Engelbart

An Agent-Based Interface for Exploring Beyond the Social Media Filter Bubble

Eric Duong*
Garrett Vercoe*
Ehsan Baharlou
University of Virginia School of Architecture

*Authors contributed equally to the research

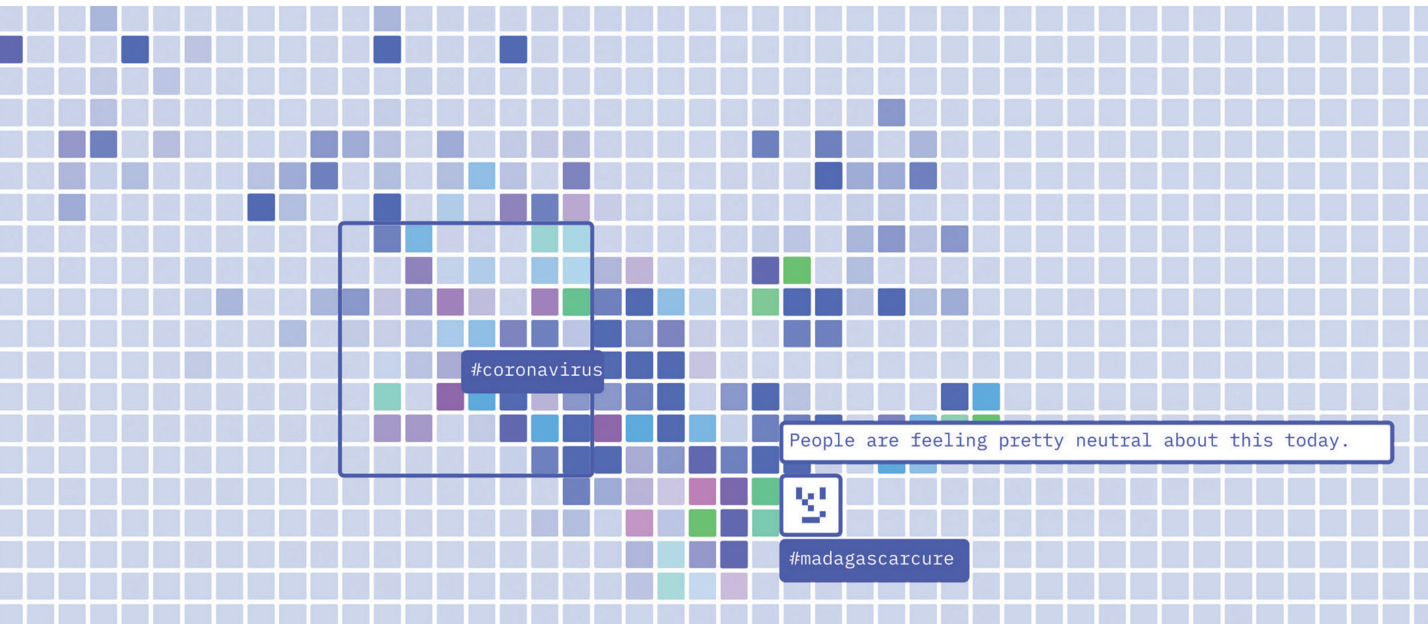

1

1 Engelbart, the autonomous agent, approaches a topic discussion and summarizes what people are saying and feeling about it.

ABSTRACT

The internet has long been viewed as a cyberspace of free and collective information, allowing for an increase in the diversity of ideas and viewpoints available to the general public. However, critics argue that the emergence of personalization algorithms on social media and other internet platforms instead reduces information diversity by forming "filter bubbles" of viewpoints similar to the user's own. The adoption of these personalization algorithms is due in part to advancements in natural language processing, which allow for textual analysis at unprecedented scales.

This paper aims to utilize natural language processing and architectural spatial principles to present social media from a collective viewpoint rather than a personalized one. To accomplish this, the paper introduces Engelbart, a data-driven agent-based system, where real-time Twitter conversations are visualized within a two-dimensional environment. This environment is interacted with by the artificial intelligence (AI) agent, Engelbart, which summarizes crowdsourced thoughts and feelings about current trending topics. The functionality of this web application comes from the natural language processing of thousands of tweets per minute throughout several layers of operations, including sentiment analysis and word embeddings. Presented as an understandable interface, it incorporates the values of cybernetics, cyberspace, agent-based modeling, and data ethics to show the potential for social media to become a more transparent space for collective discussion.

INTRODUCTION

Personalization algorithms are increasingly deciding what content users see online (Hannak et al. 2013). Without their approval, users are becoming isolated in ideologically segregated "filter bubbles" (Pariser 2011) that selectively expose content intended to keep their attention (Bakshy, Messing, and Adamic 2015). These filter bubbles have garnered concern as they form echo chambers that amplify viewpoints and limit the breadth of discussion (Sunstein 2017). Social media now resembles a "Daily Me" (Negroponte 1995) platform where users live in a private, personalized network of information. In this age of social isolation, there is a growing need for a "Daily Us" (Negroponte 1995), a place to engage with the distributed public network. Progress in natural language processing (NLP) (Joulin et al. 2017) allows for this unrestricted public view of information, yet there are inherent difficulties in visualizing and interfacing with the scale of these networks (Rafiei and Curial 2005). The predominant social media interface of the linear feed is not able to show this holistic view of the wider network, and while conducted research has explored methods of visualizing and communicating at the network scale (Correa and Ma 2011; Kitchin 1998), development is lacking in translating these networks into understandable interfaces.

This paper introduces Engelbart,[1] a real-time web application that presents the Twitter content network through an agent-based system. The goal of this spatial interface is to create a digital public space that promotes discovery of new and diverse content, providing an alternative to the informational echo chambers that often occur in linear feeds.

BACKGROUND

Understanding Cyberspace

Engelbart views Twitter as a "cyberspace" (Gibson 1984) unconstrained by the algorithmic content recommendations that form filter bubbles. At the beginning of the internet, this metaphor of cyberspace emerged as a conceptualization of a new information-based space that transcended physical boundaries (Benedikt 1991). Composed of the flow of data between computers, cyberspace is simultaneously virtual and real as a construct of thought (Novak 1991). For instance, a message sent from one user to another is said to exist in cyberspace but cannot be located as a tangible object. This ambiguity does not mean that cyberspace is unable to be scrutinized—quite the opposite. Because data is the fundamental building block of cyberspace, its information can be processed quickly and efficiently. In doing so, it is possible to uncover invisible relations through the mapping from data to representation (Novak 1991).

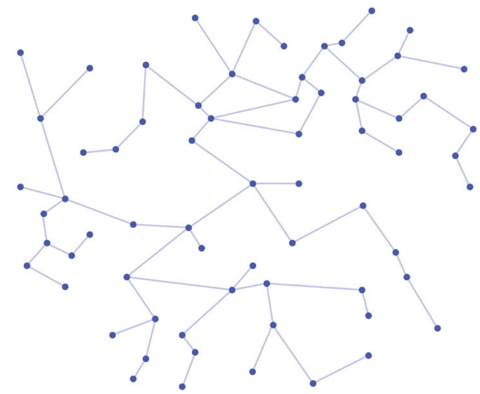

2 Node-edge graph.

Advances in NLP (Joulin et al. 2017) have accelerated efforts to analyze human language data at scale. Recent studies have explored ways to represent larger cyberspatial networks built from rich data (Chen et al. 2018; Correa and Ma 2011; Ma and Muelder 2013), but they often resemble the established form of a node-edge graph (Fig. 2) (Battiston et al. 2020). Within the context of cyberspace, nodes represent individual units such as computers or users while the edges that connect them represent the transfer of data (Boccaletti 2006). Projects like Social Mirror (Gillani et al. 2018) have been able to use this layout to effectively show the exacerbation of political polarization online, but are unable to show higher-order interactions. Because edges can only connect pairs of nodes, the node-edge graph cannot represent real social systems where interactions regularly occur in groups of three or more nodes (Battiston et al. 2020). This feature of node-edge graphs makes them difficult to scale for the representation of larger networks (Correa and Ma 2011).

The computational landscape today has shifted toward enabling computers to solve problems in their own way. As data becomes more abundant and rich, this approach has become not only more effective but even necessary. Just as architects are beginning to employ machine intelligence for form-finding (Carpo 2017), this research uses it to build a dynamic representation of cyberspace that shows high-order interactions of Twitter as a whole. This allows for it to act as a space of its own, and presents the opportunity to introduce an occupant that offers more detailed information from cyberspace. This translation of cyberspace into an occupiable space becomes an agent-based model—a series of simulated interactions between an autonomous entity and its environment (Jennings 2000).

Agent-Environment Interaction

Agent-based modeling is the computational simulation of the interactions between one or more autonomous entities and their rule-based environment (Jennings 2000). Engelbart uses this approach to create an interface that delivers information from two scales: the individual agent and the collective environment. As a social media alternative, both the agent and the environment are designed to be understandable for the end user.

For the agent, this project looks to cybernetics to inform the way it negotiates the environment and relates to the user. As the study of communication between human and machine, the field of cybernetics has long been experimenting with developing and using AI to solve complex design problems. Yet, despite their focus on the leading edge of technology, cyberneticists hold to a philosophy of humanism through machines—believing that machines should not only solve problems but also rationalize and relate to them (Negroponte 1973). Despite decades of technological advances since the first cyberneticists, the field has remained relevant because of this central tenet of humanism. The early cyberneticist Gordon Pask (1962) maintained that the act of conversation is enough to be a meaningful cybernetic interaction, which has become a point of inspiration for the interaction between agent and user.

For the environment, psychogeography is foundational to visualize the high-order network of cyberspace. In the same way that city plans do not reveal the social complexities of a community, the objective mapping of networks in node-edge graphs limits the view of collective interactions. In the Naked City (Debord 1955), Guy Debord explores how the city plan can be deconstructed and layered with subjective interpretations of space. The result is a map that views Paris from an experiential lens that focuses on the social centers of the city. With NLP, this same approach can be applied to understand the social dynamics of the web and create a human-readable subjective mapping.

The Ethics of Cyberspace

Cyberspace is not only a metaphor for the internet but also an ideology. Alongside the idea of cyberspace came ideas for how it should be governed. At its inception, the complete freedom of the internet encouraged ideals of equal access to information. The 1996 Declaration of the Independence of Cyberspace reflected this belief, calling for the civilization of cyberspace to be "more humane and fair than the world your governments have made before" (Barlow 1996, 7). Yet the modern internet has only become more locked down as companies consolidate control over

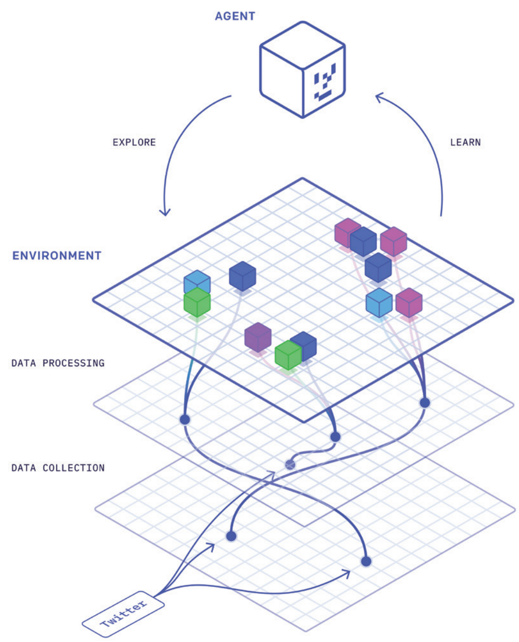

3 The framework for construction of the agent-based model.

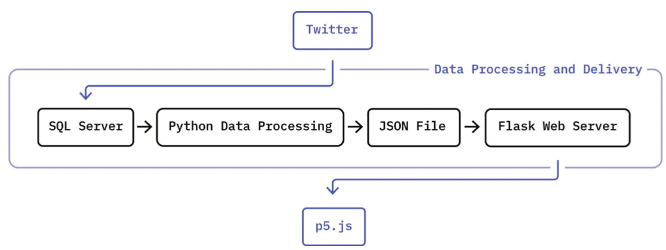

4 Technical system diagram of the web application.

communication, commerce, and search (Patelis 2013). In the social media sector, projects like Civic Signals (Pariser 2019) are working to rebuild platforms as true public spaces. Principles of urban planning are central to their proposals and include ideas like online parks or digital cities (Pariser 2019) free from the algorithms that lock users into filter bubbles. This project upholds the original promise of a humane cyberspace while also looking toward these spatial models for transparent and open online spaces. In combination with agent-based modeling, Engelbart presents a more equitable and understandable cyberspace.

METHODS

Design Framework

Personalization algorithms cause users on social platforms such as Twitter to engage with the app for longer, with the downside of forming narrow filter bubbles of information. To provide an alternative to a user's filter bubble, the Engelbart web application shows the collective content on the Twitter platform rather than a select

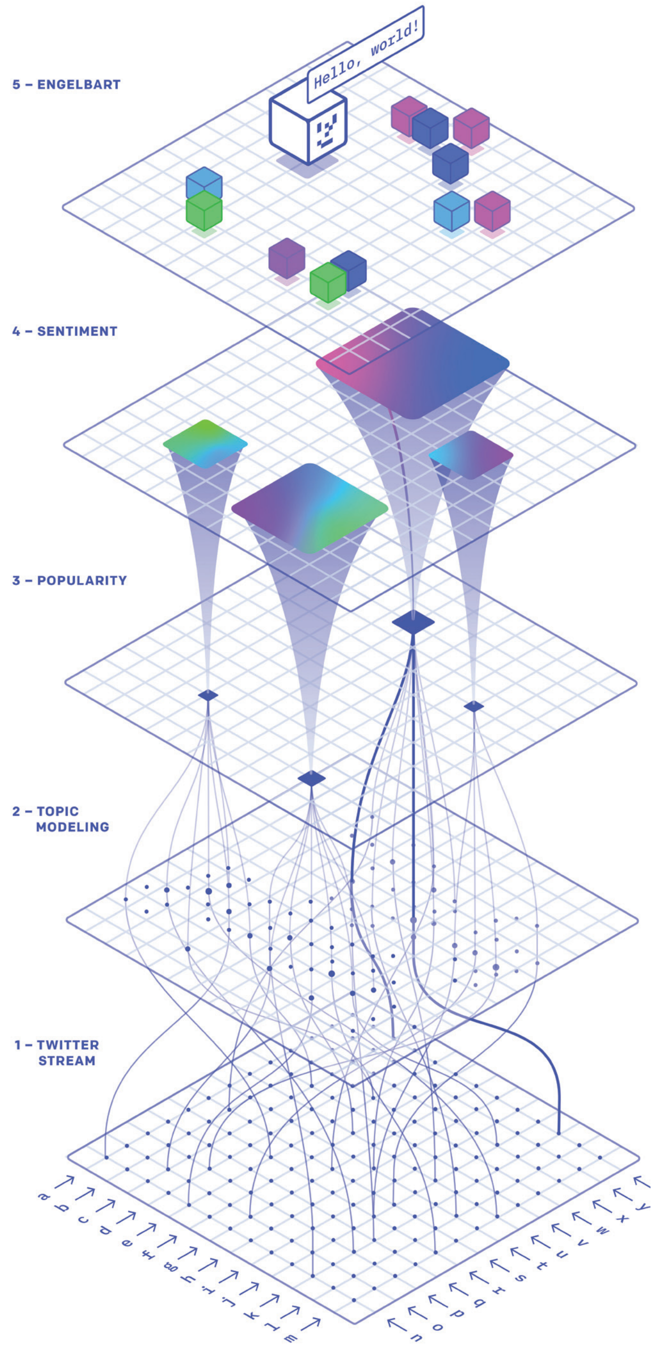

5 (1) The bottom layer represents the raw data capture from a live Twitter stream. On the way to the next layer, sentences are divided into words in preparation for further processing. (2) Words are vectorized and rearranged within multidimensional space based on topic similarity. (3) The vector space is flattened to reveal dense clusters of varying size. (4) Sentiment analysis visualizes topic opinions through color. (5) Topic clusters become squares in the grid interface that reflect the sentiment of the tweets that make them up. Engelbart studies each topic and relays what it has learned to the user.

slice. However, the collective content of a platform as large as Twitter becomes quickly incomprehensible. To aid in human understanding, the system is structured as an agent-environment model to reflect the mental model of an in-person communication network (Fig. 3).

As this research was conducted during COVID-19 in 2020, the project is designed as an open-source web application[2] so that it can be accessible to distributed audiences and reflect the medium of social media that it pulls from. Because of the dynamic nature of social media, a primary design consideration is to show emerging conversations in real time.[3] The interface is built in Processing and Javascript, while the data collection and processing uses Python and NLP packages to aggregate and summarize the data (Fig. 4).

Aesthetics and Experience

The primary context for understanding this collective view of Twitter for this project is the mapping of multidimensional Twitter data onto a two-dimensional rectangular grid. This superimposition is influenced by the von Neumann neighborhood model for cellular automata (von Neumann 1966) and the traditional checkerboard visualization often used for social dynamic modeling (Hegselmann 2017). This allows for the drawing of higher-order relationships simply through the proximity of cells to one another. The model is also conceptually influenced by the architectural public square (Mattson 1999), providing a structured yet open space for public discourse to be viewed. Within this environment, the agent personifies the Twitter data to form a 1:1 connection between the user and the collective data. This creates a flow of information that is done conversationally by the agent over time, helping to reduce the cognitive load of the user.

Layering System

The agent-environment model is composed of five separate layers: (1) collection and cleaning of Twitter data; (2) topic modeling of categorical clusters; (3) popularity; (4) sentiment; and (5) Engelbart, the agent that visualizes and interfaces with the topical categories. Figure 5 is a process diagram of this system.

Layer 1: Collecting Twitter Data

The first layer and foundation of the system is the data collection layer. This layer utilizes Tweepy (Roesslein 2020), an API to access Twitter, and captures the stream of tweets as they are posted in real time. Tweets are collected in a 10% random sample of overall Twitter activity. Three fields are captured: (1) tweet message;

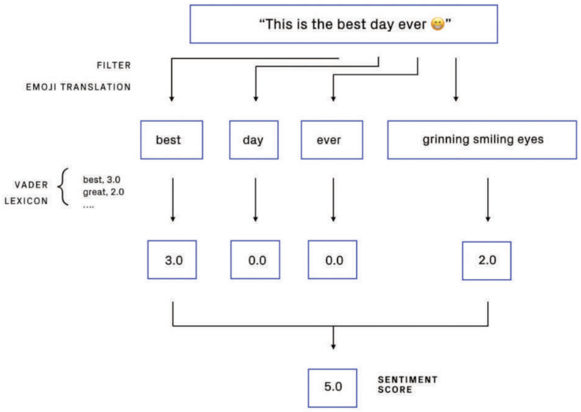

6

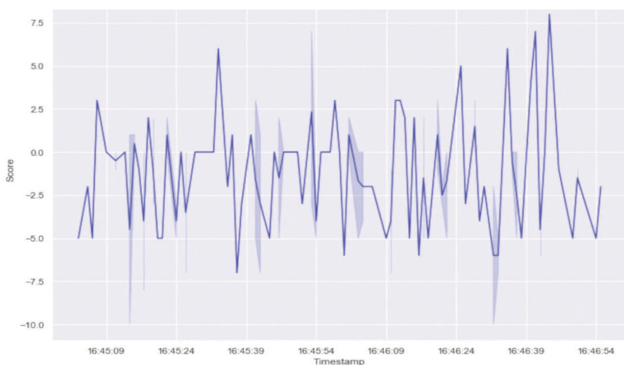

7

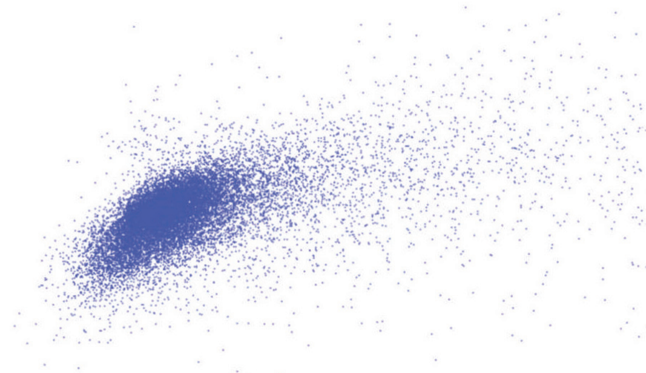

8

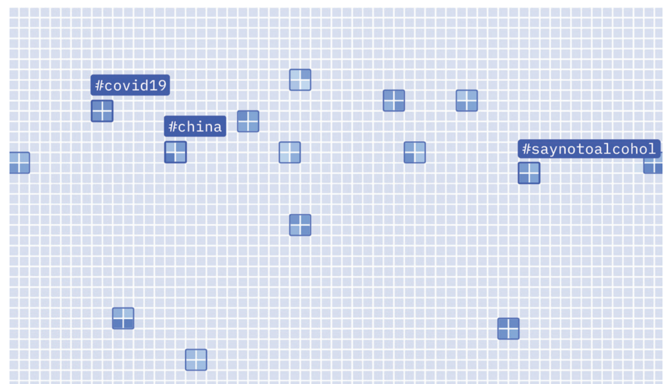

9

6 Tweet sentiment deconstruction by the VADER algorithm.

7 Sentiment score variation of collected tweets over a two-minute period.

8 Vectorized 3D word embeddings from 10,000 tweets, captured in March 2020 in a 20-minute window.

9 Topic clusters placed within the 2D grid, where closer distance equates to closer topic similarity. #covid19 and #china are shown as neighbors while #saynotoalcohol is farther away.

(2) post time; and (3) sentiment, which is calculated by the VADER algorithm (Hutto and Gilbert 2014). The VADER algorithm is a sentiment lexicon library specifically attuned to social media (Fig. 6). No other content is recorded, to protect the privacy of users.

Before this data is piped to the rest of the system, the tweets are tokenized and stemmed with the Natural Language Toolkit (NLTK) package so that they can work with NLP libraries in the following layers. As the stream is taking in tweets globally, tweets are filtered to include only messages written in English. The database is sent to the rest of the system every 15 minutes for processing.

Layer 2: Topic Clustering with Deep-Word Embeddings

The content discussed on Twitter is varied and broad. Spatializing tweets on a 1:1 basis quickly becomes incomprehensible. To help make sense of the larger picture, tweets are aggregated into groups based on their hashtag and keyword usage. These groups are then given a multidimensional vector based on their content similarity to other tweets and clustered by FastText (Mikolov et al. 2013), a nearest-neighbors word embedding algorithm. The result is a series of topic clusters positioned in space relative to how similar they are to each other (Fig. 8).

The nature of these deep-word embeddings of topic clusters is that their vectors are highly dimensional (n = 100) and impossible to visualize in their original form. To make them understandable to the user, singular value decomposition is used to translate this high-dimensional space into a 2D field that is overlaid and scaled responsively onto the environment's grid system (Fig. 9). This 2D coordinate system allows for a clearer organization of topic clusters to the end user compared to a 3D coordinate system, which creates difficulty in measuring the distance, and thus the similarity, between clusters.

Layer 3: Topic Popularity

Alongside topic distribution, the size of each topic cluster corresponds to the amount of user activity within the topic, rather than tokenized measures such as likes or retweets. The more tweets that mention a particular topic, the wider the area of squares covered on the grid environment (Fig.

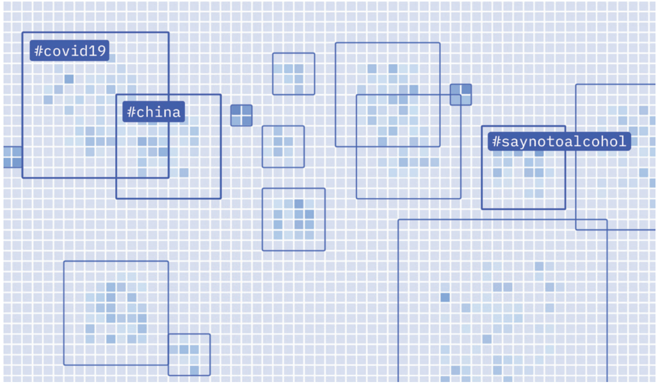

10

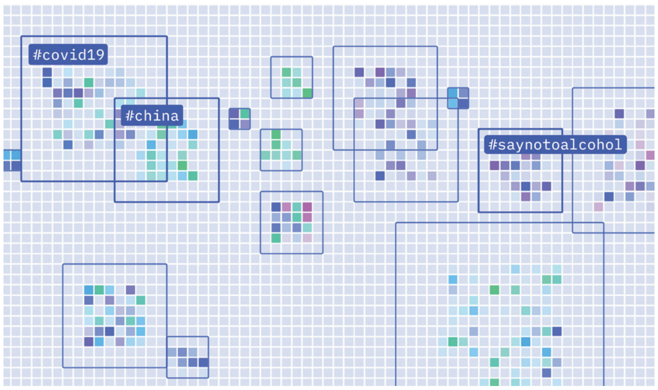

11

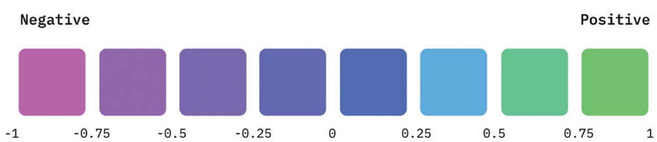

12

10 Topic clusters are scaled based on their popularity. Due to #covid19 being a popular discussion, it is one of the largest topic areas.

11 Topics with sentiment data mapped to their squares. Sentiment regarding #covid19 is mixed, while #saynotoalcohol is entirely negative.

12 The sentiment color scale from most negatively charged tweets to most positively charged. Sentiment scores of tweets in topic are normalized and bucketed onto a -1 to 1 scale.

10). To show the dynamism of conversations happening in real time, squares pulsate and reposition within the topic cluster area.

Layer 4: Sentiment Diversity

Though many people may be conversing about the same topic online, these individual voices often form a diverse array of opinions and feelings in the discussion. To account for these individual voices, when a new conversation topic appears in the application, its sentiment variety is shown through the proportion of colored squares within the topic cluster (Fig. 11). This differentiation is determined by a colored scale of sentiment polarity for the given topic (Fig. 12).

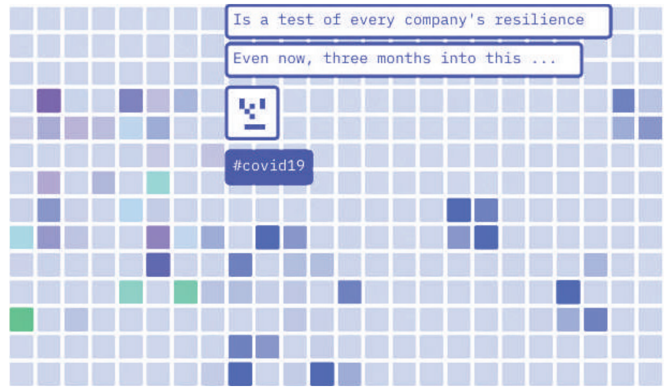

13

14

13 Engelbart reveals the topic cluster of #covid19 with diverse sentiment. He comments on company resilience during COVID-19.

14 User hovers over topic #harrypotterathome. Upon clicking, they are directed to the Twitter page of that discussion.

Layer 5: Interactive Agent

Engelbart, the intelligent agent, learns what people are talking about in the environment and returns a concise summary (Fig. 13). This provides a direct connection between the individual user and the collective user base of Twitter.

On initial load, the only information visible to the user is the location of topic clusters and their popularity. Engelbart must explore the environment to reveal further information about the Twitter topic clusters. To do this, Engelbart approaches the closest unvisited clusters one at a time based on a nearest neighbor search algorithm. Once Engelbart arrives at a cluster, it reveals the summary, sentiment, and hashtag of the topic.

Engelbart's summaries are constructed by utilizing a term frequency-inverse document frequency (TF-IDF) summary algorithm (Bun and Ishizuka 2001), which strings together the most frequently used phrases for the given topic. During its summary, Engelbart interjects emotive language and changes its facial expression in reaction to the overall sentiment of the topic cluster.

After summarizing, Engelbart reveals the sentiment and color diversity of the topic to the user. The user can then

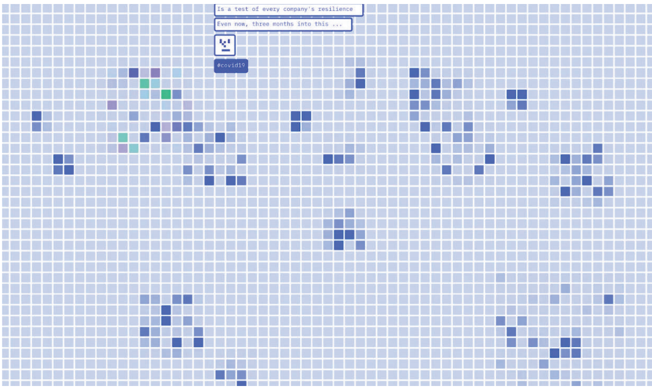

15

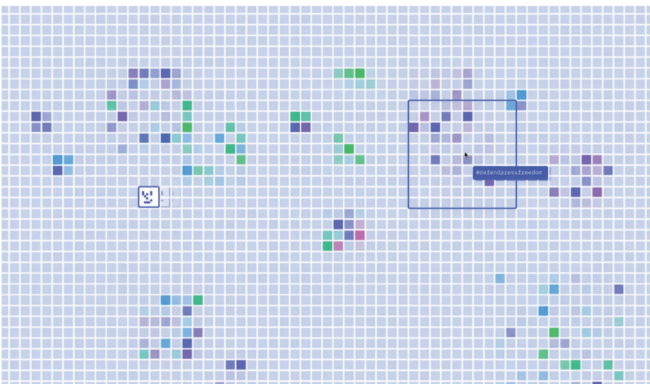

16

15 May 5, 2020, 11:00 am ET. Engelbart reveals the first topic cluster, #covid19, in the environment.

16 May 5, 2020, 11:05 am ET. Engelbart after revealing all of the topic clusters, with #defendpressfreedom highlighted.

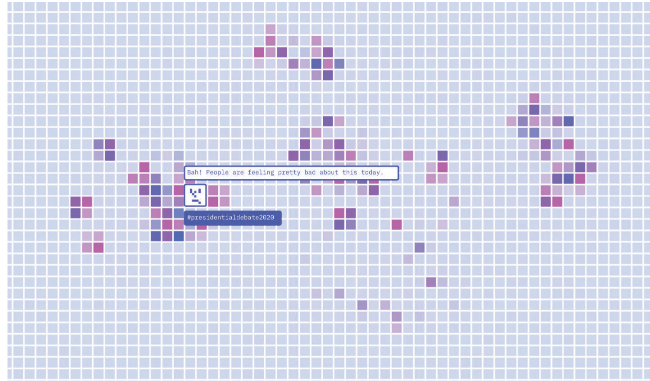

17

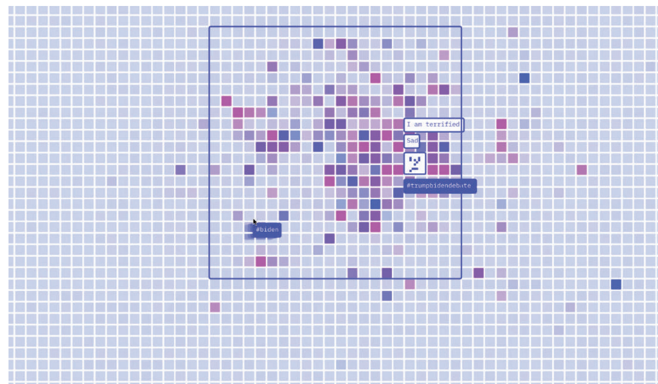

18

17 Sept 29, 2020, 3:30 pm ET. Engelbart with revealed environment, hours before a major presidential debate between Donald Trump and Joe Biden. Several topic clusters of low relativity are apparent, with #presidentialdebate2020 exhibiting mostly negative sentiment.

18 Sept 29, 2020, 9:00 pm ET. Engelbart with revealed environment during the presidential debate. Topic clusters show most all tweets are negatively charged and center around one category: the debate. Engelbart comments on #trumpbidendebate, "Sad. I am terrified."

hover over the cluster to reveal the associated hashtag. Upon clicking the topic cluster, the browser opens a Twitter page for the hashtag in another tab, allowing the user to directly engage in the discussion (Fig. 14).

Engelbart then moves through the environment to the next topic cluster, where it will again reveal hidden information about the topic. This process repeats until all topics have been visited.

RESULTS

Reflection

In its holistic view of Twitter, Engelbart helps users see past their filter bubbles by presenting the collective discussion in its entirety. Its indiscriminate approach to data collection ensures that the diverse opinions of Twitter users are represented in both the agent and the environment. Though the scale of the collected data would be unreadable on its own, visual abstraction and layering of information allows for the interface to embrace the topical complexity of Twitter without misrepresenting its data.

To break the filter bubble, users must be able to listen and relate to the views of others. However, to a user who has been within a filter bubble, the collective discussion is likely to be unfamiliar. Engelbart walks the user through the environment and delivers in-depth summaries after processing topic information (Fig. 15). Though Engelbart is unfamiliar with its environment at first, it eventually reveals all topic clusters to present a comprehensive view of the current discussion. Within this view, users are able to revisit topics of interest (Fig. 16) and navigate directly to Twitter in order to become a part of the conversation.

By disrupting the conventional interface of a linear feed, Engelbart presents a different view of social media. Zooming out to a large scale of information removes references to any one Twitter user and promotes a slower-paced browsing of the site. The translation from data to space becomes a visual language that comes closer to the vision of cyberspace: a virtual public square where topics occupy space and take the shape of the discussion going on within them. When observed at different times, each environment captures a moment on Twitter that can reveal shifting conversations when compared to other moments in time (Figs. 17–18).

Limitations

Engelbart should not be used as a representation of online discussion and society as a whole. Though Engelbart treats data fairly in the collection process, the inherent biases of Twitter itself are reflected in Engelbart and the environment. For instance, the character limit does not allow for long-form expressions of opinion, while the political leanings of the platform may not accurately reflect public viewpoints.

As a hybrid between an agent-based model and user interface, Engelbart should be evaluated in the context of both. The model itself is only able to make a new data pull on a new load, making it unable to change over the course of a single instance. In addition to the limited ways in which Engelbart can change the environment, the agent-environment interaction has been simplified compared to other agent-based modeling research.

Engelbart also has its limitations as an interface. In its current state, the conversational component is a one-way flow of information from computer to user. Users have no part in the dialogue-driven interaction, limiting the depth of the information to the summaries that Engelbart provides. By opening this interaction to user inquiry, Engelbart could provide more specific information relating to the topic and become a more robust tool for learning.

The meshing of agent-based model and interface has allowed for a view past the filter bubble, but there is room for the application to develop as a more comprehensive interactive system. In the exploration of these possibilities, continued research can explore emerging modalities of interaction between model and interface.

CONCLUSION

Building on research of cybernetics, agent-based modeling, and cyberspace, Engelbart posits a transparency of social media through its spatial data interface. The Engelbart web application integrates several layers of data processing of the real-time collective discussion of Twitter and presents it as a 2D agent-environment model. The conceptual framework serves as an early prototype for applying architectural spatial principles towards datacentric online spaces. With Engelbart, users can discover topics and opinions that are not represented in the filter bubbles of their own personal feeds, and be pointed to conversations to which they would not have otherwise contributed.

Future Directions

Further iterations could begin to provide users with more direct control of information retrieval. The ability to manipulate the time of the environment would let users observe changes as topics form and change, acting as a quasi-historical method for understanding tipping points of larger movements or societal thought patterns (Fig. 19). Additionally, incorporating a flexible zoom functionality would allow users to view the data of each cluster with more granularity (Fig. 20). The project demonstrates the potential at a grander scale for self-learning and emergent spatial systems that can catalog and adapt their structure to fit online social media content. Further developments of the system could see Engelbart acting on the form of the environment as it simultaneously changes with the data,

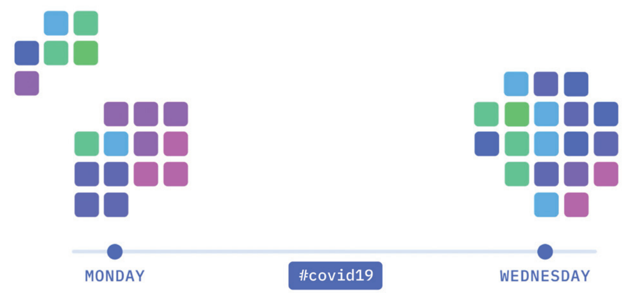

19

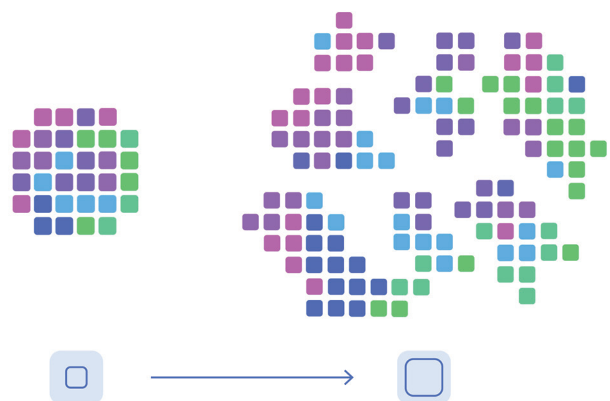

20

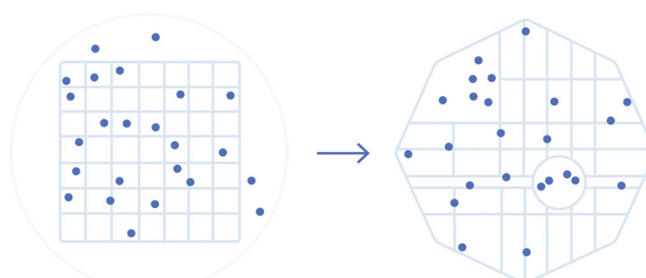

21 In addition to improving the learning potential of the agent in the system for better summarizing of collective topics, there is potential for the spatial environment to be architecturally malleable based on the manifold structure of the word embeddings. This would result in cyberspatial cities of content.

allowing for the emergence of entirely new self-organizing or predictive compositions of cyberspace (Fig. 21). As more of the world goes online, distributed information networks will simultaneously grow in scale and accessibility. This agent-environment model for interfacing with scaling networks allows for an unprecedented ability to discover information from multiple viewpoints, helping to promote equal and ethical online spaces not constrained by the algorithms within which they operate.

ACKNOWLEDGMENTS

This research was completed under the Advanced Technologies undergraduate thesis studio during spring 2020 at the University of Virginia School of Architecture. The primary aim of this studio is to explore different agencies, such as design agency, material agency, and fabrication agency.

The research team would like to thank the following people for helping to improve this research: John Comazzi, Leena Cho, Robin Dripps, Ali Fard, Melissa Goldman, Zaneta Hong, Ryan Shih, and Katelyn Stenger.

NOTES

1. Engelbart is named after Douglas Engelbart, a pioneer in the field of human-computer interaction (Markoff 2013).
2. Open-source code repository available for research at https://github.com/garrettvercoe/Engelbart-Public.
3. A two-minute demo of the application is viewable at https://vimeo.com/411441564.
4. The real-time nature of the application is a primary constraint for the technology stack and allowed data complexity. In all cases, the favored methods and technologies are those that have the quickest processing time. This is why TF-IDF is implemented over GPT-3, FastText over more modern transformers like BERT or ELMo, and singular value decomposition over deep manifold learning.

REFERENCES

Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. "Exposure to Ideologically Diverse News and Opinion on Facebook." In *Science* 348 (6239): 1130–1132.

Barlow, John Perry. 1996. "A Declaration of the Independence of Cyberspace." *Duke Law and Technology Review* 18: 5–7. https://doi.org/10.7551/mitpress/2229.003.0006.

Battiston, Federico, Giulia Cencetti, Iacopo Iacopini, Vito Latora, Maxime Lucas, Alice Patania, Jean-Gabriel Young, and Giovanni Petri. 2020. "Networks beyond Pairwise Interactions: Structure and Dynamics." *Physics Reports* 874 (August): 1–92.

Benedikt, Michael. 1991. "Introduction." In *Cyberspace: First Steps*. Cambridge, MA: MIT Press.

Boccaletti, Stefano, Vito Latora, Yamir Moreno, Dong-Uk Hwang, and Mario Chavez. 2006. "Complex Networks: Structure and Dynamics." *Physics Reports* 424 (4–5): 175–305.

Bun, Khoo Khyou, and Mitsuru Ishizuka. 2001. "Emerging Topic Tracking System." In *Proceedings Third International Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems (WECWIS)*, San Juan, CA, 21–22 June 2001, 2–11. IEEE.

Camp, Jean, and Y.T.- Chien. 2000. "The Internet as Public Space: Concepts, Issues, and Implications in Public Policy." *ACM SIGCAS Computers and Society 30 (3):* 13-19. DOI:10.1145/572241.572244.

Carpo, Mario. 2017. *The Second Digital Turn: Design Beyond Intelligence*. Cambridge, MA: MIT Press.

Chen, Wei, Fangzhou Guo, Dongming Han, Jacheng Pan, Xiaotao Nie, Jiazhi Xia, and Xiaolong Zhang. 2018. "Structure-Based Suggestive Exploration: A New Approach for Effective Exploration of Large Networks." *IEEE Transactions on Visualization and Computer Graphics* 25 (1): 555–565.

Correa, Carlos, and Kwan-Liu Ma. 2011. "Visualizing Social Networks". In *Social Network Data Analytics*, edited C. Aggarwal, 307–326. Boston: Springer.

Debord, Guy. 1955. "Introduction to a Critique of Urban Geography." Translated by Ken Knabb. *Les Lèvres Nues 6* (September): 95–109. https://www.cddc.vt.edu/sionline/presitu/geography.html.

Gibson, William. 1984. *Neuromancer*. New York: Berkley.

Gillani, Nabeel, Ann Yuan, Martin Saveski, Soroush Vosoughi, and Deb Roy. 2018. "Me, My Echo Chamber, and I: Introspection on Social Media Polarization." In *WWW 2018 [Proceedings of the 2018 World Wide Web Conference]*, Lyon, France, 23–27 April 2018, 823–831. International World Wide Web Conferences Steering Committee.

Graham, Stephen. 1998. "The End of Geography or the Explosion of Place? Conceptualizing Space, Place and Information Technology." *Progress in Human Geography* 22 (2): 165–185.

Hannak, Aniko, Piotr Sapiezynski, Arash Molavi Kakhki, Balachander Krishnamurthy, David Lazer, Alan Mislove, and Christo Wilson. 2013. "Measuring Personalization of Web Search." In *WWW 2013 [Proceedings of the 22nd international conference on World Wide Web]*, Rio de Janeiro, Brazil, 13–17 May 2013, 527–538. Association for Computing Machinery.

Hegselmann, Rainer. 2017. "Thomas C. Schelling and James M. Sakoda: The Intellectual, Technical, and Social History of a Model." *Journal of Artificial Societies and Social Simulation* 20 (3): 15.

Hutto, Clayton, and Eric Gilbert. 2014. "Vader: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text." In *Proceedings of the International AAAI Conference on Weblogs and Social Media (ICWSM),* Ann Arbor, Michigan, 1–4 July 2014. AAAI.

Jain, Shikha, and Krishna Asawa. 2015. "EMIA: Emotion Model for Intelligent Agent." *Journal of Intelligent Systems* 24 (4): 449–465.

Jennings, Nicholas. 2000. "On Agent-Based Software Engineering." *Artificial Intelligence* 117 (2): 277–296.

Joulin, Armand, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. "Bag of Tricks for Efficient Text Classification." In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics,* Valencia, Spain, April 2017, 427–431. Association for Computational Linguistics.

Kitchin, Robert. 1998. "Towards Geographies of Cyberspace." *Progress in Human Geography* 22 (3): 385–406.

Ma, Kwan-Liu, and Chris Muelder. 2013. "Large-Scale Graph Visualization and Analytics." *Computer* 46 (7): 39–46.

Markoff, John. 2013. "Computer Visionary Who Invented the Mouse." New York Times, July 3, 2013. https://www.nytimes.com/2013/07/04/technology/douglas-c-engelbart-inventor-of-the-computer-mouse-dies-at-88.html.

Mattson, K. 1999. "Reclaiming and Remaking Public Space: Toward an Architecture for American Democracy." *Nat Civic Rev* 88: 133–144.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." In *Workshop Track Proceedings of the 1st International Conference on Learning Representations*, Scottsdale, Arizona, 2–4 May 2013. ICLR.

Negroponte, Nicholas. 1973. *The Architecture Machine*. Cambridge, MA: MIT Press.

Negroponte, Nicholas. 1995. *Being Digital*. New York: Knopf.

Novak, Marcos. 1991. "Liquid Architectures in Cyberspaces." In *Cyberspace: First Steps*, edited by Michael Benedikt, 225–254. Cambridge, MA: MIT Press.

Pariser, Eli. 2011. *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. New York: Penguin Publishing Group.

Pariser, Eli. 2019. "'Filter Bubble' Author Eli Pariser on Why We Need Publicly Owned Social Networks." *The Verge*, November 12. https://www.theverge.com/interface/2019/11/12/20959479/eli-pariser-civic-signals-filter-bubble-q-a.

Pask, Gordon. 1962. *An Approach to Cybernetics*. London: Hutchinson & Co.

Patelis, Korinna. 2013. "Political Economy and Monopoly Abstractions: What Social Media Demand." In *Unlike Us Reader: Social Media Monopolies and Their Alternatives*, edited by Geert Lovnik and Miriam Rasch, 117–127. Amsterdam: Institute of Network Cultures.

Rafiei, Davood, and Stephen Curial. 2005. "Effectively Visualizing Large Networks Through Sampling." In *VIS 05. IEEE Visualization*, 375–382. Minneapolis.

Roesslein, J., 2020. *Tweepy: Twitter for Python!* https://github.com/tweepy/tweepy.

Slovic, Paul. 2007. "'If I Look at the Mass I Will Never Act': Psychic Numbing and Genocide." *Judgment and Decision Making* 2 (2): 79–95.

Sunstein, Cass. 2017. *#Republic: Divided Democracy in the Age of Social Media*. Princeton, NJ: Princeton University Press.

Von Neumann, John. 1966. *Theory of Self-Reproducing Automata*. Urbana: University of Illinois Press.

Eric Duong is a computational designer currently working on research related to COVID-19 and shared spaces. Duong recently completed undergraduate studies at the University of Virginia School of Architecture with a secondary major in computer science.

Garrett Vercoe's work specializes in the building of humane technology that expands what people can think and do. Vercoe recently completed undergraduate studies at the University of Virginia School of Architecture with a secondary major in cognitive science, and an additional computer science concentration.

Ehsan Baharlou is an Assistant Professor of Architecture at the University of Virginia. He also worked as a postdoctoral associate at the Department of Architecture at the Massachusetts Institute of Technology (MIT). Ehsan holds a doctoral title from the Institute of Computational Design and Construction (ICD) at the University of Stuttgart. His work focuses on advanced technology to develop novel approaches to design computation and digital fabrication.