

Using machine-coded event data for the micro-level study of political violence

Research and Politics
July-September 2014: 1–8
© The Author(s) 2014
DOI: 10.1177/2053168014539924
rap.sagepub.com


Jesse Hammond^{1,2} and Nils B Weidmann²

Abstract

Machine-coded datasets likely represent the future of event data analysis. We assess the use of one of these datasets—Global Database of Events, Language and Tone (GDELT)—for the micro-level study of political violence by comparing it to two hand-coded conflict event datasets. Our findings indicate that GDELT should be used with caution for geo-spatial analyses at the subnational level: its overall correlation with hand-coded data is mediocre, and at the local level major issues of geographic bias exist in how events are reported. Overall, our findings suggest that due to these issues, researchers studying local conflict processes may want to wait for a more reliable geocoding method before relying too heavily on this set of machine-coded data.

Keywords

Event data, political violence, micro-level analysis, Global Database of Events, Language and Tone, geographic information system

The introduction last year of the Global Database of Events, Language and Tone (GDELT; Leetaru and Schrod, 2013) has caused a stir in academic and policy communities alike. With a quarter of a billion observations, a 35-year temporal span, and daily updates through automated coding, the advantages are manifold. Using automatic geo-referencing routines, event data not only come with temporal coordinates, but are also tagged with geographic coordinates. As compared to the previous generation of machine-coded event datasets, this makes the new generation of event data suitable for the micro-level, geo-spatial analysis of political events. In this paper, we assess the use of these datasets for micro-level studies, focusing in particular on political violence where spatial analysis has become a widely used approach. While there has been earlier work attesting to the validity of machine-coded events, our focus in this paper is the quality of geolocalization. In other words, we ask whether machine-coded event datasets can approximate the spatial pattern of a conflict to a reasonable extent.

To answer this question, we correlate the spatial patterns of violence as coded by two established human-coded event datasets to those given by GDELT. Since these datasets attempt to code the same type of event, and use many of the same sources to code event data, we expect the correlation to be high. Yet, this is not what we find: our spatial-temporal

analysis shows that GDELT does correlate significantly with the other event datasets, although these correlations remain fairly low. However, when we collapse our data to a time series—eliminating the geographic dimension—correlations become much stronger. This suggests that spatial disaggregation accounts for much of the disagreement between machine- and human-coded data. To further explore this, we analyze how spatial remoteness accounts for the mismatch between the data. The results confirm a geographic bias; whereas GDELT over-reports violence close to the capital, the opposite applies in remote locations. This is problematic, since we risk falsely associating civil war violence with urban areas (Kalyvas, 2004). Despite the proven reliability of machine-coding techniques, our findings indicate that significant work is required in optimizing the automated geolocalization of event data. In the next section, we briefly sketch the development of event data in political analysis, before turning to our analysis and results, which focus on political violence more narrowly.

¹University of California at Davis, USA

²University of Konstanz, Germany

Corresponding author:

Jesse Hammond, University of California at Davis, 1 Shields Avenue, Davis, CA 95616, USA.

Email: jrhammond@ucdavis.edu



Automatic and non-automatic approaches to event coding

The recent surge of academic interest in political violence has spurred the development of event data. An event dataset is one that lists individual acts or interactions along with precise coordinates. Event data have a history in political science; earlier approaches originated mostly in international relations, focusing on actions between states. Examples date back a few decades, and include the World Event Interaction Survey (WEIS; McClelland, 1976) and the Conflict and Peace Data Bank (COPDAB; Azar, 1980). Later, Schrodt et al. (1994) developed machine-assisted approaches to generate the same kind of data from news sources, which have been used in the creation of large international datasets such as the Kansas Event Data System (KEDS; Gerner et al., 1994). Validity checks indicate that machine-coded events can match human-coded data with a high degree of accuracy regarding the substance of events and the actors involved (Best et al., 2013; Schrodt and Gerner, 1994).

A recent generation of event datasets combines new data with a sharper substantive focus. Most importantly, location plays a major role. Event datasets such as the Armed Conflict Location and Event Dataset (ACLED; Raleigh et al., 2010), the Geo-referenced Event Dataset (GED) released by the Uppsala Conflict Data Program (Sundberg and Melander, 2013), or the Social Conflict in Africa Database (SCAD; Salehyan et al., 2012) all list events with precise spatial coordinates, making it possible to study patterns of violence as well cross-link the events to other spatial data (Gleditsch and Weidmann, 2012). This was done to allow both cross-national and subnational analysis of political violence, as for example in Raleigh and Hegre (2009) or Weidmann and Ward (2010). All these datasets rely largely on human coding of news reports, and thus require significant effort and time.¹

Further developments in machine coding attempt to generate a similar type of data. For example, the recent GDELT project applies automated coding techniques to the generation of event data from English language media reports, with the key addition of automated geocoding. This results in a data product that largely relies on the same sources as the human-coded datasets (local and international sources from newspapers and newswires, usually provided through repositories such as LexisNexis and Factiva), but whose creation is fully automated. The machine-coding approach is to define lists of actors and events, as well as the words identifying them in natural language. News reports are fed into a parser that identifies events based on the actor and event keywords given in the dictionaries.² GDELT relies on the CAMEO event classification scheme to code events (Gerner et al., 2002). Similarly, a simple lookup of place names is used for each event to extract location information from media reports (Leetaru and Schrodt, 2013).

There are several reasons why machine-coded events could diverge from human-coded ones. Firstly, there may be differences regarding the events that are detected: the automated coding process may fail to pick up events, or code others that the human coder chose not to include. This could be due to ambiguous or non-standard language patterns describing events, but could also be due to the fact that human-coded datasets require background information on civil wars when coding events: the GED, for example, requires events to be explicitly linked to an ongoing civil war in order to be included in the dataset. Machine-coding routines do not have this background knowledge, and rely entirely on information and keywords supplied in a single article. This latter effect should lead machine-coding algorithms to be much less selective when it comes to the events included. Secondly, even if events are included in both data collections, these coding approaches may differ when it comes to the nature of the event. For example, discrepancies could exist in the coding of the attacker or the target of violence. Lastly, even with identical source data and coding techniques, locational discrepancies could exist due to issues with the geo-referencing process. Automated geo-referencing is far from straightforward; oftentimes, reports mention different locations, and it is difficult to discern which locational information in a report actually describes the event's location. For example, due to frequent mentions of the capital in news reports, a "capital bias" could occur that places many events wrongly in the country's most prominent city.

This information, however, is key for the rapidly evolving micro-level analysis of violence (Kalyvas, 2012), which oftentimes relies on geographic information. How well does GDELT track variation in violence at the local level? Earlier work has analyzed the validity of machine coding more generally (Schrodt and Gerner, 1994), and there have been a few attempts to conduct similar analyses of GDELT in particular. A paper by Arva et al. (2013) compares GDELT to another event dataset—the commercial product ICEWS³—that relies on machine coding. However, this comparison is done at the country level, and therefore is not particularly insightful for micro-level researchers. Another GDELT validation by Ward et al. (2013) also uses ICEWS for comparison, but focuses on temporal trajectories at a fixed location, which again cannot tell us much about correspondence across space. Therefore, our analysis is the first attempt to probe GDELT's usefulness for micro-level research. The next section presents our approach and the results.

Data

Our analysis compares GDELT to two other established datasets on political violence, the ACLED (Raleigh et al., 2010) and GED (Sundberg and Melander, 2013).⁴ To ensure overlap between these datasets, we focus on African violence

between 1997 and 2008. Within this range, we further subset the data into only country-years during which at least one civil conflict was ongoing, as coded by the UCDP-PRIO armed conflict dataset (Themner and Wallenstein, 2013). The resulting coverage for all three datasets includes 25 countries and 136 conflict-years. Within this range, we aggregate the three datasets at the level of cell-months, using uniform grid cells of approximately 55 kilometers on each side from the PRIO-GRID dataset (Tollefsen et al., 2012). Our dataset consists of 5897 unique cells, with a total of 547,812 monthly observations.

Our comparison focuses on the presence or absence of conflict events at the cell-month level between these three datasets. As they ostensibly cover the same type of event, we are interested in seeing how similar subnational GDELT coverage of civil war violence is to the more established hand-coded datasets. Since GDELT covers a much wider range of political events, we filter out those that fit the definition of civil war violence. Specifically, a GDELT event is kept in the dataset if it meets the following criteria for information:

1. the event is classified as direct armed clash or battle;
2. the event can be geolocated to at least the administrative district or city level;
3. the initiator is identified generally as a government body, military organization, or insurgent group.

This subsetting only requires event type, location, and actor, a bare minimum for event data analysis. For the hand-coded datasets, we also subset by event type (violent conflict event), and geographic precision (events coded at the city or district level). For a more in-depth explanation of our coding rules see Appendix A, and for the results of parallel analyses using alternative coding rules, see Appendix B.

We also realize that these datasets, even when coding the same type of event, use different general rules in converting raw data to event records. The GED coding method attempts to create one record “per event”, meaning that events lasting more than one day only count as one observation, and it differs from the other two datasets in that it only codes events where at least one fatality is recorded, leading to a lower overall event count. The ACLED coding, on the other hand, codes what are essentially “event-days”, with one observation per day that an event was ongoing. Finally, GDELT’s automated coding system can vary as to how many observations correspond to one “event”, as it relies on specific pieces of information in a news article to determine whether two sources describe the same event. To control for this issue, we dichotomize monthly event-counts to a conflict variable (0/1), with 1 corresponding to cases where at least one event was

reported in a given cell-month, and 0 otherwise. We also run our set of models on the raw event-count data, and find very similar results (Appendix C).

A final potential issue is that the three datasets use partly different sources. ACLED and GED do not have specific rules regarding what sources are used to code events: they rely primarily on newswire sources, but also use a wide variety of media and non-media sources to maximize coverage. GDELT, on the other hand, relies on a much more limited set of media sources, primarily four news streams available on LexisNexis: Agence France-Presse, Associated Press, BBC Monitoring Service, and Xinhua News Service (and Google News from 2002 onwards). In order to rule out the possibility that differences in the sources used account for potential discrepancies, we run a set of models on a subset of ACLED and GED that only includes events recorded by the four main media sources used by GDELT. We find virtually identical results (Appendix D).

Correlations across datasets

Our first step is to compare time trends in event reporting, similar to the approach followed in other GDELT validation exercises. Figure 1 displays aggregated time series for GDELT, ACLED, and GED for the period we examine, showing the total number of grid cells experiencing one or more conflict events in a given month.

When looking at the data as a time series (ignoring spatial variation), GDELT correlates with ACLED at 0.64, and with GED at 0.33. Overall, this is in line with previous findings that automated event coding can generally track hand-coded data (Best et al., 2013), but the overall strength of the correlations remains modest. However, these correlations drop dramatically once we add the spatial dimension. At the cell-month level, GDELT correlates with ACLED at 0.26, and with GED at 0.20. A closer look at the confusion tables (Table 1) reinforces this impression. While all datasets agree on the absence of conflict in the vast number of cases (top left cells), the majority of grid cells that GDELT codes as experiencing conflict are not classified as such by either of the two datasets (bottom left cells), nor does GDELT pick up many of the cases that ACLED or the GED code as conflict (top right cells). Taken together, these figures provide evidence that the automated coding process GDELT uses to geolocate events may differ from human coders: even though it tends to code a similar set of events as ACLED and GED (as confirmed by the time trends above), it seems to have issues in placing those events in the same locations as human coders.

To further study how the correlations between the datasets vary over time and space, we visualize “overlap” between the datasets in Figures 2 and 3. The bars indicate the number of conflict cases identified by each dataset over time (Figure 2) and space (Figure 3). In each plot the

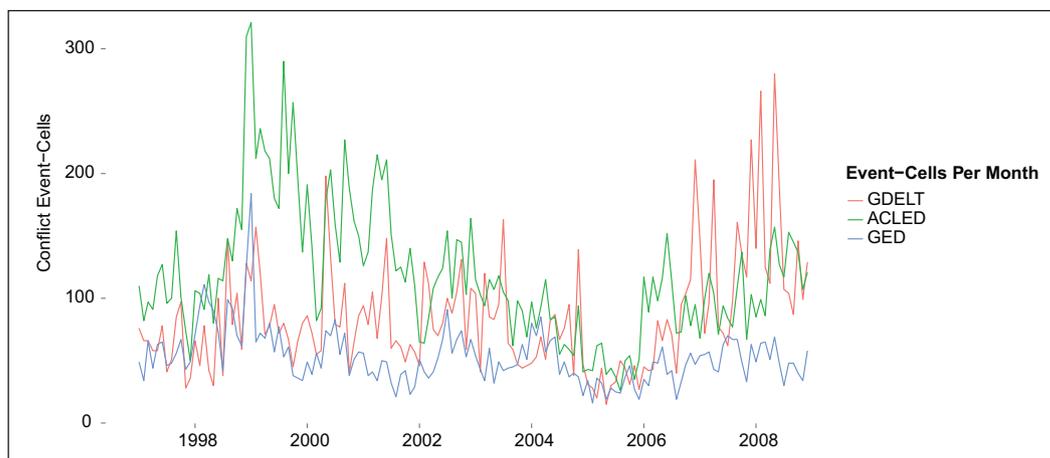


Figure 1. Total events over time.

Table 1. Confusion matrices.

	ACLED = 0	ACLED = 1	GED = 0	GED = 1
GDELT = 0	538,541	5552	540,494	3589
GDELT = 1	2398	1331	2906	823

N = 547,812 cell-months.

ACLED: Armed Conflict Location and Event Dataset; GED: Geo-referenced Event Dataset.

overlapping area in the center represents the number of cases where both datasets code conflict (the true positives). Figure 2 pools all observations for a given month to create the plot. Figure 3 pools observations by cell, and orders the cells by distance from capital. Due to the large number of grid cells (nearly 6000), we subset the sample visualized in Figure 3 to only include grid cells where both datasets report at least one conflict event in the same month.

The non-overlapping parts of the bars show that there are a large number of cases that are coded as conflict by ACLED or GED, but that are not captured by GDELT (and, as Appendix D shows, these cases are not due to the wider range of sources used by the hand-coded datasets). No clear trend can be identified when it comes to variation over time (Figure 2). Visualizing conflict reports over space shows higher reporting by all datasets in grid cells close to the capital city, but there is no visually discernible trend in how distance to the capital affects the likelihood of both datasets reporting an event in the same cell-month. We further explore spatial variation in disagreement between the datasets using regression analysis. We create a variable measuring incidences where a cell-month was coded as conflict by one of the human-coded datasets, but not GDELT, or by GDELT but not by the human-coded dataset. We regress this variable on the distance from the capital (logged and normalized to the 0–1 interval by country), controlling for

logged population and the amount of rugged terrain within a given cell. The results are shown in Table 2.

Controlling for population, we find clear evidence for geographic bias in GDELT. Distance from the capital decreases GDELT coverage as compared to the human coded datasets (Models 2 and 4), whereas the opposite is true if we move closer to the capital (Models 1 and 3). If we assume that human geo-referencing is more accurate (which we believe is a reasonable assumption), these results are consistent with GDELT not being able to geo-reference events accurately, and wrongly placing them near the capital.

Even if capital bias exists in GDELT, is there reason to worry? In other words, do these misallocated events in GDELT fundamentally change the results we obtain based on these data? To analyze this, we estimated simple structural models using the different datasets in our study. Again, we stick to a binary dependent variable (violence in cell-month). Following standard methodology, we include both the spatial and the temporal lag of violence, both measured in the month before, as well as control variables for population and mountainous terrain. Of particular interest here is the coefficient linking geographic remoteness, as measured by distance to the capital, to the presence of conflict in a given cell-month. Table 3 shows the results.

Most coefficients behave as expected and are consistent across the datasets. For example, violence is more likely to happen in cells with high population or mountainous terrain, and in those that experienced conflict in the month before. However, different datasets seem to give different answers when it comes to the effect of remoteness on violence. The human-coded datasets (ACLED and GED, Models 5 and 6) confirm the frequent finding that violence is more likely in remote areas of a country. This finding has been established with a number of data collections other than ACLED and GED (Buhaug and Rød, 2006; Buhaug

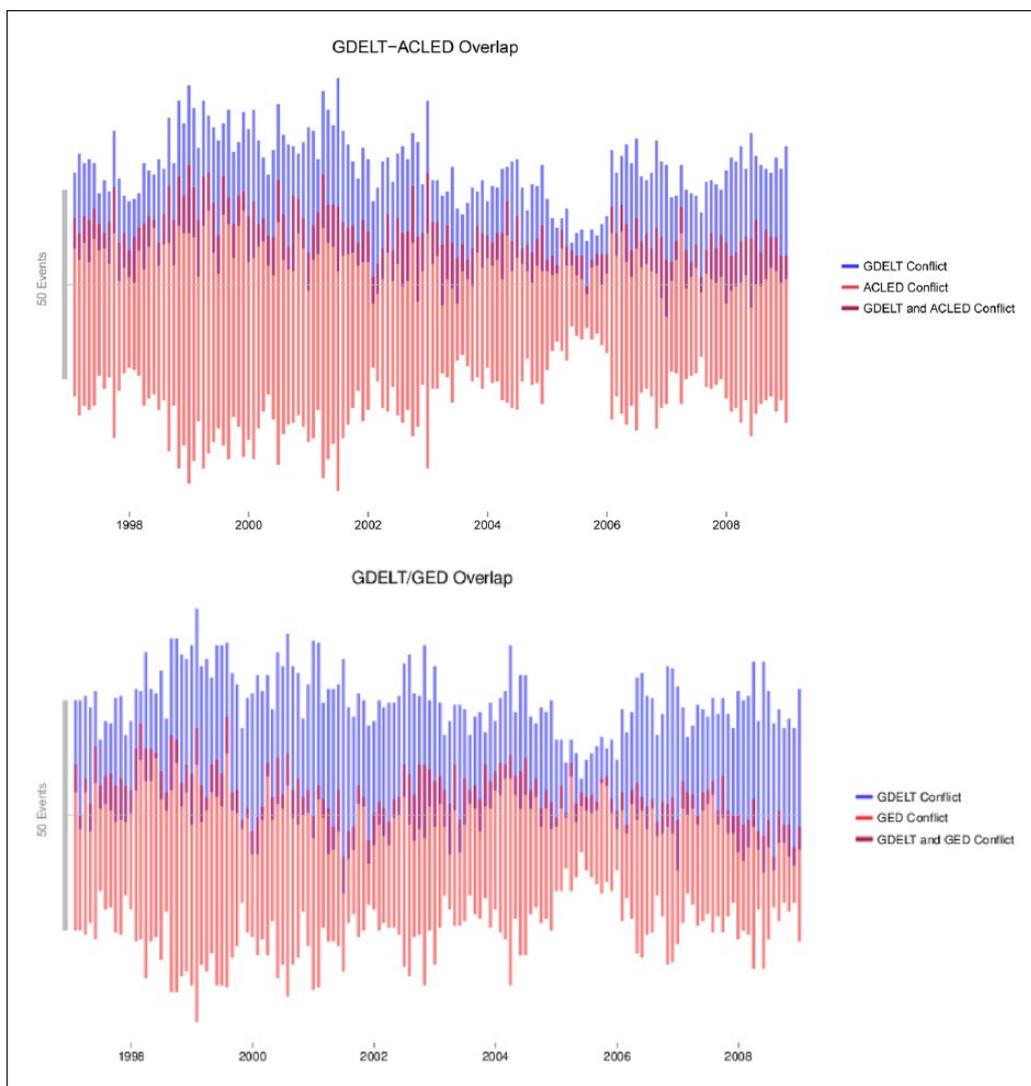


Figure 2. Grid-cell events over time.

et al., 2008). Results based on GDELT, however, suggest exactly the opposite (Model 7), and show that violence according to GDELT is more likely to occur close to the capital. We take this as clear evidence for a capital-centric geocoding pattern in GDELT, and this geographic bias, as compared to existing datasets on violence, is something that must be taken into account. As we have shown above, an analysis of the location of violence based on GDELT can lead to findings going squarely against many existing works, and this gives significant reason for caution when considering GDELT for micro-level studies of civil war.

Conclusions

Machine coding of event datasets can process a large number of reports in little time, and thus have some advantages over their human-coded counterparts. In this short article,

we have scrutinized the use of one of these event datasets—GDELT—for geo-spatial analysis at the subnational level. Previously, machine-coded event data have been used mostly for the study of international relations, and have been shown to be valid and reliable for this purpose. However, so far there have been few attempts to find out whether relying on machine coding produces data that is equally suitable for analyses at the subnational level. Our analysis reveals a considerable lack of agreement between human-coded and machine-coded data. We show that this is largely due to problems in geo-localization. While GDELT seems to track temporal ups and downs in violence as identified by the human-coded datasets, it places a disproportionately high number of events closer to a country's capital, undercounting events in more remote areas. For geo-spatial analyses of violence, this may be reason to worry. If we cannot be sure that the spatial accuracy of

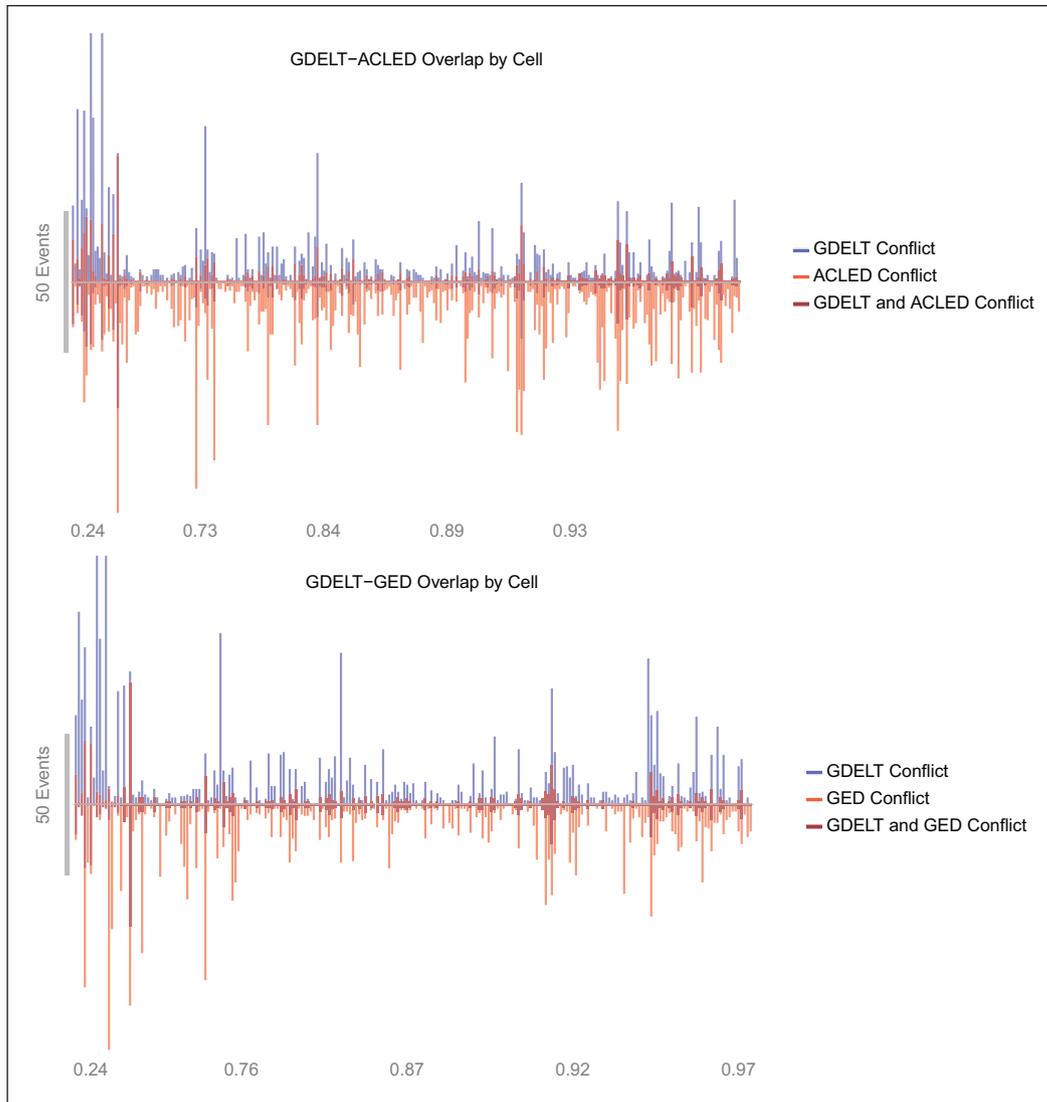


Figure 3. Grid-cell events over logged and normalized capital distance.

events is within reasonable limits (Weidmann, 2015), this can make machine-coded event datasets difficult to use for fine-grained analyses of the dynamics of violence on the ground.

We believe, however, that further work will be able to address these difficulties. As a first step towards more transparency in the machine-coding process, datasets should include pointers to the original articles used to code an event. This will enable more thorough validation studies not only of automatic geocoding, but of event content coding as well. Despite the fact that GDELТ is positively correlated with trends identified in other datasets, the overlap is still far from perfect. With traceback information in the dataset, users can go back to the original articles and find out whether, for example, “protest”

as coded by GDELТ corresponds to the type of “protest” they are interested in. Overall, our current recommendation is that GDELТ should be viewed at best as a complement, rather than a substitute, for existing event data. The high level of noise in the GDELТ data, coupled with the geographic accuracy issues we identify, suggests that using GDELТ instead of a more detailed hand-coded dataset may lead to distorted or incorrect inferences when explaining spatial dynamics in civil war violence. However, with more work on the refinement of automatic coding, this may well be where event data collection is (and should be) going.

Declaration of conflicting interest

The authors declare that there is no conflict of interest.

Table 2. Logistic regression results: event record disagreement by cell-month.

	GDELT-ACLED		GDELT-GED	
	GDELT = 1, ACLED = 0	GDELT = 0, ACLED = 1	GDELT = 1, GED = 0	GDELT = 0, GED = 1
	(1)	(2)	(3)	(4)
Distance to capital	-3.74*** (0.20)	1.08*** (0.16)	-3.86*** (0.18)	1.52*** (0.20)
Population	0.68*** (0.02)	0.66*** (0.01)	0.73*** (0.02)	0.71*** (0.02)
% Mountainous	0.69*** (0.09)	0.66*** (0.06)	0.63*** (0.09)	0.57*** (0.07)
Constant	-8.32*** (0.53)	-11.22*** (0.36)	-9.82*** (0.68)	-13.25*** (0.52)
Observations	547,812	547,812	547,812	547,812

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Country-level fixed effects not shown.

GDELT: Global Database of Events, Language and Tone; ACLED: Armed Conflict Location and Event Dataset; GED: Geo-referenced Event Dataset.

Table 3. Logistic regression results. Dependent variable: occurrence of violence in cell/month.

	Dependent variable:		
	ACLED Conflict	GED Conflict	GDELT Conflict
	(5)	(6)	(7)
Conflict ^(t-1)	2.60*** (0.04)	2.63*** (0.05)	2.86*** (0.05)
Spatial lag ^(t-1)	0.06*** (0.003)	0.13*** (0.01)	0.004 (0.004)
Distance to capital	0.14 (0.15)	1.18*** (0.18)	-2.41*** (0.18)
Population	0.60*** (0.01)	0.65*** (0.02)	0.68*** (0.02)
% Mountainous	0.44 (0.06)	0.31*** (0.07)	0.48*** (0.08)
Constant	-10.24*** (0.36)	-12.05*** (0.44)	-9.80*** (0.51)
Observations	547,812	547,812	547,812
Log likelihood	-26,121.86	-19,047.40	-14,483.13
Akaike inf. crit.	52,303.71	38,154.80	29,026.27

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Country-level fixed effects not shown.

GDELT: Global Database of Events, Language and Tone; ACLED: Armed Conflict Location and Event Dataset; GED: Geo-referenced Event Dataset.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Supplementary material

The online appendix is available at: <http://rap.sagepub.com/content/1/2/2053168014539924/suppl/DC1>

Notes

1. Some new data initiatives still in development, such as the Social, Political, and Economic Event Data Project (SPEED; <http://www.clinceneter.illinois.edu/data/speed/>) and the Event Data on Armed Conflict and Security Project (EDACS; <http://www.conflict-data.org/edacs/index.html>), attempt to utilize the speed of machine coding with direct human oversight. We eagerly look forward to learning more about the advantages of this hybrid approach as these projects mature.
2. See <http://eventdata.parusanalytics.com/data.html> for a more detailed description.
3. See http://www.lockheedmartin.com/us/products/W-ICEWS/W-ICEWS_overview.html for more information.
4. For an in-depth discussion of the differences and similarities between GED and ACLED, see Eck (2012).

References

- Arva J, et al. (2013) Improving forecasts of international events of interest. In: *annual meeting of the European Political Science Association*.
- Azar EE (1980) The Conflict and Peace Data Bank (COPDAB) project. *Journal of Conflict Resolution* 24(1): 143–152.
- Best RH, Carpino C and Crescenzi MJC (2013) An analysis of the TABARI coding system. *Conflict Management and Peace Science* 30(4): 335.
- Buhaug H, Cederman L-E and Rød JK (2008) Disaggregating ethnic conflict: a dyadic model of exclusion theory. *International Organization* 62(3): 531–551.
- Buhaug H and Rød JK (2006) Local determinants of African civil wars, 1970–2001. *Political Geography* 25(3): 315–335.
- Eck K (2012) In data we trust: a comparison of UCDP GED and ACLED conflict events datasets. *Cooperation and Conflict* 47(1): 124.
- Gerner D, et al. (2002) The creation of CAMEO (Conflict and Mediation Event Observations): An event data framework

- for a post cold war world. In: *annual meeting of the American Political Science Association*.
- Gerner DJ, et al. (1994) Machine coding of event data using regional and international sources. *International Studies Quarterly* 38(1): 91–119.
- Gleditsch KS and Weidmann NB (2012) Richardson in the information age: Geographic information systems and spatial data in international studies. *Annual Review of Political Science* 15(1): 461–481.
- Kalyvas SN (2004) The urban bias in research on civil wars. *Security Studies* 13(3): 160–190.
- Kalyvas SN (2012) Micro-level studies of violence in civil war: Refining and extending the control-collaboration model. *Terrorism and Political Violence* 24(4): 658–668.
- Leetaru K and Schrodt PA (2013) GDELT: Global data on events, language, and tone, 1979–2012. In: International Studies Association annual conference, San Francisco, CA.
- McClelland C (1976). *World Event/Interaction Survey*. Codebook. Ann Arbor.
- Raleigh C and Hegre H (2009) Population size, concentration and civil war: A geographically disaggregated analysis. *Political Geography* 28(4): 224–238.
- Raleigh C, Linke A, et al. (2010) Introducing ACLED: An armed conflict location and event dataset. *Journal of Peace Research* 47(5): 651–660.
- Salehyan I, et al. (2012) Social conflict in Africa: A new database. *International Interactions* 38(4): 503–511.
- Schrodt PA, Davis SG and Weddle JL (1994) Political science: KEDS - A program for the machine coding of event data. *Social Science Computer Review* 12(4): 561–587.
- Schrodt PA and Gerner DJ (1994) Validity assessment of a machine-coded event data set for the Middle East, 1982–92. *American Journal of Political Science* 38(3): 825–854.
- Sundberg R and Melander E (2013) Introducing the UCDP georeferenced event dataset. *Journal of Peace Research* 50(4): 523–532.
- Themnér L and Wallensteen P (2013) Armed conflicts, 1946–2012. *Journal of Peace Research* 50(4): 509–521.
- Tollefsen AF, Strand H and Buhaug H (2012) PRIO-GRID: A unified spatial data structure. *Journal of Peace Research* 49(2): 363–374.
- Ward MD, et al. (2013) *Comparing GDELT and ICEWS event data*. Working paper, Duke University. Available at: http://mdwardlab.com/sites/default/files/GDELTICEWS_0.pdf. (accessed 12 December 2013).
- Weidmann NB (2014) On the accuracy of media-based conflict event data. *Journal of Conflict Resolution* 58(5).
- Weidmann NB and Ward MD (2010) Predicting conflict in space and time. *Journal of Conflict Resolution* 54(6): 883–901.