Modeling the Influence of Narratives on Collective Behavior

Case Study: Using social media to predict the outbreak of violence in the 2011 London Riots.

James Houghton[1], Michael Siegel[2], Daniel Goldsmith[3]

[1] Research Associate, Sloan School of Management, Massachusetts Institute of Technology
[2] Principal Research Scientist, Sloan School of Management, Massachusetts Institute of Technology
[3] Principal Consultant, PA Consulting Group

77 Massachusetts Avenue, E70-676
Cambridge MA 02139

281.728.6848

Houghton@mit.edu
MSiegel@mit.edu
Daniel.Goldsmith@paconsulting.com

Published in The Proceedings of The 31st International Conference of the System Dynamics Society, Cambridge, Massachusetts USA, July 21 – July 25, 2013

Abstract

This paper considers the problem of understanding the influences of narratives or stories on individual and group behavior. Narrative theory describes how stories help people make sense of the world, and is being used to explain behavior in domains such as security, health care, and consumer behavior. We are interested in using narrative theory to develop better predictions of behavior and have developed a multi-methodology approach to combine narrative influence with system dynamics modeling of group behavior. Our model quantifies how individuals use narratives to understand current events and make decisions. We model the time-varying strength of cultural narratives as a degree of belief in the narrative's explanatory power, updated heuristically in response to observations about similarity between cultural narratives and current events. We use Twitter posts to measure narrative-significant observations in the real world. Using this approach, we investigate a case study of the violent riots in London in 2011 and demonstrate how relevant narratives can be identified, monitored, and included in behavior models to predict violent activity.

Introduction

The study of narratives is providing exciting new opportunities to better understand the role stories play in human psychology and sociology and help explain human behavior in many domains. These opportunities are especially notable in security contexts, where the inclusion of narrative impacts can help explain the factors that contribute to radicalization, violent social mobilization, and insurgency, among other challenges. For example, in 2007, the US House of Representatives Committee on Armed Services held a hearing to discuss the “Battle of Ideas” in the war on terrorism, and the task of “winning [our adversaries’] hearts and minds”. The hearing underscored the fact that often times violent extremism is driven by a specific underlying worldview and that it is important to understand the narratives which make up this worldview. (House Armed Services Committee, 2007) Just as conventional conflict requires an understanding of the contested physical terrain, a winning a battle of ideas will require understanding of the narrative landscape. .

This paper considers the problem of understanding the narratives of inter-group conflict. It develops a mixed methodology approach that combines system dynamics models with social media data in order to better explain and predict outcomes of inter-group conflicts, particularly the escalation of group behavior towards violence. In addition to drawing upon previous system dynamics modeling research, our work incorporates emerging research on narrative networks that give insight on how stories exert powerful influences on human thoughts and behavior.

The theory of ‘Narrative’ influence on collective behavior is gaining traction in social psychological circles as an alternate to methods assuming rational economic actors; and with the potential to “serve as a barometer for public views” and attitudes. (Monitor 360, 2012). Our work formalizes critical dynamics relating to the way stories exert powerful influences on group behavior in a system dynamics model and utilizes social media data for parameterization. We believe that incorporating narrative perspectives in formal simulation models, and drawing upon social data-sets, is an exciting new opportunity for system dynamics research to help develop responses to better predict and mitigate inter-group conflicts.

The paper is organized as follows. The first section provides a brief overview of the social psychology research into narrative influence that provides the theoretical basis for the quantitative model. The second section gives a detailed formulation of the system dynamics model. Subsequently the paper discusses the use of social media streams as a proxy for important model components. The third section is a case study in which the methods developed in the paper are applied to understand an incident of rioting in London in 2011. Finally, the paper concludes with discussion of lessons learned from the case study and their implications for theory and practice.

Brief overview of narrative theory

Narrative theory, broadly defined, is concerned with how narratives influence cognition and behavior. Over the last several decades, researchers have investigated the origin, transformation, and behavioral impact of narratives. A central premise of narrative theory is that human beings understand their lives in terms of stories. While narratives can come from many sources, most of the narratives an individual maintains “derive from public narratives conveyed… by other people; things… intentionally taught” and come from a pool of cultural stories which cover “religion, politics, popular culture, regional identity, racial and ethnic identity, attitudes toward other members of the culture and toward minority members, attitudes toward outsiders” (Beach, 2010, p. 30).

Narratives have an important role in consolidating memories, shaping emotions, and providing group distinctions, among other impacts. Other authors have described the impact of narrative on behavior, by arguing that: "humans use… narratives to understand the world and their place in it, to frame events, and to plan and justify their actions" (Halverson, Goodall, & Corman, 2011, p. 181) , and that narratives are how humans learn " what is right and wrong and why, how and why things happen, how to perform tasks to produce desired effects… who we are, where we fit in the scheme of things, and what our rights and responsibilities are." (Beach, 2010, p. 29) A well-cited phrase that describes the role narratives play in behavior—"I can only answer the question 'What am I to do?' if I can answer the prior question 'Of what story or stories do I find myself a part?'" (MacIntyre, 1981)—helps capture our central interest in using narratives to better understand human behavior.

Identifying relevant narratives

The first step in utilizing narrative theory in our simulation models is to identify the narratives relevant for a target population If we can understand the narratives present in a population, and measure their shifting influences on that population's attitudes and decision making, we gain insight into that population's potential collective behavior and can prepare an informed response. The task of identifying narratives present in a population has been addressed by several sociologists in a limited number of case studies. For example: researchers have investigated narratives of Islamic Extremism (Halverson, Goodall, & Corman, 2011); narratives relevant to the Palestinian/Israeli Conflict (Hammack, 2009); narratives related to the 2012 US presidential election (Rosenstiel & Jurkowitz, 2012); and narratives with influence on the behavior of spouses in long distance relationships (Bergen, 2010). There are many narrative researchers expanding the range of population narratives, and we anticipate that as research progresses there will be increasingly better mappings between populations and narratives.

We believe that investigating historical precedent, news articles, cultural literature, and social media can serve as input for us to achieve a rough understanding of the narratives related to a problem of interest.

Building a quantitative model of narrative influence

Our goal is to increase the predictive capacity of narrative theory, and we therefore need to move beyond only mapping narratives to populations.   Our goal is to establish a quantitative method to determine the strength and influence of narratives, and ultimately their impact on behavior. This section describes the key pieces of the model and our approach to model development that will allow us to start evaluating real world cases.

A metric of narrative strength

Cultures maintain multiple narratives in order to "assure a closer correspondence between the narrative and the underlying facts in any given case." (Skeel, 2009, p. 1203) Narratives with the most fidelity then have the capacity to direct behavior. Fisher explains that narratives which "represent accurate assertions about social reality… constitute good reasons for belief or action." (Fisher, 1987, p. 105) Narrative strength refers to power of each narrative as they compete among multiple narratives to explain the world as it is observed— the strongest narrative will win out in this competition. We model narrative strength as the probability, that an individual sees themselves as participants in a particular cultural narrative.

Narrative building and decay

Our next step is to model how the strength of each narrative is built over time. According to Dehghani et al, “the rate of retrieval of cultural narratives [is dependent] upon the degree of surface and structural similarity with the presented [situation]” (Dehghani, Sachdeva, Ekhtiari, Gentner, & Forbus, 2009). In this manner, individuals update their beliefs about narratives by assessing the similarity of current events with their cultural narrative, with greater fidelity leading to greater narrative strength, as seen in Figure 1.

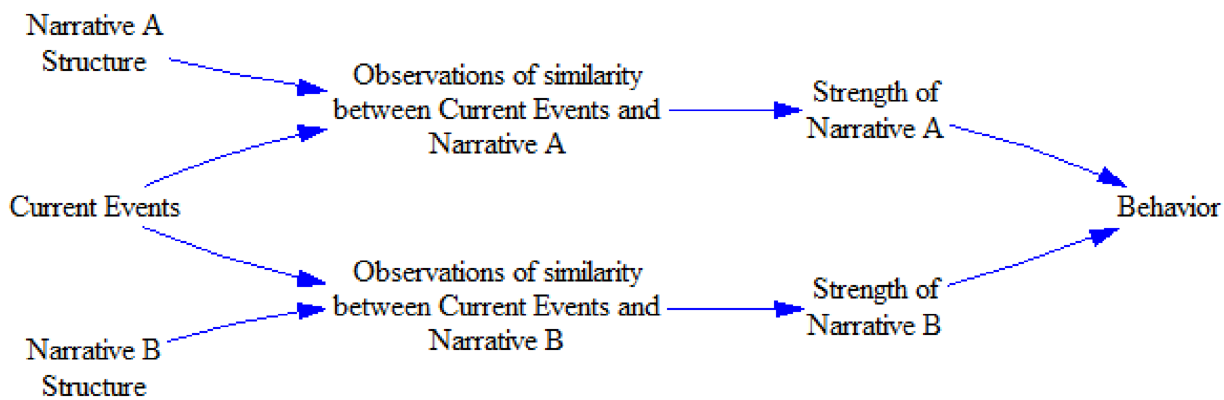

Figure 1: Narratives compete to influence behavior.

To simplify our model, we start with the assumption that each observation of narrative fidelity contributes equally to the strength of that narrative. This simplification is appropriate at the level of the population with the joint assumption that more influential events are more readily observed. Our focus on ‘observations’ of events captures much of this effect. This formulation is in line with previous work that models the role of messaging in the growth of insurgency. (Choucri, Goldsmith, Madnick, Morrison, & Siegel, 2007)

Next, we focus on the strength of narrative strength over time. Consistent with system dynamics modeling of decision making, we make the assumption that new observations are more relevant to an individual’s assessment of narrative fidelity than are older ones, and as such have more influence over narrative strength.(Sterman, 2000) Similarly, as an observation of narrative fidelity is succeeded by others, its influence over the individual’s assessment of their situation declines.

We next formulate an approach to measure the relative strength of a new observation versus an older observation. We model the strength of each observation as it decays with the addition of subsequent observations. Figure 2 illustrates this formulation. This formulation is consistent with prior modeling work on advertising. (Brady, 2009)

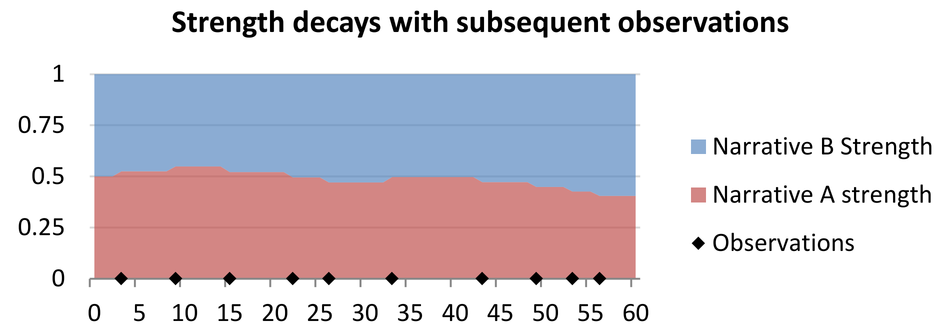

Figure 2: The influence of observations B upon narrative strength of A as it decays with subsequent message volume (Hypothetical Observations)

As more and more observations are made, the strength of the oldest observations approach zero as they become less and less relevant to the present situation.

To formalize this relation, we assume that the 'strength' of each observation decays exponentially with the number of subsequent observations. If the initial strength of an observation is '*a*', and '*n*' subsequent observations have been made, then we can describe the residual strength of an individual observation as:

$$Strength_{observation} = a \cdot (1-a)^n$$

The strength of a narrative (say 'Narrative A') is the sum of the residual strengths of all the observations which contribute to that narrative:

$$Strength_A = \sum_i a \cdot (1-a)^{n_i}$$

We can model this by tracking a level for each narrative, and when a new observation is made, multiplying all narrative levels by (1-a) before adding the new observation's strength to the relevant narrative. Say that observation 'i' contributes to Narrative A, then upon its arrival, all narratives would update as:

$$\begin{aligned} Strength_A[i] &= Strength_A[i-1] \cdot (1-a) + a \\ Strength_B[i] &= Strength_B[i-1] \cdot (1-a) \\ &\ldots \end{aligned}$$

We need to re-map this to the time domain for inclusion in a System Dynamics model. In the time domain, especially when we discretize our models, we may have multiple observations occurring in the same time period, and contributing to different narratives. We lose some information about the order in which observations are made, and so we need to add their contributions to the strengths of each narrative in a way that does not preference one over the other without cause.

If we knew the order that '*m*' messages arrived, we would know that at the end of the time period, those messages would have contributed a certain fraction of total strength of all narratives equal to:

$$\sum_{i=1}^{m} a \cdot (1-a)^{i-1}$$

We also know that the initial value of each narrative will have shrunk by a factor of:

$$(1-a)^m$$

Now of course, for the sum of all narrative strengths, the previous two equations sum to 1; so we can speed our calculation by avoiding the summation. The final equations which describe building and decay for Narrative A are:

$$Building_A = [1-(1-a)^m] \cdot \frac{Observaions_A}{m}$$
$$Decay_A = [1-(1-a)^m] \cdot Strength_A$$

A general representation of this structure using the System Dynamics modeling paradigm is described in Figure 3.

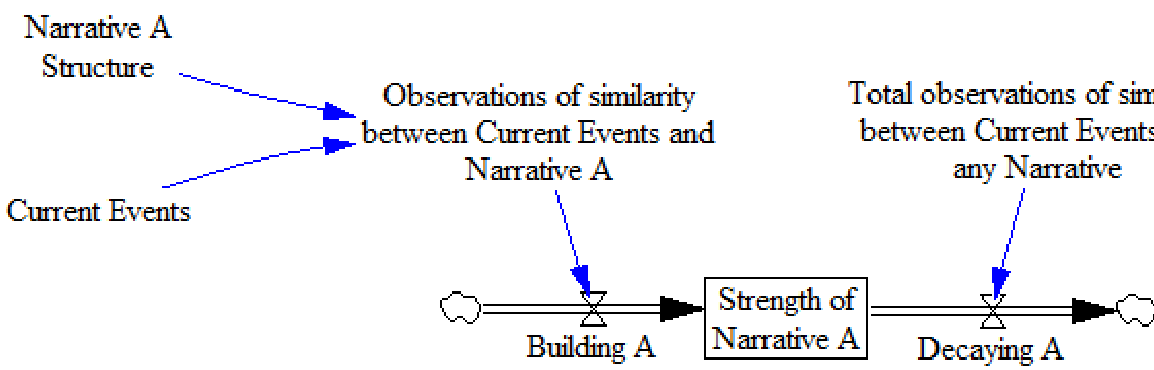

Figure 3: Stock and Flow of Narrative Strength

Defining initial narrative strength and decay time:

Finally, we must determine the initial 'strength' of an observation, the parameter 'a' in the preceding equations. While the initial strength of an observation is a hard parameter to estimate, we can utilize research that suggests that an individual can hold in short term memory about 10 pieces of information, which we may call the individual's "observation capacity". As an initial estimate, we'll set a value of 'a' such that the half-life of an individual observation is approximately ten observations per person in the population.

Setting our equation above for the residual strength of an individual observation to half of its initial value, and the exponent 'n' to the 'observation capacity' of the population, approximately $10 \cdot p$ we can estimate 'a':

$$\frac{a}{2} = a \cdot (1-a)^n$$

$$a = 1 - \left(\frac{1}{2}\right)^{\frac{1}{10 \cdot p}}$$

Further research is needed to improve this value, but for our initial purposes we believe our approach provides a reasonable estimate of the magnitude of the parameter.

Addressing variance within the population

As mentioned previously, we define 'narrative strength' as a measure of how much a member of a population understands their situation to be described by a particular narrative. As a first approximation, if an individual views their situation as being 100% likely to be caused by 'Narrative A', then the strength of that narrative for him is 1. If he sees his situation as 50% likely to be consistent with Narrative A, and 50% likely to be consistent with Narrative B, then the strengths of each are .5. For the total population, we look at the average strengths of the narratives over all individuals, such that the total strengths of all relevant narratives is equal to 1.

Using this approach, our model deals in terms of population average narrative strength, which omits information about the distribution of opinions within the population. These differences could be significant, both mathematically and in terms of their impact on the overall behavior. For instance, if Narrative A has an average strength of .3 convinced-person-equivalents per person, then the population can be composed of 30% complete believers in narrative A and 70% who give it no credit. Alternately, the population could be a completely homogeneous group that is 30% convinced of the veracity of narrative A, or a skew distribution with some function in between, as seen in Figure 4.

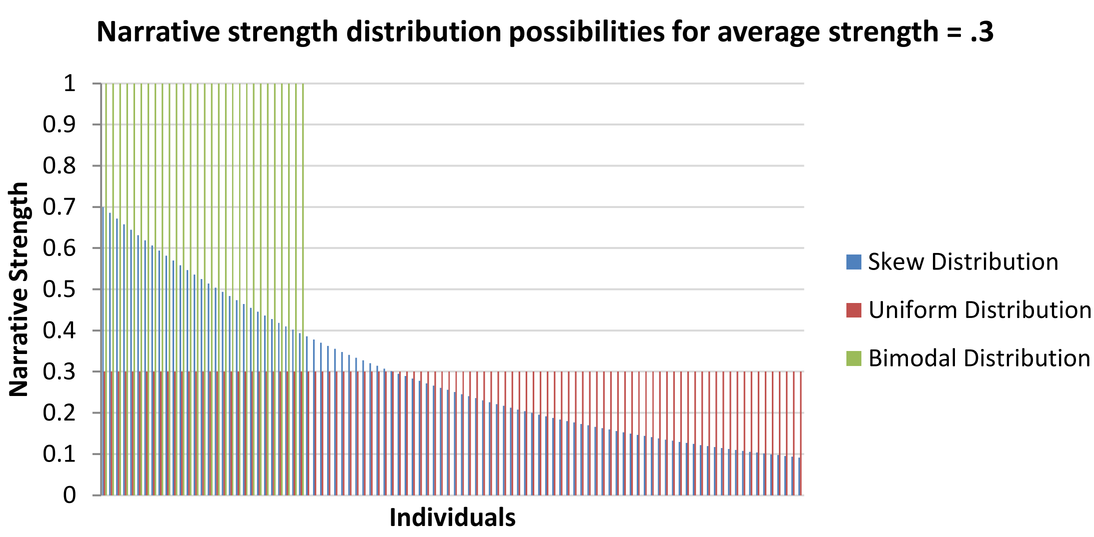

Figure 4: Many distributions of narrative strength over the population can yield the same average narrative strength

For the purpose of our analysis, we choose to conflate these two sources of variance into a single experimental parameter, included in the threshold model of narrative influence on behavior.

Modeling narrative influence:

In looking at how narratives influence collective behavior, we base our influence model on prior research on collective behavior thresholds. (Granovetter, 1978)

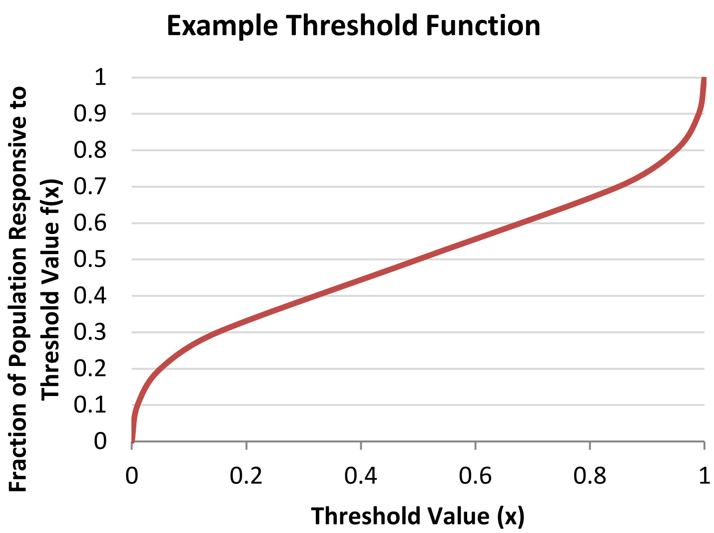

Figure 5: A threshold function maps a particular value of a parameter to the fraction of the population that responds with behavior change to that level

Granovetter describes a 'threshold' as the value of a quantity that is necessary for a given individual to participate in a behavior. Figure 5 illustrates how these individual thresholds combine to create a cumulative population threshold function. In a simple two-narrative model with no external factors, a narrative threshold can be expressed simply as a function of one of the narrative strengths:

$$\%\, participation = f(Strength)$$

When more than two narratives compete to explain the behavior, or when non-narrative factors also work to influence behavior, we may add additional dimensions to our 'decision space'.

$$\%\, participation = f(Strength_A, Strength_B, Factor_1 \ldots)$$

We can make some general statements about the shape of each curve. We will assume that when a single narrative increases in strength, the likelihood of participation in behavior advocated by that narrative increases monotonically. As all narratives must sum to one, for models with 3 or more narratives, there are regions of the 'decision space' which are inaccessible, bounded by the surface $Strength_A + Strength_B + \cdots = 1$. Beyond these constraints, the shape of the curve is case specific. An area for extension is to examine how individuals come to have their various individual thresholds.

Tying the model together

We expect to see feedback between narrative strength and current events. For example, if a narrative becomes strong enough to drive behavior, then resulting events have an increased likelihood of being consistent with that narrative, leading to more observations of similarity, and so on.

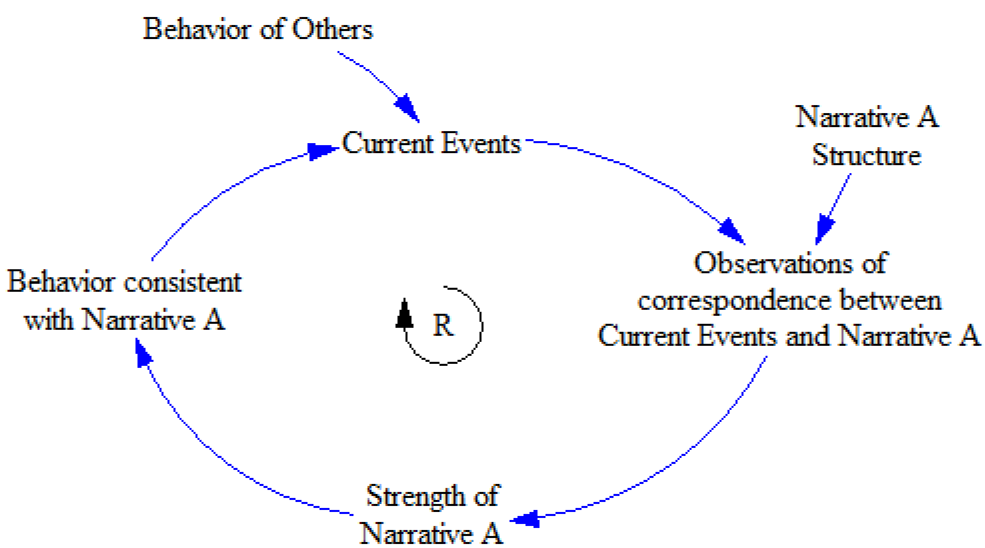

Figure 6: Narratives may drive behavior to creates a reinforcing feedback loop

The self-reinforcing power of narratives (illustrated in Figure 6) suggests an ability to lock-in to a particular behavior pattern once a dominant narrative is identified. This increases the importance of being able to monitor narrative strength as events unfold, in order to understand what eventualities are possible. In the next section we will propose one method for measuring narrative strength using social media data.

Monitoring narratives with social media

We believe social media provides exciting new opportunities to monitor narratives. While it is difficult to measure 'observations' in real time, we think this challenge can be overcome by using social media messages which point out an observed similarity between current events and a relevant narrative as a proxy variable

Social Media "has created a new way for us to collect, extract and utilize the wisdom of crowds in an objective manner with low cost and high efficiency". (Yu & Subhash, 2012) Our research suggests that we can quantify 'observations' by measuring the volume of social media messages which claim similarity between current events and cultural narratives.

For our initial work, we assume that the relative frequencies of messages claiming similarity between current events and each narrative gives a reasonable analog for the relative frequencies of narrative observations being made in the population. In the future, we can relax these assumptions and/or improve our mapping by using additional variables (i.e. demographics.)

With the above assumption, we can model the narrative building influence of social media messages in the same way that we modeled the influence of 'observations', a parallel pathway illustrated in Figure 7. We can now estimate the strength of each narrative, and its influence on behavior.

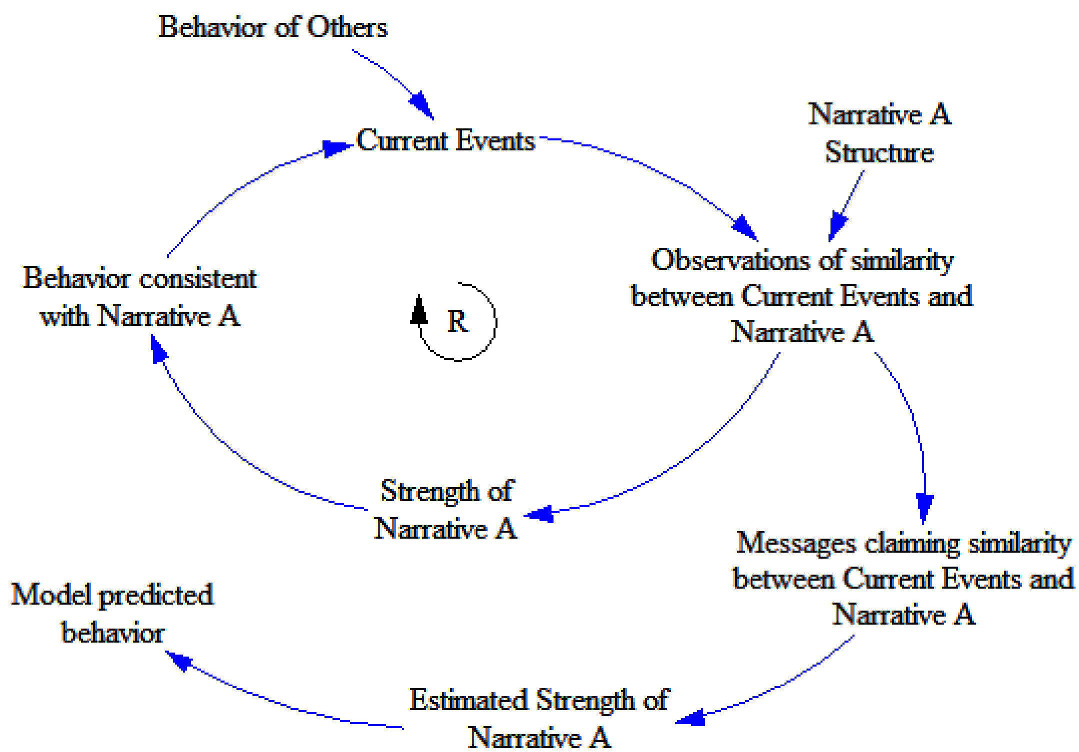

Figure 7: Social Media provides a parallel measurement pathway that may give insight into Narrative Strength

There are many challenges associated with collecting social media data and mapping it to the model we have constructed. Here we will touch here upon the challenges faced in constructing the case study that follows this discussion, and one method for dealing with them.

The first primary challenge is to identify how narratives are being discussed in relation to current events. In the case study that follows, we did not have the advantage of published narrative analyses of the population. Instead, we drew heavily on historical events which mirrored the behavior we wished to evaluate, and upon primary source reports of the behavior of interest. We assumed that a reference to these historical events constituted a reference to the overall narrative.

The second primary challenge is to measure the volume of messages that claim similarity between current events and relevant narratives. This data collection, storage, and analysis task is primarily difficult because of the sheer volume of messages to be processed. In preparing the case study that follows, we tapped into the archive of Crimson Hexagon's Foresight tool, and filtered data using keywords suggested by the narratives and by news events. We specifically looked for instances in which a narrative keyword appeared in the same message as a current event keyword.

The third primary challenge is to estimate the relevant initial-strength parameter which we labeled 'a' above. Our description of the parameter depends on what we called the 'observation capacity' of the population In mapping this parameter to social media, we need to account for the fact that not all observations result in messages, and that the size of the population changes as new groups are exposed to the same ideas and so we measured the total number of people making claims about a particular set of narratives. For the purpose of the case study, we assumed a constant relationship between observations and messages.

Case study - rioting in London

To demonstrate how this model can be used to understand and predict instances of collective social behavior, we have constructed a simple model of the rioting that occurred in London, England in August of 2011. This case study is not intended to be a complete description of how the riots progressed, so much as an example of how to incorporate quantitative narrative assessment into a larger model of collective behavior.

Riots are a very visible example of narrative-driven collective behavior, and the London 2011 riots in particular have several traits which simplify the analysis significantly. The first of these traits is that narratives and social media messages use the English language, simplifying data collection for the English speaking modelers. The second is that the riots were part of a larger historical pattern of violent behavior, with close analog situations in recent cultural memory. This gives us a starting point for identifying language used to reference the narrative.

Overview of the case study

On August 4, 2011, a resident of Tottenham named Mark Duggan was fatally shot by police. Initial reports were unclear as to whether Duggan was armed or aggressive toward police, and initial police response to complaints by Duggan's family were seen by the local population as inadequate.

On August 6th, 2011, a march began on Tottenham police station, in protest of the perceived injustice of Duggan's death. At about 8:30 PM, violence began with throwing of missiles at the police station and burning of two police cars. Violent behavior, arson, and looting spread throughout the night.

Prior to the disturbances, London Police used a 'Community Engagement Model' to understand and respond to potential outbreaks of violence. According to the official police report of the disturbances, "The engagement model did not… enable the disorder to be predicted and did not help… [during] a crucial turning point in the dynamics of the crowd." In this case study, we will attempt to determine if our model of narrative influence could have been helpful in prediction of violence. (Metropolitan Police Service, 2012)

Historical background:

London has a history of violent protest, notably including the race riots of the 1950s and 1970s. Several of these protests follow a similar pattern to that of the 2011 riots: A death or injury in police custody leads to an ostensibly peaceful march on the local police station, and the protest becomes violent.

Table 1: Narrative Relevant Rioting in London

Year	Catalyst	March?	Delay	Duration
1981	Michael Bailey dies of prior stab wounds in police custody	N	~24 hrs	1 night
1985	Cherry Groce paralyzed by police bullet	Y	~12 hrs	2 nights
1985	Cynthia Jarrett dies of heart failure in police custody	Y	~24 hrs	1 night
1995	Wayne Douglass dies of heart failure in police custody	Y	~12 hrs	1 night

(The Guardian, 1981) (Tompson, 1988) (Reuters, 1995)

These examples show how one narrative maintained by the population of Tottenham has played out in the past. They give the modeler insight into the structure of the narrative, and suggest keywords for identifying when the narrative is being referenced.

Narratives:

For this case study we have identified two narratives relevant to the rioting behavior, which we summarize here. We have chosen to call the first narrative "Oppression by Elites":

> *In this conflict-ridden world, there are regular people, and there are the powerful elite.* (Tompson, 1988) *The elite use their connections to the corrupt government and police to economically disenfranchise the regular people and actively repress their humanity. Police brutality represents an acute manifestation of this pervasive oppression, and galvanizes the people to take action against their oppressors. (Muir, 2011) If peaceful demonstration fails, the people exact violent retribution to make the powerful elite pay for the suffering they have caused the people, and recover some measure of self-empowerment. (Reuters, 1995)*

We call the second narrative 'Law-and-Order':

> *In this generally peaceful world, there are upstanding citizens, and there are criminals. When a criminal commits a crime against a citizen, the police work to catch them and punish them, sometimes at personal risk, so that the citizens can go on with their lives. The police are not perfect, but their mistakes are not directional or intended, and they will admit errors and make amends.*

Members of the population of Tottenham could use either of these narratives to understand the death of Mark Dugan, each narrative suggesting a very different response. As observations of events surrounding the death contribute to each narrative's strength, the dominant narrative will in general drive the behavior of the population.

Data collection:

We collected approximately 60,000 Twitter and Facebook messages over 11 days that claimed similarity between the identified narratives and the ongoing relationship between police and the community, as seen in Table 2. Examples of these messages can be seen in Table 4 and Table 5.

Table 2: Daily Messaging Volume

Analysis Date	Oppression Message Volume	Law-and-Order Message Volume	Authors (Est)
7/28/2011	249	1177	1257.1
7/29/2011	442	565	904.7
7/30/2011	356	643	927.6
7/31/2011	353	503	780.4
8/1/2011	294	588	763.6
8/2/2011	375	590	846.4
8/3/2011	334	598	859.6
8/4/2011	386	708	968.7
8/5/2011	426	610	946.6
8/6/2011	5173	1347	3814.7
8/7/2011	22041	22033	32454.9

Additionally, we collected data for the date on which rioting actually began (August 6) at hourly resolution, as seen in Table 3.

Table 3: Hourly Messaging Volume

Hour	Oppression Message Volume	Law-and-Order Message Volume	Authors (Est)
0	16	29	17.1

1	15	27	15.9
2	14	22	12.9
3	14	34	20.0
4	8	38	22.4
5	7	39	22.9
6	9	37	21.8
7	8	26	15.3
8	16	31	18.2
9	18	20	11.8
10	21	17	10.0
11	16	19	11.2
12	20	25	14.7
13	25	22	12.9
14	25	12	7.1
15	16	17	10.0
16	5	19	11.2
17	7	37	21.8
18	17	44	25.9
19	35	81	47.6
20	994	400	235.3
21	1044	1088	640.0
22	1798	555	326.5
23	931	559	328.8

Table 4: Examples of "Oppression by Elites" messages

Timestamp	Author	Message
8/6/2011 12:05	DJPIONEER	R.I.P Mark Duggan. Justice will be served one day....
8/6/2011 16:42	Ms_Benjamin	Protest at tottenham police station :(RIP Mark Duggan xx
8/6/2011 16:53	ShanikaBeee	We lock off tottenham high road!!! WE WANT JUSTICE!!!! RIP MARK DUGGAN MY BRUDDA ♥
8/6/2011 18:28	STYLERGFAM	Outside tottenham police station R.I.P mark duggan
8/6/2011 20:13	SkeamzMusic	I hear der rioting in tottenham over the death of mark . F[**]k feds . R.I.P mark . Family dat !
8/6/2011 20:15	terridab[***]h	Wooooow its all kicking off in totty police car getting blowen up that aint having it R.I.P mark

Table 5: Examples of "Law-and-Order" messages

Timestamp	Author	Message
8/5/2011 10:57	Hughes_Mark	Police officer shot by 'gangster' Mark Duggan in north London was saved by radio.
8/5/2011 15:29	tpsjones	Tottenham Hale man, named locally as Mark Duggan, was shot by CO19 officers in Ferry Lane after arrest attempt
8/5/2011 20:55	Cartwitt09	Officers from Operation Trident claimed there was an illegal firemarm inside the taxi which was transporting Mark Duggan.
8/6/2011 12:04	SachinChokshi	Gangster' Mark Duggan shot by police in back of London cab after shootout

Model building:

The components of our model relevant to narrative building and decay in the London riots case study are shown in In Figure 8.

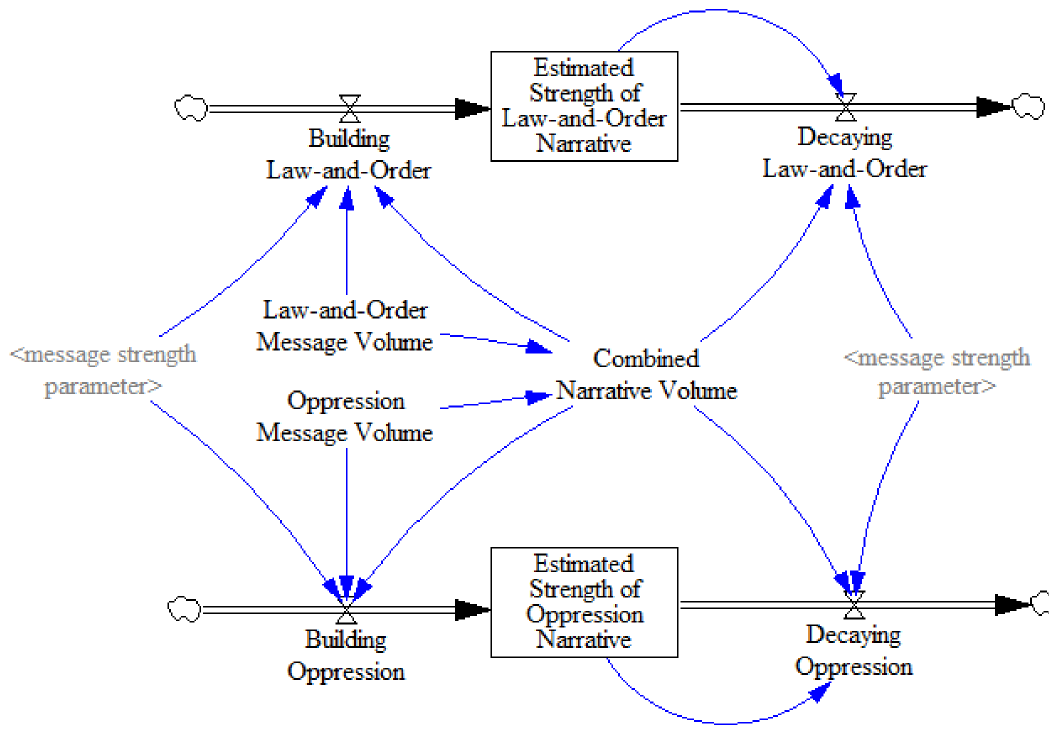

Figure 8: The narrative model incorporates message volume to estimate the strength of each narrative

We are primarily interested in running the model over the day during which the riot began. However, to establish initial values of the narrative strengths at the beginning of that day, we need to look at data from the preceding days. Figure 9 shows the sensitivity of the narrative strengths on August 5th, (the start of our analysis and a day before the riots began) to a range of narrative strength estimates on July 28th (a date sufficiently in advance of the riots to let us calculate a baseline). Including measurement data with our narrative strength model allows us to select an initial value for the strength of the opposition narrative to 0.37.

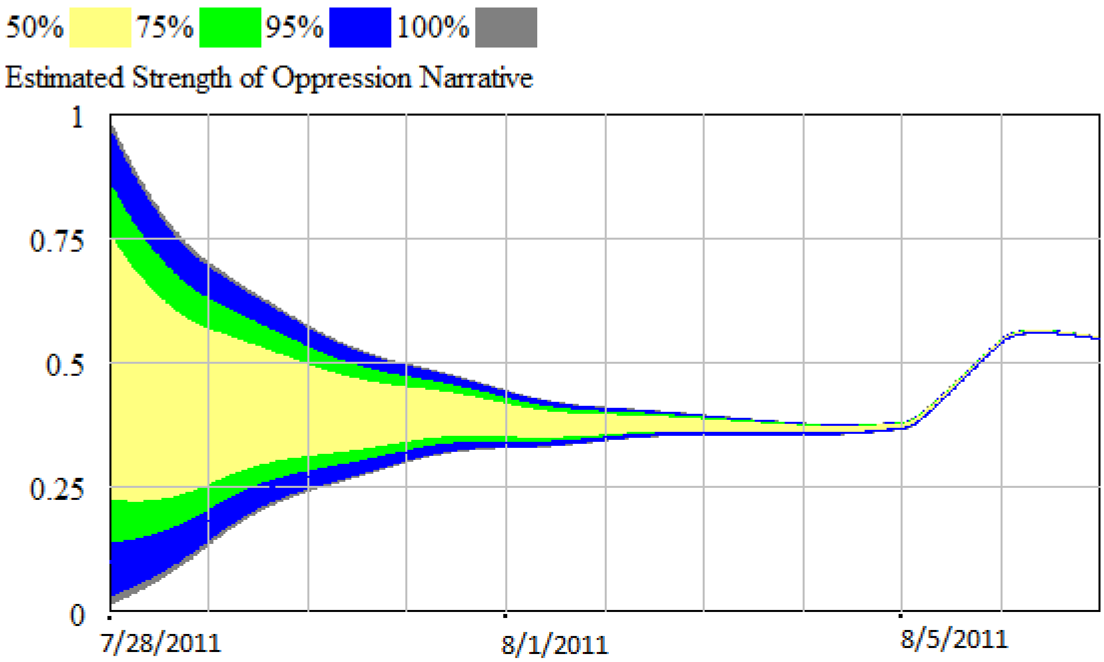

Figure 9: This sensitivity analysis shows that the Narrative Model converges to an estimate of narrative strength without any estimate of initial conditions

We chose to estimate the narrative strength parameter 'a' as a function of the measured number of authors participating in discussion about the narrative in a given day, some contributing multiple messages. This estimate of the relevant 'digital population' is a rough estimate, as it is certainly not constrained to the population of interest, and future work will involve methods for identifying populations of authors within social media. For this analysis, a rough estimate of population size is sufficient for modeling narrative strength, as shown by the sensitivity analysis presented in Figure 10.

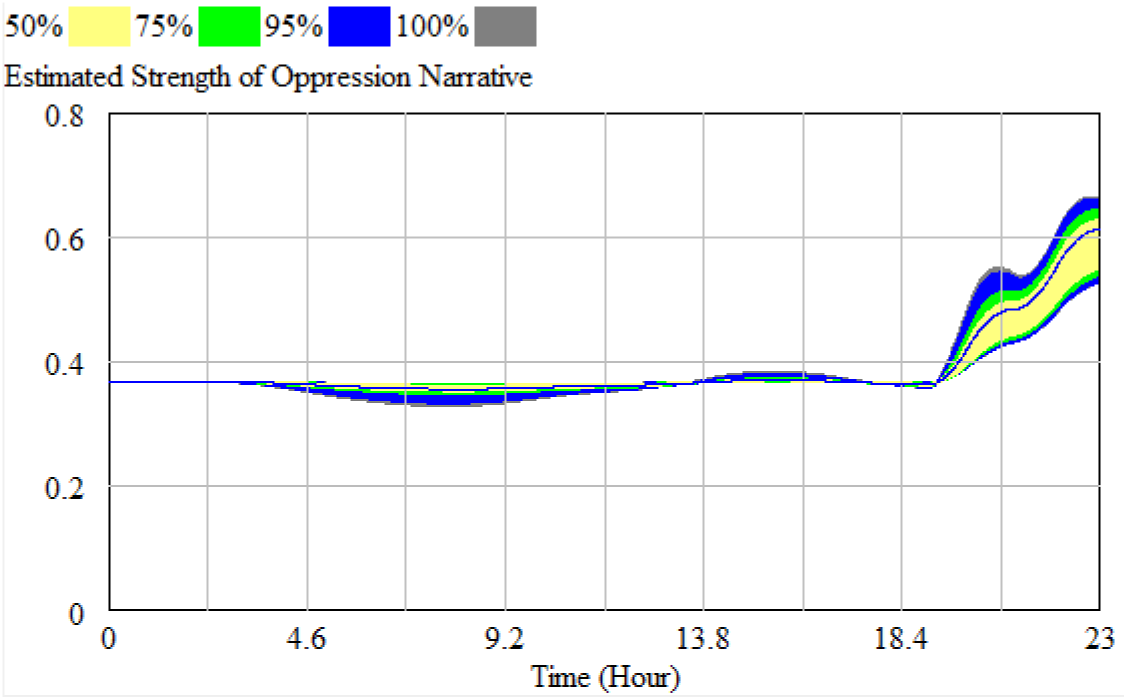

Figure 10: This sensitivity analysis shows that the general behavior of narrative strength is consistent over a tenfold change in the number of message authors

The component of this case study that deals with actual riot behavior as influenced by narrative strength is intentionally simplified to its core elements. We represent the willingness of individuals to join a riot with a threshold function, mapping the percent willing to join a riot as a function of the strength of the "Oppression by Elites" narrative and the fraction of the population already participating in violent behavior.

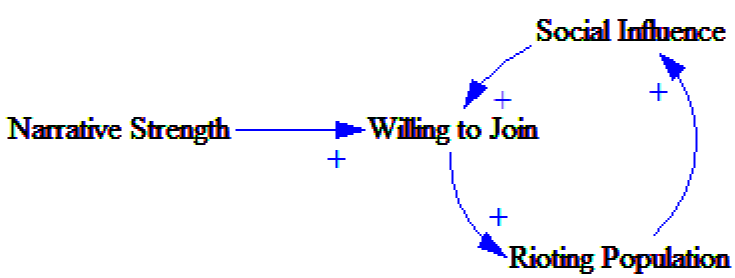

Figure 11: This causal loop diagram of the behavioral component of the model shows how both narrative strength and the size of the rioting population influence those willing to participate in violent behavior

As seen in Figure 11 the behavioral component of the model shows that the participation of individuals in violent behavior increases with the strength of the violence-driving narrative until a tipping point is reached and social influence takes over and rapidly increases participation. Figure 12 shows how these loops are constructed in a complete system dynamics model, with 'Oppression Narrative Influence' and 'Social Influence' together representing the threshold function which describes willingness to participate.

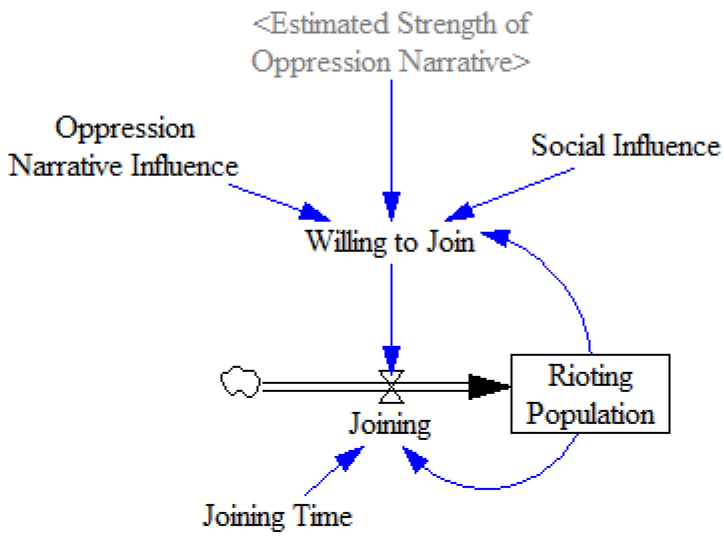

Figure 12: The time dynamics of the rioting population are driven by the lag between identifying the current size of the riot and deciding to participate

We model the influence of the oppression narrative upon the rioting population as an order of magnitude increase per quartile in narrative strength, and the social influence as a linear function that takes effect after a minimum fraction of the population has joined in with the riot. Figure 13 shows how these components combine to drive the fraction of the population willing to join the riot at any given time.

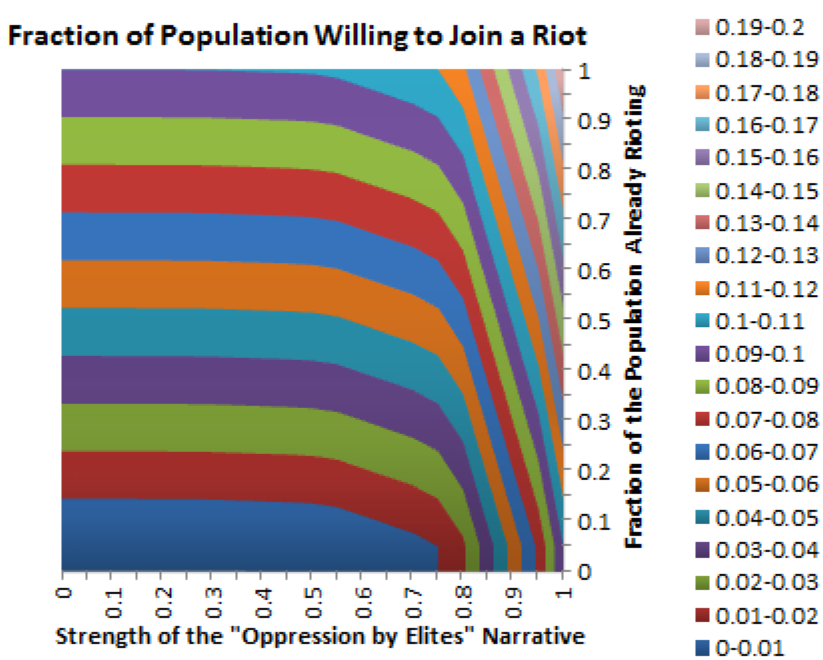

Figure 13: The fraction of a population that is willing to join in a riot is modeled as a function of the strength of the "Oppression by Elites" narrative and the existing size of the riot

Result:

We plot the development of narrative strength over the course of the day in Figure 14. As the volume of messaging regarding the narratives picks up in the evening, we see a substantial change in narrative strength for the first time in the course of the surveyed data.

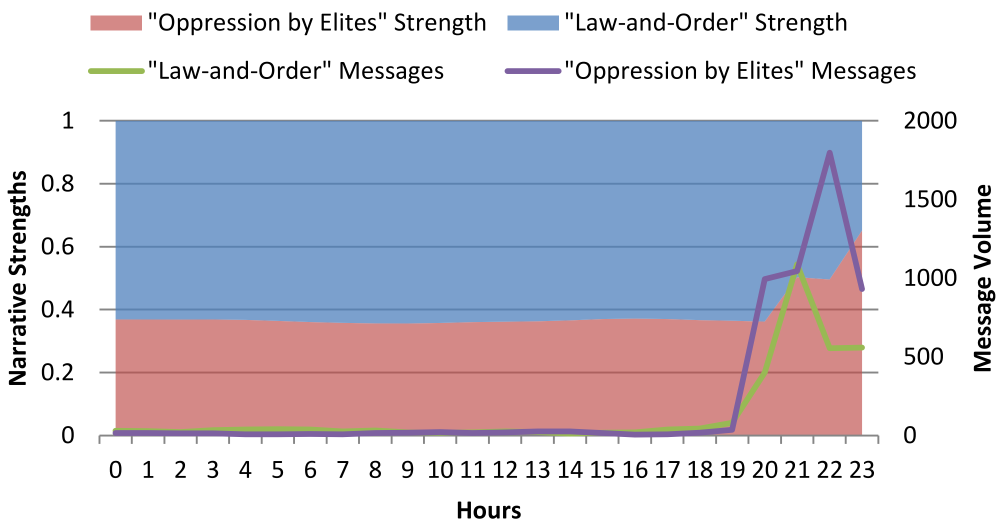

Figure 14: Narrative strength evolves quickly with the addition of an influx of new messages (and new observations) in the hour preceding violence and continues to vary as the night continues.

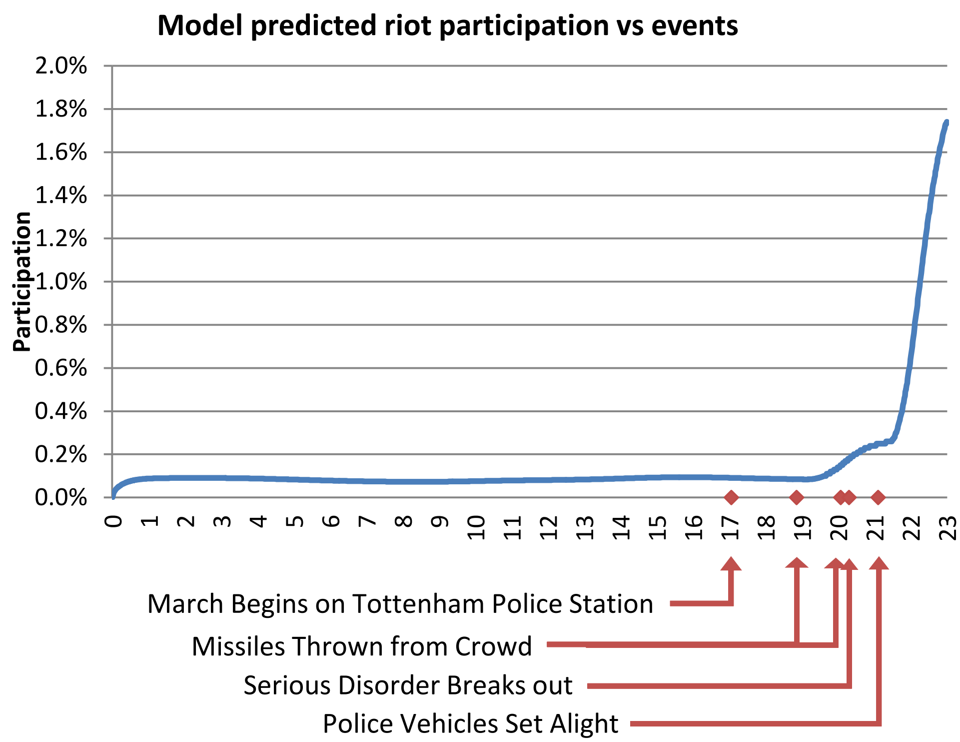

Figure 15: Predicted % Participation in Riot Behavior and Major Initial Events in Riot

According to the London metropolitan police service, generally nonviolent protests began at approximately 17:05, and violence began at approximately 20:30.

Rioting spread over the city for the next two nights until increased police presence and a wider sense of popular outrage worked to break the reinforcing social loop and change the narrative conversation to one of Law-and-Order.

Discussion and recommendations for future work

Using lessons from narrative theory we have developed a model which quantitatively represents the strength of narratives in a population, and maps their ability to influence collective behavior. We have postulated a method for estimating narrative strength using social media messages as a proxy variable for narrative-building observations. We have demonstrated in a case study how this model and data collection process can be applied to real-world situations to predict results consistent with observed behavior. We have shown that narrative strength models informed by social-media analysis may be a viable indicator of collective behavior.

Narratives are a powerful tool for either inciting or defusing violent behavior, and our ultimate goal in the study and modeling of narrative influence is not merely to predict incidents of violent behavior but to allow for policy decisions which can prevent them. Figure 16 suggests how a sense of narrative strengths can influence policy decisions in order to promote narratives of peace and stability.

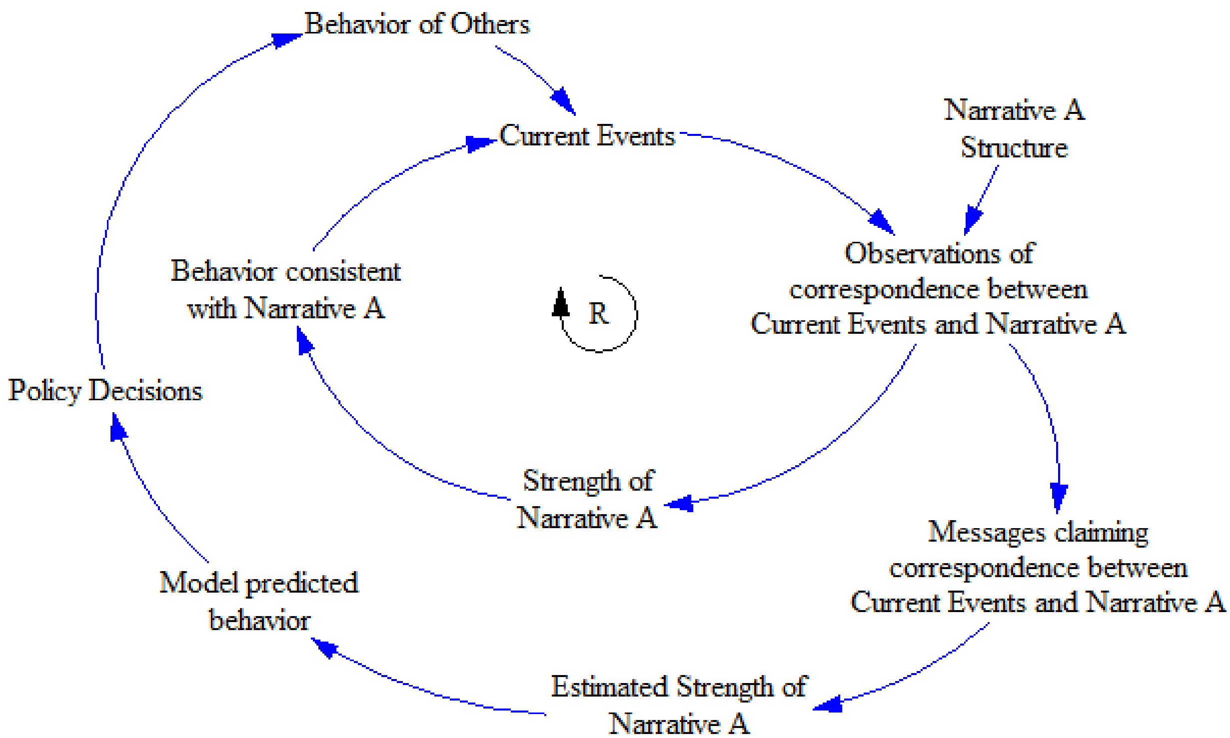

Figure 16: In operation, narrative models of collective behavior have the potential to shape policy decisions. These policy decisions can act to weaken narratives of violence by creating events consistent with other available narratives.

We have noted throughout this paper opportunities for future work to strengthen parts of our analysis. First, we can improve our method for measuring the relative strengths of various narratives within an individual in a controlled setting. Sociological research can then test the hypothesis of narrative strength decay. Work similar to that of Dehghani et al. which looks at narrative recall could be performed over time to verify that observations of similarity between narratives and suggested events declines exponentially with the number of subsequent messages. Ideally, controlled experiments to map threshold functions relating narrative strength to behavioral proclivity would be added to strengthen our analysis.

Additional work would also help identify the bound of populations within social media spheres. This work could take two paths, the first looking at geographical distributions of individuals, the second looking at linguistic differences between groups, measuring their degree of self-identification through the use of in-group language.

Finally, it would be useful to incorporate other work being performed regarding the identification of narratives and narrative activity “in the wild.”

While the technology for sharing ideas and messages will change, narratives themselves have a certain sticking power. We have confidence that the work we do today to understand and measure narrative activity will continue to be relevant despite changes in these social technologies.

Funding Acknowledgement

This work is funded by the Office of Naval Research under award number N00014-09-1-0597. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

The authors wish to acknowledge Crimson Hexagon for provision of ForSight services and support.

Bibliography

Beach, R. L. (2010). *The Psychology of Narrative Thought: How the Stories We Tell Ourselves Shape Our Lives.* USA: Xlibris Corporation.

Bergen, K. M. (2010). Accounting for difference: Commuter wives and the master narrative of marriage. *Journal of Applied Communication Research, 38*(1), 47-64.

Dehghani, M., Sachdeva, S., Ekhtiari, H., Gentner, D., & Forbus, K. (2009). The role of Cultural Narratives in Moral Decision Making. *Proceedings of the 31st Annual Conference of the Cognitive Sciencce Society.* Washington, D.C.

Fisher, W. R. (1987). *Human Communication as Narration: Toward a Philosophy of Reason, Value, and Action.* Columbia, South Carolina: University of South Carolina Press.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 1420-1443.

Halverson, J. R., Goodall, L., & Corman, S. R. (2011). *Master Narratives of Islamist Extremism.* Palgrave Macmillan.

Hammack, P. L. (2009). Exploring the Reproduction of Conflict Through Narrative: Israeli Youth Motivated to Participate in a Coexistence Program. *Peace and Conflict*, 49-74.

House Armed Services Committee. (2007, July 11). Strategic Communications and the Battle of Ideas: Winning the Hearts and Minds in the Global War Against Terrorists. *Hearing before the Terrorism, Unconventional Threats and Capabilities Subcommitttee*. U.S. Government Printing Office.

MacIntyre, A. (1981). *After Virtue: A study in moral theory.* London: Duckworth.

Metropolitan Police Service. (2012). *Four Days in August: Strategic Review into the Disorder of August 2011.* London.

Monitor 360. (2012). *Master Narratives Audience Segment Report: Muslim Communities of France.*

Muir, H. (2011, September 5). Tottenham riots: missteps in the dance of police and a frustrated community. *The Guardian*.

Reuters. (1995, December 15). Man's Death Sets Off Riot In London. *The New York Times*.

Rosenstiel, T., & Jurkowitz, M. (2012). *The Master Character Narratives in Campaign 2012.* Washington D.C.: Pew Research Center: Project for Excellence in Journalism.

Skeel, D. A. (2009). Competing Narratives in Corporate Bankruptcy: Debtor in Control vs. No Time to Spare. *Michigan State Law Review*, 1187.

Slater, T. (2011). From 'criminality' to marginality: Rioting against a broken state. *Human Geography, 4*(3), 106-115.

The Guardian. (1981, April 13). How smouldering tension erupted to set Brixton aflame. *The Guardian*.

Tompson, K. (1988). *Under Siege: Racism and Violence in Brittain Today.* Penguin.

Yu, S., & Subhash, K. (2012). A survey of Prediction Using Social Media. *arXiv preprint*, 1203.1647.