

Framing The Invisible Failure In International Press Coverage Of The 2024 CrowdStrike Crisis Through Computational Topic Modeling And Sentiment Analysis

Luis Eduardo Muñoz Guerrero

Universidad Tecnológica de Pereira, lemunozg@utp.edu.co <https://orcid.org/0000-0002-9414-618>

Received: 10/03/2024

Accepted: 21/05/2024

Published: 15/06/2024

ABSTRACT

On 19 July 2024, a faulty update to CrowdStrike's Falcon sensor triggered a global technology outage. This study examines how the international press framed the crisis through a computational approach combining topic modeling (LDA) and sentiment analysis (VADER) over a corpus of 250 GDELT articles published between 19 and 30 July 2024. The model identified four dominant frames with asymmetric affective loads. The results reveal an inverse-scale paradox: the smallest frame concentrates the strongest negativity and causal attribution, whereas the larger frames dilute responsibility into an ambivalent neutrality. Read through Entman's (1993) four framing functions, the findings show how the media redefined a technical failure as a financial and litigation crisis, delegating corrective treatment to markets and courts rather than public regulation.

KEYWORDS: media framing; CrowdStrike; topic modeling (LDA); sentiment analysis; risk communication.

1 Introduction

ON 19 JULY 2024, a faulty update to the Falcon sensor of the cybersecurity firm CrowdStrike triggered one of the largest technology-outage incidents recorded to date. An error in a channel configuration file (channel file 291) set off cascading failures across approximately 8.5 million Windows devices worldwide, temporarily paralyzing operations in critical sectors such as aviation, healthcare, financial services, broadcasting, and emergency services. The economic magnitude of a failure of this nature is not an isolated anomaly but a manifestation of contemporary technological dependence: research on the cost of information-system disruptions had already documented that critical-infrastructure outages generate large-scale cascading losses, with a particularly severe impact on the aviation sector (Chen et al., 2017; Dynes et al., 2007; Lis & Mendel, 2019). In the CrowdStrike case, direct losses among Fortune 500 firms were estimated at 5.4 billion dollars. Delta Air Lines, whose operational recovery took several days longer than that of its competitors owing to the age of its crew-management systems, filed a 500-million-dollar lawsuit against the company in October 2024, alleging gross negligence.

The incident laid bare a paradox characteristic of contemporary digital infrastructure: its invisibility under normal conditions and its catastrophic hypervisibility upon failure. The literature on technological-monoculture risk warns that system homogeneity and dependence on a single provider introduce structural fragilities that remain latent until a failure materializes them simultaneously across millions of systems (Lala & Schneider, 2009); this pattern of dependence has been documented specifically in the healthcare sector, where the transition toward single technology vendors replaces local control with a relationship of dependency (Brunner et al., 2023). The abstract nature of this risk prevents general audiences from grasping the real magnitude of their dependence

until the disruption materializes the harm in tangible form, which makes the media the primary agent of translation between the technical event and the public's understanding of the risk.

From the perspective of framing theory, the media do not report the world as it is; rather, they select aspects of perceived reality and make them more salient, promoting particular problem definitions, causal interpretations, moral evaluations, and treatment recommendations (Entman, 1993). Framing constitutes a multiparadigmatic research program whose productivity lies precisely in the plurality of theoretical approaches it accommodates (D'Angelo, 2002), and its exercise is never neutral: selection and salience distribute symbolic power by privileging some interpretations over others (Entman, 2007). In the CrowdStrike crisis, this framing function takes on special relevance: the immediate cause—a 40-kilobyte configuration file—is radically inaccessible to the non-specialist public. The question guiding this research is therefore: how did the international news media construct the narrative frames of the CrowdStrike failure between 19 and 30 July 2024, and what affective load does sentiment analysis reveal in those frames?

To answer it, we adopt a Computational Data Science (CDS) approach that combines topic modeling via Latent Dirichlet Allocation (LDA) with lexicon-based sentiment analysis (VADER), applied to a corpus of $N = 250$ articles extracted from the GDELT API. This approach makes it possible to process large volumes of text with replicable validity (Maier et al., 2018) and complements qualitative approaches to framing analysis in digital media with empirical scale (Walter & Ophir, 2019).

2 METHODOLOGICAL RATIONALE

The application of computational methods to media framing analysis has accelerated markedly over the past decade. Maier et al. (2018) establish the methodological foundations for the valid and reliable use of LDA in communication research, identifying four critical challenges: text preprocessing, selecting the optimal number of topics, assessing model reliability, and the valid interpretation of emergent topics. Their methodology, with some extensions, directly informs the pipeline implemented in this study.

LDA is a generative probabilistic model that assumes each document is a mixture of latent topics and that each topic is a distribution over the corpus vocabulary (Blei et al., 2003). Formally, each document is generated by drawing a topic distribution θ from a Dirichlet with parameter α and, for each word, a topic z according to θ and then the word from the topic's word distribution ϕ (with Dirichlet prior β); the joint probability of a document factorizes as in Equation 1.

$$P(\theta_d, z_d, w_d \mid \alpha, \beta) = p(\theta_d \mid \alpha) \cdot \prod_n p(z_{d,n} \mid \theta_d) \cdot p(w_{d,n} \mid z_{d,n}, \beta) \quad (1)$$

In the context of framing analysis, the emergent topics are interpreted as the dominant media frames—that is, the packages of meaning that journalists recurrently use to organize coverage of an event (Entman, 1993). This functional equivalence between an LDA topic and a journalistic frame has solid roots in cultural sociology, which early on demonstrated the affinity between topic modeling and the analysis of news coverage (DiMaggio et al., 2013), and it has been systematically operationalized through approaches that treat the topics of an LDA model as networks of frames (Walter & Ophir, 2019). Its application has been validated in the analysis of cybersecurity news (Ghasiya & Okamura, 2021) and in precedents of topic modeling applied to high-profile cyber incidents such as the 2014 Sony Pictures hack (Afful-Dadzie et al., 2016).


The corpus to which this approach was applied comes from the Global Database of Events, Language, and Tone (GDELT), the publicly available real-time news repository with the broadest geographic coverage. For the purposes of this study, 250 English-language articles published between 19 and 30 July 2024 with explicit mention of the terms CrowdStrike, outage, or Falcon

sensor were extracted via the GDELT API. The corpus covers publications from North America, Europe, and Asia-Pacific, ensuring a transnational perspective on the coverage.

The pipeline implemented over this corpus comprises three stages: (1) standard preprocessing (tokenization, stopword removal, lemmatization); (2) LDA modeling with sklearn over a grid of $K \in \{3, 4, 5\}$ topics, selecting $K = 4$ for semantic coherence and minimum perplexity; and (3) article-by-article sentiment analysis using the VADER lexicon (Hutto & Gilbert, 2014) from the NLTK package, extracting the compound score in the range $[-1, +1]$. The complete pipeline results, including the topic–document distributions and word–topic probabilities, are documented in the file `resultados_crowdstrike_2024.md` of this repository.

Figure 1 provides a methodological overview of the study, integrating the corpus-construction pipeline, the LDA generative model that produces the frames, and the VADER step that scores their affective load.

FIGURE 1. Methodological overview: corpus pipeline, LDA generative model, and VADER sentiment scoring



The choice of K was guided by the model's perplexity over the corpus —the inverse of the geometric mean likelihood per word, where lower values indicate better predictive fit (Equation 2)—

$$\text{Perplexity}(D) = \exp\left(-\frac{\sum_d \log p(w_d)}{\sum_d N_d}\right) \quad (2)$$

Once the model was fitted, each article was assigned to its dominant topic, the one with the highest probability in its topic–document distribution (Equation 3); this rule determines the per-topic counts in Table 1 and defines the subsets over which sentiment is aggregated.

$$\text{topic}(d) = \arg \max_k \theta_{d,k} \quad (3)$$

In the sentiment stage, VADER sums the lexical valences of each article's tokens —adjusted by its heuristic rules— and normalizes them to a compound score bounded in $[-1, +1]$ (Equation 4), where $a = 15$ is VADER's normalization constant, distinct from the Dirichlet prior α of Equation 1.

$$\text{compound} = \frac{x}{\sqrt{x^2 + a}}, \quad x = \sum_i \text{valence}_i, \quad a = 15 \quad (4)$$

Finally, the sentiment of each topic is summarized as the mean of the compound scores of the articles that have it as their dominant topic (Equation 5), a value that corresponds to the "Mean score" column of Table 2.

$$M_k = \frac{1}{n_k} \sum_{d \in D_k} \text{compound}(d), \quad n_k = |D_k| \quad (5)$$

3 RESULTS

The LDA model with $K = 4$ topics identified the following dominant frames over the corpus of 250 articles, assigning each article to its highest-probability topic according to Equation 3 (Table 1).

TABLE 1. Identified topics and keywords (LDA)

#	TOPIC	KEYWORDS (TOP 10)	N ARTICLES
T0	Global operational impact	crowdstrike, delta, outage, seek, damages, microsoft, hires, report, boies, david	32
T1	Technical response and recovery	crowdstrike, outage, microsoft, says, ceo, devices, testify, million, affected, called	49
T2	Financial and legal consequences	crowdstrike, outage, global, caused, update, tech, provides, microsoft, users, days	82
T3	Media coverage and public reaction	outage, crowdstrike, microsoft, global, compensation, bug, services, tech, hundreds, experts	87

Note. Model: LDA sklearn · 4 topics · corpus: 250 articles (GDELT 19–30 Jul 2024)

Beyond the volume and vocabulary of each topic, VADER sentiment analysis reveals markedly different affective loads across the four frames (Table 2).

TABLE 2. Sentiment distribution by topic


#	TOPIC	POSITIVE	NEUTRAL	NEGATIVE	MEAN SCORE	INTERPRETATION
T0	Global operational impact	2	7	23	−0.257	Predominantly negative coverage
T1	Technical response and recovery	2	28	19	−0.053	Neutral / ambivalent tone
T2	Financial and legal	17	34	31	−0.112	Neutral / ambivalent tone

#	TOPIC	POSITIVE	NEUTRAL	NEGATIVE	MEAN SCORE	INTERPRETATION
	consequences					
T3	Media coverage and public reaction	18	35	34	-0.116	Neutral / ambivalent tone

Note. Method: VADER (NLTK) · Compound scale: -1 (very negative) to +1 (very positive)

Figure 2 represents, in proportional terms, the sentiment composition (positive, neutral, negative) of each of the four topics, allowing one to appreciate at a glance the negative predominance of T0 against the relative balance of the higher-volume topics.

FIGURE 2. Diverging sentiment composition by topic, with per-topic mean compound score



By way of illustration, Table 3 collects a sample of representative headlines for each topic together with their sentiment score.

TABLE 3. Representative headlines by topic

TOPIC	HEADLINE	SENTIMENT	SCORE
T0	Delta is preparing a major damages lawsuit against CrowdStrike and Microsoft	Negative	-0.586
T0	Microsoft issues report on what caused the huge CrowdStrike crash	Negative	-0.103
T0	Delta planning to pursue damages after IT outage, CrowdStrike shares slide	Negative	-0.178
T1	Microsoft: Windows 'Clearly' Needs Better Resilience After CrowdStrike Outage	Positive	+0.681
T1	Microsoft says millions more devices could have been hit by CrowdStrike outage	Neutral	+0.000

TOPIC	HEADLINE	SENTIMENT	SCORE
T2	CrowdStrike Will Take Hit From Outage But 'Limit' Longer-Term Damage: Analyst	Negative	−0.649
T2	Why did it take days for Delta to restore normal service after outage?	Positive	+0.296
T3	CrowdStrike down after report Delta Air to seek compensation over IT outage	Neutral	+0.000
T3	CrowdStrike scams on the rise after global outage	Negative	−0.586

4 DISCUSSION AND CONCLUSIONS

The results of the computational analysis reveal a framing pattern that transcends mere description of the incident to configure a crisis narrative with systemic implications. Reading this pattern through the four framing functions that Entman (1993, p. 52) identifies —problem definition, causal interpretation, moral evaluation, and treatment recommendation— reveals the logic by which the international media turned a software failure into an event of broad public relevance.

Table 4 summarizes the correspondence between these four functions and the empirical frames identified by the model, which the following paragraphs develop in detail.

TABLE 4. Entman's four framing functions applied to the corpus

FRAMING FUNCTION (ENTMAN, 1993)	DOMINANT ARTICULATION IN THE CORPUS	TOPIC(S)	LEXICAL AND AFFECTIVE EVIDENCE
Problem definition	Shift from the technical to the financial and legal	T0, T2, T3	Core lexeme outage articulated with damages and compensation; only T1 introduces technical vocabulary (devices, resilience), barely 19.6% of the corpus
Causal interpretation	Concentrated responsibility (T0) versus passivized dilution (T2)	T0, T2	boies, david (Boies Schiller Flexner law firm) in T0; caused, update with no agent subject in T2
Moral evaluation	Neutralization: the absence of evaluation equals the absence of symbolic sanction	T1	Mean score $M = -0.053$; the most positive headline in the corpus (+0.681) belongs to this topic
Treatment recommendation	The market and the courts as arbiters, not public regulation	T2, T3	67.6% of the corpus (169/250); absence of regulatory vocabulary (regulation, legislation, oversight)

Note. Interpretive synthesis · Framework: Entman (1993) · Corpus: 250 articles (GDELT 19–30 Jul 2024)

The first framing function, problem definition, operates here through a semantic shift of capital importance. The four topics identified by LDA share the core lexeme outage, but each articulates it with a different semantic field. T0 associates it with damages and seek (claim for damages); T2, with caused, global, and update (technical causality at global scale); T3, with compensation and experts (expert legitimization of the compensation claim). Only T1 introduces genuinely technical vocabulary (devices, resilience), and it is precisely the topic with the smallest presence in the corpus relative to the total number of articles (19.6%).

This finding is consistent with what Boholm (2021) documents in her longitudinal analysis of cyber-threat coverage: when the technological substrate fails, the media lack a narrative repertoire for describing abstract technical causes and resort to the available repertoire of tangible consequences —financial losses, cancelled flights, lawsuits— to anchor the story in the public's experience. The most negatively loaded phrase in the corpus (VADER score = -0.649) refers not to the failure itself but to its market consequences: "CrowdStrike Will Take Hit From Outage But 'Limit' Longer-Term Damage: Analyst". The problem is thus defined not as a vulnerability in software-update architecture but as a risk to shareholder value.

The causal-interpretation function is distributed asymmetrically across the topics. T0 —the one with the greatest affective negativity ($M = -0.257$)— concentrates causal responsibility with almost judicial precision. The words boies and david in its high-probability vocabulary are not generic terms: they refer directly to the law firm Boies Schiller Flexner, retained by Delta Air Lines to represent its lawsuit against CrowdStrike. The presence of this legal lexicon in the most negative framing topic indicates that causality, when articulated precisely, takes the form of civil liability.

By contrast, T2 —the largest topic in the corpus, with 82 articles (32.8%)— dilutes causality into passivized formulations: caused, update, provided. There is no clear agent subject; the failure appears to have occurred without an actor. This pattern is consistent with what has been documented in studies on the media framing of cybersecurity incidents: coverage tends to present failures as systemic or inevitable phenomena, which reduces public pressure on the responsible corporate actors (Anderson et al., 2016). The quantitative paradox is revealing: the frame that most clearly names the responsible party (T0) is the smallest and the most negative; the one that occupies the most space (T2, T3) tends toward a neutrality that functions as causal dilution.

The moral-evaluation function of framing presents its most counterintuitive finding in T1, the topic of technical response and recovery. With a mean sentiment score of -0.053 , T1 is statistically neutral: 28 of its 49 articles receive a neutral categorization from VADER, and the highest-scoring positive headline in the entire corpus ($+0.681$) belongs precisely to this topic ("Microsoft: Windows 'Clearly' Needs Better Resilience After CrowdStrike Outage").

This neutrality of the corporate response is not innocuous: in framing terms, the absence of moral evaluation amounts to an absence of symbolic sanction. The media report that CrowdStrike and Microsoft respond, that the CEO speaks out, that devices are restored, but they pass no judgment on the adequacy of that response. Situational crisis communication theory holds that the strategy and timing of the organizational response decisively condition how the media attribute responsibility and reputation (Coombs, 2007; Mason, 2019); in the CrowdStrike case, executives took hours to address the problem publicly. Yet this communicative failure is not reflected in the affective load of T1. The press, ultimately, assimilates the corporate response within a frame of technical normalization —"the problem was identified and is being resolved"— that strips the corporate action of moral responsibility.

This phenomenon can be interpreted in light of what the literature calls the risk of "technological monoculture": the systemic dependence of millions of organizations on a single provider —conceptualized as a security risk inherent to system homogeneity (Lala & Schneider, 2009) and empirically documented in the healthcare sector's transition toward single technology vendors

(Brunner et al., 2023)—introduces a power asymmetry so pronounced that the media treat it as a given structural fact, not as an anomaly requiring critical evaluation.

Entman's fourth framing function—treatment recommendation—emerges implicitly in the thematic distribution of T2 and T3. Together, these two topics concentrate 67.6% of the corpus (169 of 250 articles) and organize the story around two corrective vectors: private litigation (T2: damages, lawsuit, legal) and the market's reaction (T3: compensation, shares, experts). In none of the topics identified by LDA does a significant regulatory vocabulary emerge (regulation, legislation, government oversight), suggesting that the recommended treatment—implicit in the thematic structure of the corpus—is the market and the courts, not public regulation.

This finding has implications that go beyond the CrowdStrike case. Studies on the media framing of data and artificial-intelligence risks document a similar tendency: the media frame technological failures within logics of civil liability and market correction, while regulatory frames remain marginal except in contexts of explicit political debate (Nguyen, 2023). The CrowdStrike crisis, in this sense, reproduces the structural pattern of techno-economic framing: harm is quantified in dollars, responsibility is processed in the courts, and the underlying infrastructure remains outside public scrutiny.

Alongside the reading of Entman's four functions, the data reveal a structural finding of particular relevance: the inverse relationship between the size of the thematic clusters and their negative affective load. T0, the smallest topic ($n = 32$, 12.8% of the corpus), is significantly the most negative ($M = -0.257$), whereas the larger topics—T2 ($n = 82$) and T3 ($n = 87$)—exhibit ambivalent tones close to neutrality ($M = -0.112$ and -0.116 , respectively).

Table 5 contrasts, for each topic, its relative position by volume and by negativity together with the proportion of negative articles, showing that the smallest frame (T0) concentrates the greatest negative load—71.9% of its articles—while the higher-volume frames (T2, T3) tend toward neutrality.


TABLE 5. The paradox of scale and negativity

TOPIC	N (% OF CORPUS)	RANK BY VOLUME	MEAN SCORE (VADER)	RANK BY NEGATIVITY	% NEGATIVE
T0 Global operational impact	32 (12.8%)	4th (smallest)	-0.257	1st (most negative)	71.9%
T1 Technical response and recovery	49 (19.6%)	3rd	-0.053	4th (least negative)	38.8%
T2 Financial and legal consequences	82 (32.8%)	2nd	-0.112	3rd	37.8%
T3 Media coverage and public reaction	87 (34.8%)	1st (largest)	-0.116	2nd	39.1%

Note. Derived from Tables 1 and 2 · % negative = negative articles / N of the topic

Figure 3 places the four topics on a size–sentiment plane, making visible the inverse relationship between a frame's volume and its negative load: T0, the smallest, alone occupies the region of greatest negativity.

FIGURE 3. Topic volume, mean sentiment, and negative-article share across the four frames, illustrating the inverse-scale paradox



This paradox has a coherent interpretation within Entman's model: the headlines that most precisely articulate the harm, the responsibility, and the legal claim —those that constitute T0— generate more intense affective responses precisely because they close the narrative circle of the crisis: there is an identified victim (Delta Air Lines), a named perpetrator (CrowdStrike), quantified harm (500 million dollars), and corrective action under way (the lawsuit). This narrative completeness correlates directly with greater affective negativity.

Conversely, the volume-dominant frames —T2 and T3— operate through what we might call an "opening narrative": the failure was global, the cause was technical, experts assess, compensation is negotiated. The absence of narrative closure dilutes the affective load and produces the statistical neutrality that VADER detects. This dynamic is consistent with the social amplification of risk framework, according to which the media not only intensify but also attenuate a risk's signal depending on how it fits their conventional narrative frames of responsibility attribution (Kasperson et al., 1988), and with the evidence that coverage of technological risks tends to adopt a globally moderate tone (Brantner, 2021).

Taken together, the results of this study have several theoretical and methodological implications for the field of risk communication and the computational analysis of frames.

First, they confirm the productivity of the LDA–VADER pipeline as a method for analyzing frames at scale for real-time news corpora. The methodology proposed by Maier et al. (2018) for the valid application of LDA in communication research proves operationalizable in acute-crisis contexts, with the advantage that the GDELT corpus makes it possible to capture the temporal evolution of framing within a window of days without depending on manual media selection.

Second, the study illustrates the relevance of topic-differentiated sentiment analysis. Computational framing studies tend to report sentiment for the corpus as a whole; disaggregation by thematic cluster reveals dynamics that the aggregate figure would conceal. In this corpus, the overall corpus sentiment would be approximately -0.13 , a moderate figure that does not communicate the polarization that exists between T0 (high negativity, litigation) and T1 (neutrality, corporate response).

Third, the finding of causal dilution in the dominant frames raises questions about the role of the media in corporate accountability during technological crises. If most articles frame the failure in terms of diffuse consequences (T2, T3) rather than concrete responsibility (T0), media pressure on the causal actors is structurally weaker than the raw number of published articles would suggest. This observation is relevant for policymakers interested in improving accountability in the cybersecurity sector.

These findings must nevertheless be interpreted in light of several limitations of the study. The corpus, although of adequate size for computational analysis ($N = 250$), is limited to English-language articles, which biases the sample toward Anglophone media and may not reflect the framing patterns of the press in Romance, Germanic, or Asian languages. Additionally, the 11-day temporal window (19–30 July 2024) captures the acute phase of the crisis but does not allow analysis of the evolution of framing in the medium and long term, when legal and regulatory frames tend to gain greater presence. Finally, VADER was designed and validated on social-media text (Hutto & Gilbert, 2014) and may underestimate the affective load of news headlines containing irony or institutional euphemisms, a limitation consistent with the evidence that sentiment-analysis performance varies according to the type of text to which it is applied (Castellanos et al., 2021) and with the documented limits of lexicon-based approaches relative to machine-learning-based ones (Qi & Shabrina, 2023).

In sum, this study has shown that international media coverage of the CrowdStrike crisis of 19 July 2024 was organized into four distinct frames with asymmetric affective loads. The media's framing function, in Entman's (1993) sense, operated through: (1) the redefinition of the technical failure as a financial and litigation crisis; (2) the concentration of causal responsibility in the smallest and most negative frame of the corpus (T0); (3) the neutralization of the moral evaluation of the corporate response (T1); and (4) the implicit delegation of correct treatment to the market and the courts (T2, T3).

The central paradox—that the most voluminous frames are the most neutral, while the smallest frame concentrates negativity and responsibility—illuminates an inverse-scale mechanism in the symbolic management of technological crises: the greater the complexity of the event, the greater the pressure toward frames that dilute causal attribution and moderate moral judgment. This mechanism is functionally convenient for the corporate actors involved but epistemologically impoverishing for the publics that depend on the media to understand the risks of the digital infrastructure on which they themselves depend.

Future research should extend the corpus to publications in multiple languages, broaden the temporal window to six or twelve months to capture the evolution of framing during the judicial process, and incorporate framing-analysis techniques based on topic model networks, whose greater sensitivity to the relational structure of frames complements the strengths and offsets the limitations of LDA (Walter & Ophir, 2020).

REFERENCES

1. Afful-Dadzie, E., Nabareseh, S., Oplatková, Z. K., & Klímek, P. (2016). Framing media coverage of the 2014 Sony Pictures Entertainment hack: A topic modelling approach. *Proceedings of the European Conference on Cyber Warfare and Security*.
2. Anderson, R., Böhme, R., Clayton, R., & Moore, T. (2016). Examining the costs and causes of cyber incidents. *Journal of Cybersecurity*, 2(2), 121–135. <https://doi.org/10.1093/cybsec/tyw001>
3. Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
4. Boholm, M. (2021). Twenty-five years of cyber threats in the news: A study of Swedish newspaper coverage (1995–2019). *Journal of Cybersecurity*, 7(1), tyab016. <https://doi.org/10.1093/cybsec/tyab016>

5. Brantner, C. (2021). Covering technology risks and responsibility: Automation, artificial intelligence, robotics, and algorithms in the media. *International Journal of Communication*, 15, 5266–5288.
6. Brunner, J., Anderson, E., Mohr, D. C., Cohen-Bearak, A., & Rinne, S. T. (2023). From "local control" to "dependency": Transitions to single-vendor integrated electronic health record systems and their implications for the EHR workforce. *Journal of General Internal Medicine*, 38(Suppl. 4), 1023. <https://doi.org/10.1007/s11606-023-08281-6>
7. Castellanos, E., Xie, H., & Brenner, P. (2021). Global news sentiment analysis. In *Springer Proceedings in Complexity* (pp. 105–117). Springer. https://doi.org/10.1007/978-3-030-77517-9_9
8. Chen, Z., Rose, A. Z., Prager, F., & Chatterjee, S. (2017). Economic consequences of aviation system disruptions: A reduced-form computable general equilibrium analysis. *Transportation Research Part A: Policy and Practice*, 95, 207–226. <https://doi.org/10.1016/j.tra.2016.09.027>
9. Coombs, W. T. (2007). Protecting organization reputations during a crisis: The development and application of situational crisis communication theory. *Corporate Reputation Review*, 10(3), 163–176. <https://doi.org/10.1057/palgrave.crr.1550049>
10. D'Angelo, P. (2002). News framing as a multiparadigmatic research program: A response to Entman. *Journal of Communication*, 52(4), 870–888. <https://doi.org/10.1111/j.1460-2466.2002.tb02578.x>
11. DiMaggio, P., Nag, M., & Blei, D. (2013). Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of U.S. government arts funding. *Poetics*, 41(6), 570–606. <https://doi.org/10.1016/j.poetic.2013.08.004>
12. Dynes, S., Johnson, M. E., Andrijcic, E., & Horowitz, B. (2007). Economic costs of firm-level information infrastructure failures: Estimates from field studies in manufacturing supply chains. *The International Journal of Logistics Management*, 18(3), 420–442. <https://doi.org/10.1108/09574090710835147>
13. Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4), 51–58. <https://doi.org/10.1111/j.1460-2466.1993.tb01304.x>
14. Entman, R. M. (2007). Framing bias: Media in the distribution of power. *Journal of Communication*, 57(1), 163–173. <https://doi.org/10.1111/j.1460-2466.2006.00336.x>
15. Ghasiya, P., & Okamura, K. (2021). Investigating cybersecurity news articles by applying topic modeling method. 2021 International Conference on Information Networking (ICOIN), 432–438. <https://doi.org/10.1109/icoin50884.2021.9333952>
16. Hutto, C., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1), 216–225. <https://doi.org/10.1609/icwsm.v8i1.14550>
17. Kasperson, R. E., Renn, O., Slovic, P., Brown, H. S., Emel, J., Goble, R., Kasperson, J. X., & Ratick, S. (1988). The social amplification of risk: A conceptual framework. *Risk Analysis*, 8(2), 177–187. <https://doi.org/10.1111/j.1539-6924.1988.tb01168.x>
18. Lala, J. H., & Schneider, F. B. (2009). IT monoculture security risks and defenses. *IEEE Security & Privacy*, 7(1), 12–15. <https://doi.org/10.1109/msp.2009.11>
19. Lis, P., & Mendel, J. (2019). Cyberattacks on critical infrastructure: An economic perspective. *Economics and Business Review*, 5(2), 24–47. <https://doi.org/10.18559/ebv.2019.2.2>
20. Maier, D., Waldherr, A., Miltner, P., Wiedemann, G., Niekler, A., Keinert, A., Pfetsch, B., Heyer, G., Reber, U., Häussler, T., Schmid-Petri, H., & Adam, S. (2018). Applying LDA topic modeling in communication research: Toward a valid and reliable methodology. *Communication Methods and Measures*, 12(2–3), 93–118. <https://doi.org/10.1080/19312458.2018.1430754>
21. Mason, A. (2019). Media frames and crisis events: Understanding the impact on corporate reputations, responsibility attributions, and negative affect. *International Journal of Business Communication*, 56(3), 414–436. <https://doi.org/10.1177/2329488416648951>

22. Nguyen, D. (2023). How news media frame data risks in their coverage of big data and AI. *Internet Policy Review*, 12(2). <https://doi.org/10.14763/2023.2.1708>
23. Qi, Y., & Shabrina, Z. (2023). Sentiment analysis using Twitter data: A comparative application of lexicon- and machine-learning-based approaches. *Social Network Analysis and Mining*, 13(1), 31. <https://doi.org/10.1007/s13278-023-01030-x>
24. Walter, D., & Ophir, Y. (2019). News frame analysis: An inductive mixed-method computational approach. *Communication Methods and Measures*, 13(4), 248–266. <https://doi.org/10.1080/19312458.2019.1639145>
25. Walter, D., & Ophir, Y. (2020). Strategy framing in news coverage and electoral success: An analysis of topic model networks approach. *Political Communication*, 38(6), 707–730. <https://doi.org/10.1080/10584609.2020.1858379>