


## LETTER

# Recurrent neural network reveals overwhelming sentiment against 2017 review of US monuments from humans and bots

Caitlin McDonough MacKenzie PhD<sup>1</sup>  | Tony Chang PhD<sup>2</sup> |  
 Mallika A. Nocco PhD<sup>3</sup> | Rebecca S. Barak PhD<sup>4</sup> | Molly C. Bletz PhD<sup>5</sup> |  
 Sara E. Kuebbing PhD<sup>6</sup> | Michael Dombeck PhD<sup>7</sup>

<sup>1</sup> Climate Change Institute University of Maine, Orono, Maine

<sup>2</sup> Conservation Science Partners, Truckee, California

<sup>3</sup> Department of Land, Air, and Water Resources, University of California-Davis, Davis, California

<sup>4</sup> Department of Plant Science and Conservation, Chicago Botanic Garden, Glencoe, Illinois

<sup>5</sup> Department of Biology, University of Massachusetts-Boston, Boston, Massachusetts

<sup>6</sup> Department of Biological Sciences, University of Pittsburgh, Pittsburgh, Pennsylvania

<sup>7</sup> David H. Smith Conservation Research Fellows, Society for Conservation Biology, Washington, DC

## Correspondence

Caitlin McDonough MacKenzie, Climate Change Institute, University of Maine, 210 Sawyer, Orono, ME 04469, USA.  
 Email: [caitlin.mcdonough@gmail.com](mailto:caitlin.mcdonough@gmail.com)

Caitlin McDonough MacKenzie, Tony Chang, and Mallika A. Nocco are cofirst authors.

## Funding information

David H. Smith Conservation Research Fellowship

## Abstract

In the United States, the conservation of federal lands reflects a social history of public advocacy, public policy, and public comments. US federal agencies solicit public comments to scope for ideas, solve problems, and use the best available science for policy-making, legislation, and management. Online comment submission has led to staggering numbers of comments that are challenging to summarize. Here, we analyze comments received by the Department of the Interior in response to the proposed executive review of 27 national monuments designated and expanded between 1996 and 2016. We used a deep recurrent neural network (AWD-LSTM) to classify sentiment of 754,707 comments with higher precision and recall (F1-score = 0.98) than support vector machine and Naïve Bayes approaches. Over 97% of unique comments opposed the executive review, suggesting overwhelming support for maintaining national monument designations. Using cosine similarity, we also found that duplicates or potential automated software bots comprised over two-thirds of comments. We offer recommendations for comment submission, collection, and analysis in the current techno-political climate.

## KEYWORDS

machine learning, national monument, public comment, recurrent neural network, sentiment analysis

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. Conservation Letters published by Wiley Periodicals, Inc.

## 1 | INTRODUCTION

US federal agencies solicit comments to scope for ideas, solve problems, and gain substantive information, including the best available science. Administrative laws require acknowledgement, summation, and response to comments in their totality. In recent decades, the shift from postal to online comment submissions with minimal cyber security has created challenges for efficient and transparent review. For example, “point and click” submissions from advocacy groups can lead to overwhelming numbers of similar or duplicate comments (Shulman, 2009). Without quantitative analysis of the full comment data, the summary of comments can be manipulated to support either side of controversial actions. Automated software bots developed to mimic human participation in online activities are a growing manipulation used to disrupt the public comment process. For example, during the Federal Communications Commission’s net neutrality comment period, researchers determined that 94% of comments were submitted by bots (Hitlin, Olmstead, & Toor, 2017). Sentiment and text similarity analyses are culturomics tools well-suited for not just identifying bots, but also tackling and understanding the totality of large public comment datasets (Lennox, Verissimo, Twardek, Davis, & Jarić, 2020; Wang, Wu, Zheng, & Wang, 2018).

Public comments, bots, and “point and click” advocacy intersected with conservation in the 2017 proposed review of US national monuments. In April 2017, US Executive Order 13792 (EO 13792) directed then-Secretary of the Interior, Ryan Zinke, to review 27 national monuments designated or expanded since 1996. National monuments are created by the executive branch of the US government under the authority of the 1906 Antiquities Act and preserve areas of historical, cultural, and biological significance. For example, the 27 monuments proposed for review include native cultural sites, endemic and endangered plants and animals, and also support local economies (for a summary of cultural and biological resources at each of the 27 monuments proposed for review, please see the David H. Smith Fellows’ public comment, archived at [https://conbio.org/images/content\\_policy/PUBLICCOMMENTSsmithNationalMonuments.pdf](https://conbio.org/images/content_policy/PUBLICCOMMENTSsmithNationalMonuments.pdf), further, a list of all 158 national monuments, is available at <https://www.nps.gov/archeology/sites/antiquities/monumentslist.htm>).

While there was no precedent for a President to revoke or downsize a national monument under the Antiquities Act, EO 13792 falls under a general pattern of Protected Area Downgrading, Downsizing, and Degazettement (PADDD) (Mascia et al., 2010). PADDD is a long-standing and widespread but largely overlooked practice in both the United States (e.g., Hetch Hetchy Valley in

Yosemite) and globally (Qin et al., 2019). Recent spikes in PADDD proposals in the United States are overwhelmingly associated with industrial scale development, usually oil, gas, or mineral extraction (Golden Kroner et al., 2019). EO 13792 was a large-scale politically motivated PADDD effort; many of the targeted national monuments were created by President Obama in states led by President Trump’s allies (e.g., Bear’s Ears in Utah, Katahdin Woods and Waters in Maine). The executive order represented a flashpoint in conservation policy, raising questions like: how permanent are protected areas? Who determines the future of public lands? As a result, the public comment period following EO 13792 was ripe for attracting many comments, including “point and click” submissions from vocal advocacy groups on all sides of the issue.

Following EO 13792, Zinke opened a public comment period on the Review of Designations under the Antiquities Act (RDUAA). This vague call differed from typical Department of the Interior (DOI) comment periods that solicit feedback on rule-making, land management planning processes, or regulatory issues (P. Hanceford per. comm.). Most calls for public comments explicitly exclude sentiment, yet in this call the DOI was, in part, requesting public sentiment toward the monuments. The DOI issued a response on August 24, 2017 to the submitted comments, calling the high volume of comments in opposition to the review “a well-orchestrated national campaign organized by multiple groups,” effectively dismissing most comments, and offering no methods or quantitative breakdown of sentiment (Zinke, 2017).

Examining public sentiment in the national monument case study is valuable to conservation research specifically, and to agencies soliciting and responding to public comments more broadly. In response to RDUAA, conservation organizations have attempted to analyze the sentiment toward national monuments captured in public comments. Understanding public attitudes toward protected areas like national monuments is increasingly recognized as important for sustaining conservation (Kotowicz, Richmond, & Hospital, 2017). These previous analyses employed both analog (i.e., reading and classifying a subsample of the comments to extrapolate sentiment) and digital (i.e., traditional machine learning) techniques to summarize public sentiment. If adopted by federal agencies, machine learning (ML) could ease agency personnel workloads, while improving transparency and accountability to the public. Agencies like the Department of Defense, Food and Drug Administration, and the Securities and Exchange Commission have explored using ML, and have found that these tools increase the efficiency of agency reviews of applications and filing (Bauguess, 2017; Onyshkevych, 2020; Rocca, 2017). If adopted as a routine part of the public comment process, ML would likely ease

personnel workloads while improving transparency and accountability to the public.

Here, we employ ML to analyze public sentiment in the full database of RDUAA public comments and demonstrate the potential for recurrent neural networks (RNN) to categorize complex language with high precision. A previous analysis of public sentiment in the RDUAA public comments found high levels of opposition (99.2%) to the national monuments review using a ML support vector machine (SVM) approach to estimate sentiment across the full comment dataset (Wang, Phillips, Beavers, & Stoner, 2017). However, this SVM approach displayed low skill at distinguishing comments that supported the review, which creates more uncertainty in a politically charged debate where it is particularly important to correctly classify unpopular sentiments with reasonable certainty (Wang et al., 2017). RNN have demonstrated high utility in sentiment analysis in recent years for interpretation of human emotion in text, even for particularly challenging tasks such as movie reviews and social media posts, where similar words are often used to express opposing sentiments (Ali, El Hamid, Mostafa, & Youssif, 2019; Porwal, Ostwal, Phadtare, Pandey, & Marathe, 2018). The RNN method can offer better precision and recall in heavily skewed datasets like RDUAA, as they factor the sequence and relative position of words in addition to their meaning when estimating sentiment.

The DOI summary of public comments about the proposed review of national monuments failed to assess and report public sentiment with clarity and transparency. Our objectives were: (1) compare performance of a deep learning approach to traditional machine learning methods for natural language processing in the review of public comments; (2) classify and summarize the sentiment of all 754,707 RDUAA comments; (3) evaluate the number duplicate/form letter comments or potential “bots” within the submitted comments. We address gaps in the DOI summary by (1) classifying comments into opposing, supporting, or unknown/neutral sentiment toward the review; and (2) sorting classified comments into three distinct groups: *human* (unique comment), *form letter* (individual comment drafted by nongovernmental organizations and customized for submission by humans), or *bot* (identical comments submitted in bulk).

## 2 | METHODS

### 2.1 | Data

We downloaded all 754,707 comments submitted during the open comment period from the US Public Comments Registry using a custom python web data extrac-

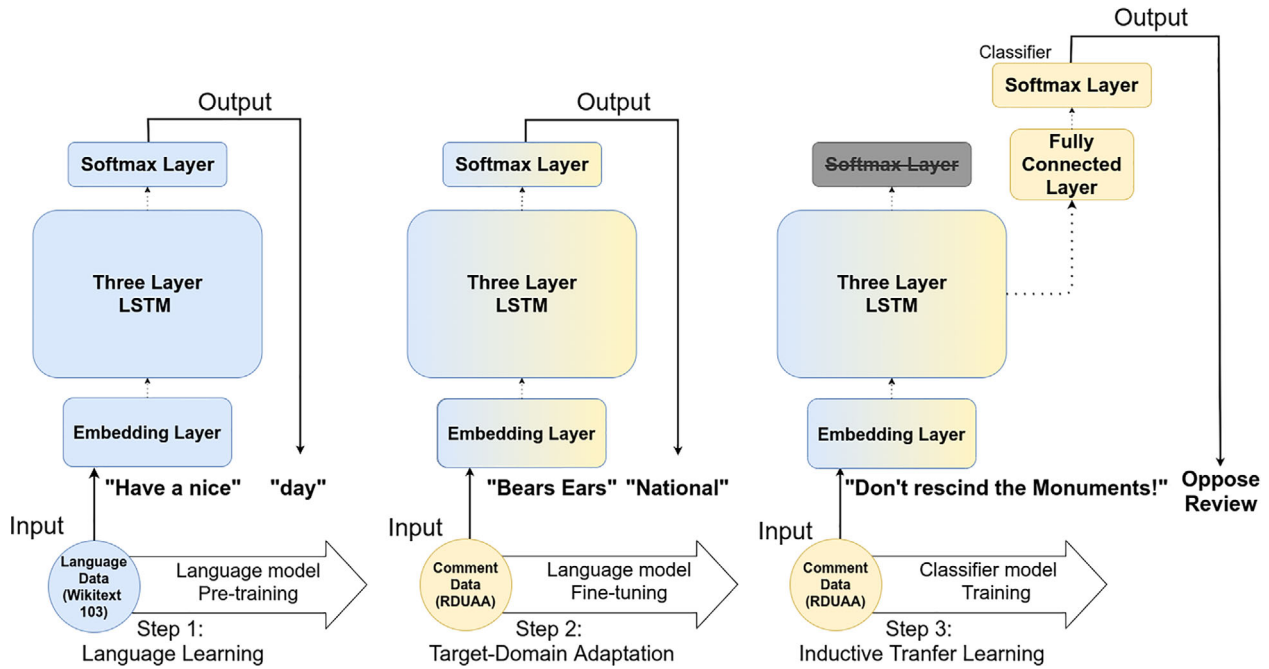
tion script to create the RDUAA dataset. A subset of 10,708 public comments was manually coded into three sentiment classifications regarding the national monument review: “unknown” or unknown/neutral sentiment; “oppose”; and “support.” These unique comments were used for training and model evaluation. We used a 90%–10% split of the data (9,637 comments for training and 1,071 comments for model evaluation). Training and validation datasets had similar distributions of all three sentiment classifications.

### 2.2 | Transfer learning

We used the Universal Language Model Fine-tuning for Text (ULMFiT) classification transfer learning method on a pretrained Averaged Stochastic Gradient Descent Weight Dropped Long Short-Term Memory model (AWD-LSTM) (Merity, Keskar, & Socher, 2017) to encode sentiment within the RDUAA language domain (Figure 1). This approach regularizes and trains the RNN while avoiding overfitting the data (See Supplemental Materials for details). The ULMFiT transfer learning method fine-tunes the learned-word-embedding and LSTM layers of the AWD-LSTM model by adjusting prefitted parameters to the linguistic properties of the RDUAA language domain. Initial word embeddings for the model were taken from the Wikitext-103 dataset, which consists of 28,595 preprocessed Wikipedia articles and 103 million words (Figure 1, Step 1).

After language learning, we fine-tuned the model with the RDUAA dataset because language used in public comments may differ from Wikipedia articles (Figure 1, Step 2). We fine-tuned for 15 epochs using a slanted triangular learning rate, which oscillates between fast and slow rates to optimize parameter fitting. Learned parameters within the LSTM network were held constant, or “frozen” to prevent complete loss of prefitted weights during the initial epochs. Learned parameters were gradually unfrozen during later epochs as word embedding parameters were adjusted for RDUAA. Following target-domain adaptation, we modified the AWD-LSTM model for sentiment classification using an inductive transfer learning approach (Figure 1, Step 3). Two additional linear fully connected layers were appended to the last LSTM layer and then trained using the 9,637 “training” comments in the labeled RDUAA sentiment dataset.

We ran the trained AWD-LSTM model on the labeled RDUAA model evaluation dataset (1,071 comments). Then we used the same training and validation data to build two additional models using the multinomial Naïve Bayes and support vector machine approaches from the



**FIGURE 1** Schematic of the AWD-LSTM model layered, modular architecture that includes language learning, target-domain adaptation, and inductive transfer learning steps. Blue indicates input processing, learning, and training, while yellow indicates AWD-LSTM outputs associated with each step adapted from (Howard & Ruder, 2018)

scikit-learn machine learning library in Python 3.6. We assessed performance of AWD-LSTM compared to these other contemporary machine learning approaches. All three machine learning simulations were performed using a Microsoft Azure NC6 machine, NC-series using an Intel Xeon E5-2690 v3 2.60 GHz v3 (Haswell) processor as well as Nvidia K80 GPU 12 GB GDDR memory. Machine learning simulations were not time (<24 hours/run) or cost (<\$10/run) prohibitive.

### 2.3 | Bot and form letter designation

After sentiment-classification, we further sorted comments into three groups to understand their likely sources: *human*, *form letter*, and *bot* (Table S1). *Human* comments were defined by their complete uniqueness from other comments. *Form letter* comments were collections of very similar comments that contained small differences from one another, typically the addition of a submitter's name or a custom sentence. We recognize that form letters are also likely submitted by humans, but use the categories *human* and *form letter* to separate uniquely written comments from form letters provided by different organizations (e.g., Patagonia, American Motorcycle Association). We designated *form letters* using the cosine similarity metric (Huang, 2008), which is a common method for comparing text and plays a crucial role in tasks such as bot or plagiarism detection (Wang, Wu, Zheng, and Wang 2018).

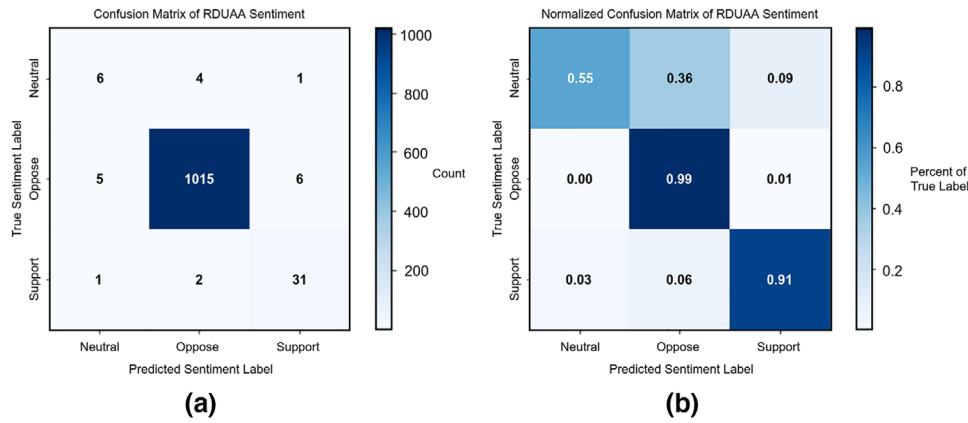
We designated *bots* as complete duplicates of comment text (nonunique) within the full RDUA dataset.

The cosine similarity was calculated for each comment in its vectorized form (after mapping through word embedding) compared to all other comments. After iteratively changing thresholds at 5% intervals, we found that 0.90 was the optimized threshold to allow for form letters with different signatures and addresses. We set 0.90 as a cosine similarity threshold between unique, individual *human* comments (<0.90) and *form letters* (>0.90). We recognize the limitations of this application of the cosine similarity approach, as more sophisticated bots could potentially mimic form letters or even unique, human comments. However, we posit that without any bot security measures, there was no reason for bot designers to create more sophisticated bots. To further illustrate our application of cosine similarity, we share example comments that display the range of differences at varying cosine similarity levels surrounding the 0.90 threshold in supplemental materials (Table S1).

## 3 | RESULTS AND DISCUSSION

### 3.1 | Model performance and validation

The AWD-LSTM model had high overall classification accuracy (98%) for the test dataset of public comments. The confusion matrices (Figure 2) show that inaccurate



**FIGURE 2** The (a) confusion matrix and (b) and normalized confusion matrix for sentiment classification of the RDUAA public comment dataset using an inductive transfer learning approach with a trained, AWD-LSTM model. Labels are “Neutral” for unknown/neutral sentiment, “Oppose” for comments that oppose the national monument review, and “Support” for comments that support the national monument review

**TABLE 1** The AWD-LSTM model presents the strongest ability to accurately classify “Oppose” and “Support” sentiment within the validation dataset ( $n = 1,071$ ) compared to Naïve Bayes (NB) and support vector machine (SVM) approaches, which tend to bias toward the “Oppose” sentiment class label. The best performing model in each class label across Precision, Recall, and F1-score is noted in bold text

Class Label	Precision			Recall			F1-score			$n_{\text{class}}$
	AWD-LSTM	NB	SVM	AWD-LSTM	NB	SVM	AWD-LSTM	NB	SVM	
Unknown sentiment	<b>0.67</b>	<b>0.99</b>	0.83	<b>0.55</b>	0.16	0.26	<b>0.59</b>	0.27	0.4	11
Oppose review	<b>0.99</b>	0.94	0.98	<b>0.99</b>	0.99	0.99	<b>0.99</b>	0.97	0.99	1,026
Support review	<b>0.9</b>	0.88	0.87	<b>0.92</b>	0.12	0.77	<b>0.91</b>	0.21	0.81	34
Weighted average	<b>0.98</b>	0.94	0.97	<b>0.98</b>	0.94	0.97	<b>0.98</b>	0.91	0.97	1,071

predictions tended to assign “oppose” sentiment to comments with “unknown” or neutral sentiment. This is not surprising, as “unknown” comments were often long and confusing for human trainers to decipher as well. The AWD-LSTM model’s accuracy rate at classifying the small minority of “support” comments was 91% and the primary confusion came from the model assigning “oppose” to “support” comments. A previous SVM approach produced only 58% accuracy in “support” classification, and was unable to classify unknown or neutral sentiments (Wang et al., 2017).

Additionally, in our own comparison using the same test data, the AWD-LSTM model displayed superior ability to accurately classify sentiment compared to Naïve Bayes and SVM approaches, especially for the small number of comments supporting the national monuments review (“Support Review,” Table 1). The Naïve Bayes model tended to predict false negatives on the validation dataset; AWD-LSTM outperformed both contemporary ML approaches (Table 1). The structure of the AWD-LSTM model creates a more complex mapping of feature space which allows for deeper, more subtle connections between words that are not reliant on the order in which those words are

expressed. This subtlety is reflected in the AWD-LSTM model’s F1-score of 0.99 and 0.91 for the “oppose” and “support” classes, respectively. We also share six specific comments across the three sentiment classes to illustrate the subtlety and accuracy of AWD-LSTM compared to the Naïve Bayes and SVM approaches (Table 2). For example, Comment #1 uses negative words like “displeasure,” “remove,” and “fear” to express support for RDUAA, while Comment #6 uses similar words like “rescind” and “destroy” to express opposition to RDUAA (Table 2). The AWD-LSTM correctly classified sentiment in both of these cases. Additionally, Comment #4 illustrates the vague language present in “unknown” comments correctly classified by the AWD-LSTM approach. Future work in ML comment analyses could explore deeper analytical insights using comment narratives with an RNN approach.

### 3.2 | Vast majority of comments oppose national monument review

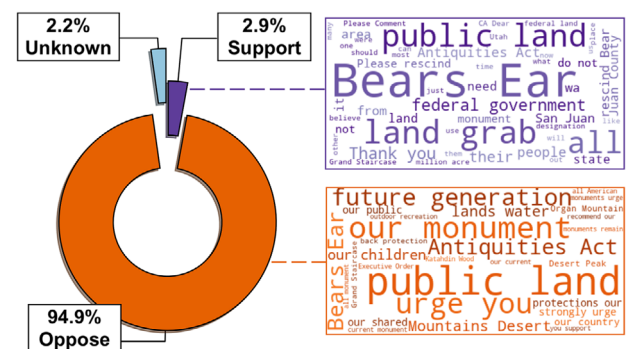
The overall RDUAA sentiment opposed the 2017 national monuments review; comments were extremely positive



**TABLE 2** Example of model classification for comments with different sentiment labels. Purple, blue, and orange text have true labels of support (“+”), unknown or neutral (“?”), and oppose (“-”), respectively. Sentiment classification by AWD-LSTM, support vector machine (SVM), and Naïve Bayes (NB) approaches are also listed for each comment

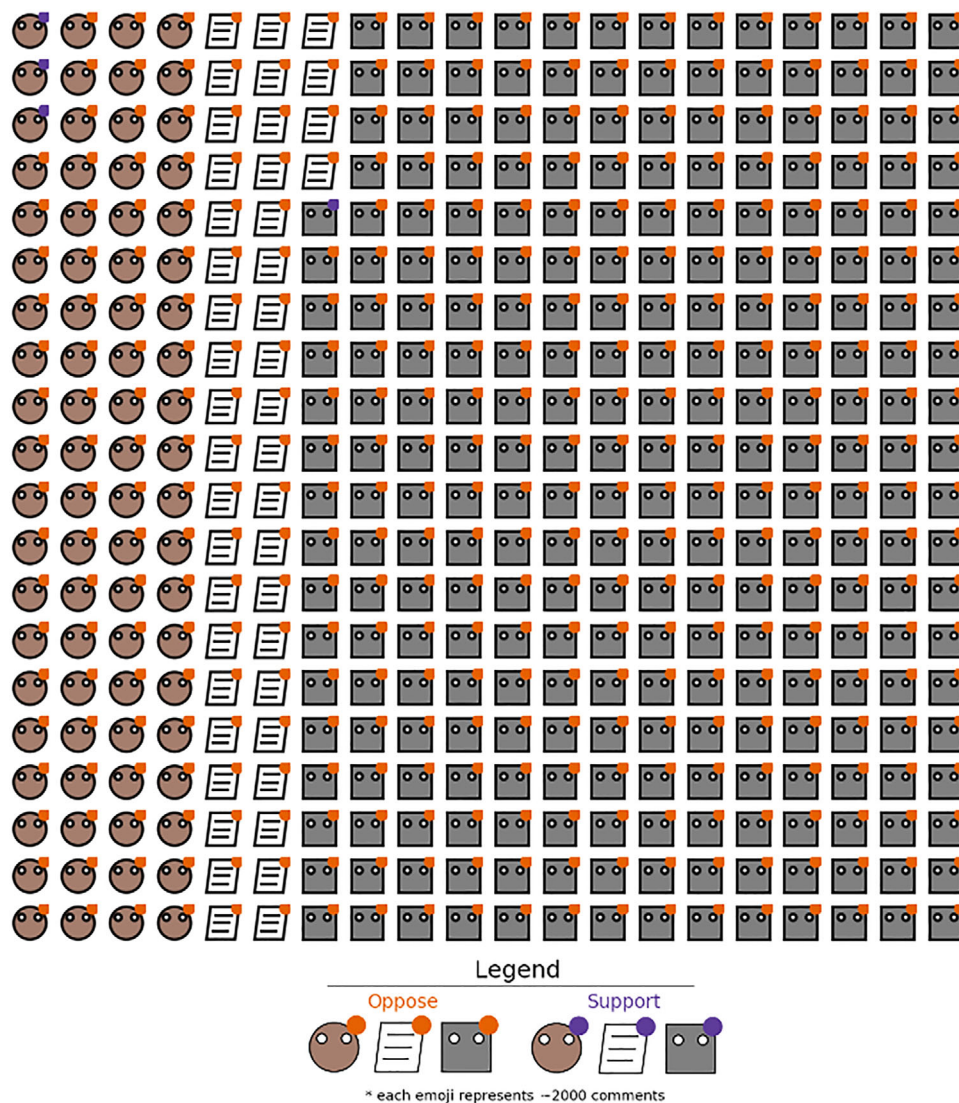
Comment	True label	SVM label	NB label	RNN label
1. “I am writing you [sic] to express my displeasure with the designation of an area of Maine as Katahdin Woods and Waters National Monument. This designation was an abuse of the Antiquities Act, as the area does not meet the outlined requirements. The designation was promoted by a wealthy individual via backdoor government connections and plenty of money to grease the wheels. The area has a diverse and beneficial history of productive private land use for recreation and economic activity. Those of us in the region fear the endless expansion of this monument and the iron fist of federal land ownership in the region. Please remove this designation and return the land to private ownership or state oversight.”	+	-	-	+
2. “No monument!!! We are wasting money on stuff no one wants. It stops the people from their land.”	+	+	-	+
3. “Dear Secretary Ryan Zenke, let the locals decide, they know best, not the big city folks from new York, S.F. etc. tks”	?	?	-	?
4. “I wish this would stop, but I’m fully cognizant that an oligarch is behind this. It still doesn’t stop me from pointing out how greedy and self centered things have become.”	?	-	-	?
5. “Please, do not rescind the monuments. Especially Bears Ears and Grand Staircase Escalante., I am an Utah residents [sic] and enjoy the beautiful state I live in. I believe if taken away the Utah Government will destroy it. Please listen to the people that are for these monument [sic], not just those who oppose it.”	-	+	-	-
6. “I am against the removal of national monuments from their designation.”	-	+	-	-

toward sustained public lands protection. We found 94.9% of comments opposed the review, 2.9% of comments supported the review, and 2.2% of comments were classified as unknown or neutral sentiment (Figure 3). We further categorized results into humans (146,595, 20% of RDUAA), form letters (84,745, 11%), and bots (515,193, 69%). Among humans, form letters, and bots, sentiment was overwhelmingly opposed to the review (Figure 4; individuals: 97.4%, form letters: 96.4%, bots: 99.6%). We highlight that although most comments were submitted by bots, human public sentiment overwhelmingly opposed the national monuments review. The official DOI summary response obscured the overwhelming negative sentiment of human comments by labeling all comments opposed to the review as “a well-orchestrated national campaign organized by multiple groups” (Zinke, 2017). In contrast, our analysis demonstrates that unique human comments were also negative toward the review, along with form letters from NGOs. Additionally, the DOI failed to mention the potential disruptive use of automated soft-



**FIGURE 3** The pie chart (left) categorizes all 754,707 Review of Designations under the Antiquities Act (RDUAA) comments by sentiment. Word clouds (right) contain the most prevalent words used in comments that either supported (purple) or opposed (orange) the monument review

were bots in the comments, which we were able to detect in our analysis.



**FIGURE 4** The Review of Designations under the Antiquities Act (RDUA) comments classified by sentiment and categorized as human (brown circle emoji), form letter (black and white letter emoji), or bot (gray square emoji). Each emoji represents ~2,000 comments and is color-coded with a small circle in the top right corner to represent the aggregate support (purple) or opposition (orange) to the 2017 review. No form letter emojis in support of the review are present as they are too few (<2,000) to be represented

### 3.3 | “Bots” comprise 2/3 of comments

*“Dear President Trump, Please don’t rescind or alter Bears Ears National Monument. This culturally rich and recreationally spectacular place is part of our national legacy and the legacy of future generations. It’s one of the important wild places where we go to run, hike, camp, ski, fish, climb and spend time with our friends and families. These public lands are not just beautiful but economically beneficial to our local communities and our nation as a whole.”*

The above text was submitted as 99,748 separate, duplicate comments comprising 13.4% of all RDUA comments. Though bot interference in political elections is notorious, they have received little attention from the scien-

tific community in regards to public comments. Bots are automated programs developed to mimic human participation in online activities. They are perpetuated by a few individuals and represent a prevalent problem that dilutes civic engagement. For example, bot-run Twitter pages were found active leading up to the 2016 US presidential election (Bessi & Ferrara 2016). Also, a Pew Research Center study of public comments regarding net neutrality found 94% of submitted comments were bots (Hitlin et al., 2017). These challenges threaten the ability of agencies to respond to unique human comments, especially in the current political climate where reduction in federal agency staff limits expert capacity to handle increasing numbers of comments.

The DOI's statement, "The DOI received approximately 2.6 million form comments associated with NGO-organized campaigns, which far outnumbered individual comments" (Zinke, 2017) did not align with our results. We contend that bots, not form letters, overshadowed humans. The DOI's estimate includes comments submitted after the solicitation closed and asserts that form letters comprised 93% of all comments (Zinke, 2017). In contrast, we found only 11% of RDUA comments were form letters, while individual, unique comments comprised 20%. Furthermore, the DOI summary does not acknowledge that sentiment of individual comments mirrors sentiment in form letters. In this case, excluding form letters and bots does not change public sentiment, because unique, individual comments resoundingly opposed the national monuments review. Thus, the administration, although not mandated to act on comments, generally ignored public sentiment. In addition, this case study highlights how bots hamper public participation in comment periods, reducing the impact of individuals and obfuscating the best available science.

On December 4, 2017, through Presidential Proclamation, President Trump reduced the size of two national monuments: Bears Ears National Monument was reduced by 87% from 1.35 million to 201,876 acres (<https://www.whitehouse.gov/presidential-actions/presidential-proclamation-modifying-bears-ears-national-monument/>) and Grand Staircase Escalante National Monument by 42%, from 1.7 million to 1,003,863 acres (<https://www.whitehouse.gov/presidential-actions/presidential-proclamation-modifying-grand-staircase-escalante-national-monument/>), marking the only time national monuments have been reduced in size (<https://www.nps.gov/archeology/sites/antiquities/monumentslist.htm>). Land removed from the national monument designation was made available for mineral and geothermal leasing, mining, and sale.

## 4 | CONCLUSION AND RECOMMENDATIONS

The digitization of public comment submission increases access and democratizes the process of scoping for ideas, solving problems, and using the best available science in policy-making, legislation, and management. However, this opportunity presents new challenges, highlighted by the large volume of comments. We demonstrate that deep learning and natural language processing algorithms can automate the difficult process of sentiment analysis to efficiently classify and summarize large comment datasets. Our approach addresses a key challenge for machine learning techniques—very unbalanced datasets—through the

use of a recurrent neural network with a more refined feature mapping approach. Though this is one of the first applications of AWD-LSTM modeling for conservation, it is prevalent in the computation and language community (Akinwande & Remy, 2017; Howard & Ruder, 2018; Krause, Lu, Murray, & Renals, 2017; Merity et al., 2017; Yang, Dai, Salakhutdinov, & Cohen, 2017). Additionally, we were able to run the AWD-LSTM model for the entire RDUA in less than a day for about ten dollars using a virtual machine. Therefore, we argue that available expertise, hardware, and software do not pose knowledge-action barriers for implementing culturomics tools like AWD-LSTM and text similarity into federal operations.

Our study revealed the prevalence of bots in this particular public comment period, which matches reports of high bot activity both within other comment calls and across social media. Without concerted intervention, bots will continue dominating public comments and overwhelming human voices. To keep pace with the digital age, we offer guidance for agencies on managing bots and improving transparency surrounding public comments. These recommendations should be considered in view of the requirements of the Administrative Procedure Act or other relevant laws; changes in these statutes could also be put forth by Congress which may impose certain restrictions on the ways that agencies deter, filter, and quantify bot submissions:

1. *Deter Bots*: Implementation of Completely Automated Public Turing tests to tell Computers and Humans Apart (CAPTCHAs; e.g. asking users to retype a written code) can deter bot submissions. Though these tests are a standard industry practice, passing even a simple CAPTCHA is not currently required to submit a public comment. Deterrence must coevolve with rapid bot development. At minimum, CAPTCHAs should be part of federal register submission portals. Additionally, timestamping comment submission to the second will aid in the detection of instantaneous bot submissions using time-series approaches. Timestamping beyond the date was not present in RDUA data.
2. *Filter and Quantify Bots*: Bot deterrence is unlikely to prevent all bot comments. ML approaches (e.g., Botometer, <https://botometer.iuni.iu.edu/#/>)—as used in this study—can enable agencies to define, count, and filter bot comments. We recommend agencies apply bot filtration to all public comment databases.
3. *Increase Transparency*: Currently, agencies are not explicitly required to report or summarize bot deterrence and filtration efforts. As agencies grapple with unprecedented volumes of online submissions, bots quietly undermine public participation and trust in decision-making. We recommend agencies report the



breakdown of duplicate comments (bots and form letters) relative to unique comments and describe how these categories inform decision-making (e.g., are bot comments ignored?).

Public comments have been overlooked in recent calls for science advocacy (Lubchenco, 2017; Toomey, Knight, & Barlow, 2017; Young et al., 2014). Public comment periods are designed to be a meaningful platform for participation in policy-making and regulatory decisions, and offer an opportunity to share expertise and, in some cases, sentiment, with federal agencies. Scientists can use this system to provide the best available research, share knowledge on impacts of proposed regulations, and add their expertise to decision-making processes. While bots may make writing public comments seem futile, we argue that if agencies implement these recommendations, human-submitted comments will continue to be an important avenue for promoting science-based policy-making. As conservation scientists, we implore federal agencies to confront bots. In turn, we will work on producing substantive, plain language public comments and sharing the best available science with policy-makers.

## ACKNOWLEDGMENTS

We thank the following scientists for coding the test set of public comments: Stephanie Borelle, Emily Darling, Paul Elsen, Meredith Holgerson, Kurt Ingeman, Jonathan Koch, Bonnie McGill, Morgan Tingley, and Grace Wu. We thank Mike Anderson and Phil Hanceford for providing valuable input on the public comment process, and Kay Havens, Susan Jewell, Sarah Reed, Kristin Floress, Bonnie McGill, Kurt Ingeman, and Stephanie Borrelle for their helpful comments on earlier drafts of this manuscript. Thank you also to two anonymous reviewers and the editors for incisive, constructive feedback. Much of the collaborative work for this paper took place at Schoodic Institute at Acadia National Park and Iowa Lakeside Laboratory. RSB, MCB, TC, SK, CMM, and MAN were supported by the David H Smith Conservation Research Fellowship. We thank Shonda Foster, Smith Fellowship program director, for her support.

## AUTHOR CONTRIBUTIONS

Conceived and designed the study, revised drafts: all. Created and applied deep learning techniques: TC. Wrote initial draft: TC, MAN, MCB, RSB, SEK, CMM. Contributed advice, experiences, and policy perspective: MD. Led project: CMM.

## DATA AVAILABILITY STATEMENT


All public comments are available at <https://www.regulations.gov/docketBrowser?>

`rpp = 25&po = 0&dct = PS&D = DOI-2017-0002&refD = DOI-2017-0002-0001` and the code is available at <https://gitlab.com/tonychang/smith-monuments-sentiment>.

## CONFLICT OF INTEREST

The authors declare no conflict of interest.

## ORCID

Caitlin McDonough MacKenzie PhD  <https://orcid.org/0000-0002-2775-5628>

## REFERENCES

- Akinwande, V., & Remy, S. L. (2017). Characterizing the hyperparameter space of LSTM language models for mixed context applications. *arXiv preprint arXiv:1712.03199*, 1–4.
- Ali, N. M., El Hamid, A., Mostafa, M., & Youssif, A. (2019). Sentiment analysis for movies reviews dataset using deep learning models. *International Journal of Data Mining & Knowledge Management Process (IJDKP)*, 9, 19–27. (June 14, 2019).
- Baugess, S. W. (2017). The role of big data, machine learning, and AI in assessing risks: A regulatory perspective. *Champagne Keynote Address: OpRisk North America 2017*, New York, New York. June 21, 2017. Retrieved from <https://www.sec.gov/news/speech/baugess-big-data-ai>
- Bessi, A., & Ferrara, E. (2016). Social bots distort the 2016 US Presidential election online discussion. *First Monday*, 21, 1–14.
- Golden Kroner, R. E., Qin, S., Cook, C. N., Krithivasan, R., Pack, S. M., Bonilla, O. D., ... Mascia, M. B. (2019). The uncertain future of protected lands and waters. *Science*, 364, 881–886.
- Hitlin, P., Olmstead, K., & Toor, S. (2017). Public comments to the federal communications commission about net neutrality contain many inaccuracies and duplicates. *Pew Research Center*, 1–18.
- Howard, J., & Ruder, S. (2018). Universal language model fine-tuning for text classification. *arXiv Prepr. arXiv:1801.06146*, 1–12.
- Huang, A. (2008). Similarity measures for text document clustering. *Proceedings of the Sixth New Zealand Computer Science Research Student Conference (NZCSRSC2008)*. Christchurch, New Zeal. pp. 9–56.
- Kotowicz, D. M., Richmond, L., & Hospital, J. (2017). Exploring public knowledge, attitudes, and perceptions of the marianas trench marine national monument. *Coastal Management*, 45, 452–469.
- Krause, B., Lu, L., Murray, I., & Renals, S. (2016). Multiplicative LSTM for sequence modelling. *arXiv preprint arXiv:1609.07959*, 1–12.
- Lennox, R. J., Verissimo, D., Twardek, W. M., Davis, C. R., & Jarić, I. (2020). Sentiment analysis as a measure of conservation culture in scientific literature. *Conservation Biology*, 34(2), 462–471.
- Lubchenco, J. (2017). Environmental science in a post-truth world. *Frontiers in Ecology and the Environment*, 15, 3.
- Mascia, M. B., & Pailler, S. (2010). Protected area downgrading, downsizing, and degazettement (PADDD) and its conservation implications. *Conservation Letters*, 4, 9–20.
- Merity, S., Keskar, N. S., & Socher, R. (2017). Regularizing and optimizing LSTM language models. *arXiv Prepr. arXiv:1708.02182*, 1–9.
- Onyshkevych, Boyan Deep Exploration and Filtering of Text (DEFT) (2020). Defense Advanced Research Projects Agency. Retrieved from <https://www.darpa.mil/program/deep-exploration-and-filtering-of-text>.

- Porwal, S., Ostwal, G., Phadtare, A., Pandey, M., & Marathe, M. V. (2018). Sarcasm detection using recurrent neural network. *Proceedings of the 2018 Second International Conference on Intelligent Computing and Control Systems*. IEEE, Mandurai, India. Pp. 746–748.
- Qin, S., Golden Kroner, R. E., Cook, C., Tesfaw, A. T., Braybrook, R., Rodriguez, C. M., ... Mascia, M. B. (2019). Protected area downgrading, downsizing, and degazettement as a threat to iconic protected areas. *Conservation Letters*, 33, 1275–1285.
- Rocca, Mitra (2017). Lessons learned from NLP implementations at FDA. NSLP Workshop June 15, 2017. Retrieved from <https://pharm.ucsf.edu/sites/pharm.ucsf.edu/files/cersi/media-browser/Rocca.pdf>
- Shulman, S. W. (2009). The case against mass e-mails: Perverse incentives and low quality public participation in US federal rulemaking. *Policy & Internet*, 1, 23–53.
- Toomey, A. H., Knight, A. T., & Barlow, J. (2017). Navigating the space between research and implementation in conservation. *Conservation Letters*, 10, 619–625.
- Wang, S., Phillips, S., Beavers, O., & Stoner, J. (2017). *Public support for public lands: Analysis of comments regarding review and potential loss of protection for America's national monuments*. Key-Log Economics report for the Wilderness Society, Charlottesville, Virginia. Pp. 1–24.
- Wang, Y., Wu, C., Zheng, K., & Wang, X. (2018, August). Social bot detection using tweets similarity. *International Conference on Security and Privacy in Communication Systems* (pp. 63–78). Springer, Cham.
- Yang, Z., Dai, Z., Salakhutdinov, R., & Cohen, W. W. (2017). Breaking the softmax bottleneck: A high-rank RNN language model. arXiv preprint arXiv:1711.03953, 1–18.
- Young, J. C., Waylen, K. A., Sarkki, S., Albon, S., Bainbridge, I., Balian, E., ... Margerison, C. (2014). Improving the science-policy dialogue to meet the challenges of biodiversity conservation: Having conversations rather than talking at one-another. *Biodiversity and Conservation*, 23, 387–404.
- Zinke, R. K. (2017). *Final report summarizing findings of the review of designations under the Antiquities Act*. Washington, D.C.: Memo. Press.

## SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

**How to cite this article:** McDonough MacKenzie C, Chang T, Nocco MA, et al. Recurrent neural network reveals overwhelming sentiment against 2017 review of US monuments from humans and bots. *Conservation Letters*. 2020;e12747. <https://doi.org/10.1111/conl.12747>