“Is my Internet down?”: Sifting through User-Affecting Outages with Google Trends

Ege Cem Kirci
ETH Zürich

Martin Vahlensieck
ETH Zürich

Laurent Vanbever
ETH Zürich

ABSTRACT

What are the worst outages for Internet users? How long do they last, and how wide are they? Such questions are hard to answer via traditional outage detection and analysis techniques, as they conventionally rely on network-level signals and do not necessarily represent *users' perceptions of connectivity*.

We present SIFT, a detection and analysis tool for capturing user-affecting Internet outages. SIFT leverages users' aggregated web search activity to detect outages. Specifically, SIFT starts by building a timeline of users' interests in outage-related search queries. It then analyzes this timeline looking for spikes of user interest. Finally, SIFT characterizes these spikes in duration, geographical extent, and simultaneously trending search terms which may help understand root causes, such as power outages or associated ISPs.

We use SIFT to collect more than 49 000 Internet outages in the United States over the last two years. Among others, SIFT reveals that user-affecting outages: *(i)* do not happen uniformly: half of them originate from 10 states only; *(ii)* can affect users for a long time: 10% of them last at least 3 hours; and *(iii)* can have a broad impact: 11% of them simultaneously affect at least 10 distinct states. SIFT annotations also reveal a perhaps overlooked fact: outages are often caused by climate and/or power-related issues.

CCS CONCEPTS

• **Networks → Network reliability**.

KEYWORDS

Internet outages, Anomaly detection, Data mining, Google Trends

ACM Reference Format:
Ege Cem Kirci, Martin Vahlensieck, and Laurent Vanbever. 2022. “Is my Internet down?”: Sifting through User-Affecting Outages with Google Trends. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22), October 25–27, 2022, Nice, France.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3517745.3561428

1 INTRODUCTION

What were the most user-affecting outages over the last two years? This question is tough to answer with traditional techniques for at least three reasons: *(i)* it involves measuring the impact on users, *(ii)* it inherently depends on the users' location, and *(iii)* it requires comparing independent outages. Eventually, this question is hard to answer because it concerns the users and not a particular service.

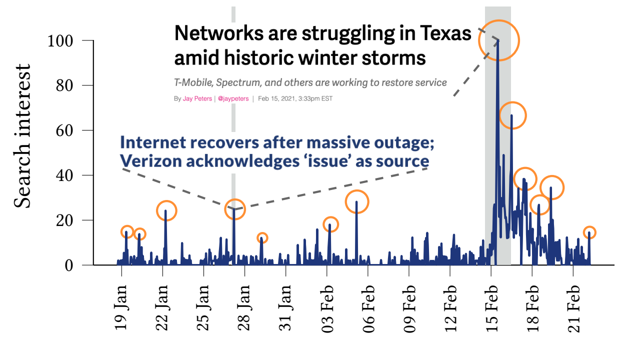

Figure 1: The <Internet outage> popularity index in Texas. The news verify circled spikes such as Verizon [44] and winter storm outages [40].

Can we formalize and extract a *user indicator* to study Internet outages from users' perspectives? This *users-as-sensors* approach can potentially enable observability of a broad range of failures along with extensive and ubiquitous visibility into diverse parts of the Internet. We argue that tracking aggregated web search activity over outage-related search queries can estimate actual Internet outages, similar to past work on web search activity in health [15, 18, 25, 38, 41] and economics [7, 8, 10, 19].

We present SIFT – a detection and analysis tool for capturing and analyzing *users' perceptions* of Internet outages. SIFT discovers outages by identifying spikes in aggregated user interest around outage-related search queries. Initially, SIFT constructs a continuous and normalized timeline for a selected search term (e.g., <Internet outage>) by stitching partial and limited responses crawled from the search aggregation service. Then, SIFT detects spikes – trends in user interest – and identifies these spikes' start, peak, and end time. Finally, SIFT analyzes the spike's magnitude, duration, and timing from the extracted spike metadata. SIFT also attaches context attributes to each spike by looking for simultaneously trending search terms coinciding with the spike's time and location. These search terms range from causal relations such as <Power outage> or <Thunderstorm>; to company names of network and application providers such as <Verizon>, <T-mobile>, and <Youtube>.

Fig. 1 demonstrates SIFT in action. The timeline represents the user interest for aggregated search queries semantically clustered into the <Internet outage> search term in the scope of Texas. The x-axis illustrates a cut view of our two-year-long timeline in the winter of 2021. The y-axis indicates the normalized popularity index of the search term, meaning its proportion among all the searches for the given time and geographical area. The orange circles mark

IMC '22, October 25–27, 2022, Nice, France

ACM ISBN 978-1-4503-9259-4/22/10…$15.00
https://doi.org/10.1145/3517745.3561428

the spikes for which we could find a corresponding news article. The figure includes two such news captions – a Verizon outage and a general power outage – with gray stripes indicating the outage duration that SIFT infers. SIFT analyzes the detected spikes in their *(i)* impact, *(ii)* area, and *(iii)* context. First, SIFT reveals that the impact and the duration of the power outage is more significant than the Verizon one. The news validate this finding: the winter storm power outage impacted various network providers for multiple days, whereas the Verizon outage was specific to an ISP and lasted for several hours. Second, SIFT analyzes the outage area by matching concurrent spikes from distinct states. For example, although the winter storm power outage is significant in Texas, SIFT detects that the Verizon outage exhibits spikes in 27 states concurrent in time, validating the news [42]. Third, SIFT analyzes the simultaneously trending search terms. For instance, the aggregation engine suggests the Verizon search term for that spike, whereas multiple ISP names for the winter storm. The winter storm spike also returns the <Power outage> and <February 13-17, 2021 North American winter storm> search terms as trending, enabling profound insights into the underlying causes of the outage.

We use SIFT to conduct a two-year study in the United States. Our preliminary results suggest that SIFT complements existing techniques and distinguishes itself with its visibility and observability capabilities. First, the *users-as-sensors* approach enables visibility into otherwise hard-to-observe domains. For example, SIFT shows that application-level issues such as DNS misconfigurations or backend/ mobile network problems that are invisible/ unreachable to ping probing reveal themselves to users. Second, SIFT unlocks unique observability by formalizing the extraction of a publicly available, semantically rich, and extensive data set. For instance, without SIFT, insights such as the relationship between natural disasters and the most impactful outages or the ubiquity of power-related Internet outages would have required dedicated studies a-priori targeted at exploring these relationships.

What were the most user-affecting outages? Contrasting the spikes SIFT identified in 2020 and 2021 uncovers a key theme in user-affecting outages. We discover spike occurrences between two years are similar (i.e., 25,494 for 2020 and 23,695 for 2021), with similar distributions at the state and monthly levels. However, the long-lasting spikes (i.e., those with *at least* five hours in user interest) are 50% greater in 2020 than in 2021. The outliers are clear: California in August and September 2020 and Texas in January and February 2021. As we initially discussed for Texas in Fig. 1, SIFT's context analysis reveals why: power outages resulting from wildfires and winter storms. Our analysis shows climate disasters have dictated the outliers for the last two years.

2 BACKGROUND: GOOGLE TRENDS

SIFT uses Google Trends (GT) [51], an aggregation service, to analyze Google users' search query statistics. We choose Google based on its overwhelming dominance in the online search market in the United States [14]. Analytical use of GT requires a rigorous understanding of the underlying mechanisms. In the following, we explain the relevant details of what GT does, the data it produces, its object types, and its rising suggestions feature.

Google Trends. Google Trends is a service that analyzes the popularity of a search term over a selected time frame and geographical area. For a given time frame, the data points of a search term represents its proportion of all searches on all topics over that time frame and geographical area [47]. This normalization allows users to detect spikes in search interest by comparing a search term's popularity to itself over time. We define a *spike* as a sudden acceleration of interest in a search term compared to what it typically occupies among all searches.

Data points. Users can request GT data over different *time frames*, meaning the requested data's time range; and *time blocks*, which is the frequency of data points over the specified time frame. GT poses some limitations on the user: a single time frame can only extend to one week for the frequency of hourly time blocks. When a request arrives, GT draws an unbiased random sample of Google search data for the given time frame and geographical area. Next, to provide anonymity for the search engine users, GT rounds off tiny search volumes below a certain threshold to 0. Then, GT indexes the data points of a time frame on a scale from 1 to 100, where 100 represents the data point with the maximum search interest for the selected time frame [34].

Objects types. GT distinguishes between *search queries* and *search topics*, collectively referred to as *search terms*. Search queries represent the direct user inputs to the search engine. On the other hand, search topics are the more reliable abstractions that pull in and group a set of semantically similar phrases, misspellings, and acronyms, covering all languages. In our study, we keep track of the search topic <Internet outage> to leverage Google's semantic clustering and aggregation directly. GT keeps track of both object types with the same statistical procedures [34].

Rising suggestions. Beyond indexing search interests, GT also suggests related queries and topics for a given input term. In more detail, the *rising terms* represent the search terms that see the most significant increase in their search interests over the selected time frame and geographical area of the input term. GT assigns weights to these suggestions proportional to their percent increase in the selected time frame [34].

3 SYSTEM DESIGN

This section starts with an illustration of the SIFT workflow by following an example depicted in Fig. 2. Then, we delve into the technical details of each workflow component, namely, *(i)* processing, *(ii)* detection, and *(iii)* analysis.

3.1 Overview

The SIFT workflow begins with a user specifying a time range, geographical area, and search term(s) as an input (1). SIFT partitions the selected time range into consecutive and overlapping weekly time frames (2) to construct an hourly extended time series. We repeat this procedure for every state of the United States for the <Internet outage> search term covering the last two years.

As Fig. 2 depicts, GT responds to SIFT requests by weekly time frames (3) consisting of 168 data points (i.e., one per hour) and rising terms for that week (4). As illustrated by the value of the

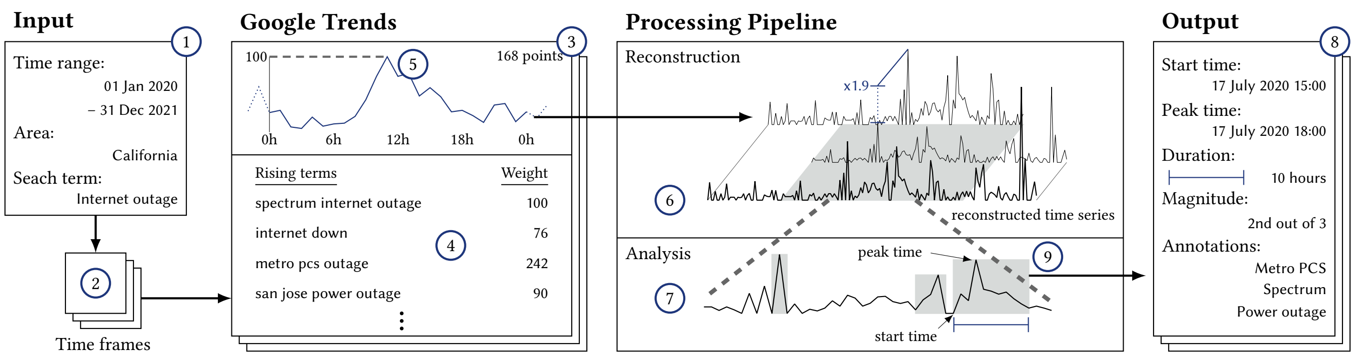

Figure 2: The SIFT workflow.

highest data point 5, GT normalizes each time frame in its own scope. In the depicted response, the <spectrum internet outage> and <san jose power outage> are among the rising search terms. SIFT repeats this process for daily time frames on spike days to capture more targeted and fine-grained rising terms for each spike.

The SIFT processing pipeline involves time series reconstruction 6 and spike analysis 7. The reconstruction part illustrates SIFT stitching and renormalizing two consecutive and overlapping time frames into a continuous time series. At each stitching step, SIFT extends the time series with the next time frame after normalizing it against the overlap. The analysis part illustrates SIFT’s peak detection and annotation procedure. In the depicted example, SIFT detects three spikes with varying magnitudes and durations.

Finally, SIFT reports each spike’s start, peak, and end times as well as their magnitude, duration, and context annotations for evaluations 8. In this running example the rightmost spike 9 lasted for 10 hours, was the second greatest in magnitude for California (i.e., for the figure), and was annotated with three search terms: <Spectrum>, <Metro PCS>, and <Power outage>, possibly indicating a power outage impacting multiple network providers.

3.2 Processing

The processing pipeline begins with the stitching of consecutive and overlapping time frames. We require this step due to two challenges introduced by GT: random sampling and piecewise normalization.

Random sampling. GT draws an unbiased random sample from the Google search database and normalizes the resulting data upon a request. This methodology results in random error between the returned sample proportion (i.e., the search interest index) and the undisclosed population proportion. The normal approximation dictates that the standard error of an unbiased sample proportion is inversely related to the sample size, which means the confidence interval improves with larger sample set. We mitigate the sampling error with an iterative method. First, we build a time series from a single set of time frames and detect the resulting spikes. Then, we repeat this procedure but instead take the average of two time frames to reduce the sampling error at each time frame position. We follow this procedure until the set of spikes we detect converge. This averaging process takes six rounds of re-fetches to conclude.

Piecewise normalization. GT scales each time frame relative to its highest data point. This piecewise indexing prevents SIFT from identifying spikes’ global magnitude. This restricted approach keeps us from ranking outages and eliminating local noise. SIFT reconstructs a continuous time series from piecewise time frames by initially fetching consecutive and overlapping time frames. Then, SIFT uses the intersecting regions to identify the scaling ratio between the consecutive time frames. Finally, SIFT rescales the right-adjacent time frame by this ratio and appends it sequentially to the preceding time series. This procedure enables SIFT to recalibrate piecewise time frames to build continuous global time series. SIFT then renormalizes the global time series and indexes the data points on a scale from 0 to 100.

3.3 Detection

After reconstructing a stitched and calibrated time series, SIFT detects spikes in the <Internet outage> search topic. Spike detection involves two challenges. First, search interests rise over multiple time blocks for a single spike. Therefore, SIFT has a mechanism to avoid recounting multiple successive peaks as multiple distinct spikes. Second, SIFT needs to identify each spike’s start and end times to measure their duration.

Spike. Existing spike detection techniques typically require knowing the underlying distribution of events, which is not the case with Internet outages. Therefore, SIFT characterizes spikes by leveraging a notion called *topographic prominence* [27]. The SIFT detection algorithm starts at the highest peak, then continues forward in time block by block until the current time block’s value is less than half of the value in the previous block (or zero). This point marks the ending of the spike. The start point is determined by stepping backward in time starting from the peak, either until the current block’s value is zero or the endpoint of another spike.

Duration. SIFT defines the duration of spikes as the time elapsed between their start and end times. This interval represents the duration of the user interest for the given spike. Since the magnitudes of spikes are not directly comparable between different states due to geographical normalization, in this work, we use duration as the *directional indicator* for the impact of the outages.

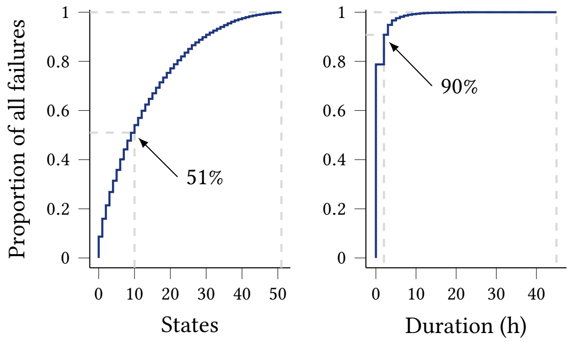

Figure 3: Characteristics of all spikes in 2020–2021.

3.4 Analysis

SIFT uses GT's rising suggestions feature to extract the most relevant search terms around peak times. We use a series of heuristics to cluster and rank suggested search terms most effectively.

Heavy-hitters. SIFT distinguishes interesting search terms from random correlations by superimposing all the suggestions from all the spikes and checking their frequency. Our data set reveals that only 33 of the 6655 search terms suggested comprise half of the overall suggestions. Thus, we prioritize the heavy-hitters in the annotation ranking. These heavy-hitters include but are not limited to: <Power outage>, <Xfinity>, <Spectrum>, <Comcast>, <AT&T>, <Cox Communications>, <Verizon>, and <Electric power>.

Ranking. For a given spike, SIFT applies several transformations on the rising suggestions set. First, it ranks the rising suggestions according to their weights (i.e., their percentage increase in the time frame and geographical area). Then, it prioritizes heavy-hitter terms over the others. Finally, SIFT applies a natural language processing library with pre-trained word vectors to cluster semantically similar phrases such as <is Verizon down> and <Verizon outage>.

4 EVALUATION

This section demonstrates SIFT's novel insights over three dimensions: *(i)* impact, *(ii)* area, and *(iii)* context. First, SIFT provides two unique directional indicators for quantifying the impact on the users: The spike magnitude (i.e., how aggressive the search interest is) and the spike duration (i.e., for how long the outage interests the users). We show that the impact of spikes varies greatly, and the most impactful ones are indeed newsworthy. Second, we introduce the concept of area, meaning the region where spikes co-occur. This indicator illustrates the extent and geography of the outages from the user's perspective. Third, we present the context where semantic integration of other user searches brings unique insights.

Implementation. We implement SIFT in Python. The system consists of a backend database, a data extraction and collection module, a processing pipeline, and a running web interface to display the requested data to the SIFT user. As the data collection module's primary bottleneck is GT's IP-based rate-limiting, the collection module first maps the queued workload into fetcher units hosted behind separate IP addresses. The collection module then merges the responses gathered from the fetchers into a unified database. SIFT requested 160 238 time frames from GT throughout our study.

Table 1: The most impactful spikes based on their durations.

Spike time	State	Duration (h)	Outage
15 Feb. 2021–10h	TX	45	Winter storm [35]
09 Nov. 2021–04h	CA	23	Xfinity [13]
08 Jun. 2021–09h	CA	22	Fastly [53]
26 Dec. 2020–12h	TN	21	AT&T [49]
29 Oct. 2020–09h	GA	20	Comcast [9]
15 Jun. 2020–14h	CA	19	T-Mobile [3]
13 Apr. 2020–11h	NC	18	CenturyLink [17]

The ANT outages data set. Throughout our evaluation, we use a state-of-the-art active probing data set (i.e., ANT outages data set [52]) to compare it against SIFT's findings where possible. The ANT data set includes eleven-minute time slots of active probing measurements collected from six distinct locations in the world [43]. It reports IP subnets, the start time of outages, and their durations based on the reachability of the probed end nodes in each subnet. We augment this data set with geolocations based on the Maxmind IP-geolocation data set [29].

4.1 Impact

User interest is quantifiable over two metrics: magnitude and duration of spikes. Since GT normalizes search interest over all queries in a selected geographical area, magnitude fits well with temporal comparisons on a fixed geography. However, duration is more stable for inter-state comparisons as it does not get scaled on a per-state basis. We use duration as the quantifying metric in the following analysis.

As illustrated in Fig. 3, SIFT spikes exhibit a skewed distribution in locations and durations. The left-side cumulative frequency plot demonstrates the distribution of the 49 189 spikes over the geographical states. The x-axis ranks the states in decreasing order. The y-axis illustrates that the top ten states, including California, Texas, Florida, and New York, host 51% of the total spikes. Due to state-level normalization, we can not simply explain this imbalance by state populations. In §4.3, we discuss how the contextual analysis reveal the root cause of the imbalance for the top states: California and Texas. The right-side cumulative frequency plot shows the distribution of the overall spikes in their duration. SIFT finds that 10% of spikes attract user interest for at least three hours.

Table 1 demonstrates the most impactful outages SIFT discovers regarding spike duration. The Texas winter storm power outage on 15 Feb. 2021 appears to be the most long-lasting outage in our data set. The <Power outage> annotation confirms the news. Another interesting data point is the California T-Mobile outage on 15 Jun. 2020. Although this highly impactful outage is observable by SIFT, we failed to trace the outage in the ANT data set. That could be due to mobile nodes not responding to probes and escaping the ANT's detection methodology. Lastly, Fig. 4 demonstrates the daily distribution of all spikes. The figure implies that the Internet sees fewer outages during weekends. We conjecture that this is probably due to less human error on the service side.

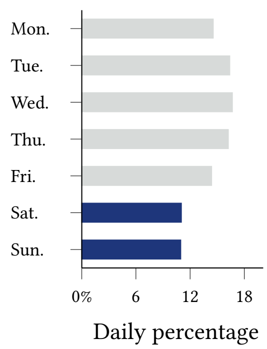

Figure 4: Daily distribution of all spikes.

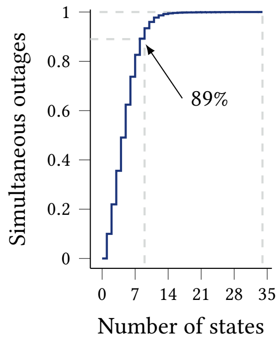

Figure 5: Distribution of spikes over outages.

Table 2: Most extensive spikes based on their geographical footprint.

Spike time	States	Outage
22 Jul. 2021–14h	34	Akamai [32]
17 Jul. 2020–19h	30	Cloudflare [21]
04 Oct. 2021–15h	29	Facebook [26]
26 Jan. 2021–16h	27	Verizon [42]
11 Nov. 2020–23h	27	Youtube [39]
15 Dec. 2021–14h	26	AWS [45]
08 Jun. 2021–09h	26	Fastly [53]
23 Jan. 2020–18h	25	Comcast [50]
30 Aug. 2020–09h	24	CenturyLink [20]

4.2 Area

One unique capability of SIFT is its immersive vantage points – Internet users. SIFT can analyze partial outages and detect their spatial and temporal patterns over a selected region. Fig. 5 demonstrates the number of distinct states simultaneously observing a spike. The figure illustrates that 11% of all the outages include 10 or more states, meaning there are spikes simultaneously occurring in 10 or more states for that particular time.

Table 2 illustrates the most extensive spikes SIFT identifies, representing the skewed data points in Fig. 5. The Akamai outage on 22 Jul. 2021 appears to be the most extensive outage in our data set, with spikes occurring in 34 states simultaneously. We fail to trace the Akamai and Youtube outages from Table 2 in the ANT data set. That is likely because the applications were unavailable but responsive to pings: the Akamai outage was a DNS misconfiguration [32], whereas the Youtube outage was a video buffering problem [39].

We also investigate why Table 2 does not include an outage with all the states, as news report that some of these outages are country-wide. We discover a substantial spike in all the states for the Facebook outage [26], but with certain lags for the remaining 22 states. We suspect the local time difference causes lagged spikes, especially with leisure applications like Facebook. We leave further analysis resulting from user behavior to future work.

4.3 Context

As SIFT runs on aggregated web search activity, it can draw insights from outside the Internet by linking search term correlations and similarities. The ninth most popular suggestion for <Internet outage> search term is <Power outage>. The dependency between power and Internet outages is well-known from past work studying Internet outages during natural disasters [28, 30, 37, 48]. However, to the best of our knowledge, no research exists on investigating how prevalent this dependency is among high-profile Internet outages. This question deserves further investigation since Internet availability relies on the power infrastructure.

Fig. 6 illustrates the monthly distribution of <Internet outage> spikes with <Power outage> annotation that lasts at least five hours. Note that the overall Internet outage spikes lasting at least five hours only comprise the top 3.5% of all the spikes. Power outages are responsible for 73% of the spikes that last at least five hours, showing how prevalent the power outages are among the top spikes. The state-level analysis of the monthly data quickly reveals two outliers: Texas in January and February 2021 and California in August and September 2020. We validate the individual spikes with news indicating the power outages on these days from winter storms and wildfires, respectively.

Table 3 demonstrates several high-profile power outages resulting in <Internet outage> spikes from various states discovered by SIFT. Our manual verifications of the news for the span of our study reveal a grim yet perhaps overlooked observation: the climate disasters seem to be a significant reason behind long-lasting Internet outages. We leave a more in-depth and comprehensive analysis of this particular finding for future work.

5 RELATED WORK

Internet outage detection and analysis is a broad research area. We suggest a recent survey by Aceto *et al.* that thoroughly navigates the field for interested readers [1]. Due to space considerations, we focus on the two approaches we know exist in the wild as professional services. We consider this property a sign of the solution’s generality and evidence of its practical soundness.

Complaint-based approach. Downdetector [12] is the most prominent example of complaint-based approaches. Downdetector monitors online service outages through a proprietary user-driven mechanism. The official resources suggest that the system analyzes user-entered complaints submitted on their website and social media profiles, thus automatically detecting problems based on unusual amounts of complaints [36].

One of SIFT’s contributions is to formalize the extraction of a publicly available, extensive data set that is search engine queries and their semantic aggregations. The Downdetector data set is unavailable by the time of this writing. Besides, Downdetector does not offer suggestions for the possible root causes of the failures.

Past studies investigated Twitter as a potential data source for passive crowdsourcing of Internet-related problems, such as popular service unavailabilities [33] and DDoS event forecasting [54]. Unlike SIFT, these systems do not offer geographical insights and a comparative approach between detected events. We consider other data sources, such as Twitter, complementary to our approach.

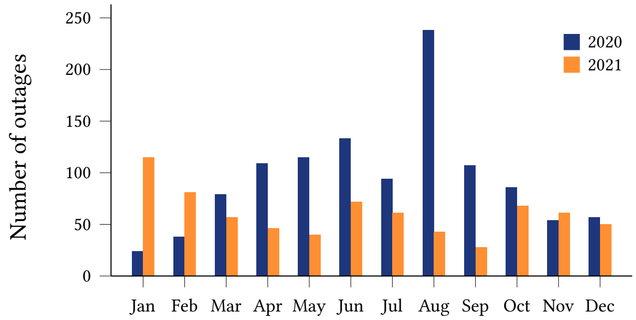

Figure 6: The number of power annotated spikes with a duration of at least 5 hours. 2020 California wildfires and 2021 Texas winter storms are the highlights.

Active probing approach. Active probing is a popular approach in both the industry [4, 31] and academia [43, 46, 48]. Active probing tests the reachability of a set of addresses by periodically contacting the corresponding end nodes. Ultimately, the methodology interprets unresponsiveness as a proxy for the livelihood of the respondent end nodes and, implicitly, the hosting network.

The responsiveness-as-livelihood assumption introduces several weaknesses as traffic may be absent for reasons other than an outage. First, a past study argues that only a tiny portion (i.e., 3.6%) of the possible IPv4 address space responds to pings [23]. Internet trends such as IPv6, mobile devices, network address translation, and firewalls suggest the ever-increasing difficulty of finding responsive end nodes. Second, active probing is inherently costly as it requires continuous traffic generation and analysis. Address scanning is also open to misinterpretation as anomalous behavior or abusive usage. Finally, the gathered data is hard to analyze as it gives little information about the problem's extent, impact, or root causes.

SIFT does not replace these methodologies but rather complements or augments them. For example, researchers do not have access to private user-complaint data. In this case, search engine query statistics is a viable proxy for extracting user's perception. As for the active probing approach, SIFT assesses the impact of the detected outages and potentially reveal their underlying causes.

6 CONCLUSION AND FUTURE WORK

This work presents SIFT – an outage detection and analysis tool for discovering user-affecting Internet outages by tracking users' aggregated web search activities. SIFT formalizes the extraction and analysis of user search queries by generating calibrated time series of user interest with statistical rigor and specifying directional indicators to compare detected outages quantitatively. Besides, SIFT takes advantage of this unique data source by decorating the detected events with simultaneously trending user interests such as associated service names or potential root causes.

This work is only the beginning of what SIFT has to offer. First, SIFT is a good fit for studying trends over more extended periods. What effect has the climate crisis had on the Internet over the past ten years – has the rise in wildfires impacted the Internet's reliability? Can tracking Internet security-related search terms uncover the trending attacks? Second, SIFT enables multidisciplinary connections thanks to its overarching data source. What patterns exist between the Internet and seemingly unrelated areas, such as economics or political events? We share SIFT's source code to help answer these and further questions.[1]

Although SIFT formalizes the data extraction and analysis steps, our preliminary study validates SIFT manually and mainly for newsworthy outages. The validation step is particularly challenging for SIFT since traditional approaches do not directly provide ground truth. SIFT detects what users sense, and to users, an outage is a far broader experience than network-level probes and monitors can capture, such as malfunctioning applications, unresponsive micro-services, and corrupted data. That is to say; we believe a more detailed validation study can unfold two promising research directions. First, SIFT can help researchers put traditional data sets in context. In the ANT outages data set, which outages are identified as impactful by SIFT, and what are their characteristics in network-level signals? Second, SIFT can help characterize which types of outages traditional data sets do not capture. What separates the outages SIFT detects but ANT outages data set does not possess? In future work, we will cross-validate the two approaches, namely the ANT and SIFT outages data sets, and characterize their respective advantages, disadvantages, and complementary features.

Table 3: Most impactful power outages for various states detected by SIFT.

Spike time	State	Duration (h)	Outage
15 Feb. 2021–10h	TX	45	Winter storm [16]
06 Sep. 2020–18h	CA	18	Heat wave [6]
11 Aug. 2021–09h	MI	15	Heavy rain and storm [11]
24 Oct. 2021–18h	WA	13	Storm [22]
22 Jul. 2021–14h	CO	9	Severed power line [24]
12 Aug. 2021–20h	OH	7	Storm [5]
11 Dec. 2021–23h	KY	7	Tornado [2]

ACKNOWLEDGMENTS

The authors would like to thank the reviewers and our shepherd, Romain Fontugne, for their insightful comments. The authors also thank Carsten Peters and Alex Schröder for their support in obtaining research permission with Google Trends, John Heidemann for sharing the ANT outage data set, and Stefano Vissicchio, Coralie Busse-Grawitz, and Muoi Tran for their valuable feedback.

[1]The SIFT repository is available at https://github.com/nsg-ethz/sift.

REFERENCES

[1] Giuseppe Aceto, Alessio Botta, Pietro Marchetta, Valerio Persico, and Antonio Pescapé. 2018. A Comprehensive Survey on Internet Outages. *Journal of Network and Computer Applications* 113 (2018), 36–63. https://doi.org/10.1016/j.jnca.2018.03.026

[2] Lucas Aulbach. 2021. Thousands Still without Power in Kentucky Following Devastating Tornado Outbreak. https://www.courier-journal.com/story/news/2021/12/11/kentucky-tornado-widespread-power-outages-amid-damage/6475197001/

[3] Benton Institute. 2020. June 15, 2020 T-Mobile Network Outage Report. https://www.benton.org/headlines/june-15-2020-t-mobile-network-outage-report

[4] Daniele Besana. 2017. ThousandEyes Review: Outage Detection. https://www.routerfreak.com/thousandeyes-review-outage-detection/

[5] Boggs. 2021. Several Schools Closed as Thousands Remain without Power. https://spectrumnews1.com/oh/columbus/news/2021/08/12/several-schools-closed-as-thousands-remain-without-power

[6] Matt Charnock. 2020. Charge Your Power Banks: Rotating Blackouts and Power Shutoffs Possible In Parts of Bay Area This Week. https://sfist.com/2020/09/06/charge-your-power-banks-rotating-bay-area-blackouts-possible-tuesday-and-wednesday/

[7] Hyunyoung Choi. 2009. *Predicting Initial Claims for Unemployment Benefits*. SSRN Scholarly Paper 1659307. Social Science Research Network, Rochester, NY. https://papers.ssrn.com/abstract=1659307

[8] Hyunyoung Choi and Hal Varian. 2012. Predicting the Present with Google Trends. *Economic Record* 88, s1 (2012), 2–9. https://doi.org/10.1111/j.1475-4932.2012.00809.x

[9] Crisis24. 2020. US: Tropical Storm Zeta Causes Disruptions in Georgia October 29 /Update 4. https://crisis24.garda.com/alerts/2020/10/us-tropical-storm-zeta-causes-disruptions-in-georgia-october-29-update-4

[10] Zhi Da, Joseph Engelberg, and Pengjie Gao. 2011. In Search of Attention. *The Journal of Finance* 66, 5 (Oct. 2011), 1461–1499. https://doi.org/10.1111/j.1540-6261.2011.01679.x

[11] Kyle Davidson. 2021. Storms Leave 600,000+ Michiganders without Power; Flash Flood Warning in Wayne County. https://www.freep.com/story/news/local/michigan/2021/08/11/storms-cause-power-outages-michigan/8100403002/

[12] Downdetector. 2022. About Us | We Detect When Technology Fails. https://downdetector.com/

[13] Clare Duffy. 2021. Comcast Xfinity Internet Outage Hits Customers across the US. https://www.cnn.com/2021/11/09/tech/comcast-xfinity-outage/index.html

[14] Jennifer Elias. 2020. Google 'overwhelmingly' Dominates Search Market, Antitrust Committee States. https://www.cnbc.com/2020/10/06/google-overwhelmingly-dominates-search-market-house-committee-finds.html

[15] Gunther Eysenbach. 2006. Infodemiology: Tracking Flu-Related Searches on the Web for Syndromic Surveillance. *AMIA Annual Symposium Proceedings* 2006 (2006), 244–248. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1839505/

[16] Mandi Cai Ferman, Erin Douglas and Mitchell. 2022. How Texas' Power Grid Failed in 2021 — and Who's Responsible for Preventing a Repeat. https://www.texastribune.org/2022/02/15/texas-power-grid-winter-storm-2021/

[17] Kevin Fogarty. 2020. Outages Spike in Late April as COVID-19 Trends Strain Internet. https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/outages-spike-in-late-april-as-covid-19-trends-strain-internet-58412923

[18] Jeremy Ginsberg, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, and Larry Brilliant. 2009. Detecting Influenza Epidemics Using Search Engine Query Data. *Nature* 457, 7232 (Feb. 2009), 1012–1014. https://doi.org/10.1038/nature07634

[19] Sharad Goel, Jake M. Hofman, Sébastien Lahaie, David M. Pennock, and Duncan J. Watts. 2010. Predicting Consumer Behavior with Web Search. *Proceedings of the National Academy of Sciences* 107, 41 (Oct. 2010), 17486–17490. https://doi.org/10.1073/pnas.1005962107

[20] Jazmin Goodwin. 2020. Major Internet Outage: Dozens of Websites and Apps Were Down. https://www.cnn.com/2020/08/30/tech/internet-outage-cloudflare/index.html

[21] John Graham-Cumming. 2020. Cloudflare Outage on July 17, 2020. http://blog.cloudflare.com/cloudflare-outage-on-july-17-2020/

[22] Alex Hasenstab. 2021. Massive Pacific Northwest Storm Causes Power Outages, Downed Trees. https://www.opb.org/article/2021/10/24/storm-to-cause-winds-up-to-65-mph-along-pacific-northwest-coast/

[23] John Heidemann, Yuri Pradkin, Ramesh Govindan, Christos Papadopoulos, Genevieve Bartlett, and Joseph Bannister. 2008. Census and Survey of the Visible Internet. In *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement (IMC '08)*. Association for Computing Machinery, New York, NY, USA, 169–182. https://doi.org/10.1145/1452520.1452542

[24] Zach Hillstrom. 2021. Severed Power Line Causing Water Outages and Issues in Colorado City. https://www.chieftain.com/story/news/2021/07/22/colorado-city-dealing-water-outages-due-severed-power-line/8055846002/

[25] Anette Hulth, Gustaf Rydevik, and Annika Linde. 2009. Web Queries as a Source for Syndromic Surveillance. *PLoS ONE* 4, 2 (Feb. 2009), e4378. https://doi.org/10.1371/journal.pone.0004378

[26] Santosh Janardhan. 2021. Update about the October 4th Outage. https://engineering.fb.com/2021/10/04/networking-traffic/outage/

[27] Andrew Kirmse. 2020, January. Andrew Kirmse's Page - Topographic Prominence. https://www.andrewkirmse.com/prominence

[28] Christian Köpp. 2013. Controlled Internet Outage Monitoring. *Network (Bristol, England)* 11 (2013).

[29] Maxmind. 2022. GeoIP® Databases & Services: Industry Leading IP Intelligence. https://www.maxmind.com/en/geoip2-services-and-databases

[30] Juno Mayer, Valerie Sahakian, Emilie Hooft, Douglas Toomey, and Ramakrishnan Durairajan. 2021. On the Resilience of Internet Infrastructures in Pacific Northwest to Earthquakes. In *Passive and Active Measurement*, Oliver Hohlfeld, Andra Lutu, and Dave Levin (Eds.). Springer International Publishing, Cham, 247–265.

[31] Carlo Medas. 2019. Building a Real-Time Internet Outage Detection System with Flink @ Fing | LinkedIn. https://www.linkedin.com/pulse/building-real-time-internet-outage-detection-system-flink-carlo-medas/

[32] Shikhar Mehrotra. 2021. What Led to Internet Outage That Took down Some Major Websites on July 22? Check out Why. https://www.republicworld.com/technology-news/other-tech-news/what-led-to-internet-outage-that-took-down-some-major-websites-on-july-22-check-out-why.html

[33] Marti Motoyama, Brendan Meeder, Kirill Levchenko, Geoffrey M Voelker, and Stefan Savage. 2010. Measuring Online Service Availability Using Twitter. In *3rd Workshop on Online Social Networks (WOSN 2010)*.

[34] Google News Initiative. 2022. Fundamentals Course. https://newsinitiative.withgoogle.com/training/course/fundamentals

[35] Matthew Odam and Peter Blackstock. 2021. Severe Weather Causing Cellular and Internet Outages in Austin Area. https://www.austin360.com/story/news/2021/02/15/severe-weather-cellular-internet-outages-austin/6752081002/

[36] Ookla. 2020. How Downdetector Works. https://www.ookla.com/articles/how-downdetector-works

[37] Ramakrishna Padmanabhan, Aaron Schulman, Dave Levin, and Neil Spring. 2019. Residential Links under the Weather. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM '19)*. Association for Computing Machinery, New York, NY, USA, 145–158. https://doi.org/10.1145/3341302.3342084

[38] Camille Pelat, Clément Turbelin, Avner Bar-Hen, Antoine Flahault, and Alain-Jacques Valleron. 2009. More Diseases Tracked by Using Google Trends. *Emerging Infectious Diseases* 15, 8 (Aug. 2009), 1327–1328. https://doi.org/10.3201/eid1508.090299

[39] Jay Peters. 2020. YouTube Went down around the World, but It's Now Fixed. https://www.theverge.com/2020/11/11/21561764/youtube-down-outage-loading-videos

[40] Jay Peters. 2021. Networks Are Struggling in Texas amid Historic Winter Storms. https://www.theverge.com/2021/2/15/22284309/texas-winter-storms-internet-cell-network-issues-t-mobile-att-spectrum

[41] Philip M. Polgreen, Yiling Chen, David M. Pennock, and Forrest D. Nelson. 2008. Using Internet Searches for Influenza Surveillance. *Clinical Infectious Diseases: An Official Publication of the Infectious Diseases Society of America* 47, 11 (Dec. 2008), 1443–1448. https://doi.org/10.1086/593098

[42] Associated Press. 2021. Thousands Hit by Internet Outage on East Coast. https://www.mercurynews.com/2021/01/26/internet-outage-strikes-many-on-east-coast

[43] Lin Quan, John Heidemann, and Yuri Pradkin. 2013. Trinocular: Understanding Internet Reliability through Adaptive Probing. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM (SIGCOMM '13)*. Association for Computing Machinery, New York, NY, USA, 255–266. https://doi.org/10.1145/2486001.2486017

[44] Wire Reports. 2021. Internet Recovers after Massive Outage; Verizon Acknowledges 'Issue' as Source | WRAL TechWire. https://wraltechwire.com/2021/01/27/internet-recovers-after-massive-outage-verizon-acknowledges-issue-as-source/

[45] Reuters. 2021. Amazon Cloud Unit Recovers from Brief Outage Affecting Third-Party Services. https://www.reuters.com/markets/commodities/amazon-owned-twitch-down-many-users-2021-12-15/

[46] Philipp Richter, Ramakrishna Padmanabhan, Neil Spring, Arthur Berger, and David Clark. 2018. Advancing the Art of Internet Edge Outage Detection. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. Association for Computing Machinery, New York, NY, USA, 350–363. https://doi.org/10.1145/3278532.3278563

[47] Simon Rogers. 2016. What Is Google Trends Data — and What Does It Mean? https://medium.com/google-news-lab/what-is-google-trends-data-and-what-does-it-mean-b48f07342ee8

[48] Aaron Schulman and Neil Spring. 2011. Pingin' in the Rain. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (IMC '11)*. Association for Computing Machinery, New York, NY, USA, 19–28. https://doi.org/10.1145/2068816.2068819

[49] Staff reports. 2020. AT&T Outage Sunday Updates: 'Progress Continues' in Restoring Service after Nashville Bombing. https://www.tennessean.com/story/

news/crime/2020/12/26/att-outage-internet-911-phones-nasvhille-updates-saturday/4046849001/
[50] Michael Tanenbaum. 2020-01-23T15:09:11.484449-05:00. Comcast Experienced a Nationwide Internet Outage on Thursday. https://www.phillyvoice.com/comcast-nationwide-internet-outage-2020-january-23-philadelphia/
[51] Google Trends. 2022. Google Trends. https://trends.google.com/trends/?geo=CH
[52] USC/LANDER Project. 2022. The ANT Lab: Analysis of Network Traffic | ANT Datasets | Internet Outage Data. https://ant.isi.edu/
[53] Jordan Valinsky and David Goldman. 2021. Massive Internet Outage: Websites and Apps around the World Go Dark. https://www.cnn.com/2021/06/08/tech/internet-outage-fastly/index.html
[54] Zhongqing Wang and Yue Zhang. 2017. DDoS Event Forecasting Using Twitter Data. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Melbourne, Australia, 4151–4157. https://doi.org/10.24963/ijcai.2017/580