Unpack that Tweet: A Traceable and Interpretable Cognitive Modeling System

Upol Ehsan, Christopher Purdy, Christina Kelley, Lalith Polepeddi, Nicholas Davis
School of Interactive Computing
Georgia Institute of Technology
Atlanta GA, 303332 USA
{ehsanu; cpurdy3; christina.kelley; lpolepeddi; ndavis35}@gatech.edu

Abstract

The rhetoric of world leaders has considerable influence on civic engagement and policy. Twitter, in particular, has become a consequential means of communication for politicians. However, the mechanisms by which these politicians use Twitter to communicate with the public are not well-understood from a computational perspective. This paper describes an analytic method and system for examining the political tweet-making process. We offer three main contributions: a traceable and interpretable system with a model that simulates tweets produced by world leaders, a framework which allows users to simulate "what-if" scenarios and unpack the underlying mechanisms of the system, and a transferable model-generation process contextualized with a foundational use case centered on U.S. President Donald Trump's Twitter account @realDonaldTrump. We end with a discussion of the strengths and limitations of our system and a plan for assessment.

Keywords: social justice, cognitive model, artificial intelligence, interpretable AI, Twitter, social media, political rhetoric, communication, computational creativity

Introduction

Words construct our social world. They empower us, embolden our souls, make us laugh and cry like few things can. Words from a person of influence, power, and privilege have magnified and compounded implications on our social fabric. We can gain important insights into the minds of leaders by analyzing their communication methods. These insights have the potential to catalyze informed and proactive civic engagement about social justice issues. If we can create computation tools that provide insights into the process, we can foster informed participation in social commentary and civic engagement.

Twitter, a 140-character-limited microblogging social networking service, has emerged as a communication means for political leaders. Particularly, it has become the primary communication method for one of the most influential people in the world: the 45th U.S. President, Donald J. Trump. As a result, using Twitter as a means of official governmental communication has been elevated to an unforeseen level with real economic and policy consequences (Mathew 2017; Noguchi 2017). In the past, historians and sociologists have studied political speeches and their impact on policy and diplomacy, but little research examines the particular impact that tweets have on the social landscape. To build a foundation for this, we must first understand *how* these tweets are formed. Schwarz (2005) defines inquiry as the process used by learners to investigate a phenomenon, explore topics, and utilize them to understand the phenomenon. One way to facilitate the process of inquiry is to provide an analytic tool that allows access to the decision-making processes of simulated realistic what-if situations. Contextual and data-driven simulation of tweets can empower users to improve their conceptual models and understanding of the tweeting behavior (Piaget 1950; Papert 1980; Joyner et al. 2014; Joyner and Goel 2015). With this said, a computational understanding of tweets from a cognitive perspective of the tweeting behavior is an understudied area that requires rigorous investigation.

Interpretability and *traceability* are desirable features in AI systems, as they afford explanation and inference through inspection of inner components. Researchers are attempting to make AI systems more explainable by introducing elements of interpretability, rationalization, and transparency (Harrison et al. 2017). *Explainable AI* refers to AI systems that empower humans to appropriately understand, trust, and operate: it attempts to reduce the black-boxed nature of AI systems. A *black-boxed* system inherently lacks explanatory power as it can only be inspected with respect to inputs and outputs. This lack of interpretability makes it difficult for users, especially non-experts, to trust the system's decisions. For instance, neural networks are typically considered black-boxed and uninterpretable without a visualization layer or explanation (Zeiler and Fergus 2014; Yosinski et al. 2015). A *clear-boxed* system is interpretable, as users can *trace* through inner components and inspect to understand "behind the scenes" processes.

Inspired by the interpretability side of explainable AI, we present in this paper a *traceable* and *interpretable* system that uses a cognitive model to provide insights into the potential rationale behind simulated tweets from world leaders. The system uses data-driven belief bases to provide these insights; it allows users to simulate "what-if" scenarios and allows "behind the scenes" access to the decision-making processes of generated tweets. For instance, given a realistic simulated tweet, the user can examine the modeled beliefs and strategies used to produce that tweet. The generated tweet anchors the exploratory process while the ability to unpack modules in the system can facilitate interpretation of the system's actions. Unlike black-boxed connectionist systems, ours is an analytic tool that allows users to trace through decisions made by the system, thus facilitating inquiry. By transforming the implicit processes to explicit ones and providing a deeper insight into what is happening "behind the scenes", our system can facilitate informed and proactive civic participation in society.

To contextualize and appropriately ground the system with an example of consequence, we model the behavior of @realDonaldTrump, the personal Twitter account of U.S. President Donald Trump. His predilection for tweeting has upended norms of political communication, and there is often confusion around the interpretation of his tweets ("President Trump's Tweets, Annotated" 2017). With political polarization on the rise (Gentzkow 2016) and reflected through social media (Conover et al. 2011), it is incumbent on us to systematically understand the key factors that influence his tweeting behavior. Despite the tangible impact of President Trump's tweets and countless automated Trump tweetbots, there is a dearth of research on clear-boxed computational models of this behavior. As a result, his tweeting behavior is an appropriate case study for our proposed system.

Our contributions are threefold. First, we present a system that allows simulation of the tweeting behavior of a world leader with a traceable and interpretable model. Second, we describe a framework which enables users to unpack the tweet generation process as they inspect the model. Simulation coupled with traceability in the decision-making pipeline can help users develop better conceptual models of the phenomenon. Third, we contextualize the model with a use case and offer a transferable process of model-generation, so that others can utilize our work to analyze the communicative behavior of other leaders.

We begin this paper by reviewing related work on the impact of rhetoric from leaders and use of Twitter to understand social phenomenon, underscoring the need for traceable and clear-boxed computational cognitive models. Next, we describe a hybrid methodology that grounds the process of building our cognitive system. We proceed to delineate different components of the model, justifying the basis of each part, before proposing evaluation plans for the system that focus on open cycles of inquiry by the user. Finally, we conclude with the potentials and limitations of the current system and a plan for future work.

Related Work

Rhetoric, Power, and Communication

Rhetoric and communication are core elements of politics, and they have an important impact on social justice (Hart 1987). The rhetoric of a powerful figure has tangible effects on society, and leaders have long utilized and revolutionized novel means of communication to influence the population. Rigorous analysis of the methods and content can help us unearth potential rationale behind statements and their impact on society.

Analysis of Dr. Martin Luther King Jr.'s rhetorical influences and speeches, for instance, can help us understand his views on segregation, religion, and civil rights. It also helps us see how he was able to connect with and inspire people from all walks of life (Washington 1986). Viewing Mahatma Gandhi's political communication through the lenses of non-violence, non-cooperation, and civil-disobedience can augment our understanding of his pervasive appeal that transcends time and borders (Yamabhai 1973).

Similarly, U.S. President Franklin D. Roosevelt played a pioneering role in using radio to connect with people; in particular, "radio provided him with a direct link to his voting public and the next generation of voters" (p.90, Yu 2005). His fireside chats tangibly affected the trajectory of U.S. history: over the radio, Roosevelt "led the nation to unite behind [his] call to war" (p.89, Yu 2005). He was able to adapt this emerging medium to fulfill political needs.

Political Communication via Social Media

In recent years, social media platforms have added another new dimension and transformed political rhetoric and participation (Tumasjan et al. 2010; Rainie et al. 2012; Farrell and Drezner 2008; Wattal et al. 2010). Twitter, in particular, has played a rising role in political campaigns, allowing candidates to share information and gauge reactions in real-time (Conway et al. 2015; Stieglitz and Dang-Xuan 2012). It played an instrumental role in mobilizing citizens in the Arab Spring (Eltantawy and Wiest 2011). Twitter is not only used as a medium for political deliberation, but it also serves a reflection of political sentiment (Tumasjan et al. 2010).

Political leaders across the globe have utilized Twitter as a direct channel of communication with citizens. The dialogic characteristic of Twitter facilitates its communicative power; in fact, a vast majority of randomly sampled tweets contain an "@" sign, indicating direct communication taking place (Honey and Herring 2009). For instance, Indian Prime Minister Modi has successfully used Twitter to reach out and engage with an increasingly connected electorate of 800 million people (Kotoky 2014). In addition to domestic engagement, PM Modi has engaged global audiences by tweeting in the native languages of host nations during his state visits (Chronicle 2014).

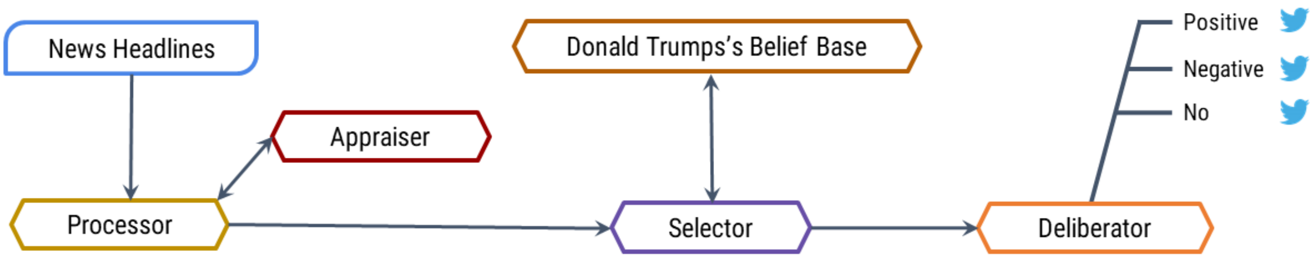

Figure 1: A diagram of the model applied to President Trump

Like Prime Minister Modi, President Trump has successfully transformed and shaped his message using Twitter, and has thus been able to engage millions of people (Philipps 2015). Today, the 140 characters of a tweet from President Trump can change the fortunes of a company by impacting stock prices (Bukhari 2017; Wieczner 2016). In fact, foreign governments now have dedicated employees monitoring President Trump's tweets (Ryall 2017). Gauging this importance, the National Archives have even requested that the White House store all presidential tweets (Braun 2017).

Existing Analytic Methods

Investigators from a diverse range of backgrounds have attempted to analyze President Trump's style: from psychologists to linguistics to AI researchers (Keohane 2016; Marshall 2016; Simms 2016; McGill 2017). Although there are tweetbots that mimic his style, there is little work--either within or outside academia--that aims to *computationally* model President Trump's tweeting behavior in a manner that is both *traceable* and *interpretable*.

A range of methods have been applied to create bots and mimic his tweets. In fact, automated Trump tweetbots outnumbered his nearest opponent Hillary Clinton by 7:1 (Kollanyi 2016). For instance, @DeepDrumpf uses a deep recurrent neural network to generate tweets (Hayes 2016). However, almost all of these bots and models are black-boxed deep neural network models. They lack the interpretability and traceability that can afford explanation and insights into the behavior.

Although clear-boxed *computational* analysis of President Trump's tweeting behavior may not exist, there are analyses of his strategies in non-computational forms. Most notably, George Lakoff (2017) has recently created a taxonomy of President Trump's tweets in an attempt to explain the process. He denotes the President's tweeting strategies as *pre-emptive framing*, in which Trump offers up a new idea; *diversion*, in which he diverts attention from one issue to another; *deflection*, in which he "attacks the messengers" and changes the conversation; and *trial balloon*, in which he tests public reaction to a topic. This approach has explanatory power, but not computational power.

Our approach, by comparison, serves to alleviate the limitations of solely analytical and solely computational approaches by affording both explanatory *and* computational power.

Methods

We used a mixed-methods approach to curate data and create our model. Our goal at each step was to maintain a grounded, evidence-based process for understanding the beliefs and behaviors of the @realDonaldTrump account.

When analyzing the existing tweets of @realDonaldTrump, we imposed a series of constraints to focus our analysis and ground our corpus. We examined only tweets with a timestamp between June 16th, 2015 (when Trump announced his candidacy for presidency) and March 31st, 2017 (when we finalized our methodology). We allowed tweets from any device, and we did not take the @POTUS account into consideration. With these constraints, we had a final corpus of 8,544 tweets from the @realDonaldTrump account.

Topic	Sentiment	Anger	Joy	Sadness	Fear	Disgust
Hillary Rodham Clinton	-0.20	0.13	0.086	0.38	0.12	0.35
Russia	0.28	0.01	0.56	0.06	0.06	0.13

Table 1: Sample beliefs from the system's belief base. Sentiment is scored on an interval of [-1, 1], and anger, joy, sadness, fear, and disgust are scored on an interval of [0, 1].

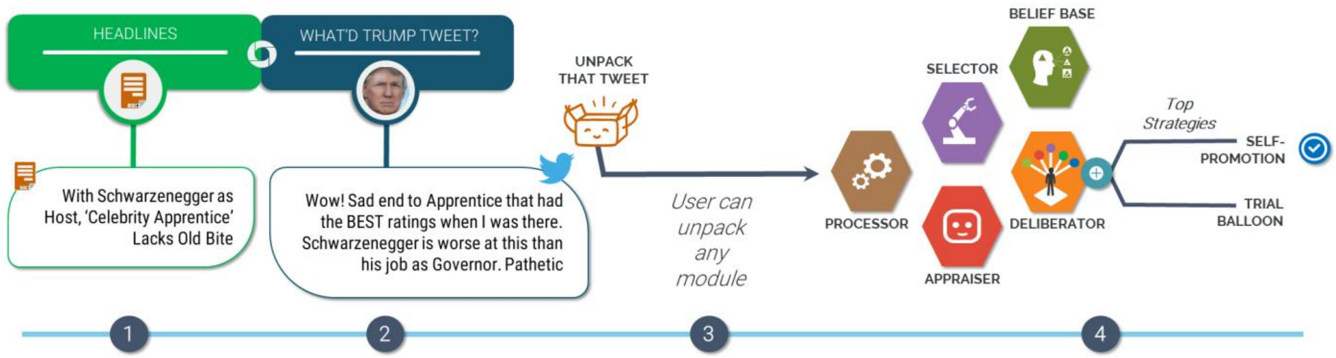

Figure 2: A sample user interaction pipeline unpacking the Deliberator of the system

Once we had identified this tweet corpus, we worked to construct a "belief base." We began by determining the policies of the 2016 Trump campaign, as stated by the campaign website (Trump 2016). Next, we used IBM's Bluemix Natural Language Understanding service to conduct topic modeling and sentiment analysis of the @realDonaldTrump tweet corpus defined above. Together, this allowed us to construct evidence-based beliefs using the words and policies of the campaign and the tweets of the candidate and president. Table 1 describes a list of beliefs stored in the system's belief base.

In parallel, we built the foundation of the Deliberator module, which decides the output of the model. This involved qualitative coding and thematic analysis of a randomly selected subset of 400 tweets from our corpus; we leveraged the aforementioned tweet taxonomy of George Lakoff to assign one or more strategies to each of the tweets in the subset. Two researchers conducted a majority of the coding, and a third reviewed it for reliability purposes. We expanded Lakoff's original taxonomy to include two additional tweet strategies: *self-promotion*, in which President Trump praises himself, his cabinet, or his policies, and *mobilization*, in which he calls for people to take some collective action.

Finally, we used these tweeting strategies to create a grammar. The grammar is a set of rules that generates synthetic sentences. In our case, these sentences are simulated tweets, which are the outputs of the model. To contextualize the grammar construction, we used thematic analysis (Aronson 1995).to create prototypical responses that exemplify each tweeting strategy. Multiple strategies, such as deflection and diversion, can trigger the same grammar rules.

The Model

Our system receives as input a news headline and returns as output a simulated tweet. The user can choose to unpack and trace the processes underlying each component of the system. At the center of this system is a generative model (Fig. 1) which is comprised of the following modules:

The Processor

Upon receiving a news headline, the model sends the text to the Processor module, which is tasked with extracting the key entities and topics from the text. This task is accomplished by the Natural Language Understanding service of IBM's Bluemix. With a list of relevant topics in hand, the Processor turns to the Appraiser module to determine the affective information contained in the news headline.

The Appraiser

The Appraiser component detects the sentiment and emotion expressed in the news headline. This, too, is accomplished by Bluemix's Natural Language Understanding service. "Sentiment" is represented as a real number on the interval [-1, 1], where -1 represents negative sentiment and 1 represents positive sentiment. "Emotion" is represented as a five-tuple of real values, where each value pertains to the strength of the following five emotional elements: anger, joy, sadness, fear, and disgust (each on the interval [0, 1]). The list of topics, the sentiment, and the emotion tuple (collectively called the "headline belief") are sent to the Selector.

The Selector

The Selector examines the topics found in the headline and compares them against the topics of the beliefs stored in the belief base (which are encoded in a list just as the headline beliefs are). It chooses the most relevant belief according the the selected topics. For example, if the input news headline talks about "Obama" and "Healthcare", the selector will search for a belief which is centered on "Obama" and "Healthcare". The headline belief and the selected belief (if there is one) are both sent to the Deliberator component.

The Deliberator

The Deliberator is tasked with performing an action according to the alignment between the headline belief and the selected belief. This alignment is determined by finding the difference between the sentiment scores and the emotion scores. If no relevant belief was chosen by the

Input Headline	Strategy Invoked	Output Simulated Tweet
“Trump was right after all about the Obama administration wiretaps”	Boasting	“I have proven to be far more correct about Obama than anybody- and it's not even close. #MAGA”
“Trump needs a better approach to immigration because bullying isn't cutting it”	Diversion	“Pathetic! The FBI is totally unable to stop the national security 'leakers'. We must put #AMERICAFIRST”
“Twitter destroys 'crack pipe' Clinton adviser who suggests the media owes Hillary an apology”	Boasting	“I have proven to be far more correct about Clinton than anybody- and it's not even close. #MAGA”
“Russia: The scandal Trump can't shake”	Deflection	“Reports about me and Russia have zero credibility. Why is failing BBC joining the 'witch hunt'? Pathetic Dems conspiracy?”

Table 2: A sample of input headlines and corresponding selected strategies and generated tweets.

Selector, the system will make no response. If the news headline "agrees" with the system's belief, a positive tweet will be generated. If the headline "disagrees" with the system's belief, a negative tweet will be generated. The tweet is generated according to the list of hand-authored grammars and the strategies developed from Lakoff’s taxonomy and our own extensions to it. Table 2 shows a collection of sample news headlines, strategies selected, and the respective tweets generated by the system.

User Interaction

The user can supply a headline to the system and request a simulated tweet to be generated about the headline. Once the tweet has been generated, the user may request to “unpack” the tweet generation process in order to find out what decisions the system had made. This exposes the mechanisms of the system that can be interpreted to gain insights.

Contextualizing the Information The data representations that drive the system, in their raw form, are not easily interpretable by human beings. Since the goal is to enable inspection and understanding of the inner workings of the system, there is a need to convert those inner workings into a human-readable format. To this end, the system supplies users with contextualized descriptions of the sentiment and emotion data that propagate through the system. For example, the system contextualizes a sentiment score of 0.75 as “positive sentiment” and a score of 0.1 as “no sentiment.”

Belief Adjustment The user may look at the relevant beliefs held in the system’s belief base and judge them for accuracy. For example, if the user sees that the President Trump model has a positive attitude for Hillary Clinton, the user may adjust the belief.

Interaction Pipeline Fig.2 outlines the pipeline of an example user interaction. In this particular example, the user is interested to know what President Trump might tweet about in response to a negative headline about Arnold Schwarzenegger’s job as host of Celebrity Apprentice, the President’s former reality show. First, the user selects that headline. Second, the system generates a tweet based on this headline. Third, the user unpacks the tweet because she wants to trace through the system’s underlying processes. Specifically, she wants to know about the strategy used to produce this tweet. Fourth, after an exploratory inspection of each component, she finds that the Deliberator reveals the chosen strategy. The chosen strategy was self-promotion because the headline sentiment and emotions matched those of its counterpart in Donald Trump’s modeled belief system.

Plans for Assessment

Now that we have outlined the process of creating our model, we provide a plan for assessing the system using open cycles of inquiry. The goal is to gain insights into the “behind the scenes” processes of a person’s simulated tweeting behavior. In this particular case, we are focusing on the Twitter account for President Trump, @readDonaldTrump. We plan to recruit 40-60 adult participants of various backgrounds for our assessment study.

In order to assess the performance of the system, we will conduct the following steps. The participant first interacts with the system with no traceability. That is, the participant selects a news headline and a tweet is generated: this is the *black-boxed* state. Next, the participant interacts with the *clear-boxed* traceable system that allows unpacking of the decision-making process behind the tweet generation. To control for ordering effects, we will randomize the order of the interactions with the black- and clear-boxed systems.

While interacting with the traceable system, the participant engages in a think-aloud protocol study while conducting the following tasks: the participant unpacks the tweet generation processes for 10 simulated tweets by “unpacking” through the Deliberator, Selector, and Appraiser. For each of 10 tweets unpacked, the participant also engages in secondary research, finds a relevant topic to update the belief base, and observes the resulting changes in system output. Next, the participant answers a questionnaire focused on three main areas: understanding the helpfulness of traceability, the section(s) where it was useful, and the efficacy of updating the belief base.

Multiple choice questions and Likert scale-based questions help us gain insights into each of the three areas.

Once the participant is done interacting with the system, we will conduct a short semi-structured interview focused on the qualitative experience of using the system. Specifically, we will try to understand three things: the experiential differences between the black- and clear-boxed systems, whether the system has enabled the participant to learn something new about the process, and any insights on the participant's expectations of the system. The assessment study will end with a quick collection of demographic data, which will be used for analysis after being adequately anonymized for personally identifiable information.

Discussion

Our system is the first, to our knowledge, to explore the workings of the "black box" of simulated Twitter rhetoric from world leaders. With Twitter playing a key role in political engagement, a systematic understanding of the tweeting behavior using a computational lens can afford us proactivity in the social discourse. Here are the main strengths of the system.

First, our system allows the user to embark on exploratory journeys by revealing the decision-making rationale behind the generation of simulated tweets. The generated tweet anchors the exploratory process, while the ability to unpack modules in the system can facilitate interpretation of the system's actions. Given that our belief base and underlying data are taken from primary sources (the President's Twitter account and stated policies), there is built-in plausibility. Also, our grammar is directly generated from portions of actual tweets from @realDonaldTrump. Therefore, the simulation generated by the model is grounded in evidence.

Second, exploration with open cycles of inquiry is likely to facilitate unexplored relations between concepts resulting in deeper understanding of a phenomenon (De Bono 1993). Moreover, it can lead a user down a novel, useful, and surprising path that catalyzes generative processing towards creative thinking and understanding of the topics (Mayer and Moreno 2003). Data-driven simulations that are traceable and interpretable can help improve conceptual understanding of rationale behind the tweeting behavior.

Third, our system has interpretive flexibility by design. We are inspired by Bijker et al.'s (1987) notion of *interpretive flexibility,* where relevant social groups give meaning and interpretation to the design of the technology. Since the model affords traceability and interpretability, each user group can tinker, simulate, and reappropriate the system to their own use cases. As a result, there is flexibility in the adoption of the system according to the meaning assigned by the relevant groups. Users of different backgrounds will each be able to find utility in their own ways. For instance, activist groups interested in a certain topic, say climate change, can simulate the tweeting behavior using appropriate headlines. As climate change activists explore through the unpacking journey, they can gain insight about the subject of the model's potential tweeting strategies. This insight can help them to be proactive rather than reactive in their civic engagement and social commentary.

Fourth, when the system unpacks the pipeline of its decision-making process, it inherently allows the user to understand any limitations by allowing update of its knowledge base. Therefore, there is some level of in-built evaluation in the interaction itself. Additionally, allowing users to update beliefs and make changes to the model enables us to leverage human computation to further improve our system.

With the strengths of the system in mind, there are limitations to acknowledge. First, with this type of model, we run the risk of ascribing intentionality where none exists. As a result, the system may "force" a response and reasoning that may be unrealistic: given a tweet, the system will inherently try to take an action, which in some cases may not be an appropriate response. Second, since we hand-coded a sample of 400 of President Trump's tweets, there may be tweeting strategies that are not covered by our classifications. However, our evaluation plan can help us to uncover some of these types of errors and make plans to correct them in the future. Third, there is a labor-intensive requirement associated with building the belief base of a particular world leader. There is no way to automate the process of grounding the belief base: researchers must manually identify the beliefs and their sources. There is additional reliance on the expertise of the researchers in terms of thematically analyzing each belief and reliably encoding it in the model.

We focus on world leaders as use cases for our system for two main reasons. First, simulating "what-if" scenarios of behavior from world leaders is likely to be consequential from a social justice perspective. Second,, the likelihood of finding accessible and cross-verifiable content describing the leaders' beliefs is higher compared to that of others in society. Our cognitive model, however, can be applied to anyone: as long there is a verifiable process of constructing a robust belief-base. Not only does the content about a person's policies, goals, and beliefs need to exist in an accessible format, but the content should also be cross-verifiable. Our transferrable model-generation process can help to cross-verify and ground the contents of the belief-base of other leaders. In addition to other things, it can be used to cluster multiple primary sources (be it tweets, speeches, policies, actions, etc.), find commonality between them, and construct the belief base. However, belief bases should be constructed acknowledging human fallibility of the leader that may result in some inconsistency between stated policies and actions. Researchers should maintain rigor in the process of construction and ensure that internal deviances are within reasonable levels. The base can also be updated and revised over time, improving the quality of the simulations and the utility of the system.

Our system has important implications for social commentary. The inner workings of a democratic society should not be a black box, and understanding the processes of leader rhetoric is a crucial step in making it more interpretable and actionable. We hope that our model can leverage computation and user interaction in order to facilitate social justice engagement. Moreover, in the future, we can transfer components of the current system to create a platform that can process past tweets from world leaders and reverse-engineer potential underlying strategies. Making the implicit explicit in these ways enables proactive participation in economic society, and this is ever more important in an increasingly polarized world.

Conclusion

In this paper, we have bridged the gap between computation and analysis by presenting a traceable and interpretable analytic tool that can help individuals gain insight into the potential rationale behind simulated tweets. We will assess the usefulness of this traceable system by comparing our clear-boxed model to a black-boxed one. We hope that our system, its interpretability, and the transferability of the process can encourage user interaction and work to increase social engagement. Understanding decision-making processes behind the communication strategies of world leaders is a crucial component of improved civic engagement in our ever-connected world.

Acknowledgements

We thank Prateek Mittal, Ashok Goel, Jose Delgado, and Joshua Killingsworth for their help in this process.

References

Aronson, J. (1995). A pragmatic view of thematic analysis. The qualitative report, 2(1), 1-3

Bijker, W. E., Hughes, T. P., Pinch, T., & Douglas, D. G. (2012). The social construction of technological systems: New directions in the sociology and history of technology. MIT press.

Braun, S. (2017). National Archives Tells the White House to Save All of President Trump's Tweets. *Time*. Retrieved 23 April 2017, from http://time.com/4723685/donald-trump-tweets-national-archives/

Bukhari, J. (2017, February 14). Why Trumponomics Could Make or Break Boeing. Retrieved April 23, 2017, from http://fortune.com/2017/02/14/boeing-stock-trump/

Conover, M., Ratkiewicz, J., Francisco, M. R., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political Polarization on Twitter. *ICWSM*, *133*, 89-96.

Conway, B. A., Kenski, K., & Wang, D. (2015). The rise of Twitter in the political campaign: Searching for intermedia agenda-setting effects in the presidential primary. Journal of Computer-Mediated Communication, 20(4), 363-380.

Chronicle, D. (2014, September 03). Meet the man behind PM Narendra Modi's strong online presence, Japanese tweets. Retrieved April 24, 2017, from http://www.deccanchronicle.com/140903/nation-current-affairs/article/meet-man-behind-pm-narendra-modi's-online-presence

Eltantawy, N., & Wiest, J. B. (2011). The Arab spring| Social media in the Egyptian revolution: reconsidering resource mobilization theory. International Journal of Communication, 5, 18.

Farrell, H. & Drezner, D. (2008) 'The power and politics of blogs', Public Choice, vol. 134, no. 1, pp. 15–30.

Gentzkow, M. (2016). Polarization in 2016. *Toulouse Network of Information Technology white paper*.

Harrison, B., Ehsan, U., & Riedl, M. O. (2017). Rationalization: A Neural Machine Translation Approach to Generating Natural Language Explanations. arXiv preprint arXiv:1702.07826.

Hart, R. P. (1987). The sound of leadership: Presidential communication in the modern age. University of Chicago Press.

Hayes, B. (2016). DeepDrumpf.com. Retrieved 23 April 2017 from http://www.deepdrumpf2016.com/

Honey, C., & Herring, S. C. (2009, January). Beyond microblogging: Conversation and collaboration via Twitter. In System Sciences, 2009. HICSS'09. 42nd Hawaii International Conference on (pp. 1-10). IEEE.

IBM. (2017). IBM Watson Bluemix. Retrieved 23 April 2017 from https://www.ibm.com/cloud-computing/bluemix/watson

Joyner, D. A., & Goel, A. K. (2015, March). Improving Inquiry-Driven Modeling in Science Education through Interaction with Intelligent Tutoring Agents. In Proceedings of the 20th International Conference on Intelligent User Interfaces (pp. 5-16). ACM.

Joyner, D. A., Goel, A. K., & Papin, N. M. (2014, February). MILA--S: generation of agent-based simulations from conceptual models of complex systems. In Proceedings of the 19th international conference on Intelligent User Interfaces (pp. 289-298). ACM.

Keohane, J. (2016). The Cry-Bully: The sad mind and evil media genius behind @realDonaldTrump. *Politico*. Retrieved 23 April 2017, from http://www.politico.com/magazine/story/2016/04/2016-donald-trump-politics-campaign-twitter-social-media-213827

Kollanyi, B., Howard, P. N., & Woolley, S. C. (2016). *Bots and automation over Twitter during the first US Presidential debate*. COMPROP Data Memo.

Kotoky, A. (2014). Modi Taps Selfie Generation to Oust Twitter-Absent Gandhis. (2014, May 19). Retrieved April 24, 2017, from https://www.bloomberg.com/news/articles/2014-05-18/modi-taps-india-selfie-generation-to-oust-twitter-absent-gandhis

Lakoff, G. (2017). Trump's Twitter Distraction. Retrieved 23 April 2017, from https://georgelakoff.com/2017/03/07/trumps-twitter-distraction/

Marshall, J. (2016). Metrical Analysis of Trump Insult Haiku. *Talking Points Memo*. Retrieved 23 April 2017, from http://talkingpointsmemo.com/edblog/metrical-analysis-of-trump-insult-haiku

Mathew, I. (2017). Here's What a Trump Tweet Does to a Company's Share Price. Fortune.com. Retrieved 23 April 2017, from http://fortune.com/2017/02/24/trump-tweet-stocks/

De Bono, E. (1993). Serious creativity. New York, N.Y.: HarperBusiness.

Mayer, R. & Moreno, R. (2003). Nine Ways to Reduce Cognitive Load in Multimedia Learning. Educational Psychologist, 38(1), 43-52. http://dx.doi.org/10.1207/s15326985ep3801_6

McGill, A. (2017, March 28). A Bot That Can Tell When It's Really Donald Trump Who's Tweeting. *The Atlantic*. Retrieved April 23, 2017, from https://www.theatlantic.com/politics/archive/2017/03/a-bot-that-detects-when-donald-trump-is-tweeting/521127/

Noguchi, Y. (2017). Some Firms Are Harnessing Trump's Tweets As A Marketing Strategy. NPR.org. Retrieved 23 April 2017, from http://www.npr.org/2017/01/24/511415043/how-businesses-are-using-trump-s-tweets-as-a-marketing-strategy

Papert, S. (1980). Mindstorms: Children, computers, and powerful ideas. New York: Basic Books.

Piaget, J. (1950). The Psychology of Intelligence. New York: Routledge.

Phillips, A. (2015, December 10). The surprising genius of Donald Trump's Twitter account. Retrieved April 24, 2017, from https://www.washingtonpost.com/news/the-fix/wp/2015/12/10/reading-6000-of-his-tweets-has-convinced-us-donald-trump-is-a-social-media-master/

President Trump's Tweets, Annotated. (2017). NPR.org. Retrieved 23 April 2017, from http://www.npr.org/2017/01/27/511216062/president-trumps-tweets-annotated

Rainie, L., Smith, A., Lehman Schlozman, K., Brady, H. & Verba, S. (2012) 'Social media and political engagement', Pew Internet & American Life Project, Washington DC. [Online] Available at: http://pewinternet.org/~/media//Files/Reports/2012/PIP_SocialMediaAndPoliticalEngagement_PDF.pdf

Ryall, J. (2017). Twitter diplomacy: South Korea appoints officer to monitor Donald Trump's tweets for foreign policy insights. *The Telegraph*. Retrieved 23 April 2017, from http://www.telegraph.co.uk/news/2017/01/05/twitter-diplomacy-south-korea-appoints-officer-monitor-donald/

Schwarz, C. V., & White, B. Y. (2005). Metamodeling knowledge: Developing students' understanding of scientific modeling. Cognition and Instruction, 23(2), 165-205.

Simms, Q. (2016). Scraping 27,000 Tweets for Presidential Sentiment Analysis. Data-Driven Journalism. Retrieved 23 April 2017, from http://datadrivenjournalism.net/news_and_analysis/scraping_27000_tweets_for_presidential_sentiment_analysis

Stieglitz, S., & Dang-Xuan, L. (2012, January). Political communication and influence through microblogging--An empirical analysis of sentiment in Twitter messages and retweet behavior. In System Science (HICSS), 2012 45th Hawaii International Conference on (pp. 3500-3509). IEEE.

Trump, D. (2016). Policies. Retrieved 23 April 2017 from https://web.archive.org/web/20161015091228/https://www.donaldjtrump.com/policies/

Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2010). Predicting elections with twitter: What 140 characters reveal about political sentiment. ICWSM, 10(1), 178-185.

Washington, J. M. (1986). The essential writings and speeches of Martin Luther King, Jr. New York: HarperOne.

Wattal, S., Schuff, D., Mandviwalla, M. & Williams, C. (2010) ‘Web 2.0 and politics:the 2008 U.S. presidential election and an e-politics research agenda’,Management Information Systems Quarterly, vol. 34, no. 4, pp. 669–688.

Wieczner, J. (2016, December 14). When Donald Trump Hate-Tweeted Lockheed Martin, Hedge Funds Were Ready. Retrieved April 23, 2017, from http://fortune.com/2016/12/13/donald-trump-twitter-lockheed-martin-stock-dump/

Woods, W. 1981. Procedural semantics as a theory of meaning. In Joshi, A.; Webber, B.; and Sag, I., eds., *Elements of Discourse Understanding*. Cambridge, UK: Cambridge University Press. 300–334.

Yamabhai, S. (1973). The rhetoric of non-violence: a critical analysis of selected speeches by MK Gandhi (Doctoral dissertation, The Ohio State University).

Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., & Lipson, H. (2015). Understanding neural networks through deep visualization. arXiv preprint arXiv:1506.06579.

Yu, L. (2005). The Great Communicator: How FDR's Radio Speeches Shaped American History. The History Teacher, 39(1), 89-106.

Zeiler, M. D., & Fergus, R. (2014, September). Visualizing and understanding convolutional networks. In European conference on computer vision (pp. 818-833). Springer International Publishing.