Divesting in Socially (Ir)responsible Internet Service Providers

Emily Huff
Colgate University
ehuff@colgate.edu

Aaron Gember-Jacobson
Colgate University
agemberjacobson@colgate.edu

ABSTRACT

Divesting in a company based on their social practices is one tactic for encouraging a company to alter their behavior. Individuals and institutions can divest in "socially irresponsible" ISPs directly by refusing to peer with or purchase transit from these ISPs, or indirectly by avoiding paths through these ISPs. While the economics of divesting in ISPs and mechanisms for exerting more control over Internet paths have been previously studied, a major challenge remains: how do individuals and institutions assess whether an ISP is socially responsible? Large institutions may have the resources to research which ISPs are socially (ir)responsible, but most individuals and institutions lack these resources, thereby limiting their ability to effect social change.

We address this challenge by designing a system for automatically assessing ISPs along various dimensions of social responsibility—e.g., environmental stewardship, censorship actions, or net neutrality practices. In particular, we gather publicly accessible web pages which discuss a particular ISP and social issue and apply stance detection—a natural language processing technique for determining a subject's attitude or commitment toward an object—to compute a "socially responsibility" score for individual ISPs and dimensions of social responsibility. We construct a data set over 600 web pages relating to net neutrality and 13 ISPs of varying sizes, and show our framework achieves 71% accuracy in determining an ISP's stance.

CCS CONCEPTS

• **Networks** → **Network properties**.

ACM Reference Format:
Emily Huff and Aaron Gember-Jacobson. 2021. Divesting in Socially (Ir)responsible Internet Service Providers. In *ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet (TAURIN '21), August 23, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3472951.3473504

1 INTRODUCTION

Globally, the Internet service provider (ISP) industry has a market size of $726 billion [15]. The potential influence of such a high-revenue industry on the equitable and responsible treatment of individuals, organizations, and the environment is substantial. But social responsibility—e.g., ISPs' environmental impact, censorship actions, and net neutrality practices—does not yet factor into how traffic is routed through the Internet. Instead, traffic is routed along paths based on availability, distance, capacity, monetary cost, or other factors unrelated to social responsibility.

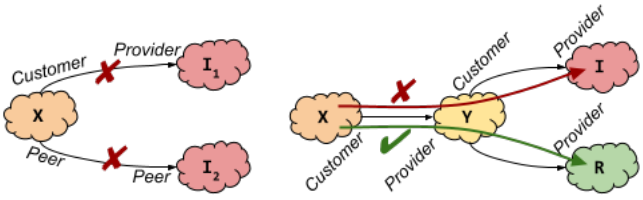

(a) Direct divestment **(b) Indirect divestment**

Figure 1: Examples of divesting in ISPs

Boycotting or divesting in a company based on their social or political practices is one tactic for encouraging a company to alter their behavior. For example, the increased distaste for tobacco, private prisons, and fossil fuels has encouraged individuals and institutions to divest in companies which support these practices [55]. These actions pressure a company to change their practices to appease investors/customers and alleviate financial losses.

Institutions and individuals can *directly* divest in "socially irresponsible" ISPs by refusing to peer with or purchase transit/access from these ISPs (Figure 1a). For example, a popular content provider can refuse to peer with a socially irresponsible ISP, thereby forcing the ISP to incur transit costs to reach the content provider. As another example, an individual or institution can avoid becoming a customer of a socially irresponsible ISP, thereby preventing the ISP from earning more revenue. Institutions and individuals can also *indirectly* divest in ISPs by avoiding paths through socially irresponsible ISPs. For example, assume institution X is a customer of ISP Y, and Y is a customer of socially irresponsible ISP I and socially responsible ISP R (Figure 1b). Whenever X's traffic transits I (or R), I (or R) will earn revenue from Y[1]. Consequently, if X avoids a path through I—using source routing [53], extensions to BGP [37, 48, 62], overlays [19, 54] or other strategies that provide more control over Internet paths [52]—X will cause I to earn less revenue from Y. Institutions and individuals can also financially *invest* in socially responsible ISPs by peering with, purchasing transit from, or preferring paths through such ISPs.

While the economics of divesting in ISPs [45, 60] and mechanisms for exerting more control over Internet paths [19, 37, 48, 52–54, 62] have been previously studied, a major challenge remains: *how do individuals and institutions assess whether an ISP is socially responsible?* There is no organization or central authority that assesses whether ISPs are socially responsible. Researchers have developed measurement tools to quantify some of an ISP's social practices—e.g., violations of net neutrality [32, 43, 57]—but with over 100K autonomous systems (ASes) in existence [4], it is impractical to

[1] Based on 95th-percentile usage, commit level, and/or destination [59]

TAURIN '21, August 23, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-8639-5/21/08…$15.00
https://doi.org/10.1145/3472951.3473504

apply these tools to every ISP. Furthermore, no measurements tools exist for certain social practices—e.g, environmental impact. While large content providers or ISPs may have the resources to research which ISPs are socially (ir)responsible, most individuals and institutions lack these resources. This limits individuals' and institutions' ability to effect social change, and makes it difficult to employ grassroots movements to create a more socially responsible Internet.[2]

We address this challenge by designing *a system for automatically assessing ISPs along various dimensions of social responsibility*. In particular, we gather publicly accessible news articles, blog posts, company policies, etc. which mention a specific ISP and social issue (e.g., net neutrality). Then we extract relevant phrases and apply stance detection—a natural language processing (NLP) technique for determining a subject's viewpoint on a specific topic—to compute "social responsibility scores" for ISPs.

We implement our framework in Python using Beautiful Soup [13], spaCy [7], and scikit-learn [5]. We apply our framework to assess the net neutrality posture of 13 ISPs of varying sizes. Our framework automatically collects over 600 documents and extracts over 250 relevant sentences. We manually label these sentences and use 10-fold cross-validation to train and test our stance detector, which achieves 71% accuracy. We also combine our inferred ISP stances with data from the Route Views Project [9] to assess how many paths from our university to popular destinations currently traverse only ISPs which support/practice net neutrality. Our implementation and dataset are publicly available [6].

The remainder of this paper is organized as follows. Section 2 illustrates the connections between ISPs' practices and individuals'/institutions' social values. Section 3 discusses our NLP-based approach for assessing ISPs' social responsibility, and Section 4 provides details on our implementation. Section 5 evaluates the feasibility and effectiveness of our approach for assessing the net neutrality posture of 13 ISPs of varying sizes. Finally, Sections 6 and 7 discuss related and future work, respectively.

2 DIMENSIONS OF SOCIAL RESPONSIBILITY

There are many ways an ISP may be socially (ir)responsible. In this section, we discuss four dimensions of social responsibility that have recently received significant attention. Our goal is to illustrate the connections between ISPs' practices and individuals'/institutions' social values. Our goal is *not* to prescribe which practices are (ir)responsible, because individuals and institutions may have different social values.

Environmental stewardship. It was projected that in 2020 ISPs would be responsible for 256 metric tons of CO_2 emissions, with two-thirds of this arising from energy usage [10]. Fortunately, Bolla et al. estimate that using energy-aware technologies and techniques in a medium-sized ISP could reduce the ISP's energy consumption by 60% [25], while Chiaraviglio et al. estimate that concentrating traffic on fewer routers and links could reduce an ISP's energy consumption by 35% [29]. Environmentally-conscious individuals/institutions may want to divest in ISPs that do not employ such energy-saving techniques. However, to the best of our knowledge, no existing tools allow third parties to measure an ISPs' energy efficiency—prior measurement studies [25, 26, 29] all rely on internal data provided by the ISPs.

Censorship. Censorship is the filtering of traffic based on factors such as content, source, or destination. Government-imposed censorship is widespread [3, 14]. However, the Internet Society's Global Internet Reports articulate numerous benefits of an open Internet [11, 12]. Consequently, human-rights-conscious individuals/institutions may want to divest in ISPs that filter network traffic. Several large-scale measurement platforms currently quantify and publicly release censorship data on a per-ISP basis [34, 56].

Net neutrality. Net neutrality is the requirement that ISPs do not discriminate among traffic in terms of performance (bandwidth, access, etc.) [30]. Net neutrality has been a popular source of debate since the early 2000s. Hahn and Wallsten [38] argue that net neutrality harms economic welfare by reducing investment incentives and discouraging innovation in the form of apps or services. Hahn and Wallste liken net neutrality to price regulations, which in the long run harm consumers by eliminating provider assurance of fast content delivery, or adding additional broadband infrastructure cost. In contrast, Cheng et al. [28], find that while broadband providers claim implementing net neutrality would reduce their incentive to expand, the increase in competition actually incentivizes bandwidth expansion, as capacity is the only option for differentiation. Equality-conscious individuals/institutions may want invest in ISPs that support/practice net neutrality, while free-market supporters may want to divest in ISPs that support/practice net neutrality.

Routing security practices. An ISP's routing security practices impact the vulnerability of users' data as well as the overall health of the Internet. For example, ISPs that validate route announcements [44] can avoid routing users' traffic along spurious paths. Alternatively, ISPs that employ source address validation (SAV) can reduce the opportunities for attackers to launch denial of service (DoS) attacks from their network, thereby improving the security of the Internet as a whole [46]. Individuals may want to divest in non-security-conscious ISPs to encourage them to change their practices. The Mutual Agreed Norms for Routing Security (MANRS) maintains a repository of ISPs that comply with certain best practices [2], but not all ISPs that employ such practices may officially join MANRS, motivating the need for an alternative mechanism to quantify ISPs commitment to routing security.

3 OUR FRAMEWORK

Our framework for assessing ISPs' social responsibility consists of three phases (Figure 2): (1) gathering publicly accessible news articles, blog posts, corporate policies, etc. which mention individual ISPs and a specific dimension of social responsibility (Section 2); (2) extracting phrases that feature an ISP and mention a specific dimension of social responsibility; and (3) using stance detection to determine the polarity of extracted phrases. We discuss each phase below. Implementation details are discussed in Section 4.

3.1 Gathering relevant documents

The first step in assessing an ISP on a specific dimension of social responsibility is to gather documents containing information about the ISP and the selected dimension. In keeping with our observation that there is no central source of information regarding ISPs'

[2]Individuals' and institutions' ability to exert control over Internet paths is also currently limited [60], but overlays provide an avenue for grassroots movements.

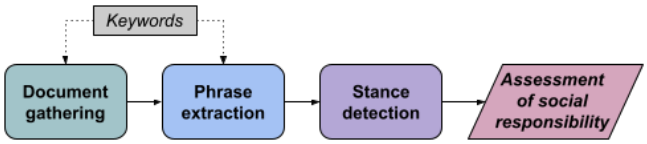

Figure 2: Overview of our framework

social responsibility (Section 1), we employ simple web searches to gather such documents. This allows us to gather publicly accessible news articles, blog posts, company policies, scientific reports, etc. which mention an ISP and a dimension of social responsibility. However, gathering documents using a simple web search raises several challenges our framework must address.

Search terms. The dimensions of social responsibility discussed in Section 2 are each complex issues with many different facets. For example, net neutrality encompasses the principles of an open Internet [61], paid prioritization [47], zero-rating [23], and blocking [36]. Consequently, we use multiple keywords to locate documents that address these various facets. We combine each keyword with the name of an ISP to locate documents related to that ISP's posture: e.g., `"Internet2" AND "net neutrality"`.

Unrelated occurrences of search terms. Documents included in search results must include both the ISP's name and the specified keyword, but there is no guarantee these two terms are related. For example, multiple documents that discuss the Internet's history [27, 49, 51] appear in search results for `"NYSERNet" AND "net neutrality"`, because the documents mention both NYSERNet—one of the Internet's first ISPs—and the issue of net neutrality. However, NYSERNet and net neutrality are discussed separately, with no connections between the ISP and dimension of social responsibility.

We remedy this issue in two ways. First, we use an advanced Google search operator (*AROUND*) that requires the search terms to occur near each other in the document [39], thereby increasing the likelihood that the document discusses the dimension of social responsibility in relation to the ISP. Second, we extract only relevant phrases from the gathered documents, as discussed in Section 3.2.

Duplicated documents. The same document may appear in whole or in part at multiple URLs. For example, an article written by a syndicated news service may appear on multiple news sites [20–22], an excerpt of a blog post may appear on a page listing all posts in a particular category, or a quote from a company spokesperson may appear in multiple documents. We filter out full duplicates by checking if two documents have the same title. We filter out "index pages" by checking if documents contain several `article` tags [16]. We do not filter out documents with common quotes, because each document may contain other novel information.

Different service offerings. Some ISPs offer several different services, including residential access, mobile access, and transit. These service offerings may differ with regards to a dimension of social responsibility: e.g., AT&T's mobile services included zero-rated streaming for a certain content provider [33, 50], but the zero-rating did not extend to AT&T's residential access or transit offerings.

At first glance, differentiating between these services seems important. For example, an individual who wants to *directly* divest in a socially irresponsible ISP (Figure 1a) cares about its residential/mobile access, whereas an individual who wants to *indirectly* divest in an ISP (Figure 1b) cares about its transit services. However, divestment is typically employed for the benefit of society, as opposed to one's own benefit. Consequently, individuals should consider *all* of an ISP's services when deciding whether to divest in the ISP. For example, a human-rights-conscious individual should divest in an ISP if it employs censorship in any of its services, regardless of whether the censorship applies to the individual.

False information. Documents may contain false/misleading information about an ISP. For example, a competitor may claim an ISP's practices harm the environment, an authoritarian government may claim an ISP does not engage in censorship, or an ISP may claim it supports net neutrality whereas third-parties claim the opposite. Fortunately, we find there are a sufficient number of document sources for large ISPs such that a few sources with false/misleading information will not change our conclusion on an ISP's stance on net neutrality (Section 5.4). However, for smaller ISPs or ISPs under authoritarian governments, false/misleading information may skew our analysis. In the future, we plan to explore whether fake news detection techniques [64] can address this issue.

3.2 Extracting relevant phrases

Although some documents focus exclusively on a single ISP and a single dimension of social responsibility—e.g., Cogent's statement on net neutrality [31]—most documents include additional information that is not specific to an ISP and/or its posture—e.g., news articles discussing AT&T's response to California's net neutrality law also provide background on the law itself [33, 40, 50]. Thus, performing stance detection on a document in its entirety may conflate other parties' postures with the ISP's posture or other unrelated issues with the target dimension of social responsibility.

Consequently, the second phase of our framework focuses on extracting relevant phrases from the gathered documents. In particular, we break each document into sentences and perform part-of-speech tagging [8] to identify the subject(s) and object(s) of a sentence. We extract all sentences whose subject(s) contain the name of the ISP of interest, and whose object(s) contain one or more of the keywords associated with the target dimension of social responsibility. For example, we extract the sentence "Cogent practices net neutrality," [31] but ignore the sentence "Pai, a former in-house lawyer for telecom giant Verizon who closed investigations into AT&T, Verizon, and other wireless providers after President Donald Trump designated him as the FCC chairman, and the FCC successfully dismantled the nation's net neutrality laws in 2018." [40].

One limitation of this approach is that it ignores sentences which do not mention an ISP and target dimension of socially responsibility yet clearly express an ISP's posture. For example, our framework ignores sentences that indirectly refer to an ISP: "'Although your company has repeatedly stated publicly that it supports legally binding net neutrality rules, this policy appears to run contrary to the essential principle that in a free and open internet, service providers may not favor content in which they have a financial interest over competitors' content,' the senators said in their letter." [33]. In the future, we plan to explore more nuanced phrase extraction which, for example, performs pronoun resolution [63].

ISP		Search results		Excluded				Final	
Name	Tier	Total	Total w/o nearby	HTTP error	Not HTML	Index page	Duplicate	Total	From ISP
AT&T	1	192	161	55	4	2	5	126	7
British Telecom	2	86	176	27	8	1	3	47	0
Cogent Comm.	2	111	172	44	11	1	4	51	14
Deutsche Telekom	1	141	173	31	13	1	3	93	5
FirstLight Fiber	3	13	77	5	0	2	0	5	0
Hurricane Electric	2	53	164	28	3	0	1	21	0
Internet2	3	96	177	37	9	2	0	48	1
KDDI	2	92	178	40	16	1	1	34	1
NYSERNet	3	25	96	11	9	0	0	5	0
Sprint	1	188	178	65	3	5	9	106	7
Telefonica	2	144	168	69	8	5	3	59	1
Telia Carrier	1	19	133	8	2	0	1	8	1
Zayo	1	43	176	16	3	4	0	20	0

Table 1: Document gathering

3.3 Performing stance detection

The final phase of our framework performs stance detection. Stance detection is a form of opinion mining which aims to infer a subject's attitude or commitment toward an object [18, 24]: e.g., whether an ISP favors or practices net neutrality.

We initially attempted to infer an ISP's posture toward a dimension of social responsibility using sentiment analysis. Sentiment analysis is another form of opinion mining which aims to infer whether a subject's opinion of an object is positive or negative. However, we found that sentences discussing ISPs' posture toward dimensions of socially responsibility rarely contained adjective-noun phrases that are the hallmark of sentiment analysis.

Our framework uses word embeddings to create a vector representation of each sentence. Word embeddings are vectors that capture the semantic similarity of words [42]. Words that are semantically similar—e.g., "Internet" and "network"—have similar vectors, whereas words with opposing meanings—e.g., "equal" and "prioritized"—have vectors that are quite different. We treat each sentence as a bag-of-words, and sums the word embeddings of all (known) words to compute a sentence's vector representation.

A sentence's vector is used as input to a support vector machine (SVM) based classifier. The classifier categories a sentence as indicating an ISP (a) *favors or practices*, (b) *opposes or ignores*, or (c) *has a mixed or unknown posture* toward the target dimension of social responsibility. The classifier is trained using a set of manually labelled sentences. Since different dimensions of social responsibility (Section 2) may be completely unrelated—e.g., environmental stewardship versus censorship—we use a separate classifier for each dimension of social responsibility.

We compute a "social responsibility score" for each ISP for each dimension of social responsibility. The score is simply the fraction of sentences that are categorized as indicating an ISP favors or practices a target dimension of social responsibility (out of the total number of sentences produced by our framework's phrase extractor (Section 3.2) for the target ISP and dimension).

4 IMPLEMENTATION

We have implemented our framework in Python. Our implementation is publicly available [6].

Data gathering (Section 3.1). We issue Google search queries for each combination of ISP name and keyword, requiring the terms to occur within 50 words of each other in the document: e.g., `"Internet2" AROUND(50) "net neutrality"`. We use Beautiful Soup [13] to extract URLs from the first 50 search results. We then fetch the webpages and extract the contents of the first `article` tag [16], if present, or all paragraph (`p`) tags. We automatically exclude pages whose title begins with the same five words as the title of an already fetched page, because these are mostly syndicated articles. We also exclude pages whose URL contains "tag" or "category", because these are mostly "index" pages.

Phrase extraction (Section 3.2). We use spaCy [7] with the `en_core_web_md` model to break each document into sentences and perform part-of-speech tagging. We extract all sentences with a nominal subject (`nsubj`) containing (an abbreviation of) an ISP's name and a direct object (`dobj`) containing one of the keywords.

Stance detection (Section 3.3). We use the word embeddings in spaCy's `en_core_web_md` model to compute a vector for each sentence. We use scikit-learn [5] to train and test a SVM-based classifier using 10-fold cross validation.

5 EVALUATION

We use our framework to assess 13 ISPs'—including Tier-1, Tier-2, and Tier-3 ISPs—posture on net neutrality and evaluate the feasibility and effectiveness of our approach.

5.1 Document gathering

We first assess whether we can gather a sufficient quantity and variety of documents that discuss ISPs' posture on net neutrality. We gather documents using the names of the 13 ISPs-and four keywords—"net neutrality," "open Internet," "zero-rating," and "paid prioritization." Our final dataset is publicly available [6].

The *Search results: Total* column in Table 1 shows the aggregate number of results returned by our search queries for each ISP. Recall that we conduct a separate search for each of the four keywords, and we limit each search to 50 results (Section 3.1). If the same URL appears in the results for multiple keywords, we only count the URL once. Unsurprisingly, we generally obtain more results for large Tier-1 ISPs and fewer results for small Tier-3 ISPs. Additionally, we

ISP		Total		Documents with			Sentences with stance		
Name	Tier	Documents gathered	Sentences extracted	0 sentences extracted	1 sentence extracted	2+ sentences extracted	Favors	Opposes	Unknown
AT&T	1	126	179	59	19	48	40	94	45
British Telecom	2	47	3	44	3	0	0	1	2
Cogent Comm.	2	51	18	38	10	3	13	0	5
Deutsche Telekom	1	93	26	71	18	4	3	17	6
FirstLight Fiber	3	5	0	5	0	0	N/A	N/A	N/A
Hurricane Electric	2	21	0	21	0	0	N/A	N/A	N/A
Internet2	3	48	3	46	1	1	1	0	2
KDDI	2	34	0	34	0	0	N/A	N/A	N/A
NYSERNet	3	5	0	5	0	0	N/A	N/A	N/A
Sprint	1	106	44	78	17	11	12	12	20
Telefonica	2	59	2	57	2	0	1	1	0
Telia Carrier	1	8	0	8	0	0	N/A	N/A	N/A
Zayo	1	20	1	19	1	0	0	0	1
TOTALS		**623**	**276**	**485**	**71**	**67**	**70**	**125**	**81**

Table 2: Phrase extraction

obtain more results for ISPs that provide a larger range of services—e.g., we obtain 192 results for AT&T, a Tier-1 ISP that provides residential access, mobile access, and transit—and fewer results for ISPs that provide a single type of service—e.g., we obtain 19 results for Telia Carrier, a Tier-1 ISP that primarily provides transit.

The *Total w/o nearby* column in Table 1 shows the aggregate number of search results when we do not require the ISP name and keyword to appear near each other in the document. This search strategy almost always produces more results (compared to *Search results: Total*), yet the number of phrases we extract from both sets of results is nearly the same (not shown). This indicates that limiting our search results based on the proximity of the search terms (Section 3.1) is useful for eliminating irrelevant documents.

The *Excluded* columns in Table 1 show the number of search results that are discarded for various reasons:

- *HTTP error*: an error occurred when trying to fetch the document—likely due to changes that occurred since the document was indexed by the search engine.
- *Not HTML*: the document is a PDF or other type of file—our prototype currently only handles HTML pages.
- *Index page*: the document is a listing of (excerpts from) other documents; we manually confirmed our URL-based heuristic (Section 4) correctly identified all index pages.
- *Duplicate*: the document is the same as another document; we manually confirmed our title-based heuristic (Section 4) correctly identified all duplicates.

The *Final: Total* column in Table 1 shows the number of documents we collect for each ISP. We collect 623 documents in total. We observe that only a few of these documents are located on the ISPs' websites (*Final: From ISP*), indicating that our quantification of social responsibility is primarily based on information from third-parties. These third-party sources include news organizations, government agencies, watchdog organizations, and personal blogs.

5.2 Phrase extraction

We next examine the sentences selected by our phrase extractor.

Our framework extracted a total of 276 sentences from within the 623 documents we gathered. The *Total: Sentences extracted* column in Table 2 shows the number of sentences we extracted for each ISP. The three ISPs for whom we gathered the most documents—126 for AT&T, 106 for Sprint, and 93 for Deutsche Telekom—were also the three ISPs for whom we extracted the most sentences—179, 44, and 26, respectively. Interestingly, we did not extract any phrases for FirstLight Fiber, Hurricane Electric, KDDI, NYSERNet, or Telia Carrier, despite having documents for these ISPs.

A similarly concerning observation is the large number of documents from which we did not extract any sentences (*Documents with 0 sentences extracted* column in Table 2). With the exception of AT&T, sentences were extracted from fewer than one-third of the documents we gathered for each ISP, and only 22% of documents overall. One possibility is that the documents we are gathering are not the right documents. Another possibility is that the our phrase extraction process is too simple/narrow. In the future, we plan to investigate and address both of these issues.

		Predicted: Favors	Predicted: Opposes
Actual	Favors	12	55
	Opposes	0	123

Table 3: Stance detection confusion matrix

5.3 Stance detection

We now assess the accuracy of our stance detection.

We manually labeled the 276 sentences output by our phrase extractor, indicating whether the sentence implies an ISP favors/practices or opposes/ignores net neutrality. Some sentences do not provide a clear indication of an ISP's stance on net neutrality or express a mixed stance, so we label these as "unknown." The *Sentences with stance* columns in Table 2 shows the number of sentences with each label for each ISP. In total, 25% of the sentences are labeled favors, 45% are labeled opposes, and the remainder are labeled unknown.

We train an SVM classifier using scikit-learn's default parameters. We perform 10-fold cross-validation to train and test our classifier. We exclude all sentences with stance "unknown" from our training and testing sets, giving us a total of 195 sentences to use for training and testing. The overall accuracy of our classifier is 71%.

The confusion matrix is shown in Table 3. Note that the number of sentences only adds to 190, as opposed to 195, because cross-validation requires a dataset whose size is a multiple of the number of folds (10). We observe that all inaccuracies stem from the classifier predicting a sentence implies an ISP opposes net neutrality when the sentence actually implies the ISP favors net neutrality. In other words, our classifier has 100% accuracy for sentences that imply an ISP opposes net neutrality, and 20% accuracy for sentences that imply an ISP favors net neutrality. We hypothesize this arises from the skew in our data: approximately two-thirds of the sentences (excluding those with stance unknown) are labeled opposed. Addressing this skew by expanding the our dataset and employing more sophisticated machine learning approaches are important areas of future work.

To help put the results in perspective: a simple majority classifier has an accuracy of 64%, whereas our classifier has an accuracy of 71%. Furthermore, even if only 20% of the sentences indicating an ISP favors net neutrality were properly labeled by our classifier, we would still draw the same conclusion for three of the four ISPs for which we have a non-negiligle number of sentences: i.e., AT&T and Deutsche Telekom would still be classified as opposing net neutrality and Cogent Communications would still be classified as supporting net neutrality; only our classification for Sprint, which is evenly balanced in its support versus opposition to net neutrality, would be impacted by our classifier's errors.

5.4 Robustness to false information

The above observation also illustrates our approach's robustness to false information. For AT&T, Cogent Communications, and Deutsche Telekom, there is a clear majority stance, with 70%, 100%, and 85%, respectively, of the sentences falling within the majority stance. Furthermore, the sentences are obtained from 67, 13, and 22, respectively, unique sources, so a single source is responsible for at most 8% of an ISP's predicted stance. Consequently, even the inclusion of a few sources of false information would not significantly change the social responsibility scores we compute.

5.5 Feasibility of divestment

Finally, we explore the feasibility of indirectly divesting in ISPs based on the socially responsibility scores we compute and Internet path data. We use real Internet path data from the Route Views Project [9], which was gathered by the route-views2.oregon-ix.net collector from April 6, 2020, 12:00 AM to 2:00 AM.

We example a small sample of paths with one of five popular destinations—netflix.com, google.com, microsoft.com, facebook.com, or doubleclick.com—which we select (based on familiarity) from the top 20 domains in the Cisco Umbrella top one million [41]. We use the BGPView API [1] to translate autonomous system (AS) numbers in the path to ISP names. We analyze eight unique paths.

We declare an ISP to be net neutral if an ISP's social responsibility score was above 75%—i.e., 75% or more of the sentences extracted from documents indicate the ISP favors/practices net neutrality. We declare a path to be socially responsible if every AS along it was net neutral. Using these criteria, we observe that 87.5% of the eight unique paths we considered were socially responsible. This implies there are a sufficient number of ISPs in the Internet that practice net neutrality to provide enough socially responsible paths.

6 RELATED WORK

A similar task is considered by Fleischmann et al. [35], wherein the authors use NLP techniques to understand what values humans express in writing. The written content in question is testimony regarding net neutrality, making the coded values (such as freedom, innovation, social order, and wealth) particularly relevant to our evaluation of social responsibility within text.

ChoiceNet [60] proposes an explicit relationship between economics and network services in order to increase competition and improve service quality. Similar to our vision, ChoiceNet makes the flow of money deliberate by allowing users to select from different service offerings, which encourages internal change among ISPs. Similarly, the alternative routing schema presented by Levin et al. [45] attempts to give users the ability to boycott networks which filter or censor. We seek to send a similarly forceful financial message while maintaining a connection to a larger social commentary.

Our vision of quantifying an ISPs' social responsibility parallels Turker's goal of measuring [58] Corporate Social Responsibility (CSR). However, rather than relying primarily on internal management surveys to evaluate social responsibility, we take an approach more similar to Abbott and Monsen [17] who explore content analysis of corporate publications. The authors convey concern for the lack of social data on a large portion of existing companies, but the plethora of information in the form of corporate publications, news articles, and blogs available via the Internet allows us to more thoroughly explore this avenue.

7 CONCLUSION AND FUTURE WORK

Assessing whether an ISP is socially responsible is a necessary, but currently difficult task, for individuals and institutions who want to (in)directly divest in ISPs to encourage social change. However, our framework for automatically computing a "social responsibility score" for many ISPs for various dimensions of social responsibility addresses this hurdle. Our evaluation demonstrates that our framework can accurately predict an ISP's stance on a target dimension of socially responsibility with 71% accuracy.

In the future, we plan to develop improved data gathering techniques, explore more nuanced phrase extraction techniques, and assess our framework's effectiveness on other dimensions of social responsibility. Additionally, we plan to develop a concrete mechanism to integrate socially responsibility scores into routing decisions. Our hope is that automated analysis of ISPs' socially responsibility will lead to a more responsible Internet.

REFERENCES

[1] [n.d.]. BGPView API Powered by Security Trails. https://bgpview.io.
[2] [n.d.]. MANRS - Mutually Agreed Norms for Routing Security. https://manrs.org.
[3] [n.d.]. OpenNet Initiative Country Profiles. https://opennet.net/country-profiles.
[4] [n.d.]. Regional Internet Registries Statistics. https://www-public.imtbs-tsp.eu/~maigron/RIR_Stats/RIR_Delegations/World/ASN-ByNb.html.
[5] [n.d.]. scikit-learn: Machine Learning in Python. https://scikit-learn.org.

[6] [n.d.]. Socially responsible routing. https://github.com/colgate-cs-research/responsible.

[7] [n.d.]. spaCy: Industrial-Strength Natural Language Processing in Python. https://spacy.io.

[8] [n.d.]. spaCy Usage Documentation: Linguistic Features. https://spacy.io/usage/linguistic-features. Access: 2021-05-31.

[9] [n.d.]. University of Oregon Route Views Project. http://www.routeviews.org/routeviews.

[10] 2008. *SMART 2020: Enabling the low carbon economy in the information age.* Technical Report. The Global eSustainability Initiative.

[11] 2014. Internet Society Global Internet Report 2014: Open and Sustainable Access for All. https://www.internetsociety.org/wp-content/uploads/2017/08/Global_Internet_Report_2014_0.pdf.

[12] 2017. 2017 Internet Society Global Internet Report: Paths to Our Digital Future. https://future.internetsociety.org/2017.

[13] 2020. Beautiful Soup Documentation. https://www.crummy.com/software/BeautifulSoup/bs4/doc.

[14] 2020. Freedom on the Net 2020: The Pandemic's Digital Shadow. *Freedom House* (23 September 2020). https://freedomhouse.org/report/freedom-net

[15] 2020. Global Internet Service Providers Industry - Market Research Report. https://www.ibisworld.com/global/market-research-reports/global-internet-service-providers-industry/. Accessed: 2021-05-21.

[16] 2021. Article - Schema.org Type. https://schema.org/Article.

[17] Walter F. Abbott and R. Joseph Monsen. 1979. On the Measurement of Corporate Social Responsibility: Self-Reported Disclosures as a Method of Measuring Corporate Social Involvement. *The Academy of Management Journal* 22, 3 (1979), 501–515.

[18] Abeer ALDayel and Walid Magdy. 2021. Stance detection on social media: State of the art and trends. *Information Processing & Management* 58, 4 (2021).

[19] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Tappan Morris. 2001. Resilient Overlay Networks. In *Proceedings of the 18th ACM Symposium on Operating System Principles (SOSP)*. 131–145.

[20] Tali Arbel. 2017. After net neutrality: Brace for internet 'fast lanes'. *The Detroit News* (20 December 2017). https://www.detroitnews.com/story/tech/2017/12/20/net-neutrality-promises/108785646 Accessed: 2021-07-01.

[21] Tali Arbel. 2017. After net neutrality: Brace for internet 'fast lanes'. *The Columbus dispatch* (20 December 2017). https://www.dispatch.com/zz/shareable/20171220/after-net-neutrality-brace-for-internet-fast-lanes Accessed: 2021-07-01.

[22] Tali Arbel. 2017. After net neutrality: Brace for internet 'fast lanes' and 'slow lanes'. *Milwaukee Journal Sentinel* (22 December 2017). https://www.jsonline.com/story/money/business/2017/12/22/after-net-neutrality-brace-internet-fast-lanes-and-slow-lanes/976693001 Accessed: 2021-07-01.

[23] Samantha Bates, Christopher Bavitz, and Kira Hessekiel. 2017. *Zero Rating & Internet Adoption: The Role of Telcos, ISPs & Technology Companies in Expanding Global Internet Access: Workshop Paper & Research Agenda.* Technical Report 2017-9. Berkman Klein Center for Internet & Society.

[24] Douglas Biber and Edward Finegan. 1988. Adverbial stance types in English. *Discourse Processes* 11, 1 (1988), 1–34.

[25] Raffaele Bolla, Roberto Bruschi, Alessandro Carrega, Franco Davoli, Diego Suino, Constantinos Vassilakis, and Anastasios Zafeiropoulos. 2012. Cutting the energy bills of Internet Service Providers and telecoms through power management: An impact analysis. *Computer Networks* 56, 10 (2012), 2320 – 2342.

[26] Edoardo Bonetto, Marco Mellia, and Michela Meo. 2012. Energy profiling of ISP points of presence. In *Proceedings of IEEE International Conference on Communications (ICC)*.

[27] Vinton Cerf. 1993. How the Internet Came to Be. *The Online User's Encyclopedia* (1993).

[28] Hsing Cheng, Subhajyoti Bandyopadhyay, and Hong Guo. 2011. The Debate on Net Neutrality: A Policy Perspective. *Information Systems Research* 22 (03 2011), 60–82.

[29] Luca Chiaraviglio, Marco Mellia, and Fabio Neri. 2012. Minimizing ISP Network Energy Cost: Formulation and Solutions. *IEEE/ACM Trans. Netw.* 20, 2 (2012), 463–476.

[30] Jay Pil Choi and Byung-Cheol Kim. 2010. Net Neutrality and Investment Incentives. *RAND Journal of Economics* 41, 3 (2010), 446–471.

[31] Cogent Communications. 2021. Net Neutrality. https://www.cogentco.com/en/net-neutrality.

[32] Marcel Dischinger, Massimiliano Marcon, Saikat Guha, P. Krishna Gummadi, Ratul Mahajan, and Stefan Saroiu. 2010. Glasnost: Enabling End Users to Detect Traffic Differentiation. In *Proceedings of the 7th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 405–418.

[33] Lauren Feiner. 2021. AT&T will start counting HBO Max against data limits, blames California net neutrality law. *CNBC* (2021). https://www.cnbc.com/2021/03/17/att-will-count-hbo-max-toward-data-caps-blames-net-neutrality-law.html Accessed: 2021-05-31.

[34] Arturo Filastò and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference. In *2nd USENIX Workshop on Free and Open Communications on the Internet FOCI*.

[35] Kenneth R. Fleischmann, Yasuhiro Takayama, An-Shou Cheng, Yoichi Tomiura, Douglas W. Oard, and Emi Ishita. 2015. Thematic Analysis of Words that Invoke Values in the Net Neutrality Debate. In *iConference 2015 Proceedings*.

[36] Hassan Habibi Gharakheili, Arun Vishwanath, and Vijay Sivaraman. 2016. Perspectives on Net Neutrality and Internet Fast-Lanes. *Comput. Commun. Rev.* 46, 1 (2016), 64–69.

[37] Brighten Godfrey, Igor Ganichev, Scott Shenker, and Ion Stoica. 2009. Pathlet routing. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. 111–122.

[38] Robert W. Hahn and Scott Wallsten. 2006. The Economics of Net Neutrality. http://dx.doi.org/10.2139/ssrn.943757 AEI-Brookings Joint Center Working Paper No. RP06-13.

[39] Joshua Hardwick. 2020. Google Search Operators: The Complete List (42 Advanced Operators). https://ahrefs.com/blog/google-advanced-search-operators.

[40] Tyler Hersko. 2021. AT&T Ignores Net Neutrality: HBO Max Won't Hit Data Caps but Competing Streamers Will. *IndieWire* (March 2021). https://www.indiewire.com/2020/06/att-net-neutrality-hbo-max-no-data-caps-1202235538 Accessed: 2021-05-31.

[41] Dan Hubbard. 2020. Cisco Umbrella 1 Million. https://umbrella.cisco.com/blog/cisco-umbrella-1-million. Accessed: 2020-02-16.

[42] D. Jurafsky and J.H. Martin. 2009. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition.* Pearson Prentice Hall.

[43] Partha Kanuparthy and Constantine Dovrolis. 2010. DiffProbe: Detecting ISP Service Discrimination. In *29th IEEE International Conference on Computer Communications (INFOCOMM)*. 1649–1657.

[44] Matt Lepinski and Kotikalapudi Sriram. 2017. BGPsec Protocol Specification. RFC 8205. https://rfc-editor.org/rfc/rfc8205.txt

[45] Dave Levin, Adam Bender, Cristian Lumezanu, Neil Spring, and Bobby Bhattacharjee. 2007. Boycotting and Extorting Nodes in an Internetwork. In *Workshop on the Economics of Networked Systems and Incentive-Based Computing (NetEcon+IBC)*.

[46] Matthew J. Luckie, Robert Beverly, Ryan Koga, Ken Keys, Joshua A. Kroll, and kc claffy. 2019. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 465–480.

[47] Richard T. B. Ma, Jingjing Wang, and Dah Ming Chiu. 2017. Paid Prioritization and Its Impact on Net Neutrality. *IEEE J. Sel. Areas Commun.* 35, 2 (2017), 367–379.

[48] Ratul Mahajan, David Wetherall, and Thomas E. Anderson. 2005. Negotiation-Based Routing Between Neighboring ISPs. In *2nd Symposium on Networked Systems Design and Implementation (NSDI)*.

[49] Chris McDonald. 2020. Internet Ascendant, Part 2: Going Private and Going Public. (22 October 2020). https://technicshistory.com/2020/10/22/internet-ascendant-part-2-going-private-and-going-public Accessed: 2021-05-31.

[50] Ben Munson. 2021. AT&T says net neutrality laws mean it has to charge for video streaming. *Fierce Video* (17 March 2021). https://www.fiercevideo.com/regulatory/at-t-says-net-neutrality-laws-mean-it-has-to-charge-for-video-streaming Accessed: 2021-05-26.

[51] Anthony J. Pennings. 2020. US Internet Policy, Part 1: The Rise OF ISPs. (15 March 2020). http://apennings.com/telecom-policy/us-internet-policy-part-1-the-rise-of-isps Accessed: 2021-05-31.

[52] Adrian Perrig, Pawel Szalachowski, Raphael M. Reischuk, and Laurent Chuat. 2017. *SCION: A Secure Internet Architecture.* Springer.

[53] Barath Raghavan and Alex C. Snoeren. 2004. A system for authenticated policy-compliant routing. In *Proceedings of the ACM SIGCOMM 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*. ACM.

[54] Stefan Savage, Thomas E. Anderson, Amit Aggarwal, David Becker, Neal Cardwell, Andy Collins, Eric Hoffman, John Snell, Amin Vahdat, Geoffrey M. Voelker, and John Zahorjan. 1999. Detour: informed Internet routing and transport. *IEEE Micro* 19, 1 (1999), 50–59.

[55] Gay W. Seidman. 2015. Divestment Dynamics: Mobilizing, Shaming, and Changing the Rules. *Social Research: An International Quarterly* 82, 4 (2015), 1015–1037.

[56] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. 2020. Censored Planet: An Internet-Wide, Longitudinal Censorship Observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS)*.

[57] Muhammad Mukarram Bin Tariq, Murtaza Motiwala, Nick Feamster, and Mostafa H. Ammar. 2009. Detecting network neutrality violations with causal inference. In *Proceedings of the 2009 ACM Conference on Emerging Networking Experiments and Technology (CoNEXT)*. 289–300.

[58] Duygu Turker. 2009. Measuring Corporate Responsibility: A Scale Development Study. *Journal of Business Ethics* 85, 4 (2009), 411–427.

[59] Vytautas Valancius, Cristian Lumezanu, Nick Feamster, Ramesh Johari, and Vijay V. Vazirani. 2011. How many tiers?: Pricing in the internet transit market. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*. 194–205.

[60] Tilman Wolf, James Griffioen, Kenneth L. Calvert, Rudra Dutta, George N. Rouskas, Ilya Baldin, and Anna Nagurney. 2014. ChoiceNet: toward an economy plane for the internet. *Computer Communication Review* 44, 3 (2014), 58–65.
[61] Tim Wu. 2003. Network neutrality, broadband discrimination. *J. of Telecommunications & High Technology Law* 2 (2003), 141.
[62] Wen Xu and Jennifer Rexford. 2006. MIRO: Multi-path interdomain routing. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*. 171–182.
[63] Hongming Zhang, Xinran Zhao, and Yangqiu Song. 2020. A Brief Survey and Comparative Study of Recent Development of Pronoun Coreference Resolution. arXiv:2009.12721
[64] Xinyi Zhou and Reza Zafarani. 2020. A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. *ACM Comput. Surv.* 53, 5, Article 109 (Sept. 2020).