of an event. These structural patterns could provide an alternative method for identifying semiotic variation and alternative histories in visual narratives. Employing this method could help researchers avoid the time-consuming process of hand coding image content. These structures should be studied further

to determine their usability as a method for detecting cultural perspective shifts on Wikipedia. Future work in this area should include exploring the presence or non-presence of these structural patterns.

Traditional media studies literature often highlights a rift between Western and Eastern societies in how they identify with and advance particular image compositions over others, thus leading to the expectation that the outsider clustering patterns would reflect the difference in perspective [13]. If this divergence of Western and Eastern attitudes were echoed in the data, we would likely see western and eastern languages in separate “blooms.” Unexpectedly, our findings appear to show a cross pollination of sorts with eastern and western “outsider” languages appearing in the same blooms. While it is not the focus of this particular study, it does raise questions about our current understanding of culturally relevant visual media with regard to crises-event based narratives.

6.1 Methodological Limitations

In our data collection, the sampling on language editions from a seed set of English Wikipedia articles likely biased topical coverage. This data collection procedure misses coups classified in other categories or lacking English language versions. Inconsistencies observed among inter-language links between topics could resolve articles to misleading topics that are either too general or too specific [45]. In the absence of reliable labels about the topical relevance of samples seeded from other language editions, we prioritized precision (ensuring all the articles in the sample were unequivocally about coups) over recall (ensuring all possible articles related to coups were analyzed). These sampling biases—as well as other temporal, societal, and user biases discussed in the following section—are endogenous to all the findings we reported. However, the differences observed nevertheless lend themselves to supporting prior empirical findings that Wikipedia editors’ contributions are influenced by their own similarity to that language’s history, culture, and politics.

6.2 Other Limitations and Future Work

The analyses in this study were largely static in nature by examining the current state of the articles at the time of data collection. The rich revision history data available for all these articles points to exciting potential to examine longitudinal dynamics in image introduction and conflict. Examining changes over time would also illuminate stronger causal processes and thicker qualitative accounts about the effects of introducing images from other languages on the collaboration as well as the diffusion of images over various co-authorship, co-citation, or co-linking relationships. As a growing body of knowledge on translation practices across Wikipedia suggests that translation of pages is still quite rare, we do not examine this possibility in the framing of this research [20, 28, 54]. However, these complex relationships suggest the need for more advanced network representations such as multi-graphs or hyper-graphs to capture the rich, multi-dimensional relationships of images, themes, languages, and articles changing over time.

Despite the multi-lingual analysis, there are endogenous systemic biases present in this analysis. There are almost certainly strong temporal biases toward more images in recent events since advancements in technology adoption may contribute to a more extensive selection of bottom-up citizen-witnessing photography while older photos may not be as widely available [1]. Second-order cultural biases over images’ value (“men on tanks” versus “women in hospitals”) likewise influence the images recorded and posted to Commons before the biases image-language-event fit we examine. The patterns of image use for coups may not generalize to other crises and conflicts where the defeated are not just disenfranchised, but extirpated [12].

We also note that many of these coups d’état take place in locations that have limited access to technology, education, or uncensored information that will necessarily limit the ability for the

local volunteers to edit these articles. Much of the editing in the Chinese Wikipedia, for example, happens from volunteers in Taiwan, Hong Kong, Singapore, international students, and expatriates living abroad [43, 76]. Additionally, the data used in these analyses cannot disambiguate editors' geographic locations or demographic characteristics; it is not always the case that a user writing in the Egyptian Arabic Wikipedia is Egyptian, speaks Egyptian Arabic, or is located in Egypt. Methods for examining the ideologies and dynamics of Wikipedians editing political topics [22] might be extended to understanding their cultural affinities and biases.

7 CONCLUSION

Wikipedia plays an important role in disseminating and documenting collective memory about historical events. However, our understanding of these collective memory processes has largely focused on textual features. The findings from this research extend this cross-cultural knowledge production literature to visual media. Comparing the visual narratives of influential political and historical moments in various languages allows a glimpse into cultural perceptions and representations of memories. Studying the visual rhetoric across languages on Wikipedia provides a unique opportunity for advancing insight into cultural differences, as culturally divergent framing appears naturally woven into its collective memory network. Evidence of cultural variation exhibited within the visual narratives of political crisis articles across Wikipedia presents a new lens through which we can examine the documentation of collective memory.

ACKNOWLEDGMENTS

We are indebted to Samuel Laurens for his assistance in applying our qualitative coding framework. Many thanks to Elissa Porter, Wiley Porter, and Deborah Keegan for copyediting and organizational suggestions. Thanks also to Leo Stewart and Alannah Oleson for their critique of this work and suggestions throughout the revision process. The completion of this study would not have been possible without all the years of support and encouragement from John Green and Bret Irvine.

REFERENCES

- [1] Kari Adén-Papadopoulos. 2014. Citizen camera-witnessing: Embodied political dissent in the age of “mediated mass self-communication”. *New Media & Society* 16, 5 (2014), 753–769.
- [2] Patti Bao, Brent Hecht, Samuel Carton, Mahmood Quaderi, Michael Horn, and Darren Gergle. 2012. Omnipedia: Bridging the Wikipedia language gap. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'12)*. ACM, New York, NY, 1075–1084. DOI: <https://doi.org/10.1145/2207676.2208553>
- [3] John Berger et al. 1977. *Ways of seeing*. 1972. London: British Broadcasting Corporation and Penguin Books (1977).
- [4] Melissa Bica, Leysia Palen, and Chris Bopp. 2017. Visual representations of disaster. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW'17)*. ACM, New York, NY, 1262–1276. DOI: <https://doi.org/10.1145/2998181.2998212>
- [5] Pasko Bilic and Luka Bulian. 2014. Lost in translation: Contexts, computing, disputing on Wikipedia. In *Proceedings of the 2014 iConference*. iSchools, 32–44. DOI: <https://doi.org/10.9776/14027>
- [6] Adda B. Bozeman. 2017. *Politics and Culture in International History: From the Ancient Near East to the Opening of the Modern Age*. Routledge.
- [7] Ewa Callahan. 2014. Crosslinguistic neutrality: Wikipedia's neutral points of view from a global perspective. In *Global Wikipedia*. Rowman & Littlefield, Lanham, Maryland, 69–84.
- [8] Ewa S. Callahan and Susan C. Herring. 2011. Cultural bias in Wikipedia content on famous persons. *Journal of the American Society for Information Science and Technology* 62, 10 (2011), 1899–1915. DOI: <https://doi.org/10.1002/asi.21577>
- [9] Wikimedia Commons. 2017. Project Scope. Retrieved from https://commons.wikimedia.org/wiki/Commons:Project_scope.
- [10] Wikimedia Commons. 2017. Welcome. Retrieved from <https://commons.wikimedia.org/wiki/Commons:Welcome>.
- [11] Lynne Cooke. 2005. A visual convergence of print, television, and the internet: Charting 40 years of design change in news presentation. *New Media & Society* 7, 1 (2005), 22–46. DOI: <https://doi.org/10.1177/1461444805049141>

- [12] Emilio Crenzel. 2011. Between the voices of the state and the human rights movement: Never Again and the memories of the disappeared in Argentina. *Journal of Social History* 44, 4 (2011), 1063–1076.
- [13] Wei Dong and Wai-Tat Fu. 2010. Cultural difference in image tagging. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'10)*. ACM, New York, NY, 981–984. DOI: <https://doi.org/10.1145/1753326.1753472>
- [14] Delia Dumitrescu. 2016. Nonverbal communication in politics: A review of research developments, 2005–2015. *American Behavioral Scientist* 60, 14 (2016), 1656–1675. DOI: <https://doi.org/10.1177/0002764216678280>
- [15] Michael Emmison, Philip Smith, Philip D. Smith, and Margery Mayall. 2012. *Researching the Visual*. Sage.
- [16] Michael-Andreas Etter and Finn Arup Nielsen. 2015. Collective remembering of organizations: Co-construction of organizational pasts in Wikipedia. *Corporate Communications* 20, 4 (2015), 431–447. DOI: <https://doi.org/10.1108/CCIJ-09-2014-0059>
- [17] Michela Ferron and Paolo Massa. 2011. Collective memory building in Wikipedia: The case of North African Uprisings. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym'11)*. ACM, New York, NY, 114–123. DOI: <https://doi.org/10.1145/2038558.2038578>
- [18] Michela Ferron and Paolo Massa. 2011. Studying collective memories in Wikipedia. *Journal of Social Theory* 3, 4 (2011), 449–466.
- [19] Michela Ferron and Paolo Massa. 2014. Beyond the encyclopedia: Collective memories in Wikipedia. *Memory Studies* 7, 1 (2014), 22–45. DOI: <https://doi.org/10.1177/1750698013490590>
- [20] Elena Filatova. 2012. Information overlap in multilingual Wikipedia and summarization. *International Journal of Cooperative Information Systems* 21, 4 (2012), 383–403. DOI: <https://doi.org/10.1142/S0218843013500019>
- [21] David Frohlich, Allan Kuchinsky, Celine Pering, Abbe Don, and Steven Ariss. 2002. Requirements for photoware. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (CSCW'02)*. ACM, New York, NY, 166–175. DOI: <https://doi.org/10.1145/587078.587102>
- [22] Shane Greenstein, Yuan Gu, and Feng Zhu. 2016. *Ideological Segregation among Online Collaborators: Evidence from Wikipedians*. Technical Report. National Bureau of Economic Research.
- [23] Angela H. Gutchess and Maya Siegel. 2011. Memory specificity across cultures. In *Memory and Political Change*, A. Assmann and L. Shortt (Eds.). Springer. Google-Books-ID: dyN9DAAAQBAJ.
- [24] Scott A. Hale. 2014. Multilinguals and Wikipedia editing. In *Proceedings of the 2014 ACM Conference on Web Science (WebSci'14)*. ACM, New York, NY, 99–108. DOI: <https://doi.org/10.1145/2615569.2615684>
- [25] Noriko Hara, Prina Shachaf, and Khe Foon Hew. 2010. Cross-cultural analysis of the Wikipedia community. *Journal of the American Society for Information Science and Technology* 61, 10 (2010), 2097–2108. DOI: <https://doi.org/10.1002/asi.21373>
- [26] Shiqing He, Allen Yilun Lin, Eytan Adar, and Brent Hecht. 2018. The_Tower_of_Babel.jpg: Diversity of visual encyclopedic knowledge across Wikipedia language editions. In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM'18)*. AAAI.
- [27] Brent Hecht and Darren Gergle. 2009. Measuring self-focus bias in community-maintained knowledge repositories. In *Proceedings of the Fourth International Conference on Communities and Technologies (C&T'09)*. ACM, New York, NY, 11–20. DOI: <https://doi.org/10.1145/1556460.1556463>
- [28] Brent Hecht and Darren Gergle. 2010. The Tower of Babel meets Web 2.0: uUser-generated content and its applications in a multilingual context. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'10)*. ACM, New York, NY, 291–300. DOI: <https://doi.org/10.1145/1753326.1753370>
- [29] Brent J. Hecht and Darren Gergle. 2010. On the “Localness” of user-generated content. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW'10)*. ACM, New York, NY, 229–232. DOI: <https://doi.org/10.1145/1718918.1718962>
- [30] Katherine Hepworth. 2016. History, power and visual communication artefacts. *Rethinking History* 20, 2 (2016), 280–302. DOI: <https://doi.org/10.1080/13642529.2014.932079>
- [31] Yuheng Hu, Lydia Manikonda, Subbarao Kambhampati, et al. 2014. What we Instagram: A first analysis of Instagram photo content and user types. In *Proceedings of the 8th International Conference on Web and Social Media (ICWSM'14)*. AAAI.
- [32] L. Rowell Huesmann. 2007. The impact of electronic media violence: Scientific theory and research. *Journal of Adolescent Health* 41, 6 (2007), S6–S13. DOI: <https://doi.org/10.1016/j.jadohealth.2007.09.005>
- [33] Dariusz Jemielniak. 2014. *Common Knowledge?: An Ethnography of Wikipedia*. Stanford University Press.
- [34] N. Kanhabua, Tu Ngoc Nguyen, and C. Niederée. 2014. What triggers human remembering of events? A large-scale analysis of catalysts for collective memory in Wikipedia. In *Proceedings of the IEEE/ACM Joint Conference on Digital Libraries. IEEE/ACM*, 341–350. DOI: <https://doi.org/10.1109/JCDL.2014.6970189>
- [35] Brian Keegan. 2013. A history of newswork on Wikipedia. In *Proceedings of the 9th International Symposium on Open Collaboration (WikiSym'13)*. ACM, New York, NY, 7:1–7:10. DOI: <https://doi.org/10.1145/2491055.2491062>

- [36] Brian Keegan. 2015. Emergent social roles in Wikipedia's breaking news collaborations. In *Roles, Trust, and Reputation in Social Media Knowledge Markets*. Springer, 57–79.
- [37] Brian Keegan and Jed R. Brubaker. 2015. "Is" to "was": Coordination and commemoration in posthumous activity on wikipedia biographies. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW'15)*. ACM, New York, NY, 533–546. DOI : <https://doi.org/10.1145/2675133.2675238>
- [38] Brian Keegan, Darren Gergle, and Noshir Contractor. 2011. Hot off the Wiki: Dynamics, practices, and structures in Wikipedia's coverage of the Tōhoku catastrophes. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym'11)*. ACM, New York, NY, 105–113. DOI : <https://doi.org/10.1145/2038558.2038577>
- [39] Brian Keegan, Darren Gergle, and Noshir Contractor. 2012. Staying in the loop: Structure and dynamics of Wikipedia's breaking news collaborations. In *Proceedings of the 8th Annual International Symposium on Wikis and Open Collaboration (WikiSym'12)*. ACM, New York, NY, 1:1–1:10. DOI : <https://doi.org/10.1145/2462932.2462934>
- [40] Brian Keegan, Darren Gergle, and Noshir Contractor. 2013. Hot off the Wiki: Structures and dynamics of Wikipedia's coverage of breaking news events. *American Behavioral Scientist* 57, 5 (2013), 595–622. DOI : <https://doi.org/10.1177/0002764212469367>
- [41] Klaus Krippendorff. 2004. *Content Analysis: An Introduction to Its Methodology*. Sage.
- [42] Paul Laufer, Claudia Wagner, Fabian Flöck, and Markus Strohmaier. 2015. Mining cross-cultural relations from Wikipedia: A study of 31 European food cultures. In *Proceedings of the ACM Web Science Conference (WebSci'15)*. ACM, New York, NY, 3:1–3:10. DOI : <https://doi.org/10.1145/2786451.2786452>
- [43] Hang-Teng Liao. 2014. Cultural politics of user-generated encyclopaedias: Comparing Chinese Wikipedia and Baidu Baike. (2014). DOI : <https://doi.org/10.2139/ssrn.2508885>
- [44] Michael D. Lieberman and Jimmy J. Lin. 2009. You are where you edit: Locating Wikipedia contributors through edit histories. In *Proceedings of the 3rd International AAAI Conference on Web and Social Media (ICWSM'15)*.
- [45] Yilun Lin, Bowen Yu, Andrew Hall, and Brent Hecht. 2017. Problematizing and addressing the article-as-concept assumption in Wikipedia. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW'17)*. ACM, New York, NY, 2052–2067. DOI : <https://doi.org/10.1145/2998181.2998274>
- [46] Brendan Luyt. 2011. The nature of historical representation on Wikipedia: Dominant or alternative historiography? *Journal of the American Society for Information Science and Technology* 62, 6 (2011), 1058–1065. DOI : <https://doi.org/10.1002/asi.21531>
- [47] Brendan Luyt. 2016. Wikipedia, collective memory, and the Vietnam War. *Journal of the Association for Information Science and Technology* 67, 8 (2016), 1956–1961. DOI : <https://doi.org/10.1002/asi.23518>
- [48] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9, Nov. (2008), 2579–2605.
- [49] Jamie Mahoney, Tom Feltwell, Obinna Ajuruchi, and Shaun Lawson. 2016. Constructing the visual online political self: An analysis of Instagram use by the scottish electorate. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 3339–3351.
- [50] Marc Miquel-Ribé and David Laniado. 2016. Cultural identities in Wikipedias. In *Proceedings of the 7th 2016 International Conference on Social Media & Society (SMSociety'16)*. ACM, New York, NY, 24:1–24:10. DOI : <https://doi.org/10.1145/2930971.2930996>
- [51] Alexander G. Nikolaev. 2009. Images of War: Content analysis of the photo coverage of the war in Kosovo. *Critical Sociology* 35, 1 (2009), 105–130. DOI : <https://doi.org/10.1177/0896920508098659>
- [52] Oded Nov, Mor Naaman, and Chen Ye. 2010. Analysis of participation in an online photo-sharing community: A multidimensional perspective. *Journal of the American Society for Information Science and Technology* 61, 3 (2010), 555–566. DOI : <https://doi.org/10.1002/asi.21278>
- [53] Jahna Otterbacher. 2014. Our news, their events? A comparison of archived current events on English and Greek Wikipedias. In *Global Wikipedia*. Rowman & Littlefield, Lanham, Maryland, 49–68.
- [54] Jahna Otterbacher. 2015. Linguistic bias in collaboratively produced biographies: Crowdsourcing social stereotypes?. In *Proceedings of the 8th International Conference on Web and Social Media (ICWSM'15)*. AAAI, 298–307.
- [55] Bulent Ozel and Han Woo Park. 2012. Online image content analysis of political figures: An exploratory study. *Quality & Quantity* 46, 4 (2012), 1013–1024. DOI : <https://doi.org/10.1007/s11135-011-9445-x>
- [56] Béatrice Parguel, Florence Benoît-Moreau, and Cristel Antonia Russell. 2015. Can evoking nature in advertising mislead consumers? The power of "executional greenwashing". *International Journal of Advertising* 34, 1 (2015), 107–134.
- [57] S. Tejaswi Peesapati, Hao-Chuan Wang, and Dan Cosley. 2010. Intercultural human-photo encounters: How cultural similarity affects perceiving and tagging photographs. In *Proceedings of the 3rd International Conference on Intercultural Collaboration (ICIC'10)*. ACM, New York, NY, 203–206. DOI : <https://doi.org/10.1145/1841853.1841887>
- [58] Christian Pentzold. 2009. Fixing the floating gap: The online encyclopaedia Wikipedia as a global memory place. *Memory Studies* 2, 2 (2009), 255–272. DOI : <https://doi.org/10.1177/1750698008102055>

- [59] Ulrike Pfeil, Panayiotis Zaphiris, and Chee Siang Ang. 2006. Cultural differences in collaborative authoring of Wikipedia. *Journal of Computer-Mediated Communication* 12, 1 (2006), 88–113. DOI : <https://doi.org/10.1111/j.1083-6101.2006.00316.x>
- [60] Jonathan M. Powell and Clayton L. Thyne. 2011. Global instances of coups from 1950 to 2010: A new dataset. *Journal of Peace Research* 48, 2 (2011), 249–259.
- [61] Thomas E. Powell, Hajo G. Boomgaarden, Knut De Swert, and Claes H. de Vreese. 2015. A clearer picture: The contribution of visuals and text to framing effects: Visual framing effects. *Journal of Communication* 65, 6 (2015), 997–1017. DOI : <https://doi.org/10.1111/jcom.12184>
- [62] Joseph Michael Reagle. 2010. *Good Faith Collaboration: The Culture of Wikipedia*. MIT Press.
- [63] Gillian Rose. 2016. *Visual Methodologies: An Introduction to Researching with Visual Materials*. Sage.
- [64] Peter J. Rousseeuw. 1987. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* 20 (1987), 53–65.
- [65] Dan Schill. 2012. The visual image and the political image: A review of visual communication research in the field of political communication. *Review of Communication* 12, 2 (2012), 118–142. DOI : <https://doi.org/10.1080/15358593.2011.653504>
- [66] David A. Shamma, Lyndon Kennedy, Jia Li, Bart Thomee, Haojian Jin, and Jeff Yuan. 2016. Finding weather photos: Community-supervised methods for editorial curation of online sources. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW'16)*. ACM, New York, NY, 86–96. DOI : <https://doi.org/10.1145/2818048.2819989>
- [67] David Shim and Dirk Nabers. 2013. Imaging North Korea: Exploring its visual representations in international politics. *International Studies Perspectives* 14, 3 (2013), 289–306. DOI : <https://doi.org/10.1111/j.1528-3585.2012.00493.x>
- [68] Susan Sontag. 2001. *On Photography*. Vol. 48. Macmillan.
- [69] Marita Sturken and Lisa Cartwright. 2001. *Practices of Looking*. Oxford University Press, Oxford.
- [70] Terri L. Towner. 2017. The infographic election: The role of visual content on social media in the 2016 presidential campaign. In *The Presidency and Social Media*. Routledge, 236–262.
- [71] Marlon Twyman, Brian Keegan, and Aaron Shaw. 2017. Black lives matter in Wikipedia: Collective memory and collaboration around online social movements. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW'17)*. ACM, New York, NY, 1400–1412. DOI : <https://doi.org/10.1145/2998181.2998232>
- [72] F. B. Viegas. 2007. The visual side of Wikipedia. In *40th Annual Hawaii International Conference on System Sciences (HICSS'07)*. 85–85. DOI : <https://doi.org/10.1109/HICSS.2007.559>
- [73] Martin Wattenberg, Fernanda Viégas, and Ian Johnson. 2016. How to use T-SNE effectively. *Distill* 1, 10 (Oct. 2016), e2. DOI : <https://doi.org/10.23915/distill.00002>
- [74] Taha Yasseri, Anselm Spoerri, Mark Graham, and Janos Keresz. 2014. The most controversial topics in Wikipedia. In *Global Wikipedia*. Rowman & Littlefield, Lanham, Maryland, 25–48.
- [75] Fatemeh “Mariam” Zahedi and Gaurav Bansal. 2011. Cultural signifiers of web site images. *Journal of Management Information Systems* 28, 1 (2011), 147–200. DOI : <https://doi.org/10.2753/MIS0742-1222280106>
- [76] Ark Fangzhou Zhang, Danielle Livneh, Ceren Budak, Lionel Jr. Robert, and Daniel Romero. 2017. Shocking the crowd: The effect of censorship shocks on Chinese Wikipedia. In *Proceedings of the 11th International AAAI Conference on Web and Social Media (ICWSM'17)*. AAAI.

Received March 2019; revised October 2019; accepted November 2019