

RESEARCH ARTICLE

Characterizing social and ecological values expressed in US Forest Service public comments using a computational approach

Sarah K. MacFarland¹  | Sonya Sachdeva² | Spencer A. Wood^{1,3} | Joshua J. Lawler¹

¹School of Environmental and Forest Sciences, University of Washington, Seattle, Washington, USA

²Northern Research Station, US Forest Service, Evanston, Illinois, USA

³eScience Institute, University of Washington, Seattle, Washington, USA

Correspondence

Sarah K. MacFarland
Email: schase25@uw.edu

Funding information

U.S. Forest Service, Grant/Award Number: 20-JV-11261985-095 and 21-CS-11132400-3

Handling Editor: Jacopo Baggio

Abstract

1. Addressing social and ecological values is a central aim of democratic environmental management and policymaking, especially during deliberative and participatory processes. Agencies responsible for managing public lands would benefit from a deepened understanding of how various publics value those lands.
2. Federal land management agencies receive millions of written comments from the public on proposed management actions annually, providing a unique source of insights into how the public assigns value to public lands. To date, little attention has been directed towards methods for analysing the public's comments to understand their expressed values, in part because the volume of comments often makes manual analysis unworkable.
3. This study introduces and applies a novel computational approach to inferring values in written text by using natural language processing and a method that combines a lexicon with semantic embedding models. We developed embedding models for four types of values that are expressed in public comments. We then fit models to 409,241 public comments on actions proposed by the United States Forest Service from 2011 to 2020 and regulated by the National Environmental Policy Act.
4. The embedding model generally outperformed the lexicon word-count, particularly for value types with shorter lexicons, and, like human evaluators, the embedding models performed better for more evident values and were less reliable for more abstract or latent values.
5. By applying the resulting model, we furthered our understanding of how the public values National Forest lands in the United States. We observed that aesthetic and moral values were expressed more often in comments for projects that received more public interest, as gauged by the number of comments a project received and in comments for projects addressing recreational management.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2025 The Author(s). *People and Nature* published by John Wiley & Sons Ltd on behalf of British Ecological Society.

KEYWORDS

content analysis, National Environmental Policy Act, natural language processing, public comments, public participation, social and ecological values, text analysis, United States Forest Service

1 | INTRODUCTION

Human and ecological values have become a pervasive concept in environmental management, receiving attention from scholars, policymakers and advocates alike. Emerging consensus supports formalizing and addressing these values as a central aim of environmental policy and management (Tadaki et al., 2017). At the individual level, a person's values influence their likelihood to adopt pro-environmental behaviours, to believe that climate is changing and to experience ecological grief (Marshall et al., 2019). When considering various management options, individuals' core values often underpin which attributes of a forest they find most valuable (Anderson et al., 2018). Consequently, decision-makers are often forced to grapple with trade-offs between multiple competing values (Wallace et al., 2021). Value disagreement between decision-makers and stakeholders (or between stakeholder groups) can contribute to intra- and inter-group conflict (Forester, 1999; Thacher & Rein, 2004), as can decision-makers' narrow focus on policy positions rather than a broader view of stakeholder emotions and values (Buijs & Lawrence, 2013; Swanson, 1994). By attending to stakeholder values, practitioners may not eliminate conflict, but they may better distinguish between conflicts stemming from fundamental and irreconcilable differences versus those that may be more readily resolved (Bengston, 1994). In cases with fundamental value misalignment, managers can craft messages that address and ameliorate value divergence (Jones et al., 2016), or avoid alienating a broad range of stakeholders (Hurst & Stern, 2020), rather than relying on vague technical and bureaucratic language that fails to recognize the worth of the public input received (Costa et al., 2019; De'Arman, 2020).

Understanding how people value public lands is crucial for agencies responsible for land management anytime they consider a proposed action. To the extent that land management agencies solicit input, it is most often a request for written comments during formal public planning processes, such as those required by the National Environmental Policy Act (NEPA) in the United States. In the United States, at least, public input received through these processes has increased substantially in recent decades, spurred by the greater ease of participation that digital platforms allow, and by an increased public interest in policymaking (Livermore et al., 2018). The written comments that are received by land management agencies offer a crucial source of information on how the public values those lands, yet these data are currently underleveraged by social scientists, who have explored how the public relates to and values natural spaces using surveys and social media but not with public comments.

Empirical research on the values people attribute to nature generally seeks to characterize particular aspects of the human-nature

relationship in ways that can aid environmental management and decision-making (Tadaki et al., 2017; Wallace et al., 2021). In this context, public comments are useful data offering unique and diverse insights into how individuals and communities perceive and interact with the natural world. Qualitative content analyses describing values in public comments find that a wide range of values are evoked in letters submitted during formal planning processes (Moyer et al., 2008; Proctor, 1998; Vining & Tyler, 1999). Core values informing the writer's policy position (e.g. integrity, respect, fairness, equity, security), values associated with preference satisfaction (e.g. preferred recreation and economic uses, preferred aesthetic characteristics), values in pursuit of overarching goals (e.g. ecosystem functioning, concern for specific species) and valued relations (e.g. places as a source of familial, community and national identity, recognition of interconnectedness) have all been observed in written comments received by the United States Forest Service (USFS), for example (Moyer et al., 2008; Vining & Tyler, 1999). However, these qualitative studies are necessarily limited in their spatial and temporal scope to specific locations and single-project case studies, making it difficult to generalize patterns in value expressions across projects.

In this study, we propose a computational approach that uses machine-learning methods to identify specific social and ecological values in a natural resource management context. We test the approach with a large dataset of written public comments received by the USFS. We focus on four social and ecological values reflecting why and how people value public lands (Bengston & Xu, 1995). Bengston and Xu's (1995) four-class typology categorizes values towards nature as either economic and utilitarian, life supporting, aesthetic or moral and spiritual. In this model, economic and utilization values stem from the usefulness of nature to achieve human ends and to satisfy our preferences. By contrast, life support values are those related to the life-sustaining functions or services provided by an ecosystem that underpin economic value and all life itself. Aesthetic values are those for which beauty and sensation are the defining conception of what is good. Lastly, a broad set of related values spanning moral values, spiritual values, attachment-orientation to nature, sense of place value and heritage value are grouped together. Although we refer to these values simply as "moral values" throughout this paper for brevity, this category is more expansive than that title implies.

To introduce our approach, we begin by giving a brief overview of the public comment analysis process currently used by federal land managers in the United States. We also describe computational approaches to text analysis, specifically semantic embeddings, a neural network-based technique that can capture semantic relationships between words and documents. We then present our embedding method for inferring value expressions, which we test and apply

to a dataset of written public comments received by the USFS from 2011 to 2020. To assess model performance, we create annotations for aesthetic, moral and spiritual, life support and economic value for a subset of comments and compare our model predictions to researcher assessments as well as to a simple lexicon word-count method. We evaluate model performance across administrative regions and project types. Finally, given model performance, we apply the method to explore the relative importance of different values elicited in responses to USFS-proposed actions across a range of project types and regions.

2 | BACKGROUND

2.1 | Public comment analysis

Federal land management agencies receive millions of letters each year from people seeking to influence decisions that affect them and the lands they value (Germain et al., 2001; Smiley et al., 2010). These submissions can range from detailed letters containing nuanced domain, legal and scientific arguments, to shorter letters written using colloquial language, to standardized form letters received via mass comment campaigns (DeArman, 2020; Shulman, 2009). Although agency staff are required to evaluate and respond to public input, personnel have discretion over what input is considered “substantive” and thus requires agency response (Council on Environmental Quality, 2024; Innes & Booher, 2004; Predmore et al., 2011). Within the USFS, human analysts are tasked with identifying sections of text that are “substantive”, and they filter out portions of public input interpreted as irrelevant or based on opinion or conjecture. With many proposed federal actions receiving tens of thousands of letters, analysing this information poses an increasingly sizable burden in terms of time and cost. These hurdles come at a time when an increasing proportion of the agency's declining budget is diverted away from detailed comment analysis and towards other purposes, such as wildfire management and suppression resulting from longer and more intense fire seasons (Fleming et al., 2015).

In addition to the time burden, manual comment analysis can pose serious reliability challenges, especially given tight timelines or insufficient training that may reduce the analysis quality or leave the agency vulnerable to litigation. Achieving high reliability between and within analysts is often challenging, especially for very large volumes of data where the likelihood of human error is increased, and when coding for latent constructs (i.e. not directly observable, but hidden or implied within the text) that require analysts' active interpretation (Riffe et al., 2019). Reliably coding latent constructs, such as social and ecological values, often relies to some extent on analysts' background and intuition (Krippendorff, 2012). This may be a notable challenge for USFS comment analysis personnel who have training in the natural sciences and may not be well-versed in the socio-psychological literature addressing values. Finally, manual analysis is fundamentally

limited by the pace of human reading and comprehension, and comments are typically analysed on an ad hoc project-by-project basis with limited standardization or consistency, limiting the ability to discern overarching trends.

2.2 | Computational approaches to text analysis

Given the challenges of manually coding and then analysing and synthesizing large volumes of text, computational methods for unstructured text data offer a potential alternative (Dokshin, 2022; Livermore et al., 2018; Scott et al., 2020). Advances in machine-learning and natural language processing (NLP) have provided powerful tools for processing and describing very large text datasets that would be prohibitive for humans to process alone. Initial efforts to leverage computing power to detect constructs in text relied on lexicons of terms related to a given theme combined with simple or more complex sets of syntactical rules to quantify that construct within a given corpora (Bengston et al., 1999; Bengston & Xu, 1995). For example, terms like “bid price”, “goods and services”, “logging”, “lumber” and “profitable” are all associated with economic values. One might count how many times these terms appear and use the count as a measure of economic values in the text. This can be refined using a set of rules (e.g. exclude the word “logging” when it appears near the words “devastating”, “indiscriminate” or “abuse”). Unlike human coders, these word-count-based methods are perfectly reliable in the sense that the computational approach always produces the same result for a given set of words while independent human coders may not. However, the range of nuance and subtlety in language which humans easily comprehend is rarely captured completely using a set of syntactical rules (Bengston & Xu, 1995).

2.3 | Semantic embedding models

The development of semantic embedding models signifies a major advancement in NLP. These models represent words not as simple counts or as clusters or networks of words, but as vectors in a dense, continuous, high-dimensional space (Mikolov et al., 2013). Semantic embeddings rapidly gained widespread interest due to their ability to capture the context and meaning of words more accurately than previous methods and to distil complex relationships between words based on their semantic similarity. In an embedding space, words located near each other tend to have similar meanings and are often synonyms or grammatical variants. A word's broader neighbourhood in the embedding space typically includes terms with closely related meanings.

These embeddings often encode a great deal of semantic and socio-cultural meaning, making them of particular interest to social science researchers. Applications of these methods have focused on a single word or set of keywords and closely related terms, such as Kozlowski et al.'s (2019) use of antonym word pairings (e.g. rich-poor, wealthy-impooverished, man-woman) to

represent cultural dimensions of class (e.g. affluence, status, education), tracking the shift in cultural markers of class over the 20th century. Similarly, Gugulica and Burghardt (2023) use a brief set of keywords to represent two cultural ecosystem services and to map the provision of those ecosystem services using social media posts. Our study builds on this work by using terms identified a priori to track latent constructs in a body of text by incorporating a robust, validated lexicon to infer social and ecological values in public comments.

3 | MATERIALS AND METHODS

In this study, we developed, evaluated and applied an approach leveraging semantic embeddings to recognize four types of values expressed in public comments submitted to the USFS as part of the NEPA planning process.

3.1 | Data

Public comments data were retrieved from the Comment Analysis and Response Application (CARA), a database the USFS uses internally to collect, analyse and archive public input on NEPA projects. Within the CARA system, each submission received by the agency is stored as a separate letter. Identical petition or form letters are disregarded. Each unique letter may contain one or more comments, which is any sections of text from within a full letter that are identified as “substantive and significant” by USFS personnel or contractors employed by the administrative unit that is responsible for managing the area where the project is proposed to take place.

We decided to use comments as a unit of analysis because (1) these data were cleaned of extraneous information, such as letter headers, salutations or descriptions of advocacy organizations, (2) comments are narrower, often focusing on a single topic, action, or objection, while letters span numerous topics and objections, (3) letters frequently include large portions of identical text with occasional unique sections, whereas individual comments are more unique and could be easily deduplicated and (4) these data reflect information USFS personnel considered relevant to a given planning process. After retrieving all comments from the CARA system between 2011 and 2020 ($n=689,552$), we excluded duplicate comments, empty entries and test comments entered by managers (excluded $n=424,260$). Given that the performance of joint word-document embeddings are sensitive to document length (Lau & Baldwin, 2016), we also excluded short comments with less than seven words ($n=11,455$) and comments with more than 2 *SD* above the mean word count ($n=3394$). This resulted our final dataset of 409,241 comments associated with 506 projects located in all USFS administrative regions, as well as projects based out of the Washington Office and national-level projects (median of 138 comments per project, max. 47,154 comments from a single project). The

raw comment texts were treated as documents for model training, except where noted explicitly below.

3.2 | Selecting social and ecological values

We chose to use the four-class value typology proposed by Bengston and Xu (1995) because it was developed specifically to capture values related to forestland areas, making it particularly well-tailored to USFS comments, and because it is one of the only value typologies that includes a rigorously evaluated list of lexical terms indicating each value type (see Table 1). Their lexicon comprises of a list of words and phrases that each indicate a given value. For example, words like “beautiful”, “scenic” and “majestic” suggest an author is evoking an aesthetic value. When developing their lexicon, Bengston and Xu (1995) assessed performance in a range of forestry and advocacy texts and found that most lexicon entries accurately indicated value expressions 90%–95% of instances; words with less than 80% accuracy were dropped from the lexicon.

3.3 | Analysis

3.3.1 | Coding for social and ecological values

Because expressions of social and ecological values in the public comments dataset are unknown, we created a subset of value-annotated comments to use as a benchmark against which to evaluate model performance. To do this, 300 comments were randomly sampled from the full dataset. Two researchers independently coded each comment for the presence or absence of expressions of each value type, taking notes on their reasoning. The values of Cohen's Kappa for the similarity of the two human coders were highest for aesthetic and economic values (0.85 and 0.71, respectively), and lower for moral and life support values (0.66 and 0.64, respectively), all of which suggest acceptable intercoder reliability (Landis & Koch, 1977). In cases of disagreement, the two coders met to review each disputed comment, discuss their notes and interpretative judgements, and assign a consensus code. Subsequent model evaluations were conducted using these consensus annotations.

3.3.2 | Representing values using a lexicon

We first measured value expressions in comments using the lexicon word-count frequency method (Bengston & Xu, 1995). Comments were scored for each value type by summing the frequency of all terms in each value lexicon in the comment text. To assess overall model performance, we sorted the comments in descending order of word-count scores and evaluated performance using the area under the curve (AUC) of the receiver operating characteristic (ROC) curve. The ROC curve tracks the true positive rate against the false positive rate at each point of the cosine similarity ranking, and computing the

TABLE 1 Bengston and Xu's (1995) typology of values, definitions and examples of lexicon keywords.

Value type	Definition	Example keywords
Aesthetic	The type of value in which beauty is the concept of what is good	aesthetically, evocatively, exhilarate, musical, rustic, scenic, sensory, symphony, untrammelled, visual
Moral & Spiritual	A broad category of related values, including moral value, indicated by regards of love, affection, reverence and respect, and spiritual value, indicated by the "experience of being related to or in touch with an 'other' that transcends one's individual sense of self and gives meaning to one's life in a deeper than intellectual level" (Bengston and Xu 1995 citing Schroeder 1992). Attachment/orientation to nature, sense of place value, and heritage value are also included in this category	cathedral, cherish, cherishing, ecocentric, immortality, mythological, sacred, spiritual, stewardship, transcendence
Life support	Values stemming from life-supporting environmental functions or services of a forest ecosystem	biospheric, complex web, carbon dioxide, degradation, ecological benefits, keystone species, potential energy, soil movement, threatened species, water quality
Economic & utilitarian	The value of a forest ecosystem stemming from its utility for achieving human ends, where the ultimate end or goal is maximizing preferences/satisfaction	commodity, crops of trees, earning, economic growth, economy, exploited for timber, industrial forest, profitable, timber harvest, utilization

area under that curve offers an aggregate measure of performance across all possible thresholds along the curve (Fawcett, 2006). An AUC score of 1.0 implies a perfect ranking of all comments containing a value expression before non-value comments, while a score of 0.5 implies no better than a random ranking. Current recommendations for AUC interpretation in psychology suggest scores between 0.50 and 0.69 to be considered low, 0.70 and 0.89 moderate, and 0.90 and 1.00 high (Streiner & Cairney, 2007).

3.3.3 | Representing values using semantic embeddings

We selected *doc2vec*, a semantic embedding model that uses a shallow two-layered neural network architecture with the objective of optimizing the prediction of words based on their shared local context with other words (Mikolov et al., 2013). Although it is a relatively simple embedding model, previous work demonstrates *doc2vec* is able to capture meaningful topics within a set of documents (Angelov, 2020). We elected to use the *dbow* variant of *doc2vec* implemented using Gensim 4.0 (Řehůřek & Sojka, 2010) given that it has performed better than *pvdM* at similar tasks despite being a simpler model (Angelov, 2020; Lau & Baldwin, 2016).

By training *doc2vec* on comments, we generated joint word-document embeddings. To represent each value within the embedding, we computed an average vector for that value type by taking the centroid mean vector of the terms in the corresponding lexicon. Although other approaches are possible (e.g. geometric mean), previous work found very little difference between methods for averaging embedding vectors (Angelov, 2020), and we therefore opted to implement the simplest approach. We scored each comment for each value type by computing cosine similarity between the document vector and the computed average vector for value terms (Lau

& Baldwin, 2016; Mikolov et al., 2013). The resulting cosine similarity scores range from -1 (indicating perfectly opposite vectors) and 0 (perfectly orthogonal vectors) to 1 (perfectly proportional vectors).

The effect of hyperparameter values on model performance is widely accepted, and optimal settings vary depending on the corpora and the type of task (Bergstra & Bengio, 2012; Caselles-Dupré et al., 2018). We therefore optimized model hyperparameters using a random search approach (Bergstra & Bengio, 2012), training a set of 1200 models to identify optimal hyperparameters (see Supporting Information). All documents (both labelled and unlabelled, see Section 3.3.1) were used to train each model. Our rationale for this was that the *doc2vec* training is completely unsupervised (i.e. the model takes only raw text and uses no supervised or annotated information), and thus there is no need to hold back the labelled data. As with the lexicon scores, we sorted comments in descending order of similarity score and evaluated model performance using AUC.

3.3.4 | Relating social and ecological values to project types

To demonstrate one potential application for these types of models, we used the final models to relate the comments' project-level value score to other variables to learn about whether various project features may affect the expression of different social and ecological values. We included projects that received at least 50 unique comments ($n=338$ projects, 406,762 comments, median of 299 comments per project) in this analysis. To represent the type of activities proposed by each project, we used the USFS project purpose labels¹ assigned to the project in the CARA system by NEPA planners, excluding eight purpose categories that were poorly represented in our

¹These are outlined by the USFS Handbook's natural resource management codes.

TABLE 2 Examples of public comment text by value type (example comments were rated highly by both the human coders as well as the embedding model). Some comments contained citations that were removed.

Value type	Comment text
Aesthetic	The Jefferson and Monongahela National Forests are close enough that in a day trip we can be within some of the most beautiful lands on the East Coast. Nothing else is as restorative for us as being in the natural world where we can regain our sense of place. It gives me hope for the well-being of future generations that the beauty and health of our national forests are protected by my fellow citizens as well as the Forest Service.
Moral & Spiritual	Please. Don't. How can you not see that this is a horrible plan? The West Coast is already getting hit daily by Fukushima radiation. The wildlife, the marine life, the plant life and the humans can't take much more without creating a domino reaction of consequences. Where is your connection to a higher truth? Where is your soul? Please find it before you rape us all.
Life Support	The section on soils (draft plan at page 60) states "...soils are vulnerable to having decreased water availability for plant growth, groundwater recharge and stream recharge...." increased temperatures and decreased precipitation predicted with climate change will likely also have significant impacts on soil composition (e.g. carbon and nitrogen) and soil biota... Changes in soil composition and biota could facilitate non-native invasive plants, result in changes in plant composition and make restoration efforts more difficult. Fire can also have considerable impact not only on soil biota but also its physical, chemical and mineralogical properties...
Economic	There also has to be studies of the effect the noise will have on the National Park experience (and economy) since the added noise levels will seriously degrade the park, which could mean of course a lot less visitors. What effect would this have then on the surrounding communities' economies that are dependent on this tourism? These social/economic impacts cannot be just brushed aside, as the Navy has tried to do. What will the socio/economic impacts be to all of the towns and cities surrounding the park?

data² (see Supplement for comment and labelled sample frequencies by region and project purpose).

For each of the four value types, we used multiple linear regression to relate the mean value score for project comments with underlying project features, namely project purposes, level of public input (comment count), project year and USFS administrative region. We used simple effect coding, which compares each group within a variable (e.g. region) to the unweighted mean of groups, for project purposes (each purpose was coded as applicable = 1, not applicable = -1) and region (national-level projects were used as the reference group). We scaled the logarithmically transformed comment counts to fall between 0 and 1 so that coefficients could be directly compared. Lastly, we recoded project year such that 2011 = 0 and 2020 = 9. To avoid issues with multicollinearity, we checked for strong correlations between predictor variables (all were <|0.54|) and high variance inflation factor values (all were <3.0), neither of which indicate cause for concern (Stevens & Pituch, 2015). We used the standardized value z-scores as the dependent variable in each of the multiple linear regression models so that coefficients could be compared across models. Although we predicted four separate dependent variables (value scores), they were predicted from the same underlying comments dataset. We therefore used a more conservative value of 0.0125 (0.05 divided by the number of related dependent variables predicted) to reduce the risk of Type I error. We measured the predictive power of our models using a coefficient of determination (R^2) metric.

²The excluded project purpose categories were facility management, grazing management, special area management, heritage resource management, regulations directives and orders, research, land ownership management, and land acquisition. Each represented either less than 3% of comments or less than 2% of projects. No project purposes were identified for 12.6% of comments.

4 | RESULTS

4.1 | Social and ecological values in public comments

Life support and moral values were the most common value types observed, expressed in 35% and 14% of the labelled comment subset, respectively. Economic and aesthetic values were less common, expressed in 9% and 8% of labelled comments, respectively.

4.2 | Embedding performance

Comparing the original lexicon word-count approach to our embedding models, our optimized models consistently performed better than the word-counts at ranking comments with researcher-labelled value expressions above comments without value expressions, with the exception of economic value, for which the two methods had comparable performances (see Tables 2 and 3).

Despite the high degree of skew in the dataset (aesthetic values were 8% of labelled comments), we were able to detect value expressions with a moderate to high degree of accuracy using the cosine similarity between document vectors and average value type vectors. Our final model performed the best for aesthetic values and the worst for economic values. For the lexicon word-count, performance was highest for economic values but performed only slightly better than random for moral values. Embedding model performance varied not only by the type of value being inferred, but also by the type of project for which the comment was submitted (Figure 1). The optimized model performed well at detecting aesthetic values across project purposes and across regions, while performance varied considerably for the other three types.

Although optimized embedding models performed better than random chance and generally performed better than the simple lexicon word-count, this was often not the case for untuned models (Figure 2). Regardless of optimization, embedding models frequently performed better than random across all value types. Models for aesthetic and moral values performed considerably better than the simple lexicon count across the 1200 models, but many of the models

for life support and economic values trained during the random search performed substantially worse than a simple word-count.

4.3 | The influence of project characteristics on social and ecological values

The project-level variables included in the multiple linear regression models explained some of the variability in project mean aesthetic, moral, life support and economic values ($R^2=0.29, 0.34, 0.20$ and 0.37 , respectively). Multiple proposed management activities were associated with the expression of one or more values (Figure 3). For example, comments for forest-product projects were less likely to express aesthetic value, but more likely than average to express life support and economic values. Comments for road management projects were less likely to contain moral, life support or economic values, whereas comments for recreation management projects were more likely to contain aesthetic and moral values and less likely to contain economic value (but see model performance discussion below).

TABLE 3 AUC scores for each value type using lexicon wordcounts compared to optimized embedding-based approach.

Value	Model performance—AUC score	
	Lexicon	Optimized embedding
Aesthetic	0.706	0.920
Moral & Spiritual	0.585	0.836
Life support	0.717	0.820
Economic	0.777	0.759

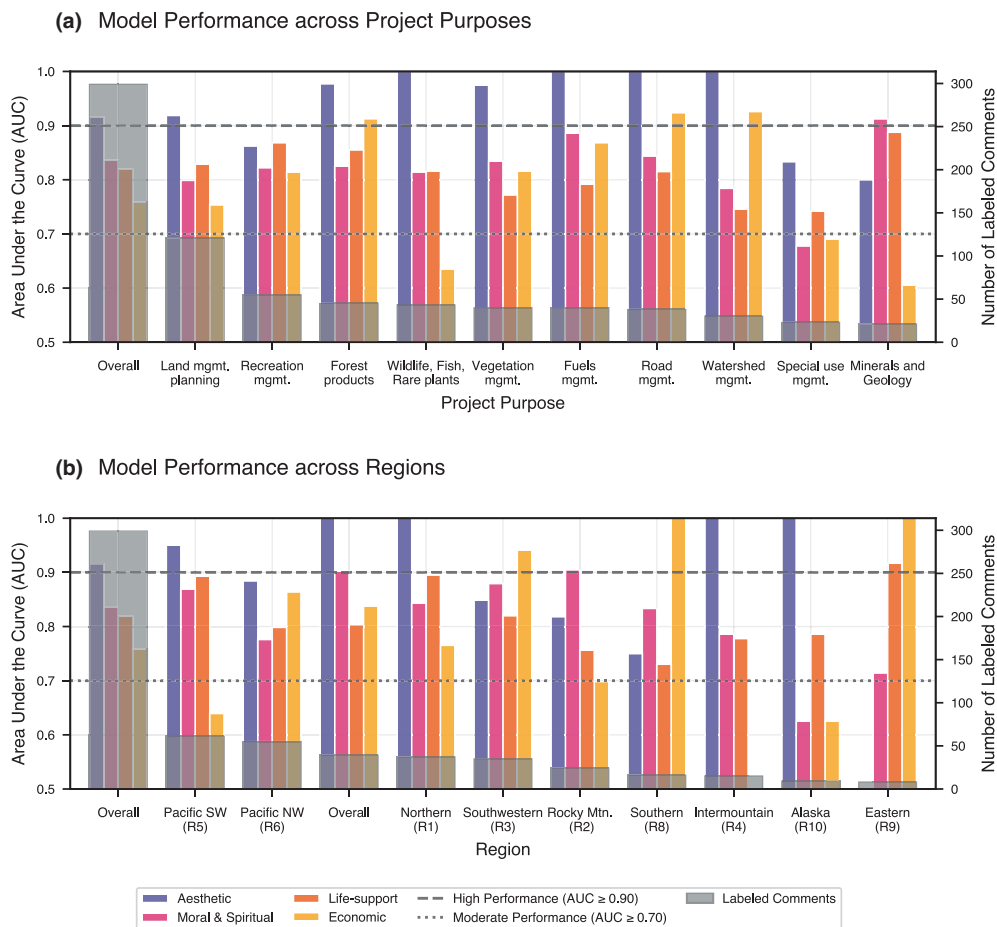


FIGURE 1 Optimized model area under the curve (AUC) score for each value type, by (a) project purpose and (b) region. Project purposes and regions are ordered by the proportion of labelled comments representing that category, indicated by the grey overlay boxes (right y-axis). AUC scores were not computed in cases where no value expressions were observed in the labelled comments. Dashed horizontal lines indicate benchmarks for model performance.

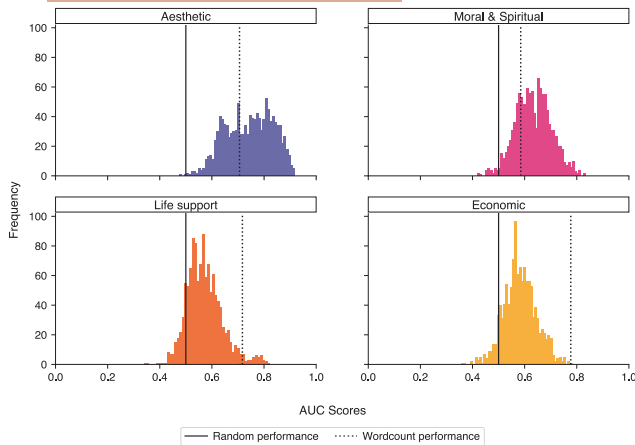


FIGURE 2 Histograms of AUC scores for the 1200 embedding models trained during the random search, by value type. Solid black lines indicate performance no better than random chance (AUC=0.5), and the dotted lines indicates the lexicon wordcount AUC score. AUC scores range from 0 to 1, with 1 indicating a perfect ranking of all labelled comments containing a value expression before non-value comments, and 0.5 indicating no better than a random ranking.




FIGURE 3 Coefficients from the models relating mean project-level value z-scores to project activities, project size and project USFS administrative region ($n=338$, Aesthetic values $R^2=0.35$, Moral values $R^2=0.26$, Life support values $R^2=0.23$ and Economic values $R^2=0.45$, national projects were used as the reference category for region). Coefficients that are insignificant at $p>0.0125$ are partially transparent.

Aesthetic and moral values were expressed more often in comments for projects that received more public interest, as gauged by the number of comments a project received, whereas there was no

association between public interest and life support or economic values. We found no effect of project year on any of the value types after controlling for other project factors. Lastly, we found that projects in some regions received comments expressing significantly different values. For example, comments from projects in the Southern Region (R8) more often expressed aesthetic values, whereas comments from projects in the Eastern Region (R9) expressed more moral values than we would expect on average, even after accounting for differences in project types within each region. Although results suggest that comments from the Alaska region (R10) expressed more economic value than we would expect on average, our results may be least reliable for this region due to the distribution of our labelled dataset (see [Supporting Information](#)).

5 | DISCUSSION

This study introduces and applies a reproducible approach for studying the expression of social and ecological values in unstructured text datasets, such as public comments, using semantic embeddings. When the approach is applied to a corpus of public comments submitted to the USFS, our embedding model consistently ranks texts expressing aesthetic, moral and spiritual, life support and economic values higher than those without such value expressions. This is among the first machine-learning efforts to infer socio-ecological values in text, and the first attempt to develop methods for analysing public comments that explicitly combine prior knowledge (i.e. lexicons) with unsupervised machine-learning methods. While demonstrating that it is possible to infer the expression of different values in public comments, we also find that there is substantial variability in model performance across comments from different types of projects that the USFS considers. We discuss below potential sources of variability in our models, stemming both from model design and from the USFS planning process that generates the experimental corpus. Despite the variability, we conclude that with careful consideration there is potential for this approach to identify where, why and how the public values public lands, providing information in an efficient, reliable manner that can inform land management. More broadly, this study serves as an example of how to develop and apply artificial intelligence to understand novel aspects of interactions within social-ecological systems.

5.1 | Modelling values in text datasets

Although the models were generally good at identifying values in comments, their performance varied substantially across project purposes and administrative regions ([Figure 1](#)). This finding has broader implications for social and ecological values research using machine-learning, as regional and contextual differences in performance highlight the importance of measuring, reporting and addressing the potential for uncertainty in artificial intelligence to produce biased results.

Inferring abstract and nuanced constructs from written text is a challenging task, even for humans, and can require substantial expertise and training to be done reliably (Krippendorff, 2012). Our models were most successful at inferring more concrete, manifest value constructs (i.e. aesthetic values) whereas they struggled to correctly infer more latent constructs that require high levels of contextual understanding and active interpretation to identify (i.e. moral and life support values). This generally mirrors the performance of human researchers, who more consistently identify concrete, manifest values but require additional discussion to develop a nuanced interpretation for latent values (see Section 3.3.1).

Embedding models successfully inferred aesthetic values across projects, but there were important regional and activity-based differences in model performance for the other values. Although discerning precise causes driving embedding performance is challenging, considering the lexicon used provides possible explanations for this variability. The lexicon was designed and evaluated to capture values expressed towards forestlands and contained words specific to forest ecosystems, and may be less accurate for text focusing on other ecosystems. Although many USFS projects occur in forest ecosystems, we suspect that lexicon–corpus (mis)alignment may account for some of the performance heterogeneity in both the lexicon and embedding models.

A second possible driver of performance difference between the models and the lexicon word-count may stem from the comprehensiveness of each value lexicon. The lexicons for aesthetic and moral values developed by Bengston and Xu (1995) are far shorter than those for life support and economic values (153 and 98 terms vs. 495 and 181 terms, respectively). For a simple word-count to accurately capture the underlying construct of interest, it needs to include an exhaustive list of relevant terms, including grammatical variants, synonyms and conceptually related words. Because embedding models cluster words with similar meanings together across embedding space, an approach leveraging embeddings may be more resilient than word-counts to omissions in a lexicon or to minor variations in terminology over time or across populations.

To investigate the embedding model's comparably poor performance on economic values, we reviewed comments that researchers labelled as expressing economic values but that the model scored as low (i.e. false negatives). Many of these comments expressed economic values related to recreation and tourism activities or leaseholders' economic interests. We also see potential evidence of model bias against these economic values in Figure 1a, where the model performs better for comments from projects more directly related to logging and worse for comments associated with other project activities. This is notable, given that terms in the economic values lexicon are chiefly related to either timber commodities (e.g. "bid price", "firewood", "harvest level", "lumber market", "stumpage"), forest utilization (e.g. "commodity", "exploit", "monetize", "utilize") or employment (e.g. "economic development", "profits", "wage", "workforce"). Adding terms representing a wider range of economic values may improve model performance for future work.

Although our optimized models performed better than word-counts, this was not true for untuned models; most models outperformed the lexicon for aesthetic and moral values, but few performed better at inferring life support and economic values (Figure 2). There is often a temptation to use generic hyperparameter values for unsupervised learning; however, our findings demonstrate the importance of rigorous model evaluation to optimize performance for specific tasks and to identify potential error and bias. In some cases, a more parsimonious approach, such as a simple word count, may outperform more complex and opaque methods.

NLP studies rarely include practices that would be recognized as rigorous in qualitative social science contexts (Tanweer et al., 2021, but see Ha & Grubert, 2023). Annotation procedures used to train and evaluate models are frequently implemented using simplistic keyword tagging or pre-existing labels rather than question-specific human evaluation, and evaluations are sometimes done by non-experts without intercoder reliability checks (Grosman et al., 2020; Lau & Baldwin, 2016; Snow et al., 2008). Incorporating rigorous qualitative practices when creating labelled training datasets or evaluating unsupervised model results can help ensure that results are both meaningful and reliable (Tanweer et al., 2021). Model evaluation is often more straightforward for supervised learning; methods for evaluating unsupervised learning are still being solidified. This work serves as one example of rigorously evaluating unsupervised models in a social science context (see also Nelson, 2020).

5.2 | Values expressed in USFS public comments

From our analysis of values in USFS projects, we gain an understanding of underlying values associated with National Forest lands in the United States. Specifically, we show how different social and ecological values relate to the types of activities being proposed, public interest in a project and regional differences. Critiques of the USFS comment analysis process have observed that public feedback framed in overtly normative terms is often dismissed or overlooked (De'Arman, 2020; Predmore et al., 2011; Vining & Tyler, 1999). So, it is worth noting that the values most likely to be dismissed as subjective preference—*aesthetic and moral values*—were expressed in 9% and 14% of labelled comments, respectively. There seems to be substantial variation in how NEPA legislation and planning guidance are interpreted and implemented by USFS planners. Some sources suggest that the USFS is required to consider the full range of values in the planning process (United States Forest Service, n.d.), while others suggest the planning rule does not "preclude the Forest Service from making decisions based on values" (Brown & Nie, 2019, p. 3), or that simple value statements should be excluded from consideration (De'Arman, 2020). This range of interpretations may partially explain the regional variation of value expressions; it may be that planning teams include wider or narrower ranges of values when identifying substantive comments.

Life support was the most prevalent value in the labelled comments dataset, expressed in roughly a third of the comments we

manually reviewed. This may reflect a common or deeply held ecological value among the public. Prior work suggests that public values towards forests have shifted over the latter 20th century away from economic values and towards ecological and non-economic values (Bengston et al., 1999, 2004). It may also be that USFS analysts are more sensitive to public feedback framed in terms of life support value than feedback framed in other ways and are thereby more likely to deem such feedback as “substantive.” The USFS 2012 Planning Rule (77 FR 21162) directs planners to place additional emphasis on ecological integrity, landscape restoration and wilderness protection, and use the best available scientific information, among other considerations. Relatedly, savvy commenters may be motivated to couch their feedback in terms of life support value and scientific or ecological justification as a strategy to influence the planning process, rather than a genuine reflection of values (Predmore et al., 2011). This is likely truer for established, engaged interest groups with experience in the USFS policymaking process, and less so for lay commenters. Given the temporal span of our dataset, we are unable to observe possible shifts in comment values pre- and post-implementation of the 2012 Planning Rule.

Similarly, we observed that comments expressing economic values were more likely to come from particular regions (e.g. USFS Region 10), even after accounting for project type. It may be that the public places greater economic value in these areas, but it is just as possible that planners in these areas are more sensitive to economic values expressed in public feedback than planners in other regions and are therefore more likely to label text containing economic values as “significant and substantive.” The dataset used here represents the body of public feedback that was formally considered in the planning process and that the agency evaluated as being relevant to decision-making. Further work is required to understand the extent to which values in this dataset are representative of the full, unfiltered body of public letters received by the agency. Additionally, it is unclear what factors affect the filtering of values in the comment analysis process, and to what extent values in substantive comments are reflected in agency response, either formal written response or in observable influence in the planning process.

5.3 | Limitations, precautions and future work

Although our work demonstrates promise for NLP and embedding models, we also illustrate several limitations. As with other research methods, considering whether a particular corpus is suitable for a given line of inquiry is crucial, as the ability to draw inferences about a given group depends on the text used in model training. As mentioned previously, our analysis used “substantive” comments, a corpus that was likely filtered for some value expressions in the raw public letters. Whether letters from the public are stored and processed in CARA is voluntary and may bias the corpus towards comments from larger and more recent projects. More fundamentally, barriers to participating in the public planning, such as cynicism and distrust in the policymaking process, inaccessibility of public meeting

locations or comment submission platforms, and a lack of cultural, language, or bureaucratic familiarity needed to meaningfully engage in the public comment process all suggest that letters are not necessarily representative of broader public values (Predmore et al., 2011; Ulibarri et al., 2022; Yakubu, 2018). So, although we cannot describe to the public's social and ecological values generally, this work can speak to the values in public feedback the agency specifically identified as important for guiding decision-making.

We trained our models on all comments in the dataset. This provided us with a dataset large enough for a model to distil subtle semantic relations such as values but required that we combine comments written by vastly different authors in vastly different social contexts. Our models thus capture broad values but obscure the heterogeneity of value expressions across individuals and stakeholder groups (e.g. industry vs. environmental interest groups vs. lay commenters). For example, the intended meaning behind economic terms used in comments appears to vary across comment writers. We observed that authors expressing moral reverence appear to use economic terms to connote negativity or scepticism towards utilitarian uses, quite different from how industry groups use these same terms. Because the model does not necessarily distinguish between different authors' use of terms, both the model and lexicon systematically mistake oppositional or sceptical uses of economic terms as genuine.

Unlike traditional lexicons, the exact analytic and algorithmic processes within embedding models are highly complex and often elude our ability to explain. Identifying sources of model error and improving performance are rarely straightforward tasks. Although the model we used (*doc2vec*) uses a two-layered neural network with a single hidden layer, more advanced deep learning models popularized in recent years (e.g. *chatGPT*) use tens or hundreds of billions of parameters, compounding opacity and interpretability issues. This increased complexity has improved performance, but it also limits researchers' understanding of how the model is trained and what distortions it is likely to produce (Kozłowski et al., 2019).

Finally, increasingly complex machine-learning methods necessitate increased technical and computing resources and expertise. Even though we used a comparably simple embedding model, model tuning nonetheless required substantial computing resources and access to computing clusters to parameterize models in a timely manner. Furthermore, these models require substantial technical and coding expertise and are often time-consuming to implement and evaluate. By contrast, word-counts can be easily and quickly implemented using widely available software or relatively simple programming and may be easier to incorporate into complex workflows. We think the embedding approach that we developed offers promise and we encourage further research investigating whether embedding models can identify values in other corpora of public comments.

6 | CONCLUSION

Formalizing and attending to the range of values expressed in public comments about land management offers planners and

decision-makers the opportunity to better understand the public input they receive, develop management strategies that better align with various publics, and to anticipate, address and mitigate value-laden sources of conflict that may arise. We recognize that there are technical, individual and institutional factors that bar land management agencies such as the USFS from considering values in the planning process, such as overwhelming workloads, negative normative beliefs about the public, fear or avoidance of conflict and a lack of leadership commitment to addressing public influence (Hoover & Stern, 2014; MacGregor & Seesholtz, 2008). Nonetheless, computational approaches leveraging emerging NLP methods, such as the one we present here, can serve as powerful tools. When care is taken to understand performance and sources of error and bias, these models can extract precise and complex semantic information efficiently and reliably. Furthermore, the psycho-social dimensions that may be observable using semantic embeddings are not limited to the four values we explored. The success of these models offers promise for their application towards other constructs of interest. Model results may also help make a broader range of social and ecological values in public comments more visible and allow researchers to better understand patterns and trends in how values are expressed at large spatial and temporal scales.

AUTHOR CONTRIBUTIONS

Sarah K. MacFarland: Conceptualization, methodology, software, validation, formal analysis, investigation, writing—original draft preparation, Writing—review and editing, Visualization, Project administration. Sonya Sachdeva: Conceptualization, methodology, validation, formal analysis, investigation, writing—review and editing, supervision, funding acquisition. Spencer A. Wood: Conceptualization, methodology, software, resources, writing—review and editing, supervision, funding acquisition. Joshua J. Lawler: Resources, writing—review and editing, supervision.

ACKNOWLEDGEMENTS

We gratefully acknowledge financial support from the USFS (agreements 20-JV-11261985-095 and 21-CS-11132400-3). Part of this work was completed using the Hyak computing cluster at the University of Washington, funded by the UW student technology fee. We also thank Dave Bengston for his helpful comments on this manuscript.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

DATA AVAILABILITY STATEMENT

These are public federal data, archived by the US Forest Service. Comments for projects are often available in planning documents associated with individual projects, and the complete dataset can be requested from the USFS.

ORCID

Sarah K. MacFarland  <https://orcid.org/0009-0005-2710-1578>

REFERENCES

- Anderson, N., Ford, R. M., Bennett, L. T., Nitschke, C., & Williams, K. J. H. (2018). Core values underpin the attributes of forests that matter to people. *Forestry: An International Journal of Forest Research*, 91(5), 629–640.
- Angelov, D. (2020). *Top2Vec: Distributed representations of topics*. (arXiv:2008.09470). *arXiv*.
- Bengston, D. N. (1994). Changing forest values and ecosystem management. *Society & Natural Resources*, 7(6), 515–533.
- Bengston, D. N., Fan, D. P., & Celarier, D. N. (1999). A new approach to monitoring the social environment for natural resource management and policy: The case of US national forest benefits and values. *Journal of Environmental Management*, 56(3), 181–193.
- Bengston, D. N., Webb, T. J., & Fan, D. P. (2004). Shifting Forest value orientations in the United States, 1980–2001: A computer content analysis. *Environmental Values*, 13(3), 373–392.
- Bengston, D. N., & Xu, Z. (1995). *Changing national forest values: A content analysis (NC-323)*. North Central Forest Experiment Station, Forest Service, U.S. Department of Agriculture.
- Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(2), 281–305.
- Brown, S. J. M., & Nie, M. (2019). Making Forest planning great again? Early implementation of the Forest Service's 2012 National Forest Planning Rule. *Natural Resources and Environment*, 33(3), 3–7.
- Buijs, A., & Lawrence, A. (2013). Emotional conflicts in rational forestry: Towards a research agenda for understanding emotions in environmental conflicts. *Forest Policy and Economics*, 33, 104–111.
- Caselles-Dupré, H., Lesaint, F., & Royo-Letelier, J. (2018). Word2vec applied to recommendation: Hyperparameters matter. Proceedings of the 12th ACM conference on recommender systems, 352–356.
- Costa, M., Desmarais, B. A., & Hird, J. A. (2019). Public comments' influence on science use in U.S. rulemaking: The case of EPA's National Emission Standards. *The American Review of Public Administration*, 49(1), 36–50.
- Council on Environmental Quality. (2024). *National Environmental Policy act Implementing Regulations* (40 CFR parts 1500–1508). <https://ceq.doe.gov/docs/laws-regulations/NEPA-Implementing-Regulations-Desk-Reference-2024.pdf>
- De'Arman, K. J. (2020). Is public participation public inclusion? The role of comments in US Forest Service decision-making. *Environmental Management*, 66(1), 91–104. <https://doi.org/10.1007/s00267-020-01278-5>
- Dokshin, F. A. (2022). The public speaks: Using large-scale public comments data in public response research. *Energy Research & Social Science*, 91, 102–689.
- Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874.
- Fleming, C. J., McCartha, E. B., & Steelman, T. A. (2015). Conflict and collaboration in wildfire management: The role of mission alignment. *Public Administration Review*, 75(3), 445–454.
- Forester, J. (1999). The consensus building handbook: A comprehensive guide to reaching agreement. In *The consensus building handbook: A comprehensive guide to reaching agreement* (pp. 463–494). SAGE Publications, Inc.
- Germain, R. H., Floyd, D. W., & Stehman, S. V. (2001). Public perceptions of the USDA Forest Service public participation process. *Forest Policy and Economics*, 3(3), 113–124.
- Grosman, J. S., Furtado, P. H. T., Rodrigues, A. M. B., Schardong, G. G., Barbosa, S. D. J., & Lopes, H. C. V. (2020). Eras: Improving the quality control in the annotation process for natural language processing tasks. *Information Systems*, 93, 101–553.
- Gugulica, M., & Burghardt, D. (2023). Mapping indicators of cultural ecosystem services use in urban green spaces based on text classification of geosocial media data. *Ecosystem Services*, 60, 101–508.

- Ha, S., & Grubert, E. (2023). Hybridizing qualitative coding with natural language processing and deep learning to assess public comments: A case study of the clean power plan. *Energy Research & Social Science*, 98, 103016.
- Hoover, K., & Stern, M. J. (2014). Constraints to public influence in US Forest Service NEPA processes. *Journal of Environmental Planning and Management*, 57(2), 173–189.
- Hurst, K., & Stern, M. J. (2020). Messaging for environmental action: The role of moral framing and message source. *Journal of Environmental Psychology*, 68, 101–394.
- Innes, J. E., & Booher, D. E. (2004). Reframing public participation: Strategies for the 21st century. *Planning Theory & Practice*, 5(4), 419–436.
- Jones, N., Shaw, S., Ross, H., Witt, K., & Pinner, B. (2016). The study of human values in understanding and managing social-ecological systems. *Ecology and Society*, 21(1), 15.
- Kozłowski, A. C., Taddy, M., & Evans, J. A. (2019). The geometry of culture: Analyzing the meanings of class through word embeddings. *American Sociological Review*, 84(5), 905–949.
- Krippendorff, K. (2012). *Content analysis: An introduction to its methodology* (3rd ed.). SAGE Publications, Inc.
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174.
- Lau, J. H., & Baldwin, T. (2016). An empirical evaluation of doc2vec with practical insights into document embedding generation. Proceedings of the 1st workshop on representation learning for NLP (pp. 78–86).
- Livermore, M., Eidelman, V., & Grom, B. (2018). Computationally assisted regulatory participation. *Notre Dame Law Review*, 93(3), 977.
- MacGregor, D. G., & Seesholtz, D. N. (2008). *Factors influencing line officers' decisions about National Environmental Policy act project design and development* General technical report PNW-GTR-766, 27 p, 766. U.S. Department of Agriculture, Forest Service, Pacific northwest Research Station.
- Marshall, N. A., Thiault, L., Beeden, A., Beeden, R., Benham, C., Curnock, M. I., Diedrich, A., Gurney, G. G., Jones, L., Marshall, P. A., Nakamura, N., & Pert, P. (2019). Our environmental value orientations influence how we respond to climate change. *Frontiers in Psychology*, 10, 938.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26, 1–9.
- Moyer, J. M., Owen, R. J., & Duinker, P. N. (2008). Forest values: A framework for old-growth Forest with implications for other Forest conditions. *The Open Forest Science Journal*, 1(1), 27–36.
- Nelson, L. K. (2020). Computational grounded theory: A methodological framework. *Sociological Methods & Research*, 49(1), 3–42.
- Predmore, S. A., Stern, M. J., & Mortimer, M. J. (2011). Constructing the public: The “substantive sieve” and personal norms in US Forest Service planning. *Journal of Environmental Planning and Management*, 54(3), 403–419.
- Proctor, J. D. (1998). Environmental values and popular conflict over environmental management: A comparative analysis of public comments on the clinton forest plan. *Environmental Management*, 22(3), 347–358.
- Řehůřek, R., & Sojka, P. (2010). Software framework for topic modeling with large corpora. Proceedings of the 2010 workshop on new challenges for NLP frameworks, 45–50.
- Riffe, D., Lacy, S., Fico, F., & Watson, B. (2019). *Analyzing media messages: Using quantitative content analysis in research* (4th ed.). Routledge.
- Scott, T. A., Ulibarri, N., & Perez Figueroa, O. (2020). NEPA and national trends in federal infrastructure siting in the United States. *Review of Policy Research*, 37(5), 605–633.
- Shulman, S. W. (2009). The case against mass E-mails: Perverse incentives and low quality public participation in U.S. Federal Rulemaking. *Policy & Internet*, 1(1), 23–53.
- Smiley, S., de Loë, R., & Kreutzwiser, R. (2010). Appropriate public involvement in local environmental governance: A framework and case study. *Society & Natural Resources*, 23(11), 1043–1059.
- Snow, R., O'Connor, B., Jurafsky, D., & Ng, A. (2008). Cheap and fast—But is it good? Evaluating non-expert annotations for natural language tasks. Proceedings of the 2008 conference on empirical methods in natural language processing, 254–263.
- Stevens, K. A., & Pituch, J. P. (2015). *Applied multivariate statistics for the social sciences: Analyses with SAS and IBM's SPSS* (6th ed.). Routledge.
- Streiner, D. L., & Cairney, J. (2007). What's under the ROC? An introduction to receiver operating characteristics curves. *The Canadian Journal of Psychiatry*, 52(2), 121–128.
- Swanson, S. (1994). Viewpoint: Integrating CRM (coordinated resource management) and NEPA (National Environmental Policy act) processes. 47.
- Tadaki, M., Sinner, J., & Chan, K. (2017). Making sense of environmental values: A typology of concepts. *Ecology and Society*, 22(1), 7.
- Tanweer, A., Gade, E. K., Krafft, P. M., & Dreier, S. K. (2021). Why the data revolution needs qualitative thinking. *Harvard Data Science Review*, 3(3), 1–32.
- Thacher, D., & Rein, M. (2004). Managing value conflict in public policy. *Governance*, 17(4), 457–486.
- Ulibarri, N., Pérez Figueroa, O., & Grant, A. (2022). Barriers and opportunities to incorporating environmental justice in the National Environmental Policy act. *Environmental Impact Assessment Review*, 97, 106–880.
- United States Forest Service. (n.d.). *FAQs on 2012 planning rule*. United States Department of Agriculture, Forest Service. https://www.fs.usda.gov/detail/planningrule/faqs#main_content
- Vining, J., & Tyler, E. (1999). Values, emotions and desired outcomes reflected in public responses to Forest management plans. *Human Ecology Review*, 6(1), 21–34.
- Wallace, K. J., Jago, M., Pannell, D. J., & Kim, M. K. (2021). Wellbeing, values, and planning in environmental management. *Journal of Environmental Management*, 277, 111–447. <https://doi.org/10.1016/j.jenvman.2020.111447>
- Yakubu, O. H. (2018). Delivering environmental justice through environmental impact assessment in the United States: The challenge of public participation. *Challenges*, 9(1), Article 1.

SUPPORTING INFORMATION

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

Figure S1. The relative frequency of comments in the full dataset (green) and the labeled sample (blue) submitted for projects associated with each project purpose. Projects may be associated with more than one purpose.

Figure S2. The relative frequency of comments in the full dataset (green) and the labeled sample (blue) submitted for projects in each USFS administrative region.

How to cite this article: MacFarland, S. K., Sachdeva, S., Wood, S. A., & Lawler, J. J. (2025). Characterizing social and ecological values expressed in US Forest Service public comments using a computational approach. *People and Nature*, 00, 1–12. <https://doi.org/10.1002/pan3.70050>