

Computational timeline reconstruction of the stories surrounding Trump: Story turbulence, narrative control, and collective chronopathy

Peter Sheridan Dodds,^{1,2,*} Joshua R. Minot,¹ Michael V. Arnold,¹ Thayer Alshaabi,¹ Jane Lydia Adams,¹ Andrew J. Reagan,³ and Christopher M. Danforth^{1,2}

¹*Computational Story Lab, Vermont Complex Systems Center,
MassMutual Center of Excellence for Complex Systems and Data Science,
Vermont Advanced Computing Core, University of Vermont, Burlington, VT 05401.*

²*Department of Mathematics & Statistics, University of Vermont, Burlington, VT 05401.*

³*MassMutual Data Science, Boston, MA 02110*

(Dated: August 18, 2020)

Measuring the specific kind, temporal ordering, diversity, and turnover rate of stories surrounding any given subject is essential to developing a complete reckoning of that subject’s historical impact. Here, we use Twitter as a distributed news and opinion aggregation source to identify and track the dynamics of the dominant day-scale stories around Donald Trump, the 45th President of the United States. Working with a data set comprising around 20 billion 1-grams, we first compare each day’s 1-gram and 2-gram usage frequencies to those of a year before, to create day- and week-scale timelines for Trump stories for 2016 onwards. We measure Trump’s narrative control, the extent to which stories have been about Trump or put forward by Trump. We then quantify story turbulence and collective chronopathy—the rate at which a population’s stories for a subject seem to change over time. We show that 2017 was the most turbulent year for Trump, and that story generation slowed dramatically during the COVID-19 pandemic in 2020. Trump story turnover for 2 months during the COVID-19 pandemic was on par with that of 3 days in September 2017. Our methods may be applied to any well-discussed phenomenon, and have potential, in particular, to enable the computational aspects of journalism, history, and biography.

I. INTRODUCTION

What happened in the world last week? What about a year ago? As individuals, it can be difficult for us to freely recall and order in time—let alone make sense of—events that have occurred at scopes running from personal and day-to-day to global and historic [1–10]. One emblematic challenge for remembering story timelines is presented by the 45th US president Donald J. Trump, our interest here. Stories revolving around Trump have been abundant and diverse in nature. Consider, for example, being able to remember and then order stories involving: North Korea, Charlottesville, kneeling in the National Football League, Confederate statues, family separation, Stormy Daniels, Space Force, and the possible purchase of Greenland.

Added to these problems of memory is that people’s perception of the passing of time is subjective and complicated [11–18]. Days can seem like months (“this week dragged on forever”) or might seem to be over in a flash (“time flies”). Story-wise, periods of time can also range from being narratively simple (“it was the only story in town”) to complicated and hard to retell (“everything happened all at once”). At the population scale, major news stories may similarly arrive at slow and fast paces, and may be coherent or disconnected. As one example, within the space of around 15 minutes after 9 pm US Eastern Standard Time on March 11, 2020, Tom Hanks

and Rita Wilson announced that they had tested positive for COVID-19, the National Basketball Association put its season on hold indefinitely due to the COVID-19 pandemic, and Trump gave an Oval Office Address during which the Dow Jones Industrial Average futures dropped. And to help illustrate the potential disconnection of co-occurring stories within the realm of US politics, at the same time as the above events were unfolding, former US vice presidential candidate Sarah Palin was appearing on the popular Fox TV show “The Masked Singer” performing Sir Mix-A-Lot’s “Baby Got Back” in a bear costume.

Here, in order to quantify story turbulence around Trump—and the collective experience of story turbulence around Trump—we develop a data-driven, computational approach to constructing a timeline of stories surrounding any given subject, with high resolution in both time and narrative (see Data and Methods, Sec. II).

For data, we use Twitter as a vast, noisy, and distributed news and opinion aggregation service [19–23]. Beyond the centrality of Twitter to Trump’s communications [24–29], a key benefit of using Twitter as “text as data” [30–33] is that popularity of story is encoded and recorded through social amplification by retweets [34]. We show that Twitter is an effective source for our treatment though our methods may be applied broadly to any temporally ordered, text-rich data sources.

We define, create, and explore week-scale timelines of the most ‘narratively dominant’ 1-grams and 2-grams in tweets containing the word Trump (Sec. III A). We supply day-scale timelines as part of the paper’s Online Appendices (compstorylab.org/trumpstoryturbulence/).

* peter.dodds@uvm.edu

Having a computationally determined timeline of n -grams then allows us to operationalize and measure a range of features of story dynamics.

First, and in a way particular to Trump, we quantify narrative control: The extent to which n -grams being used in Trump-matching tweets are due to retweets of Trump’s own tweets (Sec. IIIB). We show Trump’s narrative control varies from effectively zero (e.g., ‘coronavirus’) to high (‘Crooked Hillary’).

Second, we compute, plot, and investigate the normalized usage frequency time series for n -grams that are narratively dominant in Trump-matching tweets for three or more days (Sec. IIIC). Along with their temporal ordering, these day-scale time series provide a rich representation of the shapes of Trump-related stories including shocks, decays, and resurgence. We also incorporate our narrative control measure into these time series visualizations.

Third, we measure story turbulence at the scale of months by comparing Zipf distributions for 1-gram usage frequencies between individual days across logarithmically increasing time scales (Sec. IIID). We are able to quantify and show, for example, that story turbulence for Trump was highest in the second half of 2017 and lowest during the lockdown period of the COVID-19 pandemic in the US.

Finally, we are then able to realize a numeric measurement of ‘collective chronopathy’ which we define to be how time seems to be passing (‘time-feel’) by the rate of story turnover (Sec. IIIE). We quantify how the past can seem more distant or closer to the present as we move through time, and that it may do so nonlinearly as a function of how far back we look in time. The course of the COVID-19 pandemic in the US rendered, for example, July 2020 as being closer to April 2020 than June.

II. DATA AND METHODS

We draw on a collection of around 10% of all tweets starting in 2008. We take all English language tweets [34, 35] matching the word ‘Trump’ from 2015/01/01 on. We ignore case and accept matches of ‘Trump’ at any location of a tweet (e.g., ‘@realDonaldTrump’ matches). We break these Trump-matching tweets into 1-grams and 2-grams, and create Zipf distributions at the day scale per Coordinated Universal Time (UTC). In previous work, we have assessed the popularity of major US political figures on Twitter, finding that the median usage rank of the word ‘trump’ across all of Twitter is less than 200, tantamount to that of basic English function words (e.g., ‘say’) [28]. Consequently, our resulting data set is considerable containing around 20 billion 1-grams.

Our main collection of Trump-matching tweets thereby includes tweets about Trump and tweets by Trump. Retweets and quote retweets are naturally accounted for as they are individually recorded in our database. We further filter n -grams for simple latin character words

including hashtags and handles.

To quantify the degree to which Trump might be in control of a story, we make a second database using the subset of n -grams found in retweets of Trump’s tweets (we exclude any quote tweet matter). We then generate day-scale Zipf distributions again, in the same format as for all Trump-matching tweets.

We perform two main analyses of these time series of Zipf distributions, treating 1-gram and 2-gram distributions separately. First, we determine which n -grams are most ‘narratively dominant’ by comparing each day’s Zipf distribution with the Zipf distribution of a year before using our allotaxonomic instrument of rank-turbulence divergence (RTD) [36].

We compare the top 10,000 n -grams and set the RTD parameter α at $1/4$, a reasonable fit for Twitter data. Tuning away from $1/4$ gives similar overall results as does the use of different kinds of probability-based divergences such as Jensen-Shannon divergence. We use a rank-based divergence because of the plain-spoken interpretability and general statistical robustness that ranks confer.

We find a year to be both a stable time gap with six months to two years also producing similar results. We then create day- and week-scale timelines of keyword n -grams, filtering out occasional hashtags and user handles.

Second, we use RTD to quantify the turbulence between Zipf distributions for any pair of dates. To systematically measure change in story over time, we compare each date’s Zipf distribution to that of an approximately logarithmically increasing sequence of days δ before.

We use the one instrument of rank-turbulence divergence throughout our paper for two reasons: (1) Analytic coherence (each section builds out from the previous ones), and (2) Facilitation of quantification. While we would expect sophisticated topic modeling approaches would also help elucidate stories, our goals are more expansive regarding the experience of time.

III. ANALYSIS AND DISCUSSION

A. Computational timeline generation for dominant Trump stories

In Figs. 1 and 2, we show computational timelines of the most narratively dominant 1-grams and 2-grams in Trump-matching tweets for each week running from the start of 2016 into 2020. We coarse-grain from days to weeks by finding n -grams with the highest overall rank-turbulence divergence (RTD) sum for each week.

Our computational timelines provide n -grams as keyword hooks, and immediately give a rich overall view of the major stories surrounding Trump. Broadly, the early major chapters run through the republican nomination process, the election, and inauguration. Reflected in names of individuals and entities, places, and processes, the timelines then move through a range of US

Week	2016	2017	2018	2019	2020
1. 01/01–01/07	Hillary 34.7	hacking 28.6	Bannon 2.2	shutdown 0.0	Iran 9.6
2. 01/08–01/14	Cruz 1.0	Meryl 5.0	Mueller 0.0	shutdown 0.0	Soleimani 5.9
3. 01/15–01/21	Cruz 10.7	inauguration 0.6	DACA 6.7	Pelosi 6.8	Parnas 0.0
4. 01/22–01/28	Cruz 10.6	inauguration 3.1	Mueller 0.0	Pelosi 2.6	Ukraine 5.5
5. 01/29–02/04	Cruz 11.2	ban 2.1	Mueller 0.0	border 0.0	impeachment 0.0
6. 02/05–02/11	Cruz 5.1	Bannon 0.0	memo 2.3	Whitaker 0.0	Vindman 2.5
7. 02/12–02/18	Cruz 6.9	Flynn 0.0	Mueller 0.0	emergency 0.0	Barr 2.2
8. 02/19–02/25	Rubio 3.8	Sweden 4.9	Parkland 0.3	Jussie 0.0	Bloomberg 6.3
9. 02/26–03/04	Rubio 9.2	Russia 6.4	Mueller 0.0	Cohen 3.7	coronavirus 0.0
10. 03/05–03/11	Cruz 1.0	Russian 4.8	Mueller 0.0	Nadler 13.7	coronavirus 0.0
11. 03/12–03/18	Cruz 5.7	tax 1.8	Mueller 2.2	emergency 1.6	coronavirus 0.0
12. 03/19–03/25	Arizona 16.8	Nunes 0.0	Mueller 2.2	Barr 0.0	coronavirus 0.0
13. 03/26–04/01	women 8.3	Russia 9.9	Stormy 0.0	Schiff 5.2	coronavirus 0.5
14. 04/02–04/08	Cruz 1.5	Russia 2.8	Mueller 0.0	returns 0.0	coronavirus 0.0
15. 04/09–04/15	Cruz 1.7	Syria 0.4	Mueller 2.0	Barr 2.4	coronavirus 0.0
16. 04/16–04/22	Cruz 10.5	Russia 0.5	Mueller 0.1	Barr 0.1	coronavirus 0.0
17. 04/23–04/29	Cruz 3.0	days 0.1	Kanye 8.0	Biden 6.0	coronavirus 0.0
18. 04/30–05/06	Indiana 11.5	Trumpcare 0.0	Mueller 0.0	Barr 0.0	coronavirus 0.0
19. 05/07–05/13	Ryan 2.5	Comey 2.8	Iran 6.6	Barr 0.0	coronavirus 0.0
20. 05/14–05/20	Bernie 25.3	Comey 1.0	ZTE 4.5	Barr 0.0	coronavirus 0.0
21. 05/21–05/27	Clinton 9.5	budget 0.0	Korea 18.2	Barr 0.0	pandemic 0.0
22. 05/28–06/03	Hillary 11.9	Kathy 4.4	Roseanne 4.0	USS 3.0	Minneapolis 32.1
23. 06/04–06/10	Clinton 11.1	Comey 0.8	pardon 0.0	Mexico 27.6	police 4.2
24. 06/11–06/17	Orlando 12.4	Mueller 0.0	Kim 4.1	foreign 2.0	Tulsa 4.5
25. 06/18–06/24	Hillary 23.9	Trumpcare 0.0	children 1.0	Iran 12.9	Tulsa 2.1
26. 06/25–07/01	Clinton 13.0	Russia 5.8	Justice 8.3	Moon 29.9	bounties 0.0
27. 07/02–07/08	Crooked 80.6	CNN 0.7	toddlers 0.0	parade 0.0	Rushmore 2.3
28. 07/09–07/15	Crooked 71.5	Russian 1.2	NATO 13.0	Epstein 0.0	coronavirus 0.0
29. 07/16–07/22	Pence 2.9	Mueller 0.0	Helsinki 3.1	racist 0.8	coronavirus 0.0
30. 07/23–07/29	DNC 6.1	Scouts 0.0	Cohen 0.0	Baltimore 13.6	Portland 11.8
31. 07/30–08/05	Khan 6.5	Mueller 0.0	LeBron 0.7	Baltimore 9.4	pandemic 0.0
32. 08/06–08/12	Crooked 55.2	Korea 5.8	Omarosa 0.4	Paso 7.6	pandemic 0.0
33. 08/13–08/19	Manafort 0.0	Charlottesville 1.5	Omarosa 9.5	Greenland 6.9	
34. 08/20–08/26	Clinton 7.6	Charlottesville 3.8	Cohen 2.7	Greenland 8.0	
35. 08/27–09/02	Crooked 57.4	Harvey 0.0	Ohr 14.0	Dorian 12.2	
36. 09/03–09/09	Bondi 0.0	DACA 2.4	Kavanaugh 2.1	Dorian 12.6	
37. 09/10–09/16	deplorable 0.0	ESPN 2.7	Puerto 7.5	flavored 0.0	
38. 09/17–09/23	Clinton 6.5	Kim 4.9	Kavanaugh 1.7	Ukraine 4.5	
39. 09/24–09/30	debate 4.9	Puerto 4.7	Kavanaugh 9.5	Ukraine 6.8	
40. 10/01–10/07	Pence 4.9	Puerto 2.1	Kavanaugh 6.8	Ukraine 5.1	
41. 10/08–10/14	sexual 0.3	Puerto 1.8	Kavanaugh 4.3	Kurds 8.2	
42. 10/15–10/21	rigged 10.1	Puerto 0.2	Saudi 5.3	Kurds 3.7	
43. 10/22–10/28	star 0.0	Mueller 0.0	caravan 0.0	impeachment 0.0	
44. 10/29–11/04	FBI 5.9	Mueller 0.0	caravan 0.0	impeachment 0.0	
45. 11/05–11/11	Clinton 0.9	Gillespie 12.0	Whitaker 6.2	Ukraine 6.2	
46. 11/12–11/18	Bannon 0.0	sexual 1.7	caravan 0.0	Ukraine 5.2	
47. 11/19–11/25	Hamilton 12.4	LaVar 21.3	Saudi 1.6	Ukraine 3.5	
48. 11/26–12/02	recount 0.0	Moore 0.0	Moscow 0.1	impeachment 3.1	
49. 12/03–12/09	Taiwan 7.8	Mueller 0.0	Cohen 2.1	impeachment 0.0	
50. 12/10–12/16	Russia 2.9	Mueller 0.0	Cohen 6.9	impeachment 0.0	
51. 12/17–12/23	inauguration 11.8	Mueller 0.0	wall 9.8	impeachment 1.4	
52. 12/24–12/31	inauguration 3.2	Mueller 0.0	wall 20.4	impeachment 7.6	

FIG. 1. **Computational timeline of narratively dominant 1-grams surrounding Trump-matching tweets at the week scale.** The timelines of 1-grams and 2-grams (Fig. 2) give an overall sense of story turbulence through turnover and recurrence. For each week, we show the 1-gram with the highest sum of rank-turbulence divergence (RTD) contributions [36] relative to the year before, based on comparisons of day-scale Zipf distributions for 1-grams in Trump-matching tweets. Each 1-gram must have been the most narratively dominant for at least 1 day of the week (i.e., highest RTD contribution). The numbers and pale pink bars represent Trump’s narrative control as measured by the percentage of an n -gram appearing in retweets of Trump’s own tweets during a given week. To align weeks across years, we assign the final 8 days to Week 52, and for each leap year we include February 29 as an 8th day in week 9. See the paper’s Online Appendices (compstorylab.org/trumpstoryturbulence/) for a complete list of the top 20 1-grams around Trump at the day scale. In constructing this table and the table for 2-grams in Fig. 2, we excluded hashtags, a small set of function word n -grams, and expected-to-be surprising n -grams such as ‘of the’ and, in 2017, ‘President Trump’.

Week	2016	2017	2018	2019	2020
1. 01/01–01/07	Hillary Clinton 32.7	plant in 85.1	Steve Bannon 5.7	the government 0.0	a war 6.6
2. 01/08–01/14	Trump rally 0.0	Meryl Streep 6.6	shithole countries 0.0	the border 1.0	impeachment trial 0.0
3. 01/15–01/21	Ted Cruz 26.0	Trump's inauguration 0.0	the government 1.4	Cohen to 0.0	impeachment trial 0.0
4. 01/22–01/28	Megyn Kelly 4.9	executive order 0.0	the FBI 5.6	the government 0.0	impeachment trial 0.0
5. 01/29–02/04	Ted Cruz 19.7	travel ban 1.6	the FBI 9.4	Ralph Northam 26.0	impeachment trial 0.0
6. 02/05–02/11	New Hampshire 19.5	travel ban 1.1	military parade 0.0	El Paso 4.7	Alexander Vindman 0.0
7. 02/12–02/18	Ted Cruz 15.7	Michael Flynn 0.0	school shooting 3.1	national emergency 0.0	Roger Stone 4.0
8. 02/19–02/25	Ted Cruz 30.1	Trump administration 0.0	the NRA 0.0	Jussie Smollett 0.0	Bernie Sanders 13.6
9. 02/26–03/04	vote for 4.4	to Russia 22.0	Hope Hicks 0.0	Michael Cohen 5.3	the coronavirus 0.0
10. 03/05–03/11	Ted Cruz 2.4	travel ban 0.0	Stormy Daniels 0.0	Tim Apple 0.0	the coronavirus 0.0
11. 03/12–03/18	Trump is 0.1	Meals on 0.0	Stormy Daniels 0.0	New Zealand 17.9	the coronavirus 0.0
12. 03/19–03/25	Lyin' Ted 66.2	health care 0.0	Cambridge Analytica 0.0	Mueller report 0.0	the coronavirus 0.0
13. 03/26–04/01	Trump is 0.0	Freedom Caucus 20.8	Stormy Daniels 0.0	Mueller report 0.0	the coronavirus 0.0
14. 04/02–04/08	Ted Cruz 3.9	Susan Rice 0.3	National Guard 0.0	tax returns 0.0	the coronavirus 0.0
15. 04/09–04/15	New York 19.3	in Syria 0.2	Michael Cohen 0.0	sanctuary cities 5.3	the coronavirus 0.0
16. 04/16–04/22	Ted Cruz 28.1	turnout for 0.0	Michael Cohen 2.4	Mueller report 0.0	the coronavirus 0.0
17. 04/23–04/29	Trump rally 0.0	tax plan 0.0	the Korean 0.0	Mueller report 0.0	the coronavirus 0.0
18. 04/30–05/06	Ted Cruz 5.5	health care 0.0	Stormy Daniels 0.0	Mueller report 0.0	treated worse 0.0
19. 05/07–05/13	Paul Ryan 2.0	James Comey 6.7	the Iran 9.0	tax returns 0.0	tested positive 0.0
20. 05/14–05/20	Hillary Clinton 26.5	Saudi Arabia 12.5	are animals 0.0	Lindsey Graham 0.0	the pandemic 0.0
21. 05/21–05/27	Hillary Clinton 24.8	Saudi Arabia 8.2	the FBI 23.3	Nancy Pelosi 12.5	a mask 6.3
22. 05/28–06/03	Trump University 3.4	Kathy Griffin 5.7	Samantha Bee 4.4	John McCain 0.0	photo op 0.0
23. 06/04–06/10	Hillary Clinton 18.6	James Comey 0.2	Justin Trudeau 8.5	with Mexico 39.2	Left Democrats 75.1
24. 06/11–06/17	Trump is 0.0	obstruction of 12.6	their parents 0.0	the FBI 8.5	in Tulsa 7.4
25. 06/18–06/24	Hillary Clinton 20.6	Karen Handel 16.6	their parents 3.4	need soap 0.0	in Tulsa 2.2
26. 06/25–07/01	Hillary Clinton 20.5	Fake News 37.6	Supreme Court 3.7	Jean Carroll 0.0	American soldiers 0.0
27. 07/02–07/08	Crooked Hillary 82.8	North Korea 28.6	Trump administration 0.0	Jeffrey Epstein 0.0	Mount Rushmore 3.9
28. 07/09–07/15	Crooked Hillary 73.3	Trump Jr 0.0	Supreme Court 7.9	Jeffrey Epstein 0.0	Roger Stone 0.0
29. 07/16–07/22	Mike Pence 6.8	Secret Service 0.0	in Helsinki 1.7	a racist 0.0	in Portland 0.0
30. 07/23–07/29	Crooked Hillary 79.6	Boy Scouts 0.0	Walk of 0.0	Elijah Cummings 27.2	in Portland 8.9
31. 07/30–08/05	Khizr Khan 0.0	Maxine Waters 0.0	enemy of 22.2	El Paso 11.1	the election 3.4
32. 08/06–08/12	Hillary Clinton 10.5	North Korea 5.7	Space Force 11.1	El Paso 7.7	Social Security 0.0
33. 08/13–08/19	Trump campaign 0.0	white supremacists 0.0	security clearance 0.0	New Hampshire 26.5	
34. 08/20–08/26	Hillary Clinton 19.1	Joe Arpaio 3.5	Michael Cohen 4.3	Prime Minister 28.7	
35. 08/27–09/02	Crooked Hillary 61.8	Hurricane Harvey 0.1	John McCain 0.2	Hurricane Dorian 9.6	
36. 09/03–09/09	in Detroit 0.0	to end 0.0	Brett Kavanaugh 7.6	the Taliban 3.0	
37. 09/10–09/16	tax returns 0.0	white supremacist 0.0	Puerto Rico 8.4	Dan Bishop 37.7	
38. 09/17–09/23	Trump Jr 0.0	North Korea 12.8	Blasey Ford 0.0	a foreign 6.4	
39. 09/24–09/30	Hillary Clinton 7.5	Puerto Rico 5.2	Brett Kavanaugh 15.7	impeachment inquiry 0.0	
40. 10/01–10/07	Mike Pence 8.9	Puerto Rico 2.6	Supreme Court 6.9	Adam Schiff 13.3	
41. 10/08–10/14	sexual assault 0.0	Puerto Rico 2.2	Kanye West 0.0	the Kurds 11.3	
42. 10/15–10/21	Hillary Clinton 19.9	families of 0.0	Saudi Arabia 6.6	the Kurds 3.8	
43. 10/22–10/28	Hillary Clinton 11.7	Myeshia Johnson 0.0	the bombs 0.0	World Series 0.0	
44. 10/29–11/04	Hillary Clinton 6.5	Twitter employee 0.0	birthright citizenship 0.0	the impeachment 0.0	
45. 11/05–11/11	Trump wins 0.0	mental health 0.0	Jim Acosta 0.0	pro quo 8.1	
46. 11/12–11/18	Steve Bannon 0.0	ban on 0.0	president who 0.0	impeachment inquiry 0.0	
47. 11/19–11/25	Mike Pence 24.3	Roy Moore 0.0	Saudi Arabia 2.5	quid pro 1.3	
48. 11/26–12/02	popular vote 17.4	Native American 0.1	Trump Tower 2.5	Hong Kong 0.0	
49. 12/03–12/09	Air Force 18.2	Roy Moore 3.5	campaign finance 0.0	to impeach 7.7	
50. 12/10–12/16	of State 7.6	of sexual 0.0	Michael Cohen 7.8	articles of 0.0	
51. 12/17–12/23	Electoral College 5.8	tax bill 0.0	the wall 13.7	Christianity Today 8.1	
52. 12/24–12/31	Trump next 0.0	the FBI 0.1	Border Security 70.6	the Senate 29.1	

FIG. 2. **Computational timeline of narratively dominant 2-grams surrounding Trump at the week scale.** The timeline's construction and the plot's details are analogous to that of Fig. 2. The n -grams in both timelines match ('coronavirus' and 'the coronavirus'), expand on each other ('bounties' and 'American troops'), or point to different stories ('Mueller' and 'Stormy Daniels'). See the paper's Online Appendices for a day-scale list of the top 20 most narratively dominant 2-grams in Trump-matching tweets.

and world events (North Korea, Charlottesville, Parkland, Iran, George Floyd’s murder, Portland); US policy and systems (travel ban, Space Force, southern border wall, Supreme Court); natural disasters (hurricanes in 2017, COVID-19 in 2020); and scandals (Russia, Stormy Daniels, Mueller, impeachment, and Taliban bounties for American soldiers).

The week-scale timelines for 1-grams and 2-grams variously directly agree (e.g., ‘Epstein’ and ‘Jeffrey Epstein’ in week 27 of 2019, and ‘coronavirus’ and ‘the coronavirus’ for 9 consecutive weeks in 2020); make connections (e.g., ‘Crooked’, ‘Hillary’, and ‘Crooked Hillary’ in 2016); or point to different stories (e.g., ‘Syria’ and ‘Trump Foundation’ in week 51 of 2018).

The n -gram timelines also provide an overall qualitative sense of story turbulence. In the lead up to the 2016 election, the stories around Trump largely concern his opponents, particularly Ted Cruz and Hillary Clinton. Near the election, we see an increase in story turbulence with ‘sexual assault’, ‘rigged’, and ‘FBI’. During Trump’s presidency, the timelines give a sense of stories becoming more enduring over time. After a tumultuous first nine months of 2017, we see the first stretch of four or more weeks being ruled by the same dominant story: ‘Puerto’/‘Puerto Rico’ and ‘Mueller’. In 2018, ‘Mueller’ then dominates for 12 of 17 weeks and then later ‘Kavanaugh’ leads for 5 of 6 weeks. In 2019, ‘Barr’ is number one for 4 of 5 weeks early on. Later, ‘Ukraine’ and ‘impeachment’ reflect the main story of the last four months of 2019, broken up only by ‘Kurds’. And in 2020, COVID-19 is the main story for 13 consecutive weeks for 1-grams starting at the end of February, pushed down by the murder of George Floyd and subsequent protests, before returning again as the major story of July.

Our intention with Figs. 1 and 2 is to give one page summaries of five years of stories around Trump. Below these dominant n -grams at the week scale are n -grams at the day scale representing many other stories. For example, Tab. I shows the the top 12 most narratively dominant 1-grams for three consecutive days in late June of 2020 where the story of Russia paying Taliban soldiers bounties for killing US soldiers rose abruptly to prominence. We see that by June 27, references to the COVID-19 pandemic had dropped down the list (‘coronavirus’ would return to the top 3 on July 1), as had indicators of other stories (‘Biden’, ‘lobster’, and ‘statues’ point to Trump’s 2020 challenger, a call by Trump to subsidize the US lobster industry, and the taking down of Confederate statues in the aftermath of George Floyd’s murder). The Russian bounties story would fall away within a week, with the next non-pandemic story appearing being the arrest of Ghislaine Maxwell on July 2, 2020. As part of the Online Appendices, we provide day-scale computational timelines for the top 20 most narratively dominant 1-grams and 2-grams.

Rank	2020/06/25	2020/06/26	2020/06/27
1	coronavirus	pandemic	bounties
2	Biden	bounties	bounty
3	pandemic	coronavirus	soldiers
4	testing	Biden	militants
5	lobster	militants	Russia
6	Matter	hiring	kill
7	statues	cases	briefed
8	virus	testing	Afghanistan
9	cases	virus	intelligence
10	ad	Statues	Taliban
11	fishing	bounty	troops
12	ChinaVirus	Care	pandemic

TABLE I. Top 12 most narratively dominant 1-grams for Trump-matching tweets for an example three day period. The sharp rise of the Russian bounties for Taliban soldiers to kill US troops, along with the fall of other stories, gives an example of a microscopic detail of story turbulence we aim to quantify macroscopically.

B. Narrative control

Because we are working with Twitter, some n -grams may become narratively dominant due to high overall use by many distinct individual users because of one highly retweeted individual’s tweet, and all degrees in between. Trump is the evident potential source of many retweets in our data set, and we introduce the concept of ‘narrative control’ as part of our analysis. In Figs. 1 and 2, the underlying pale pink bars and accompanying numbers indicate our measure of Trump’s narrative control, $C_{\tau,t}$. We quantify $C_{\tau,t}$ for any given n -gram τ on Twitter for some time period t as the percentage of the occurrences of τ due to retweets of Trump tweets: $C_{\tau,t} = 100 f_{\tau,t}^{\text{rt}} / f_{\tau,t}$.

We see that Trump’s week-scale levels of narrative control vary greatly over time. A few example highs, ordered by their date of first becoming narratively dominant, are ‘Crooked Hillary’ (82.6%), ‘Fake News’ (37.6%) ‘Border Security’ (70.6%), ‘Minneapolis’ (32.1%), and ‘Left Democrats’ (75.1%).

By contrast Trump’s narrative control has been low to non-existent for narratively dominant n -grams representing many stories. For ‘Mueller’ in 2017 and 2018, $C_{\tau,t}$ ranges from 0% to 2.2% (see however Sec. III C below). The names ‘Stormy Daniels’, ‘Jeffrey Epstein’, and ‘Ghislaine Maxwell’ all register $C_{\tau,t}=0$, as does ‘the bombs’ in reference to pipe bombs mailed to leading Democrats and journalists in October, 2018.

And most strikingly, Trump’s narrative control for ‘coronavirus’ has been effectively 0. Of course, and as for all n -grams, Trump may have used other terms to refer to them. In the case of ‘coronavirus’, he has used ‘corona virus’ (with optional capitalizations), ‘invisible enemy’, and a derogatory term. But our measurement here gets at the degree to which Trump has engaged with a specific




FIG. 3. Example time series of narratively dominant 1-grams and 2-grams for Trump-matching tweets. Each time series shows the relative usage rate for n -grams (blue curve), and Trump's level of narrative control (inverted orange curve with pale pink fill). See the Online Appendices (or the ancillary files section for the paper's arXiv version) for full time series for all n -grams which have been narratively dominant on at least three days. The background shading and vertical dashed lines indicate years and quarters.

n -gram being used on Twitter in connection with him.

Transitions in narrative control also stand out. In 2020, the most striking shifts at the week scale are 0% for ‘pandemic’ to 32.1% for Minneapolis, 0% for ‘photo op’ to 75.1% for ‘Left Democrats’, and 0% for ‘coronavirus’ to 11.8% for ‘Portland’. And the start of 2020 with the US assassination of Iranian general Soleimani jumps out amid the n -grams of the impeachment hearings (‘Iran’ 9.6%, ‘Soleimani’ 5.9%, ‘a war’ 6.6%).

C. Time series for dominant Trump stories

We now take our list of narratively dominant 1-grams and 2-grams and extract day-scale time series of normalized usage rate in Trump-matching tweets. We also generate the corresponding narrative control time series.

In Fig. 3, we show a selection of n -gram time series (blue lines) with Trump’s narrative control time series inverted (orange lines with pale pink fill). Below each horizontal axis, we annotate the n -gram along with the overall narrative control percentage for section of time series visible. We normalize each time series by its maximum.

We see a range of motifs common to sociotechnical time series: spikes (‘Saudi Arabia’, ‘Syria’, ‘Cambridge Analytica’), shocks with decays (‘Puerto Rico’, ‘Kurds’), sharp drop offs (‘Ted Cruz’), episodic bursts (‘Roger Stone’, ‘Syria’, ‘sexual assault’), and noise (‘Russia’, ‘Obama’) [37–39].

While an n -gram may be ranked as the most narratively dominant for some period of time, its usage rate may fluctuate during that time. A clear example is the time series for ‘coronavirus’ which begins with an initial spike followed by a shock, and then trends linearly down until the end of June, 2020.

As it is derived from a subset’s fraction of a whole, the narrative control time series can at most exactly mirror the overall time series. Three examples of time series with high and enduring narrative control percentages are ‘Crooked Hillary’, ‘Witch Hunt’, and ‘Fake News’ (58.6%, 62.0%, and 54.6%).

More subtly, we see that Trump’s narrative control for ‘Mueller’ increases over time, though remaining modest at 4.4%.

Many of the examples we show exhibit relatively flat narrative control time series. For 2-grams, Trump did not tweet about ‘Black Lives’ (Matter) in 2016, and only a fraction in 2020 (2.5%). And we see zero narrative control for ‘sexual assault’ whose time series correlates strongly with ‘women’.

In the Online Appendices and as an ancillary file on the arXiv, we include a PDF booklet of time series for all n -grams which have been narratively dominant in Trump tweets on at least three days. We extend the time series for all n -grams to run from 2016/01/01 on, and order them by the date they first achieve narrative dominance. Moving through this ordered sequence of 274 time series

gives another qualitative experience of how the stories around Trump have unfolded in time, placing them in temporal context.

D. Story turbulence

We have so far demonstrated that the n -grams of Trump-matching tweets track major events and narratives around Trump. We now move to using our comparisons of Zipf distributions for 1-grams via rank-turbulence divergence (RTD) to operationalize and quantify two aspects of story experience: (1) Story turbulence, the rate of story turnover surrounding Trump, and (2) Collective chronopathy, the feeling of how time passes at a population scale. Chronopathy is to be distinguished from the differently defined chronesthesia [11]. Here, we use the -pathy suffix primarily to mean ‘feel’ (as in ‘empathy’) though a secondary connotation of sickness proves serviceable too (as in ‘sociopathy’).

In Fig. 4, we present visualizations of the collective chronopathy for Trump-matching tweets. We focus first on the heat map of Fig. 4A, the core distillation of our measurement of the passing of time.

For each day in each month, we compare 1-gram Zipf distributions to the 1-gram Zipf distribution for δ days before using RTD with δ varying approximately logarithmically from 1 day to 1 year (vertical axis of Fig. 4A). We generate the heat map using the median value of RTD for each month and each δ . We employ a 10 point, dark-to-light perceptually uniform color map for increasing RTD which corresponds to the speeding up of story turnover. For each δ , we indicate the five slowest months by annotation with 1 being the slowest.

Overall, we see a general trend of story turnover slowing down—the plot thickens—with several marked exceptions, leading to an extraordinary slowing down across many time scales throughout the COVID-19 pandemic in 2020.

We describe the main features of the heat map, moving from left to right through time. As we would expect, the election year of 2016 shows fast turnover at the longer time scales of $\delta \sim$ six months to a year. This turnover carries through into 2017 as the shift to being president necessarily leads to different word usage around Trump. For the shorter time scales, January 2016 has fast turnover at all time scales—Trump’s narrative was rapidly changing as he was becoming the lead contender in the republican primaries. Some of the slowest turnovers for 1 day out to 14 days occur in 2016, notably in March and August.

Story turnover generally increases through 2017 at all time scales δ , with September, 2017 being the month with days least connected to all that had come before in the previous year. As suggested by Figs. 1 and 2, the first year of Trump’s presidency burned through story, with just some of the dominant narratives concerning Flynn, Comey, Russia, Mueller, North Korea, Charlottesville, DACA, NFL, and three major hurricanes.




FIG. 4. **A.** Chronopathic equivalency heat map: Each cell represents the ‘story distance’ for a given month and a given number of δ days before. We measure story distance as the median rank-turbulence divergence (RTD) between the Zipf distributions of 1-grams used in Trump-matching tweets for each day of a month and δ days before (we use RTD parameter $\alpha=1/4$). Lighter colors on the perceptually uniform color map correspond to higher levels of story turnover. Numbers indicate the slowest five months for each value of δ . After the story turbulence of the 2016 election year and especially the first year of Trump’s presidency, there has been a general slowing down in story turnover at all time scales (the ‘plot thickens’). By July 2020, the COVID-19 pandemic has induced record slowing down of story turnover around Trump at time scales up to 91 days, punctuated by the Black Lives Matter protests in response to George Floyd’s murder. **B.** Using an example anchor of May 2020 and $\delta=56$ days (white square in panel **A**), a plot of chronopathic equivalent values of δ across time. During Trump’s presidency, the same story turnover occurred as fast as every 2.9 days (September 2017), and only July 2020 has had slower turnover (110 days). Because story turbulence is nonlinear, using a different anchor month and δ (i.e., selecting a different cell in the heatmap) potentially gives a different chronopathic equivalency plot. **C.** Derived from the inverse of time series **B**, the chronopathic speed up factor relative to the anchor of 56-days-ago in May 2020. As for Fig. 3, shading and lines give guides for years and quarters.

In 2018 and 2019, we see some of the slowest turnover at longer time scales. Trump has the consistency of being president in the previous year. Two periods where story turnover for shorter time scales slow down in the second and third years of the Trump presidency are around the government shutdown (January 2019) and the impeachment hearings (last three months of 2019).

In March 2020, the COVID-19 pandemic brings story turnover to a functional halt at shorter time scales. April is especially slow with either the slowest or second slowest turnovers for time scales ranging from 1 day to 28 days.

The same time scales for May are all still slow (top 5) but now longer time scales show less story turnover. The murder of George Floyd on May 25 and subsequent Black Lives Matter protests then leads to a sharp transition in the heat map. Of course, the pandemic had not vanished but was for a few weeks at least more of a secondary story. But in July, we return to the pandemic being the dominant narrative leading to the slowest 56 and 91 day story turnovers. The text around Trump in July 2020 is closer to that of April 2020 than for any other three month comparison.

E. Collective chronopathy

We turn finally to determining what we will call chronopathically equivalent time scales. For any given month and time scale δ —any cell in the heat map of Fig. 4A—we can estimate time scales in other months with corresponding values of RTD.

As an example, the white-bordered square in Fig. 4A marks a time scale of $\delta=56$ days (8 weeks) in May, 2020. In Fig. 4B, we show the equivalent values of δ across all months. The rapidity of story turnover in 2017 is evident, and we see episodic slowdowns in 2018 and 2019, with two example causes in 2019 being the Mueller report and the impeachment hearings. In Fig. 4C, we show the inverse of Fig. 4B normalized to be interpretable as speed up ratios.

By our analysis, our measure of story turnover across 56 days looking back from May 2020 was equivalent to around 3 days in September 2017, a speed up of $\times 19$, and 110 days in July 2020, a slowing down by a factor of $\times 2$.

Story turbulence is complicated as the speed up or slow down of story can vary nonlinearly across different memory time scales. Choosing a different month and time scale δ in Fig. 4A would change the forms of Fig. 4B and 4C. For example, in choosing 28 days instead of 56 in May 2020, we would find equivalent time scales of 2.1 and 16 in September 2017 and July 2020, with speed up factors of $\times 13$ and $\times 1.7$. So by this shorter time scale benchmark of about a month in May, story seems to have moved faster in July because of the narrative dominance of protests in June.

IV. CONCLUDING REMARKS

We have shown that interpretable, high-level timeline summaries of historical events can be derived from Twitter. Our process, which we have worked to make as rigorous as possible, is necessarily computational: In our approximately 10% sample of all of Twitter, there were from 2015/01/01 through to 2020/08/12 around 20 billion distinct 1-grams in Trump-matching tweets, with 1.5 billion of those being contained in retweets of Trump’s tweets.

Because we have shown that the narratively dominant n -grams our method find are historically sensible, we are then able to defensibly quantify narrative control, story turnover, and collective chronopathy in the context of Trump.

We observe that our focus on n -grams has resulted in timelines that are largely descriptive of major events and stories through the surfacing of the names of people, places, institutions, processes, and social phenomena. While the collective construction of what matters cannot be said to be objective—what matters socially is what matters socially [40, 41]—we find that our computationally derived timelines are not laden with overt

opinion or framing with the notable exception of certain phrases due to Trump himself.

Our study of stories around Trump suggests much future possible research.

Broader explorations of Trump’s (and others’) narrative control should be possible. We have limited ourselves here to narratively dominant n -grams and only within the same time frame. Studies across different news outlets and social media platforms could explore the extent to which Trump’s narratives are driven by external stories and vice versa, as well as their persistence in time, and certain modeling and prediction may be possible.

We have indexed stories by keywords of 1-grams and 2-grams. Connecting these n -grams to Wikipedia would possibly allow for the automatic generation of timelines augmented by links to (or summaries of) Wikipedia entries. Disconnects between Twitter and Wikipedia (and other texts) would be of interest to explore as well.

Computationally generating a taxonomy of story type would be a natural way to improve upon our work. How have the kinds of stories around Trump changed over time, which ones have persisted, and which ones have been more likely to last only a few days?

We have presented narratively dominant n -grams for the time scales of weeks, while reserving the top 20 at the day scale for the paper’s Online Appendices. While this is reasonable for a study going across now five years, for certain major events, shorter time scales of hours or even minutes may be better suited as the temporal units of analysis.

The relationship between Trump’s favorability polls with story turbulence and story kind could also be examined. As a rough observation, we note that the speeding up and then slowing down of story turnover around Trump through 2017 and into 2018 appears to be anti-correlated with Trump’s approval ratings during that same period. Having established a story taxonomy would be important for moving in this direction.

While our specific instrumental focus has involved the use of Twitter to analyze stories surrounding Trump, in doing so we have laid out a general, structured approach to quantifying and exploring story turbulence for any well-defined topic in temporally ordered, large-scale text (e.g., news outlets, online forums, and Reddit). Wherever Zipf distributions can be derived, our methods will allow for the computational construction of n -gram-anchored story timelines along with measures of story distance, turbulence, and chronopathy. For such extensions, we caution that the popularity of text must be incorporated in some fashion through measures of readership, shares, etc. We also observe that stronger variations in findings would potentially come from choosing different text sources rather than through reasonable adjustments of the methods of analysis we have laid out here.

We believe that our techniques will be of value for computationally augmented journalism and the computational social sciences [42–45]. News reporters and historians

in particular are faced with having to process ever more data, text, and media, and we see our work here as contributing to a far broader effort to develop ways to provide powerful computational support for the vital work performed by domain-knowledge experts.

ACKNOWLEDGMENTS

The authors are grateful for the computing resources provided by the Vermont Advanced Computing Core

which was supported in part by NSF award No. OAC-1827314, financial support from the Massachusetts Mutual Life Insurance Company and Google Open Source under the Open-Source Complex Ecosystems And Networks (OCEAN) project.

-
- [1] B. Gunter, *Poor reception: Misunderstanding and forgetting broadcast news* (Routledge, 1987).
 - [2] V. Price and J. Zaller, *Public opinion quarterly* **57**, 133 (1993).
 - [3] U. Neisser, R. Fivush, *et al.*, *The remembering self: Construction and accuracy in the self-narrative*, 6 (Cambridge University Press, 1994).
 - [4] A. Lang, *Journal of Communication* **50**, 46 (2000).
 - [5] E. Tulving, *Annual review of psychology* **53**, 1 (2002).
 - [6] D. L. Schacter, *Searching for memory: The brain, the mind, and the past* (Basic Books, 2008).
 - [7] R. Fivush, *Annual Review of Psychology* **62**, 559 (2011).
 - [8] W. Ocasio, M. Mauskopf, and C. W. Steele, *Academy of Management Review* **41**, 676 (2016).
 - [9] E. Garibaldi and P. Garibaldi, “Ordering history through the timeline,” (2017), cEPR Discussion Paper No. DP12508.
 - [10] E. F. Loftus, *Annual Review of Law and Social Science* **14**, 1 (2018).
 - [11] E. Tulving, in *Principles of frontal lobe function*, edited by D. T. Stuss and R. T. Knight (Oxford University Press, 2002) pp. 311–325.
 - [12] J. Edy, *Troubled pasts: News and the collective memory of social unrest* (Temple University Press, 2006).
 - [13] S. Droit-Volet and W. H. Meck, *Trends in cognitive sciences* **11**, 504 (2007).
 - [14] A. M. Sackett, T. Meyvis, L. D. Nelson, B. A. Converse, and A. L. Sackett, *Psychological Science* **21**, 111 (2010).
 - [15] I. Phillips, *European Journal of Philosophy* **18**, 176 (2010).
 - [16] S. Grondin, *Attention, Perception, & Psychophysics* **72**, 561 (2010).
 - [17] M. J. Allman and W. H. Meck, *Brain* **135**, 656 (2012).
 - [18] M. Rudd, K. D. Vohs, and J. Aaker, *Psychological science* **23**, 1130 (2012).
 - [19] P. S. Dodds, K. D. Harris, I. M. Kloumann, C. A. Bliss, and C. M. Danforth, *PLoS ONE* **6**, e26752 (2011).
 - [20] R. J. Gallagher, E. Stowell, A. G. Parker, and B. Foucault Welles, *Proceedings of the ACM on Human-Computer Interaction* **3**, 1 (2019).
 - [21] S. J. Jackson, M. Bailey, and B. Foucault Welles, *#HashtagActivism: Networks of Race and Gender Justice* (MIT Press, 2020).
 - [22] T. R. Tangherlini, S. Shahsavari, B. Shahbazi, E. Ebrahimzadeh, and V. Roychowdhury, *PLoS ONE* **15**, e0233879 (2020).
 - [23] J. Allen, B. Howland, M. Mobius, D. Rothschild, and D. J. Watts, *Science Advances* **6**, eaay3539 (2020).
 - [24] J. C. Lee and K. Quealy, *The New York Times* (2016).
 - [25] B. L. Ott, *Critical Studies in Media Communication* **34**, 59 (2017).
 - [26] A. Bovet, F. Morone, and H. A. Makse, *Scientific reports* **8**, 1 (2018).
 - [27] B. L. Ott and G. Dickinson, *The Twitter presidency: Donald J. Trump and the politics of White rage* (Routledge, 2019).
 - [28] P. S. Dodds, J. R. Minot, M. V. Arnold, T. Alshaabi, J. L. Adams, D. R. Dewhurst, A. J. Reagan, and C. M. Danforth, “Fame and Ultrafame: Measuring and comparing daily levels of ‘being talked about’ for United States’ presidents, their rivals, God, countries, and K-pop,” (2019), available online at <https://arxiv.org/abs/1910.00149>.
 - [29] Y. Ouyang and R. W. Waterman, in *Trump, Twitter, and the American Democracy* (Springer, 2020) pp. 89–129.
 - [30] J. Grimmer and B. M. Stewart, *Political analysis* **21**, 267 (2013).
 - [31] M. Gentzkow, B. Kelly, and M. Taddy, *Journal of Economic Literature* **57**, 535 (2019).
 - [32] D. Riffe, S. Lacy, F. Fico, and B. Watson, *Analyzing media messages: Using quantitative content analysis in research* (Routledge, 2019).
 - [33] R. L. Boyd, K. G. Blackburn, and J. W. Pennebaker, *Science Advances* **6**, eaba2196 (2020).
 - [34] T. Alshaabi, D. R. Dewhurst, J. R. Minot, M. V. Arnold, J. L. Adams, C. M. Danforth, and P. S. Dodds, “The growing amplification of social media: Measuring temporal and social contagion dynamics for over 150 languages on twitter for 2009–2020,” (2020), available online at <http://arxiv.org/abs/2003.03667>.
 - [35] T. Alshaabi, J. L. Adams, M. V. Arnold, J. R. Minot, D. R. Dewhurst, A. J. Reagan, C. M. Danforth, and P. S. Dodds, “Storywrangler: A massive exploratorium for sociolinguistic, cultural, socioeconomic, and political timelines using Twitter,”.
 - [36] P. S. Dodds, J. R. Minot, M. V. Arnold, T. Alshaabi, J. L. Adams, D. R. Dewhurst, T. J. Gray, M. R. Frank, A. J. Reagan, and C. M. Danforth, “Allotaxonomy and rank-turbulence divergence: A universal instrument for comparing complex systems,” (2020), available online at <https://arxiv.org/abs/2002.09770>.
 - [37] R. Crane and D. Sornette, *Prod. Nat. Acad. Sci.* **105**, 15649 (2008).

- [38] D. R. Dewhurst, T. Alshaabi, D. Kiley, M. V. Arnold, J. R. Minot, C. M. Danforth, and P. S. Dodds, EPJ Data Science **9**, 3 (2020).
- [39] M. De Domenico and E. G. Altmann, Scientific reports **10**, 1 (2020).
- [40] P. L. Berger and T. Luckmann, *The social construction of reality: A treatise in the sociology of knowledge* (Penguin UK, 1991).
- [41] A. Sandu, *Social construction of reality as communicative action* (Cambridge Scholars Publishing, 2016).
- [42] D. Lazer, A. S. Pentland, L. Adamic, S. Aral, A. L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, *et al.*, Science (New York, NY) **323**, 721 (2009).
- [43] R. Conte, N. Gilbert, G. Bonelli, C. Cioffi-Revilla, G. Deffuant, J. Kertesz, V. Loreto, S. Moat, J.-P. Nadal, A. Sanchez, *et al.*, The European Physical Journal Special Topics **214**, 325 (2012).
- [44] T. Flew, C. Spurgeon, A. Daniel, and A. Swift, Journalism Practice **6**, 157 (2012).
- [45] M. Coddington, Digital journalism **3**, 331 (2015).