

Community Life-Cycle Analysis in Multiplex Social Networks During Violent Events

Karolina Sliwa , Ema Kahr , and Mark Strembeck 

Institute for Complex Networks, Vienna University of Economics and Business (WU)

Vienna, Austria

{karolina.malgorzata.sliwa, ema.kahr, mark.strembeck}@wu.ac.at

Abstract—Tracking how online communities evolve during violent events (such as violent protests or riots) requires methods that can classify community transitions without a fixed membership-overlap threshold. We study three real-world social media datasets of violent events. Each event is represented as a multiplex network consisting of 14 layers: one interaction layer, eight emotion layers, and five moral foundation layers. In total, we analyze 27 combinations of these 14 layers. A threshold-free life-cycle framework is then applied to community transitions, allowing us to track communities hour by hour. Most communities continue from one hour to the next, with a median lineage of seven hours, but only 0.18% persist for a full day. Members mostly regroup into other communities rather than leaving the network, i.e. persistence mostly lies in the same members regrouping into successor communities rather than in the survival of a particular community. We also found that moral foundation layers regroup more users than emotion layers, but this difference does not alter the overall life-cycle dynamics. Thus, across the 27 combinations, changing the layers influences which users are grouped together, but has little effect on how long those communities last.

Index Terms—community detection, moral foundations, multiplex networks, social network

I. INTRODUCTION

In times of crisis, online social networks become the primary spaces for information seeking, emotional venting, and interpretation of the event as the situation unfolds [1]. While participating in an online discourse during crises, social media users not only generate content but also form online communities. Related studies have shown that such network communities evolve across the different phases of a crisis, revealing how different communities interpret and respond to what is happening in the physical world [2].

Violent crisis events do not only generate situational distress, but also coordinated moral responses [3]. Research has shown that communities form around a shared moral framing of an event and that morally framed emotional content especially spreads within such ideologically aligned groups [4], [5]. Users do not simply share information about what has happened, but also evaluate the event, invoke moral judgments, attribute responsibility, and demand official responses [3]. Online discussions surrounding violent events therefore become collective sense-making rather than channels for mere information sharing. These dynamics make violent events a particularly interesting setting for studying how online com-

munities form, evolve, and organize over time as a response to emotional and (often) politically charged incidents.

Incorporating message-related attributes during community detection aids in identifying communities organized by shared affective and moral framing rather than by interaction alone [6]. Studying these dynamics presents two methodological challenges that existing approaches rarely address simultaneously. The first challenge is how to incorporate these attributes into a community detection procedure. Attributes are often represented via corresponding layers in a multiplex setting [7], [8]. This opens up a challenge of which layers (or their combination) should enter community detection as this choice has a direct impact on the communities that will be detected [9].

The second methodological challenge is choosing a life-cycle tracking method that does not require a fixed membership-overlap threshold. Most existing approaches define community continuity based on the member overlap between consecutive snapshots. In these scenarios, the results are highly sensitive to the selected threshold parameter [10], [11]. Failla et al. [12] suggest addressing this issue by scoring each community transition against a set of reference transition patterns (called reference archetypes).

In this paper, we propose a framework that addresses the two challenges described above. We rely on a multiplex network, in which each layer represents one of the five moral foundations (care, loyalty, fairness, authority, and sanctity) and each of Plutchik’s eight primary emotions (joy, fear, sadness, surprise, trust, anticipation, anger, and disgust), as identified in the exchanged messages. We then apply the threshold-free life-cycle framework proposed in [12] and define a membership reassignment test against which the observed life-cycle distribution can be tested. We apply our pipeline to three real-world datasets: the 2017 Charlottesville “Unite the Right” rally, the 2017 Catalonia referendum riots, and the 2018 Gaza border protests surrounding the Jerusalem Embassy opening. We use this pipeline to answer three questions: (i) How often does each life-cycle event (e.g. emergence, continuation, death) occur, and does this exceed what would occur by chance? (ii) What impact does the choice of layers have on the detection of communities and life-cycle events? (iii) How does the community membership in each life-cycle event change over time?

The remainder of this paper is organized as follows. Section

II discusses related work. Section III describes our pipeline. Section IV reports the results of our analysis. We discuss the findings and limitations of this study in Section V. The paper concludes with Section VI and provides directions for future work.

II. RELATED WORK

Tracking Community Life-Cycles. Community detection identifies groups of users who interact more with each other than with the rest of the network [13]. In a single snapshot, this gives a useful description of who is grouped together at one particular moment, but it does not reveal how these groups change over time. Lu and Brelsford [14] tracked communities that emerged on Twitter after the 2011 earthquake in Japan and observed that crisis-affected users experienced greater community disruption than unaffected users, while content within surviving communities shifted entirely toward earthquake-related discussion. Wu et al. [15] tracked COVID-19 vaccine opinion communities on Japanese Twitter across quarterly snapshots over 2.5 years and found that community membership predicted sustained opinion shifts, with anti-vaccine community influence persisting longer than pro-vaccine effects. However, these studies work on daily to quarterly resolutions. The reorganization of communities during crisis events has rarely been studied on an hourly basis. Methodologically, this kind of analysis follows a snapshot-based design in which communities are detected independently at each time step and matched across steps by measuring membership overlap [16], [17]. Palla et al. [18] introduced the canonical formalization of community events. Greene et al. [10] proposed an event graph with six labeled transitions. These approaches require a fixed overlap threshold to decide whether two communities at adjacent steps are the same group, and results change substantially depending on the threshold [19], [11]. Failla et al. [12] address this by scoring each transition against a set of reference archetypes on a continuous scale, removing the need for a fixed threshold. We adopt their approach in the tracking stage of our pipeline. Removing the threshold, though, does not settle a second concern. Tardelli et al. [20] avoid the design entirely, since consecutive overlapping windows share members whether or not anything persists.

Multilayer Community Detection. Community detection can incorporate attributes before, during, or after detection [21]. Late-fusion methods detect communities from interaction structure first and apply attribute scores afterward. Uyheng and Carley [22] applied this strategy to daily Twitter snapshots over 75 days and characterized how hateful each community was using community-level hate speech scores after detection. In contrast, simultaneous-fusion methods incorporate user attributes directly into the partitioning step, recovering communities organized by shared attributes rather than by interaction alone. Ozer et al. [6] showed that user content and interaction structure are complementary in this respect, and also that word usage alone recovered better communities than word usage combined with hashtags and shared links. A single attribute layer, however, collapses dimensions that

may organize communities differently. A multiplex network is a specific type of multilayer network in which all layers share the same node set [23]. Applied to community detection, this allows each attribute to be represented as a separate layer rather than combined with others into one. Mucha et al. [7] formalized modularity optimization across multiplex layers. Suzuki and Tsugawa [9] showed empirically that individual layer choice substantially changes which communities are recovered in Twitter multiplex networks. Twitter multiplex networks are typically constructed from interaction layers such as retweet, reply, and mention edges. Kušen and Strembeck [24] extended this by constructing eight-layer emotion-annotated multiplex networks for different emergency events and showed that emotions shape interaction patterns that can be modeled as network motifs. Emotions motivate the choice of additional layers because they produce distinct interaction patterns during crisis events [25], and moral foundations capture a complementary dimension by predicting ideological community structure in online social networks [5], [4].

III. METHOD

This section outlines the six stages of our analysis pipeline: data collection and pre-processing (Section III-A), construction of time-ordered interaction snapshots (Section III-B), multiplex community detection (Section III-C), community tracking (Section III-D), life-cycle classification (Section III-E), and validation (Section III-F).

TABLE I

DATASET OVERVIEW. *Time frame* IS THE EXTRACTION PERIOD AND *Days* ITS LENGTH. *Tweets* IS THE NUMBER OF MESSAGES IN THE CLEANED CORPUS. *Nodes* AND *Edges* ARE THE DISTINCT ACCOUNTS AND DISTINCT DIRECTED ACCOUNT-TO-ACCOUNT PAIRS IN THE INTERACTION NETWORK ON WHICH COMMUNITIES ARE DETECTED.

Event	Time frame	Size			
		Days	Tweets	Nodes	Edges
Charlottesville	10–16 Aug 2017	7	150,784	88,424	126,639
Catalonia	28 Sep–16 Oct 2017	19	1,943,904	496,338	1,532,145
Jerusalem Embassy	11–27 May 2018	17	262,370	100,037	193,341

A. Data Collection & Preprocessing

For our study, we analyze three datasets obtained via Twitter’s Academic Search API and collected before the platform’s transition to X. All three datasets result from violent crisis events that drew sustained Twitter activity in the days following the event (see Table I). For each event, public tweets and their metadata were collected using event-related ¹ keywords and hashtags. For each message we retain the anonymized

¹Charlottesville: *charlottesville*. Catalonia: #RepublicofCatalonia, #LlibertatJordis, #CatalanReferendum, #cataloniareferendum, #CatalonianReferendum, #Catalonia, #Catalunya, #CATExit, #SPANISHPOLICE, #SpainShame, #SpainShameOnYou, #SpainIsNOTADemocracy, #SpainIsaDictatorship, #1oct catalonia, *Catalonia violence*, *Catalonia protest*, and related terms. Jerusalem Embassy: #USEmbassyJerusalem, #jerusalemembassy, #EmbassyMove, #IsraelEmbassy, #GazaMassacre, #greatmarchofreturn, #greatreturnmarch, #gazaprotests, #GazaReturnMarch, #Gaza, #Nakba, #Nakba70, #PalestineBleeds, #PalestinianLivesMatter, #GazaUnderAttack, #IStandWithIsrael, and related terms.

author ID, timestamp, and message text, along with interaction metadata covering replies, mentions, and reposts.

To identify emotions, we use TweetEval,² a transformer-based, multilabel emotion classifier fine-tuned on Twitter data [26], [27]. The model returns an independent probability score in $[0, 1]$ for each of Plutchik’s eight primary emotions [28]: anger, anticipation, disgust, fear, joy, sadness, surprise, and trust. A single message can carry multiple emotions at once. For identifying moral foundations, we use MFormer [29], a RoBERTa-based model trained on Twitter, news, and Reddit data, returning a probability in $[0, 1]$ for each of the five moral foundations: care, fairness, loyalty, authority, and sanctity [30]. Each of the thirteen classifier scores defines a separate attribute-layer in the multiplex network described in Section III-C.

B. Temporal Snapshot Construction

To capture how community structure evolves over the course of each event, we divide each dataset into overlapping temporal windows. Window length is a known methodological challenge. Short windows capture fine-grained dynamics but are sensitive to noise, while large windows smooth over rapid changes [16].

In our study, we use a sliding window of length $W = 6$ hours advancing in steps of $\Delta = 1$ hour. Within each window starting at time t , we collect all interactions (replies, mentions, reposts) whose timestamps fall in $[t, t + W)$. Repeated interactions between the same pair of users (u, v) are merged into a single directed edge with a structural weight equal to the number of interactions, treating replies, mentions, and reposts equivalently [9]. This follows standard snapshot-based practice [16]. We assigned the mean value for each attribute as edge weights per snapshot. Because consecutive windows overlap, the sequence is less sensitive to arbitrary boundary placement and produces smoother transitions between consecutive snapshots than non-overlapping windows [20].

We evaluated $W \in \{3, 6, 12, 24\}$ hours and selected $W = 6$ h for all three events, comparing candidate lengths on the fraction of nodes in the largest connected component of a snapshot [31]. Connectivity keeps rising with window length and never plateaus, so the measure alone does not identify an optimum and longer windows smooth the rapid changes we are trying to study. At six hours, the largest component already holds between 0.40 and 0.89 of the nodes, depending on the event. Six hours therefore secures most of the available connectivity at the finest resolution that our data supports (see Table II for the resulting snapshots per event).

C. Multiplex Network Construction and Community Detection

A multiplex network represents the same set of users through multiple independent layers simultaneously, where each layer encodes a different type of relationship [23], [7]. In our study, one layer encodes the structural interaction between users, while each additional layer encodes message exchanges

TABLE II
SNAPSHOT CONSTRUCTION PER EVENT. *Snapshots* IS THE NUMBER OF TIME WINDOWS. *Interactions*, *Nodes*, AND *Edges* ARE THE MEDIANS PER SNAPSHOT. ALL EVENTS USE A SIX-HOUR WINDOW ADVANCING IN ONE-HOUR STEPS. THE CHARLOTTESVILLE RIOT RUNS FOR THE SHORTEST PERIOD AND CONTRIBUTES THE FEWEST SNAPSHOTS, SO ITS SNAPSHOT-LEVEL ESTIMATES ARE CORRESPONDINGLY NOISIER.

Event	Snapshots	Interactions	Nodes	Edges
Charlottesville	141	5,516	5,100	4,933
Catalonia	456	8,580	5,602	7,613
Jerusalem	408	2,246	1,663	1,715

of a single type of emotion (e.g. anger or joy) or moral foundation dimension (e.g. loyalty or fairness). For example, two users who have interacted in a shared window are more strongly connected on the anger layer when the messages between them score high on anger. The attribute layers re-weight existing interaction edges and do not create new ones, so two users who never interacted are never connected by shared anger alone.

The structural layer enters as $\ln(1 + s_{uv})$, where s_{uv} is the number of interactions between an ordered pair of users, and the attribute layers enter as their raw scores in $[0, 1]$. All fourteen layers are defined over the same node set, so every user present in a snapshot appears in every layer. The resulting directed graphs are analyzed with *igraph* [32].

We analyze 27 layer combinations (see Figure 1). The choice of layers changes which communities are detected [9], so any finding based on a single combination could be an artifact of that choice.

	structure	emotions				moral foundations				combs.
C_{struct}	■									1
C_{emo}	■	■	■	■	■					8
C_{mf}	■					■	■	■	■	5
$C_{emo-all}$	■	■	■	■	■					1
C_{mf-all}	■					■	■	■	■	1
C_{full}	■	■	■	■	■	■	■	■	■	1
C_{emo-mf}	■	■	■	■	■	■	■	■	■	8
$C_{emo-only}$		■	■	■	■					1
$C_{mf-only}$						■	■	■	■	1
	■ layer present	■ one layer at a time	■ layer absent							

Fig. 1. **The twenty-seven layer combinations.** Each row refers to one or more combinations, and each cell represents a layer. Filled cells mark layers included in every combination of the row. Light-colored cells mark layers that enter one at a time, giving one combination per layer (e.g. C_{emo} contains structure + anger, structure + sadness, and so on). The *combs.* column shows the number of combinations per row. Rows are grouped into a structural baseline, single-attribute rows, attribute blocks, a cross-block row, and two structure-free ablations at the bottom of the table.

Communities are detected using the multiplex Leiden algorithm [33], [7]. Each layer contributes additively to the joint objective $Q = \sum_{\ell} Q_{\ell}$, with no interlayer coupling edges between the same node across layers. The standard modularity objective normalizes each layer by that layer’s own total edge weight, and we apply no further scaling, so that every layer enters with

²cardiffnlp/twitter-roberta-base-emotion-multilabel-latest

equal weight. Because the Leiden algorithm is stochastic, we run $R = 10$ independent optimizations per snapshot. Median pairwise Adjusted Rand Index (ARI)³ across runs lies between 0.863 and 0.973 for every one of the 27 combinations, so the great majority of node pairs are consistently co-clustered across runs. As the representative partition, we select the one with the highest mean ARI relative to all other runs, that is, the Rand-index medoid [34]. Communities with fewer than five members are excluded before the analysis, as groups of this size do not contain enough members to support reliable life-cycle classification [18].

D. Community Tracking

Following the threshold-free design in [12], we compute two measures for every ordered pair of communities (A_t, B_{t+1}) across consecutive snapshots. Forward inclusion $I_f(A, B) = |A \cap B|/|A|$ is the proportion of A 's members that reappear in B . Backward inclusion $I_b(A, B) = |A \cap B|/|B|$ is the proportion of B 's membership that originated in A . Both are continuous values in $[0, 1]$. For example, 100 users discuss an event together in the 6-hour window starting at hour 6. One hour later, 80 of them appear together in a new community, but that community now has 800 members because the story attracted a wider audience. Forward inclusion is 0.80 (the group largely held together) but backward inclusion is 0.10 (they are now a small minority within a much larger successor). This is an ancestor event, where original members moved together but were absorbed. Had the successor contained only those 80 users and a handful of newcomers, backward inclusion would also be high and the event would be categorized as a continuation instead. All pairwise values are recorded and passed to the life-cycle event classifier described next.

E. Life-Cycle Events

Forward and backward inclusion are computed for every pair of consecutive communities and summarized into three quantities per transition [12]. Unicity U measures whether members move together into one community in the next snapshot or split across several. Identity I measures whether those members make up most of that community's membership or arrive as a small minority within a much larger group. Outflow O is the fraction of members who do not appear in any community in the next snapshot. Each quantity is either high or low, which gives eight life-cycle events. For example, a CONTINUATION has high unicity, high identity and low outflow. A DISASSEMBLE is low on all three. Every event receives a score, computed from the three quantities, and each transition is assigned the life-cycle event with the highest one (argmax). Since we are interested in how communities evolve and dissolve, we apply the forward perspective throughout, classifying each source community A_t relative to its successors at $t + 1$ [12].

Four community evolutions are most relevant to our analysis. A CONTINUATION occurs when members stay together and

constitute the primary membership of one successor community. An ANCESTOR occurs when members move together but are absorbed into a much larger group where they are a minority (as in the 80-into-800 example above). A SPLIT occurs when members divide across several successor communities but still make up the majority of one of them, so one successor remains recognizably the original community. A DISASSEMBLE occurs when members scatter across multiple successor communities and are a minority in each of the successors, while still mostly reappearing somewhere in the network (this contrasts with the DEATH of a community, where most members do not reappear at all). In our study, the remaining types (shrink, reduction, death, dispersion) each account for under 2% of the transitions and are grouped as *other*. Building on this framework, we define a lineage as a chain of consecutive CONTINUATION events, following the dominant successor at each step. A lineage ends when another life-cycle event occurs. The successor community then starts a new lineage. For example, if a community continues for two hours and is then absorbed into a larger group, the original lineage ends when the absorption occurs. The larger group starts a new lineage.

Because consecutive windows advance in one-hour steps and overlap by $W - 1$ hours, we measure the real-time span of a lineage as $\text{end} - \text{start} + W$ in snapshots rather than as its length in steps multiplied by W , which would count the overlapping portion of each window repeatedly.

F. Validation

Evaluation is a documented problem in the community detection literature [13]. Modularity optimization in particular makes the evaluation harder because it can return many “good” communities that are structurally very different from each other [35]. Moreover, real-world networks have no ground truth to compare the detected communities against [36]. We therefore evaluate the statistical significance of detected communities via three validation tests over the structural baseline C_{struct} and the two attribute blocks $C_{\text{emo-all}}$ and $C_{\text{mf-all}}$ from Figure 1. Each test randomizes one component of the data (the edges, the attribute scores, or the community memberships) and holds the rest unchanged [16]. We repeat every test 1,000 times, record the median z-score and p-value, and consider the results statistically significant for $p < 0.001$:

- 1) Within-community connectivity: we test if users interact with others in their own community more often than expected by chance. Following [37], we randomly rewire the edges while preserving each user's in- and out-degree. We count how often two users who interact belong to the same community.
- 2) Moral foundation alignment in emotion communities: we test if users in the same community have similar moral foundation scores, even when these scores were not used to detect the communities. For each snapshot under $C_{\text{emo-all}}$, we keep the communities fixed and randomly reassign the five moral foundation scores across users [36]. We measure within-community homogeneity

³ARI measures partition agreement corrected for chance, ranging from 0 for chance-level similarity to 1 for identical partitions.

as $H = 1 - \sigma_{\text{within}}^2 / \sigma^2$,⁴ where σ_{within}^2 is the variance of the scores inside the communities, with each community contributing in proportion to its size, and σ^2 is the variance across all users.

- 3) Membership reassignment: in order to confirm that the life-cycles do not occur by chance, we keep the number of communities and their corresponding sizes fixed for each snapshot. We then randomly reassign users to communities and detect the life-cycle events (see Section III-E). We compare how often each life-cycle event occurs after reassignment with how often it occurs in the real-world network.

IV. RESULTS

A. Life-cycle across the three violent events

Most communities carry forward from one snapshot to the next. CONTINUATION accounts for between 66.9% and 76.3% of the 61,122 classified transitions, and DISASSEMBLE for between 5.5% and 10.0% (Table III). While CONTINUATION dominates in every violent event, DISASSEMBLE rates are low in the Jerusalem Embassy protests and high in the Catalonia riots.

TABLE III
DISTRIBUTION OF LIFE-CYCLE EVENTS PER VIOLENT EVENT, AS A PERCENTAGE OF ALL CLASSIFIED TRANSITIONS UNDER C_{full} . *Other* COLLECTS SHRINK, REDUCTION, DEATH, AND DISPERSION, EACH ACCOUNTING FOR UNDER 2%. N IS THE NUMBER OF CLASSIFIED TRANSITIONS IN EACH DATASET. CHARLOTTESVILLE ESTIMATES ARE NOISIER (SEE TABLE II)

Event	CONT.	ANC.	SPLIT	DISAS.	OTHER	N
Charlottesville	72.0	13.6	2.5	9.3	2.6	10,893
Catalonia	66.9	15.9	4.1	10.0	3.1	23,264
Jerusalem	76.3	11.2	4.9	5.6	2.0	26,965

The life-cycle events themselves, however, differ in how clearly they are identified. Each transition is assigned the highest-scoring of the eight life-cycle events (Section III-E). Failla et al. call this highest score *event typicality* [12]. A typicality of 1 means the transition matches a single life-cycle event exactly, and lower values mean it carries features of several. The median typicality of CONTINUATION is 1 in Charlottesville and the Jerusalem Embassy protests, and 0.82 in the Catalonia riots. For DISASSEMBLE and SPLIT it ranges from 0.36 to 0.46. The communities that carry forward therefore match a single life-cycle event, while those that break apart carry features of several.

Although CONTINUATION is both common and clearly identified at each step, few communities continue across many snapshots. A lineage follows a community through consecutive CONTINUATION events. Across the three violent events, we traced 23,091 lineages and they are short. Of these, 6,752 (29%) start at a community without any predecessor, and 16,339 (71%) start after another life-cycle event, such as an

absorption. The median lineage spans seven hours, one hourly step beyond the six-hour detection window. Of all lineages, 57% last beyond six hours, 3.7% beyond twelve hours, and 0.18% beyond a day. Persistent communities are therefore rare, as has also been reported for community structure in other real-world temporal networks [31].

B. Comparing layer combinations against the structural baseline

For each snapshot, we compare the communities found under a given combination against those found under C_{struct} and report the median ARI over snapshots (Table IV). A value of 1 means both combinations group the users in the same way. Where they differ, two users who reply to each other may be grouped together under C_{struct} but placed in separate communities under $C_{\text{mf-all}}$, because differences in their moral foundation scores weaken their connection.

TABLE IV
MEDIAN ARI BETWEEN EACH COMBINATION AND THE STRUCTURAL BASELINE, OVER SNAPSHOTS.

Combination	Layers	Charl.	Catal.	Jerus.
<i>Blocks, with structure</i>				
$C_{\text{emo-all}}$	all emotions	0.78	0.88	0.81
$C_{\text{mf-all}}$	all foundations	0.71	0.85	0.74
C_{full}	all thirteen attributes	0.76	0.88	0.78
<i>Structure-free ablations</i>				
$C_{\text{emo-only}}$	all emotions, no structure	0.77	0.87	0.80
$C_{\text{mf-only}}$	all foundations, no structure	0.68	0.83	0.69

Moral foundation layers regroup more users than emotion layers and the difference is not simply a matter of how many layers enter community detection. With all five moral foundations, the median ARI against the structural baseline ranges from 0.71 to 0.85, compared with 0.78 to 0.88 for all eight emotions. Fewer users therefore keep their structural grouping under the smaller block. The same holds when attributes enter individually, at 0.79 to 0.81 for the foundations against 0.80 to 0.84 for the emotions.

Across the five combinations and the structural baseline, the share of transitions classified as CONTINUATION varies by at most 1.0 percentage point within a violent event, and DISASSEMBLE by at most 1.5. The lineages are equally stable, with a median span of seven hours under every combination, while the number of lineages recovered varies by between 2.4% and 6.8%. Combinations therefore differ in which users they place together. They agree in the life-cycle events those communities go through and in how long they last.

C. Life-cycle events weighted by community membership

As described in Section III-E, every community receives a score for each of the eight life-cycle events, and the eight scores sum to one. Table III reports only the highest of them, counting a community of six and a community of six hundred members alike. Here we ask instead what share of members belong to communities undergoing each life-cycle event. We

⁴ H ranges from 0 (the scores inside communities vary as much as across the whole network) to 1 (all members of a community carry the same score).




Fig. 2. **Proportion of members in each life-cycle event per snapshot, shown as daily medians.** The dashed line marks the recorded time of the violent event. Charlottesville is omitted because data collection began shortly before the car attack, leaving no pre-event window.

keep all eight scores and multiply each by the community’s membership size, so that every community distributes its members across the eight life-cycle events. For example, a community of fifty users scoring 0.7 on CONTINUATION and 0.1 on ANCESTOR contributes 35 members to the former and 5 to the latter. Summed over a snapshot, the eight shares therefore account for the whole membership.

Before the referendum in Catalonia and the embassy opening in Jerusalem, the shares of members in each life-cycle event are relatively stable. The share of CONTINUATION then falls below its pre-event range 23 hours after both events. It reaches 37.1% in the Catalonia riots, against a pre-event median of 52.4%, and 30.1% in the Jerusalem Embassy protests, against 57.6%. Over the same period ANCESTOR and DISASSEMBLE both become more common, with DISASSEMBLE increasing most sharply in Jerusalem, from 7.0% to 19.0%. DEATH remains below 4% throughout, so members reorganize rather than leave the network.

The transitions also involve larger communities after the referendum and the embassy opening. Communities following the ANCESTOR evolution move into successors about $7\times$ their own size in the Catalonia riots and $6\times$ in the Jerusalem Embassy protests, against $3.5\times$ in both beforehand. Those following the DISASSEMBLE evolution have a median of 415 members after the referendum against 122 before, and 120 after the embassy opening against 42 before. Those following the CONTINUATION evolution keep the same size in both

periods. The growth is therefore limited to the communities that were absorbed or broke apart.

The change is concentrated in the largest communities. The shares return to their pre-event range after 73 hours in the Catalonia riots and 100 hours in the Jerusalem Embassy protests. Each event is therefore followed by about three to four days in which the largest communities reorganize.

D. Validation results

Within-community connectivity. Across the three violent events and all three combinations C_{struct} , $C_{\text{emo-all}}$ and $C_{\text{mf-all}}$, both users belong to the same community in between 84.8% and 96.2% of interactions. After rewiring this falls to between 4.6% and 9.2%. The difference holds in over 98% of snapshots in every event and combination ($p < 0.001$), with median z-scores between 167 and 322.⁵ The three combinations differ by less than one percentage point within each event. Users who interact are therefore far more likely to end up in the same community than if their interactions had been random.

Moral foundation alignment in emotion communities. The communities in this test are detected without considering moral foundations. We found that the moral foundation scores of users in the same community are still closer to each other than to the rest of the network. Within-community homogeneity ranges from 0.12 to 0.44 across the five foundations and

⁵The z-score gives the distance between the empirical value and the rewired values, counted in standard deviations. It is positive when the empirical value lies above the rewired values.

three events, against 0.01 to 0.06 after the scores are reassigned at random. The difference holds in over 90% of snapshots in every event ($p < 0.001$), with median z-scores between 22 and 236. Communities grouped by shared emotion are therefore also morally coherent.

Membership reassignment. After reassignment, CONTINUATION falls below 0.1% in all three violent events, against 66.9% to 76.3% in the empirical networks. The two other events that require a dominant successor move the same way. The ANCESTOR evolution falls to between 0.1% and 2.5%, against 11.2% to 15.9%, and the SPLIT evolution below 0.1%, against 2.5% to 4.9%. Only the DISASSEMBLE evolution moves in the opposite direction, rising to between 94.7% and 98.1% against 5.5% to 10.0% in the real-world network. This is expected. Every user still belongs to a community after the reassignment, but a community’s members are spread across many different ones, so they are unlikely to be grouped together again. None of the four life-cycle events came near their empirical values in any of the 1,000 repetitions ($p < 0.001$). The three layer combinations agree to within one percentage point. Members of a community therefore stay together far more often than random assignment at the same community sizes produces.

V. DISCUSSION

Violent events are followed by a period of about three to four days in which fewer communities continue into the next snapshot. Communities are absorbed into much larger successors or scatter their members across several groups, while the DEATH community evolution stays rare. The change is concentrated in the largest communities, and the shares return to their pre-event range after 73 hours in the Catalonia riots and 100 hours in the Jerusalem Embassy protests.

Crisis research offers a possible explanation. As news of a crisis event spreads, people who were not previously involved move toward the event to seek information about it [38]. On Twitter, activity tends to increase with the size and significance of the event [39]. Many of the users who join the discussion have not interacted before with one another [40]. The rise in ANCESTOR transitions after both events, with DEATH staying rare, is consistent with this reading. Communities were absorbed into much larger successors rather than ending.

Lu and Brelsford [14] studied network communities that formed on Twitter during the 2011 earthquake in Japan. They found that communities among affected users were more disrupted than among unaffected users, and that those users expanded the set of people they interacted with. In addition to observing the same disruption in our study, we provide additional insights into the reorganization of communities over their life-cycles. In other words, we show that members do not simply persist statically in their communities, but instead exhibit a highly dynamic membership turnover. In fact, only 0.18% of lineages last beyond a day, while the median spans seven hours. Emergent groups in offline disasters show the same short lifetimes, forming around an immediate need and dissolving once it has passed [41]. Our validation

test confirmed that communities carry forward far more often than they would if users were assigned to communities at random. The communities themselves are nonetheless short-lived. Persistence therefore lies in members regrouping into successor communities rather than in the survival of a particular community.

Across the 27 combinations, changing the layers influences which users are grouped together, but has little effect on how long those communities last. The relative proportions of the life-cycle events in our study remain stable, i.e., CONTINUATION is always dominant, followed by ANCESTOR, followed by DISASSEMBLE (see Table III).

A. Limitations

This study does not come without limitations. Three violent events are too few to generalize the observed patterns. The reassignment test moves users between communities but leaves the snapshots unchanged. Consecutive snapshots share five of their six hours. A community may therefore look like a CONTINUATION because its members stayed together, or because the same interactions were counted twice. Our test cannot tell these two apart. Non-overlapping windows would settle this, and we leave that to future work. The snapshot step is fixed at one hour without a sensitivity analysis. Every layer also enters the detection with equal weight. Both are choices rather than requirements. Varying the weight of the structural layer against the attribute layers would show how sensitive the detected communities are to that balance.

VI. CONCLUSION

The goal of this paper was to track how online communities emerge and evolve during violent events discussed on Twitter. We constructed a multiplex network model comprising an interaction layer, eight emotion layers, and five moral foundation layers. We applied a threshold-free life-cycle framework to identify communities and trace their evolution. This procedure was applied to three datasets, showing that persistence occurs through members regrouping into successor communities rather than through the continuation of the same community. Moreover, the three violent events in our study produced short-lived communities that reorganized rather than disappeared.

Three validation tests support these results. Membership reassignment produced under 0.1% CONTINUATION against 66.9% to 76.3% observed in our real-world datasets ($p < 0.001$). Users interact within their own community far more often than after the edges are rewired, and members of the same community carry similar moral foundation scores even where those scores did not enter community detection.

We also found that emotional and moral attributes change who belongs to communities more than they change how communities evolve. Thus, we found that emotion and moral framing influences community membership. Yet, the relative proportions of life-cycle events remain stable across layer combinations.

In our future work, we plan to apply the pipeline to a broader corpus of crisis events, including natural disasters and

sustained protests, to test if the same period of reorganization recurs across event types. Crisis research describes phases that online discourse moves through in the days after an event, from early expressions of solidarity and collective sense-making to more opinion-oriented content [3]. Aligning the life-cycle events with these phases would show if CONTINUATION, ANCESTOR, and DISASSEMBLE concentrate in certain phases of an event. Combining the life-cycle classification with the content of the exchanged messages could show what persisting and disassembling communities discuss. Linking lineage membership to individual opinion trajectories, following [15], would test if membership in a stable lineage is associated with later opinion change.

REFERENCES

- [1] E. Kušen and M. Strembeck, "Building blocks of communication networks in times of crises: Emotion-exchange motifs," *Computers in Human Behavior*, vol. 123, p. 106883, 2021.
- [2] C. Reuter and M.-A. Kaufhold, "Fifteen years of social media in emergencies: a retrospective review and future directions for crisis informatics," *Journal of Contingencies and Crisis Management*, vol. 26, no. 1, pp. 41–57, 2018.
- [3] D. Fischer-Preßler, C. Schwemmer, and K. Fischbach, "Collective sense-making in times of crisis: Connecting terror management theory with Twitter user reactions to the Berlin terrorist attack," *Computers in Human Behavior*, vol. 100, pp. 138–151, 2019.
- [4] M. Dehghani, K. Johnson, J. Hoover, E. Sagi, J. Garten, N. J. Parmar, S. Vaisey, R. Iliev, and J. Graham, "Purity homophily in social networks," *Journal of Experimental Psychology: General*, vol. 145, no. 3, p. 366, 2016.
- [5] W. J. Brady, J. A. Wills, J. T. Jost, J. A. Tucker, and J. J. Van Bavel, "Emotion shapes the diffusion of moralized content in social networks," *Proceedings of the National Academy of Sciences*, vol. 114, no. 28, pp. 7313–7318, 2017.
- [6] M. Ozer, N. Kim, and H. Davulcu, "Community detection in political Twitter networks using nonnegative matrix factorization methods," in *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 2016, pp. 81–88.
- [7] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, "Community structure in time-dependent, multiscale, and multiplex networks," *Science*, vol. 328, no. 5980, pp. 876–878, 2010.
- [8] K. Sliwa, E. Kahr, and M. Strembeck, "Using psycholinguistic features for multiplex community detection in temporal social networks," in *Companion Proceedings of the ACM Web Conference 2026*, ser. WWW Companion '26. ACM, 2026, pp. 1094–1099.
- [9] D. Suzuki and S. Tsugawa, "Evaluating community detection algorithms for multilayer networks: Effectiveness of link weights and link direction," *Complex Systems*, vol. 32, no. 2, 2023.
- [10] D. Greene, D. Doyle, and P. Cunningham, "Tracking the evolution of communities in dynamic social networks," in *2010 International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 2010, pp. 176–183.
- [11] M. Takaffoli, F. Sangi, J. Fagnan, and O. R. Zaiane, "Community evolution mining in dynamic social networks," *Procedia-Social and Behavioral Sciences*, vol. 22, pp. 49–58, 2011.
- [12] A. Failla, R. Cazabet, G. Rossetti, and S. Citraro, "Describing group evolution in temporal data using multi-faceted events," *Machine Learning*, vol. 113, no. 10, pp. 7591–7615, 2024.
- [13] S. Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, no. 3-5, pp. 75–174, 2010.
- [14] X. Lu and C. Brelsford, "Network structure and community evolution on Twitter: human behavior change in response to the 2011 Japanese earthquake and tsunami," *Scientific Reports*, vol. 4, no. 1, p. 6773, 2014.
- [15] Q. Wu, Y. Sano, H. Takayasu, S. Havlin, and M. Takayasu, "Twitter communities are associated with changing user's opinion towards COVID-19 vaccine in Japan," *Scientific Reports*, vol. 15, no. 1, p. 11716, 2025.
- [16] P. Holme and J. Saramäki, "Temporal networks," *Physics Reports*, vol. 519, no. 3, pp. 97–125, 2012.
- [17] G. Rossetti and R. Cazabet, "Community discovery in dynamic networks: a survey," *ACM Computing Surveys*, vol. 51, no. 2, pp. 1–37, 2018.
- [18] G. Palla, A.-L. Barabási, and T. Vicsek, "Quantifying social group evolution," *Nature*, vol. 446, no. 7136, pp. 664–667, 2007.
- [19] P. Bródka, S. Saganowski, and P. Kazienko, "GED: the method for group evolution discovery in social networks," *Social Network Analysis and Mining*, vol. 3, no. 1, pp. 1–14, 2013.
- [20] S. Tardelli, L. Nizzoli, M. Tesconi, M. Conti, P. Nakov, G. Da San Martino, and S. Cresci, "Temporal dynamics of coordinated online behavior: Stability, archetypes, and influence," *Proceedings of the National Academy of Sciences*, vol. 121, no. 20, p. e2307038121, 2024.
- [21] P. Chunaev, "Community detection in node-attributed social networks: a survey," *Computer Science Review*, vol. 37, p. 100286, 2020.
- [22] J. Uyheng and K. M. Carley, "Characterizing network dynamics of online hate communities around the COVID-19 pandemic," *Applied Network Science*, vol. 6, no. 1, p. 20, 2021.
- [23] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, "Multilayer networks," *Journal of Complex Networks*, vol. 2, no. 3, pp. 203–271, 2014.
- [24] E. Kušen and M. Strembeck, "An analysis of emotion-exchange motifs in multiplex networks during emergency events," *Applied Network Science*, vol. 4, no. 1, p. 8, 2019.
- [25] J. M. Salerno and L. C. Peter-Hagene, "The interactive effect of anger and disgust on moral outrage and judgments," *Psychological Science*, vol. 24, no. 10, pp. 2069–2078, 2013.
- [26] F. Barbieri, J. Camacho-Collados, L. E. Anke, and L. Neves, "TweetEval: Unified benchmark and comparative evaluation for tweet classification," in *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2020, pp. 1644–1650.
- [27] S. Mohammad, F. Bravo-Marquez, M. Salameh, and S. Kiritchenko, "SemEval-2018 Task 1: Affect in Tweets," in *Proceedings of the 12th International Workshop on Semantic Evaluation*, 2018, pp. 1–17.
- [28] R. Plutchik, "A psychoevolutionary theory of emotions," *Social Science Information*, vol. 21, no. 4, pp. 529–553, 1982.
- [29] T. D. Nguyen, Z. Chen, N. G. Carroll, A. Tran, C. Klein, and L. Xie, "Measuring moral dimensions in social media with MFormer," in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 18, 2024, pp. 1134–1147.
- [30] J. Graham, J. Haidt, S. Koleva, M. Motyl, R. Iyer, S. P. Wojcik, and P. H. Ditto, "Moral foundations theory: The pragmatic validity of moral pluralism," in *Advances in Experimental Social Psychology*. Elsevier, 2013, vol. 47, pp. 55–130.
- [31] P. Zhong, C. Ba, R. Mondragón, and R. Clegg, "Quantifying community evolution in temporal networks," *Scientific Reports*, vol. 15, no. 1, p. 45373, 2025.
- [32] G. Csardi and T. Nepusz, "The igraph software package for complex network research," *InterJournal, Complex Systems*, vol. 1695, pp. 1–9, 2006.
- [33] V. A. Traag, L. Waltman, and N. J. Van Eck, "From Louvain to Leiden: guaranteeing well-connected communities," *Scientific Reports*, vol. 9, no. 1, p. 5233, 2019.
- [34] A. Lancichinetti and S. Fortunato, "Consensus clustering in complex networks," *Scientific Reports*, vol. 2, no. 1, p. 336, 2012.
- [35] B. H. Good, Y.-A. De Montjoye, and A. Clauset, "Performance of modularity maximization in practical contexts," *Physical Review E*, vol. 81, no. 4, p. 046106, 2010.
- [36] L. Peel, D. B. Larremore, and A. Clauset, "The ground truth about metadata and community detection in networks," *Science Advances*, vol. 3, no. 5, p. e1602548, 2017.
- [37] S. Maslov and K. Sneppen, "Specificity and stability in topology of protein networks," *Science*, vol. 296, no. 5569, pp. 910–913, 2002.
- [38] C. E. Fritz and J. H. Mathewson, *Convergence behavior in disasters: A problem in social control*. National Academy of Sciences-National Research Council, 1957, no. 9.
- [39] A. L. Hughes and L. Palen, "Twitter adoption and use in mass convergence and emergency events," *International Journal of Emergency Management*, vol. 6, no. 3-4, pp. 248–260, 2009.
- [40] K. Starbird and L. Palen, "Volunteers: self-organizing by digital volunteers in times of crisis," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, pp. 1071–1080.
- [41] R. A. Stallings and E. L. Quarantelli, "Emergent citizen groups and emergency management," *Public Administration Review*, vol. 45, pp. 93–100, 1985.