Connecting Conflict Data through Wikimedia

Rex W. Douglass*

Thomas Leo Scherer†

J. Andrés Gannon‡

Erik Gartzke§

2022-03-30

Abstract

Conflict datasets are explicit about the criteria they are using, but how do we know that they accurately capture what they claim to? How can we compare across conflict datasets when they are using different lenses? We propose using new, massive, semi-structured data such as Wikipedia as a common point of connection to look across different conflict datasets. We match 11 conflict datasets to the wikidata pages that best capture their conflicts. With this, we are able to validate the datasets against each other and identify areas of overlap and possible gaps in their measurement. We discuss what these gaps mean for results based on these datasets.

1 Introduction

The proliferation of public conflict datasets, such as those from the Correlates of War project and the Uppsala Conflict Data program, has expedited quantitative conflict research by reducing or removing the time and resource intensive data collection phase. Now researchers can select which of the many conflict datasets best capture the universe of cases for their research question. In doing so, they also accept the obvious limitations of that dataset such as its temporal coverage and casualty threshold, as well as less obvious limitations such as incorrect entries and missing cases.

Massive digital corpora such as dbpedia and the common crawl allows new opportunities to explore and possibly address these limitations. They come with their own challenges however in data extraction and quality. Using them requires extensive knowledge of their ontology, structure, and creation process, along with domain expertise to connect them to the existing literature. The ultimate (and ambitious) goal of this project is to reduce these barriers by developing the ability to quickly create conflict datasets that are customizable to specific levels and features, draw upon the combined knowledge of digital corpora (informed by existing dataset), and can be instantly updatable to include new information.

As an initial step, we match the conflicts from 12 prominent datasets to wikimedia pages. We find wikimedia to have high coverage of the existing data, making it a useful tool for making connections across conflict datasets to allow for comparison and validation. This article will briefly discuss the state of conflict data and

*Department of Political Science, University of California, San Diego, CA, USA.

†Department of Political Science, University of California, San Diego, CA, USA.

‡Belfer Center for Science and International Affairs, Harvard Kennedy School, MA, USA.

§Department of Political Science, University of California, San Diego, CA, USA.

the potential (and pitfalls) of wikimedia as a source of data. We describe our process to match wikimedia to the existing data and then show descriptive statistics of the results. We next use the our matches to connect the existing datasets to each other and make comparisons between them. Finally, we discuss the potential of mining wikimedia to supplement our datasets and identify possibly fruitful areas of inquiry.

2 Overview of Existing Conflict Datasets

Conflict datasets is a loose term, often focused on episodes and instances of violence, but may refer to other datasets related to aspects of conflict. Anderton and Carter (2019) organizes their list of conflict and peace datasets substantive areas. Their are the typical categories 'Interstate Conflicts', 'Intrastate Conflicts', 'Terrorism', and 'Mass Atrocities', but also includes 'Military Spending, Armaments, and Armed Forces Personnel', 'Interstate Alliances', and 'Peace'. The different categories often represent different types of actors or motives. In an earlier piece Anderton and Carter (2011) groups datasets as 'armed conflict', 'terrorism', and 'events', which posits the level of analysis as another defining feature.

The group of 'armed conflict' datasets can be organized by actors involved, hostility level, and issue area. For example, the Correlates of War (COW) project organizes its war datasets by the actors involved: inter-state, intra-state, non-state, and extra-state (Sarkees and Wayman 2010). Anderton and Carter (2011) suggest 'war', 'sub-war conflict', and 'latent war' as a typology of hostility levels, but note that most datasets cover multiple levels; in their sorting 15 conflict datasets, the five that capture 'latent wars' capture the other types as well. Many datasets use a fatality criterion as a shorthand for the hostility level they are interested in.

For this project we connect 12 datasets to wikimedia, listed below in are described in Table 1. The Correlates of War (COW) datasets on armed conflict are divided into four types based on the actors and location of the conflict (Sarkees and Wayman 2010). The Interstate War Dataset (IWD) claims to improve on COW Interstate while using the same criteria (Reiter, Stam, and Horowitz 2016). The UCDP/PRIO Armed Conflict Dataset is similar to COW but covers a smaller time span and has a lower fatality threshold Gleditsch et al. (2002). Instead of considering states, the Divided Armies dataset collects conflicts between 'belligerents', which is any political entity with a capital, the ability to control and tax a population, and the ability to field an army (Lyall 2020).

Table 1: Conflict Datasets Matched to Wikimedia

Dataset	Focus	Unit Analysis	Dates	Fatality	State Actors	Other Actor(s)
COWExtra	Armed Conflict	Conflict	1816-2004	1000	1	Nonstate
COWInter	Armed Conflict	Conflict	1823-2003	1000	2	-
COWIntra4.1	Armed Conflict	Conflict-Actor	1818-2008	1000	1	Nonstate
COWintra5.1	Armed Conflict	Conflict-Actor	1818-2014	1000	1	Nonstate
COWNon	Armed Conflict	Conflict	1818-2005	1000	0	Nonstate
DividedArmies	Armed Conflict	Conflict-Actor-Phase	1800-2011	500	2	-
ICB	Crisis	Conflict; Conflict-Actor	1918-2016	NA	2	-
IMI	Intervention	Conflict-Dyad	1946-2006	NA	2	-
IWD	Armed Conflict	Conflict-Actor-Year	1823-2003	1000	2	-
NACH	Armed Conflict	Conflict	1518-1899	20 annum	1	Nat Am
PITF	Armed Conflict	Conflict-Year	1948-2018	1000	1	Rev or Ethnic
UCDP/PRIO	Armed Conflict	Conflict-Year	1946-2018	25 annum	1	Any

The Native American Conflict History (NACH) lists conflicts between Native American communities and colonial powers (Urlacher 2021). The Political Instability Task Force (PITF) dataset is primarily intrastate wars, classified as revolutionary wars where political groups attempt to overthrow the government or ethnic wars where groups seek major status change (Marshall, Gurr, and Harff 2009). International Crisis Behavior (ICB) differs in identifying periods of crisis between states, which may or may not escalate to conflict Brecher and Wilkenfeld (1997). The International Militarized Intervention (IMI) also slightly differs by collecting any time one country moves regular forces to another country, which means armed conflicts but also includes peacekeeping operations and evacuations Pearson and Baumann (1993). We attempted to code the Militarized International Disputes (MIDS) dataset but the success rate of our initial probe was too low to warrant further coding.

3 Validity of Wikipedia as a Data Source

Although the use of wikipedia in social science research is relatively nascent, much has been done to evaluate the accuracy of data from wikipedia (Greenstein and Zhu 2012; Jullien 2012; Mesgari et al. 2015; Benjakob and Aviram 2018). In studies of over 10,000 randomly selected articles, scholar have found Wikipedia's accuracy to be comparable to the Encyclopedia Britannica (Blumenstock 2008), with fairly consistent predictors of site accuracy such as page length, number of citations, links to other pages, and edit history. In a study on political science specifically, Wikipedia's coverage of 246 major-party gubernatorial candidates contained no errors in candidate biographies and 242 out of the 246 had the election results within 0.2% of the actual outcome (Brown 2011). Studies of historical events have similar findings, with Wikipedia's frequency and type of discrepancies being akin to sources like the Dictionary of American History and American National Biography Online, mainly concerning frequently disputed figures such as the precise size of armed forces in a battle or the number of casualties (Rector 2008). Studies also find that contrary to many university course syllabi, Wikipedia vandalism is not as common as many allege: 81% of uncited claims are flagged by Wikipedia bots using automated vandalism detection within 24 hours of being posted and fixed within three hours (Adler et al. 2011; Tramullas, Garrido-Picazo, and Sánchez-Casabón 2016).

Wikipedia's coverage, volume, and structure dwarf those of many other sources used to construct datasets in political science. Furthermore, wikipedia is accompanied by two related websites – wikidata and dbpedia – that both provide versions of the information on Wikipedia structured as Resource Description Framework (RDF) databases. As a result, each Wikipedia page has a corresponding wikidata and dbpedia page organized as 'triples' of a subject, predicate, and object. The wikidata page for Germany, for example, lists a subject (Germany), a predicate (has capital), and object (Berlin). In most cases, there are dozens to hundreds of triples for each page that can be automatically extracted using SPARQL queries (Malyshev et al. 2018). This can automate the process of identifying the universe of cases in ways that save significant time and effort compared to manual data entry.

4 Matching to Wikimedia

For each conflict dataset, coders went through every row and attempted to find the one or two wikipedia pages that best reflected the datasets information. They were instructed to focus on pages that match that entry's actors and location. They were also to consider the time-span and other rows unique to that dataset (e.g. type_of_conflict and intensity_level for UCDP/PRIO). Coders recorded the 'QCode' from the wikidata page associated with that wikipedia page. They also took notes on their matches and recorded how confident they were with that match on a 1 (low) to 3 (high) scale.

Coders were allowed to include two QCodes for each dataset entry. Most datasets were coded by two coders.[1] Coders were instructed to work independently and to not check each others codings. They were encouraged to check non-English wikipedia pages (often assisted with google translate).

Table 2 shows the number of rows in each dataset and the number that we could not match to a wikipedia page. NACH has the lowest match rate at 92%. For the rest of the datasets we were able to find wikimedia matches for over 95% of the entries. The effort resulted in around two matches on average for each entry (not necessarily unique). The average self-reported confidence levels were around 2.5, and the average maximum

[1] To date, COW-Intrastate v.5.1, IWD, and Divided Armies had one coder.

confidence level for each entry was a bit higher. Coders reported the greatest confidence in COWInterstate and IWD and the least confidence on their IMI codings.

Table 2: Summary of Results from Matching to Wikimedia

dataset	n	No Matches	Average Match Count	Average Match Confidence	Average Max Confidence
COWExtra	198	5	1.4	2.6	2.7
COWInter	337	0	2.1	2.9	3.0
COWIntra4.1	442	1	2.1	2.5	2.7
COWIntra5.1	420	14	1.0	2.4	2.4
COWNon	62	1	2.0	2.3	2.5
DividedArmies	825	17	1.3	2.5	2.8
ICB	1,065	29	1.8	2.5	2.7
IMI	1,114	89	1.2	2.2	2.3
IWD	628	1	1.0	2.9	3.0
NACH	148	6	2.0	2.5	2.7
PITF	1,598	0	1.2	2.5	2.5
UCDP/PRIO	2,384	0	2.4	2.7	2.9

For parts of the analysis it is best to have a single 'best match' for each entry of the up to four suggested matches. We select the 'best match' based on the following criteria. For each criteria, we keep any ties and continue moving down the list until there is only one QCode left. The n on each criteria indicates the number of entries with multiple QCodes after each step (out of a total n of 9,058 entries with any QCodes).

0. Take all the entires with multiple QCodes (n=5,066)
1. Take the QCode with the highest confidence (n=1,133)
2. Take the QCode with the most 'votes' (the number of coders that suggested it) (n=942)
3. Take the QCode that was often suggested across the other datasets (n=284)
4. Take the QCode that wikidata categorizes as the most 'conflictness'[1] (n=124)
5. Take the QCode that was reported first (n=22)

The result is a long-form dataset where the rows are every entry in the conflict datasets; every entry is given an identifier unique to that row and has the 'best-match' QCode and the coder confidence of that QCode. Table 3 gives a sample of the dataset. With this, we can connect entries from separate datasets based on common QCodes and compare and contrast their other variables.

Table 3: Sample of Long-Form Conflict Data with QCodes

dataset	identifier	dataset_unique_identifier	QCode	Confidence
COWExtrastate	300	COWExtrastate_300_200_-8	Q891806	3
DividedArmies	1	DividedArmies_1_16_220_5/2/1808	Q152499	3
ICB	393	ICB_393_IRQ_1990_11_29	Q37643	3
IWD	WorldWarI	IWD_WW1Ger-France_220_1914	Q361	3
NACH	685	NACH_685	Q8049168	2
PITF	133	PITF_133_2000	Q1960733	3
UCDPPRIO	249	UCDPPRIO_249_1959	Q8740	3

To get an understanding of what types of pages these QCodes represent, we can look at their 'Instance Of' property from their wikidata pages. Every wikipedia page is given an 'Instance Of' value, which designates what class of entity the wikipedia page is. For example, the wikipedia page for the Earth has a wikidata page that says Earth is an instance of the class 'Organization'. By pulling the 'instance of' values for the QCodes we matched to the conflict data, we can get a sense of what types of pages we have matched to and whether they conform with what we might expect for the different datasets.

Table 5 shows what percent of QCodes matched to each dataset have a given 'instance of'. We first see that the most common classes are what we would hope for: 'war', 'conflict', 'civil war'. We also see that the

[1] Conflictness is an ad-hoc ranking of wikidata classes based on how close they feel our understanding of armed conflict. As examples, the top five are 'war', 'civil war', 'rebellion', 'revolution', 'insurgency'. 'Military Unit' ranks towards the middle, and 'occurrance' is among the least 'conflictness'

distributions between datasets also match what we might expect; for example, COWInterstate has the high percentage of 'war', while COWIntrastate has the high percentage of 'civil war'. Looking further down, we see some 'instance of' values that we might not have expected like 'human' and 'organization'. Looking closer, we see that UN missions are classified as 'organizations', explaining why that class is more prevalent in the IMI dataset. For 'human', we find them often matched with earlier entries. For example, "Tewodros II" is matched to the COWNonstate entry on the First Ethiopian War from 1858-1861. Indeed, searches on the First Ethiopian War come up short, but within the wikipedia entry for Tewodros II we learn that he became the emporer of Ethiopia in 1855 and "Within a few years, he had forcibly brought back under direct Imperial rule the Kingdom of Shewa and the province of Gojjam. He crushed the many lords and princes of Wollo and Tigray and brought recalcitrant regions of Begemder and Simien under his direct rule".

Table 4: Top 10 QCode 'Instance Of' by Dataset

Instance_Of	All	UCDP/PRIO	COWInter	COWIntra	COWIntraV5	COWNon	COWExtra	IMI	ICB	DividedArmies	IWD	PITF	NACH
war	15.8%	18.7%	63.0%	12.7%	13.3%	24.5%	32.8%	13.9%	16.1%	46.2%	61.0%	19.4%	40.4%
conflict	14.6%	18.7%	7.0%	15.4%	15.6%	16.3%	17.5%	9.8%	7.8%	17.9%	10.0%	14.5%	27.3%
civil war	5.4%	14.3%	3.0%	22.8%	23.5%	10.2%	1.5%	6.4%	5.2%	8.4%	2.0%	26.7%	1.0%
battle	4.0%	0.6%	4.0%	1.5%	1.4%	6.1%	3.6%	0.9%	7.3%	4.4%	2.0%	-	6.1%
rebellion	3.6%	2.0%	1.0%	9.3%	9.5%	8.2%	5.1%	1.4%	0.7%	2.4%	1.0%	2.4%	4.0%
military operation	3.4%	1.5%	-	0.4%	1.0%	-	0.7%	5.9%	4.7%	-	-	-	4.0%
coup d'état	2.6%	6.7%	-	1.5%	1.0%	-	-	2.3%	1.7%	-	-	1.8%	-
human	2.5%	0.3%	1.0%	3.5%	4.1%	10.2%	4.4%	1.4%	0.5%	0.8%	1.0%	1.2%	1.0%
revolution	1.6%	2.3%	-	6.6%	5.4%	2.0%	2.2%	0.9%	1.2%	0.8%	-	4.8%	-
organization	1.4%	0.3%	-	0.4%	0.7%	-	-	3.6%	0.5%	-	-	-	-

5 Overlap of Conflict Datasets

The result is a dataset of 1,282 QCodes that are connected to 9,058 entries across 12 conflict datasets. We can use this dataset as a Rosetta Stone to connect any single entry to similar entries in other datasets. We can also use it to compare across datasets and perhaps assess their validity and completeness. Figure 1 an Upset Plot showing the overlap between the different conflict datasets. While difficult to take-in at once, some observations are interesting. First, there is very little overlap between the four COW datasets, which we would expect based on their construction. Second, COWInterstate and IWD have significant overlap, which again we would expect as IWD is a revision of COWInterstate. Third, IMI and ICB do appear to be unique in what they capture with a majority of their QCodes not overlapping with other datasets. The overlap they do have is foremost with each other. Figure 2 is the same plot with the COW datasets collapsed, clarifying some of the above observations.

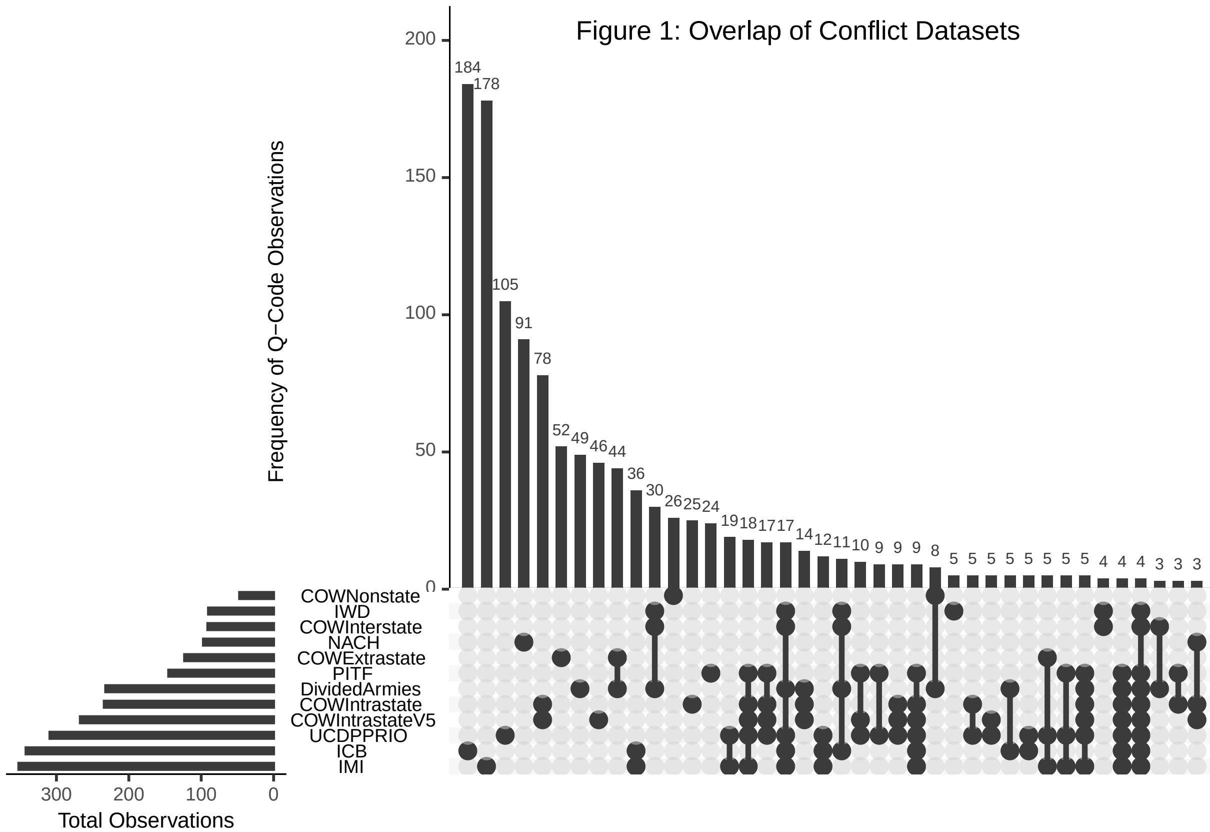

Figure 1: Overlap of Conflict Datasets

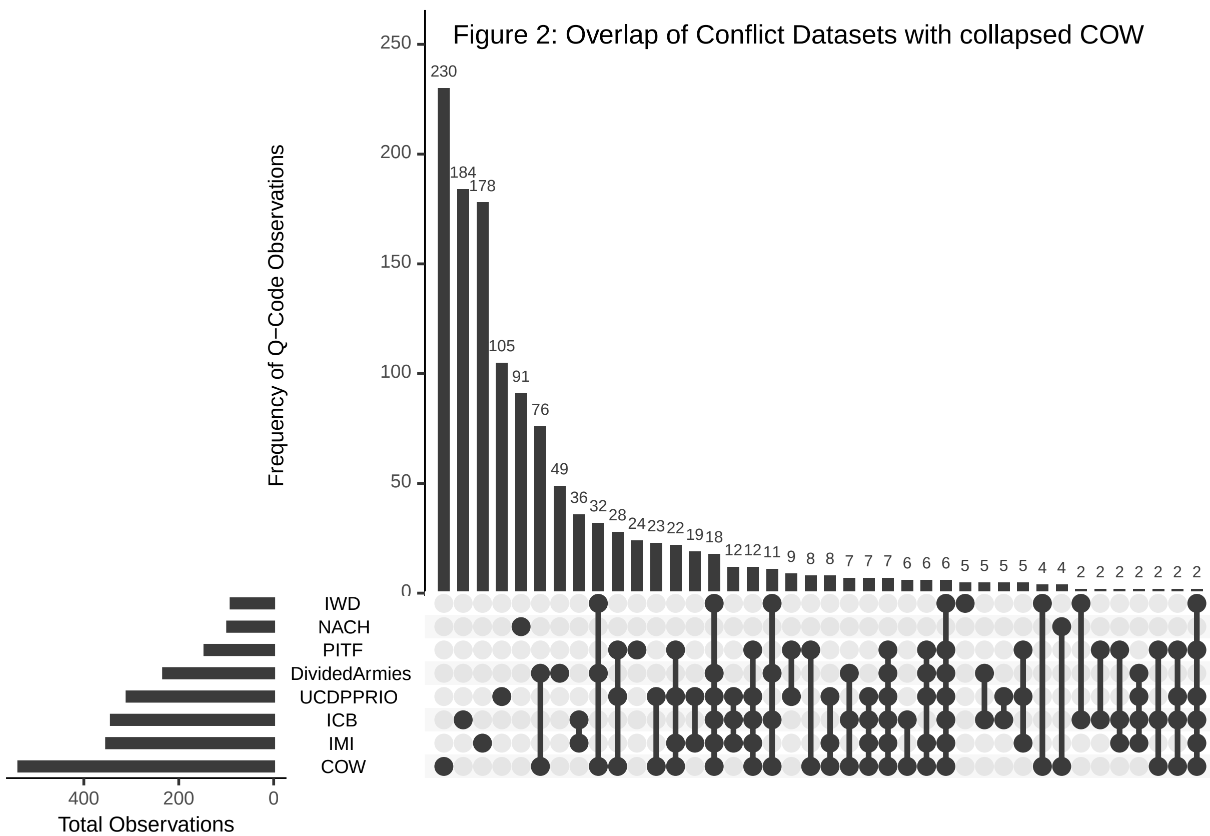

Figure 2: Overlap of Conflict Datasets with collapsed COW

Table 5 shows the pairwise overlap by showing the percent of the entries in each dataset (row) that have the same QCode as an entry in each other dataset (column). This is useful to understanding the breadth of overlap for datasets that have a format that may represent a single conflict with multiple rows, such as the conflict-year format of UCDP/PRIO. Note that the inverse cells may differ; IWD's QCodes match 26% of ICB's entries, while ICB's QCodes match 69% of IWD's entries. This makes sense given ICB's larger scope and size than IWD. This view of overlap shows the same patterns mentioned above while also highlighting the wide range in overlap across datasets.

Table 5: Percent of each dataset's entries (row) with QCodes found in the other datasets (column)

dataset	COWExtra	COWInter	COWIntra	COWIntraV5	COWNon	DividedArmies	ICB	IMI	IWD	NACH	PITF	UCDPPRIO
COWExtra	100%	3%	7%	6%	7%	46%	8%	16%	3%	0%	11%	23%
COWInter	2%	100%	14%	13%	3%	92%	59%	39%	90%	1%	13%	41%
COWIntra	3%	10%	100%	88%	7%	29%	32%	36%	8%	1%	46%	50%
COWIntraV5	4%	5%	84%	100%	5%	23%	21%	28%	4%	2%	40%	46%
COWNon	13%	0%	13%	11%	100%	30%	7%	6%	0%	2%	4%	4%
DividedArmies	18%	46%	19%	21%	9%	100%	37%	18%	43%	1%	11%	20%
ICB	3%	27%	18%	18%	3%	34%	100%	42%	26%	0%	17%	35%
IMI	5%	29%	27%	26%	4%	37%	56%	100%	28%	0%	29%	63%
IWD	1%	97%	16%	13%	1%	95%	69%	34%	100%	1%	14%	36%
NACH	0%	1%	4%	4%	8%	2%	0%	0%	1%	100%	0%	0%
PITF	8%	13%	76%	79%	5%	22%	34%	57%	10%	0%	100%	89%
UCDPPRIO	15%	10%	56%	62%	3%	17%	27%	46%	9%	0%	63%	100%

6 Validation Checks

How do we know if our matching is accurate? Our ability to connect across datasets that should overlap provides an opportunity to identify and assess anomalies. For example, every UCDP/PRIO conflict coded as having over 1,000 battle related deaths (cumulative_intensity = 1) should be included in COW, and similarly every COW conflict after 1946 should be in UCDP/PRIO. Similar overlap also occurs between PITF and UCDP/PRIO, IWD and COWInterstate, and Divided Armies and COWInterstate (not everything in Divided Armies should be in COWInterstate, but everything in COWInterstate should be in Divided Armies).

This effort continues, but an early lesson is that we do not have a good way to handle different parts of a conflict. For example, the COWInterstate conflict between Chad and Libya from 1986-1987 was matched to the 'Chadian–Libyan conflict' from 1978-1987 (Q611071), but Divided Armies did not have that QCode. Instead, the 1987 Chad-Libya conflict in Divided Armies was matched to the 'Toyota War' from 1986-1987 (Q644589), which wikidata identifies as connect to the 'Chadian-Libyan conflict.' The dates would suggest that both datasets should be matched to the Toyota War, but the longer 'Chadian-Libyan' conflict suggests there is more context to the conflict that may be relevant to understanding it but falls outside each datasets' criteria.

7 Mapping the Wikimedia Ontology

To explore the use of wikimedia as a source of 'new' information, we examine the use of wikidata's 'instance of' property and the use of the class structure. We mapped out the class structure starting with the classes matched to our conflict data and then moving up and down the class tree to identify the branches that capture most of our matched QCodes. Figure 1 takes the classes matched to existing conflict data, and maps out the classes that are in the conflict branch of wikimedia's ontology; the node size is based on frequency of each class in the existing conflict data. We can see that 'war' and 'conflict' are common classes in the existing data. We also see nodes that might seem strange, such as 'Venezuelan civil war' and 'Colombian civil war' appearing as unique nodes and that there is a node called 'legal cases'[1]

Figure 1: Conflict Data in the 'Conflict Branch' (Scaled by Conflict Data Frequency)

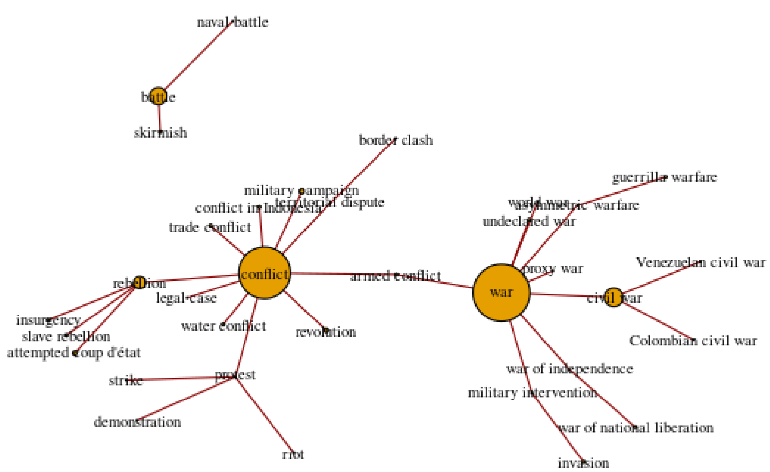

Having identified this branch as potentially useful, Figure 2 shows the complete branch; here the node size is based on the frequency of each class in wikimedia. This suggests there are many other wikimedia classes and entries that may be of use to conflict research and are worth a closer look. It also shows the that this branch should not be scraped wholesale, lest we end up with a dataset largely composed of legal cases. We are still working on ways to draw from this information, but to give a sense of scope Table 6 shows the size and the properties of the wikimedia entries for some of the more obvious classes.

[1]This class was dropped when we identify the single best QCode for each conflict.

Figure 2: Complete 'Conflict Branch' (Scaled by Wikimedia Frequency)

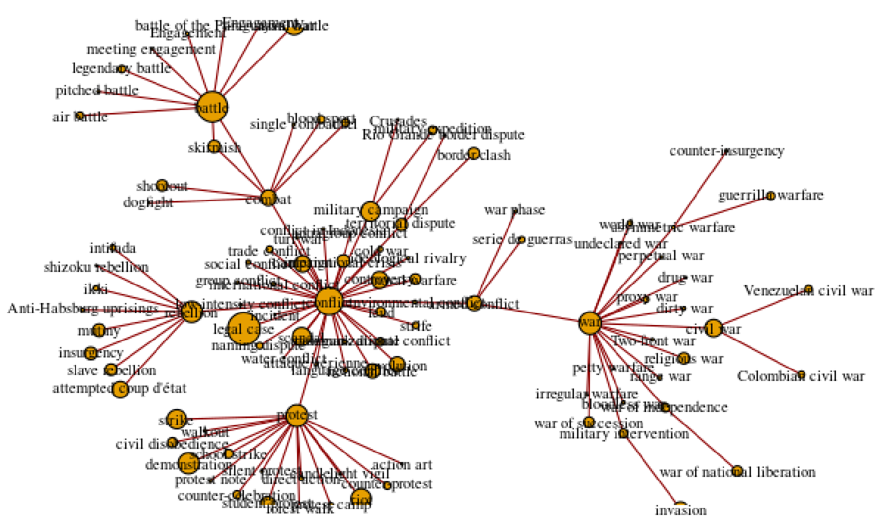

Table 6: Description of Specified Wikimedia Classes

entity type	In Existing Conflict Data	In Wikimedia	Percent with start date	Percent with end date	Percent with participants
conflict	244	1,378	1	1	0.05660377
war	259	532	1	1	0.83082707
civil war	85	134	1	1	0.43283582
battle	87	2,311	1	1	0.36088273

8 Future Promises and Pitfalls

This initial look at interacting digital corpora with exisiting conflict data highlights many challenges to this path, of which three stand out to us. First, there is the issue of data quality. We believe there is already a healthy caution around using Wikimedia, and there should continue to be. Besides missingness and miscodings, but there are also intentional structure and organization decisions that will undermine our efforts if we are not careful (i.e. legal cases as conflicts). Second, this initial step highlights the challenges around handling information at multiple levels of analysis. This challenge is already well known to anyone attempting to conduct an analysis on both COW and UCDP/PRIO wtih their respective conflict and conflict-year formats. Just as conflict datasets may organize a conflict into its constituent actors, time period, phases, and events, wikimedia also organizes conflict into some unexpected ways (i.e. the Toyota War as part of the Chadian-Libyan conflict). We often see these as additional information to be unlocked with enough effort and care.

Finally, we do not have a strong sense of how we as political scientists should approach inputting information into our source digital corpora. In our coding efforts we cautioned coders not to make any changes to wikimedia even if they found blatant errors. We were concerned that we may be accused (justifiably) of change our information sources to better suit our own ends (cooking our wikibooks). We are also wary of adding information from existing datasets to wikimedia wholesale as the information as it may bring along errors in the existing data in a cloak of legitimacy; certainly the IWD project would strongly caution against adding all of COWInterstate to wikimedia. On the other hand, we believe political scientists should be contributing to these digital corpora, both as a tool to educate and as a way to add information to data extraction efforts. The solution may be in academia adopting a version of Wikimedia's rules and standards that addresses academia's specific concerns. As a bonus, wikimedia may serve as a useful forum to identify and debate conflicting information.

9 Conclusion

This paper hopes to give a sense of the potential for interacting existing conflict datasets with large digital corpora, specifically wikimedia. While there is much work to do, with caution, our progress so suggests this will continue to be a fruitful line of inquiry.

9.0.0.1 Acknowledgments: Thanks to the Center for Peace and Security Studies and its membership for comments. Thank you to the many research assistants that contributed to this project: Nathan Cheung, Helen Chung, Janet Gleason, Daman Heer, Bhavika Kaneria, Xiangyi Kong, Anthony Limon, Ari Michelson, Christiana Moore, Prithviraj Pahwa, Jesus Rojas, Jinneva Santiesteban, Kaitlyn Willoughby, Adika Witoelar, Ruixuan Zhang. This project has benefited from the Charles Koch Foundation's support for the Center for Peace and Security Studies.

10 References

Adler, B. Thomas, Luca de Alfaro, Santiago M. Mola-Velasco, Paolo Rosso, and Andrew G. West. 2011. "Wikipedia Vandalism Detection: Combining Natural Language, Metadata, and Reputation Features." In *Computational Linguistics and Intelligent Text Processing*, 277–88. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-19437-5_23.

Anderton, Charles H., and John R. Carter. 2011. "Conflict Datasets: A Primer for Academics, Policymakers, and Practitioners." *Defence and Peace Economics* 22 (1): 21–42. https://doi.org/10.1080/10242694.2010.491677.

———, eds. 2019. "Conflict and Peace Datasets." In *Principles of Conflict Economics: The Political Economy of War, Terrorism, Genocide, and Peace*, Second, 428–37. Cambridge: Cambridge University Press. https://doi.org/10.1017/9781316875506.021.

Benjakob, Omer, and Rona Aviram. 2018. "A Clockwork Wikipedia: From a Broad Perspective to a Case Study." *Journal of Biological Rhythms*, April, 0748730418768120. https://doi.org/10.1177/0748730418768120.

Blumenstock, Joshua E. 2008. "Automatically Assessing the Quality of Wikipedia Articles," April.

Brecher, Michael, and Jonathan Wilkenfeld. 1997. *A Study of Crisis*. University of Michigan Press.

Brecher, Michael, Jonathan Wilkenfeld, Kyle Beardsley, Patrick James, and David Quinn. 2021. "International Crisis Behavior Data Codebook, Version 14." *URL: Http://Sites. Duke. Edu/Icbdata/Data-Collections*.

Brown, Adam R. 2011. "Wikipedia as a Data Source for Political Scientists: Accuracy and Completeness of Coverage." *PS: Political Science & Politics* 44 (02): 339–43. https://doi.org/10.1017/S1049096511000199.

Gleditsch, Nils Petter, Peter Wallensteen, Mikael Eriksson, Margareta Sollenberg, and Havard Strand. 2002. "Armed Conflict 1946-2001: A New Dataset." *Journal of Peace Research* 39 (5): 615–37.

Greenstein, Shane, and Feng Zhu. 2012. "Is Wikipedia Biased?" *American Economic Review* 102 (3): 343–48. https://doi.org/10.1257/aer.102.3.343.

Jullien, Nicolas. 2012. "What We Know About Wikipedia: A Review of the Literature Analyzing the Project(s)." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2053597.

Kisangani, Emizet F., and Jeffrey Pickering. 2008. *International Military Intervention, 1989-2005*. Inter-university Consortium for Political and Social Research Ann Arbor, MI.

Lyall, Jason. 2020. *Divided Armies: Inequality and Battlefield Performance in Modern War*. Vol. 164. Princeton University Press.

Malyshev, Stanislav, Markus Krötzsch, Larry González, Julius Gonsior, and Adrian Bielefeldt. 2018. "Getting the Most Out of Wikidata: Semantic Technology Usage in Wikipedia's Knowledge Graph." In *The Semantic Web – ISWC 2018*, edited by Denny Vrandečić, Kalina Bontcheva, Mari Carmen Suárez-Figueroa, Valentina Presutti, Irene Celino, Marta Sabou, Lucie-Aimée Kaffee, and Elena Simperl, 376–94. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-00668-6_23.

Marshall, Monty G., Ted Robert Gurr, and Barbara Harff. 2009. "PITF-State Failure Problem Set: Internal Wars and Failures of Governance, 1955–2008." *Dataset and Coding Guidelines*.

Mesgari, Mostafa, Chitu Okoli, Mohamad Mehdi, Finn Årup Nielsen, and Arto Lanamäki. 2015. "'The Sum of All Human Knowledge': A Systematic Review of Scholarly Research on the Content of Wikipedia: 'The Sum of All Human Knowledge': A Systematic Review of Scholarly Research on the Content of Wikipedia." *Journal of the Association for Information Science and Technology* 66 (2): 219–45. https://doi.org/10.1002/asi.23172.

Pearson, Frederic S., and Robert A. Baumann. 1993. *International Military Intervention, 1946-1988*. Inter-university Consortium for Political and Social Research Ann Arbor, MI.

Pettersson, Therese, Shawn Davies, Amber Deniz, Garoun Engström, Nanar Hawach, Stina Högbladh, and Margareta Sollenberg Magnus Öberg. 2021. "Organized Violence 1989–2020, with a Special Emphasis on Syria." *Journal of Peace Research* 58 (4): 809–25.

Rector, Lucy Holman. 2008. "Comparison of Wikipedia and Other Encyclopedias for Accuracy, Breadth, and Depth in Historical Articles." *Reference Services Review* 36 (1): 7–22. https://doi.org/10.1108/00907320810851998.

Reiter, Dan, Allan C. Stam, and Michael C. Horowitz. 2016. "A Revised Look at Interstate Wars, 1816–2007." *Journal of Conflict Resolution* 60 (5): 956–76.

Sarkees, Meredith Reid, and Frank Wayman. 2010. *Resort to War: 1816-2007*. CQ Press.

Tramullas, Jesús, Piedad Garrido-Picazo, and Ana I. Sánchez-Casabón. 2016. "Research on Wikipedia Vandalism: A Brief Literature Review." In *Proceedings of the 4th Spanish Conference on Information Retrieval*, 15:1–4. CERI '16. New York, NY, USA: ACM. https://doi.org/10.1145/2934732.2934748.

Urlacher, Brian R. 2021. "Introducing Native American Conflict History (NACH) Data." *Journal of Peace Research* 58 (5): 1117–25.