Using Coreference Resolution to Measure Sentiment in Media Coverage

Farhan Iqbal
farhan.iqbal@uga.edu

UNIVERSITY OF GEORGIA

July 30, 2021 – Academy of Management Annual Conference

The role of the media

The media are key stakeholders for organizations as:

Information Intermediaries

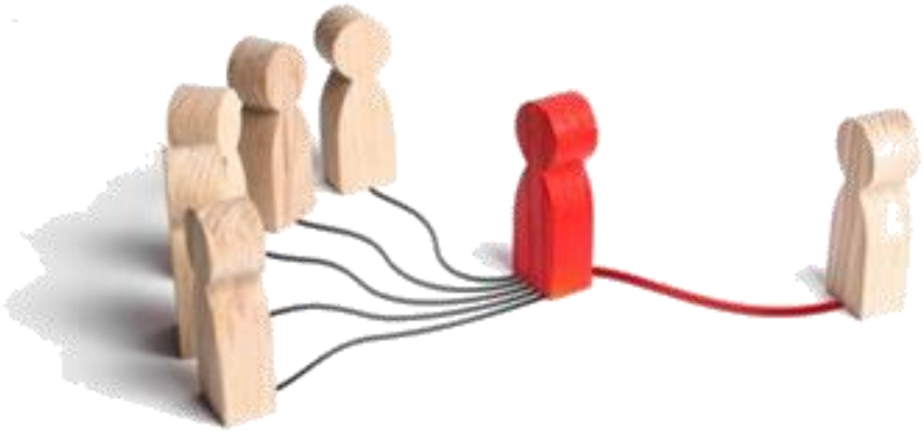

Social Arbiters

Bednar, 2012; Bednar et al., 2013; Bitektine, 2011; Deephouse, 2000; Dyck et al., 2008; Graf-Vlachy et al., 2020; Wartick, 1992

If variance in media coverage is meaningful, we should capture it

How should we measure sentiment?

The best-in-class way to measure sentiment is by using human raters

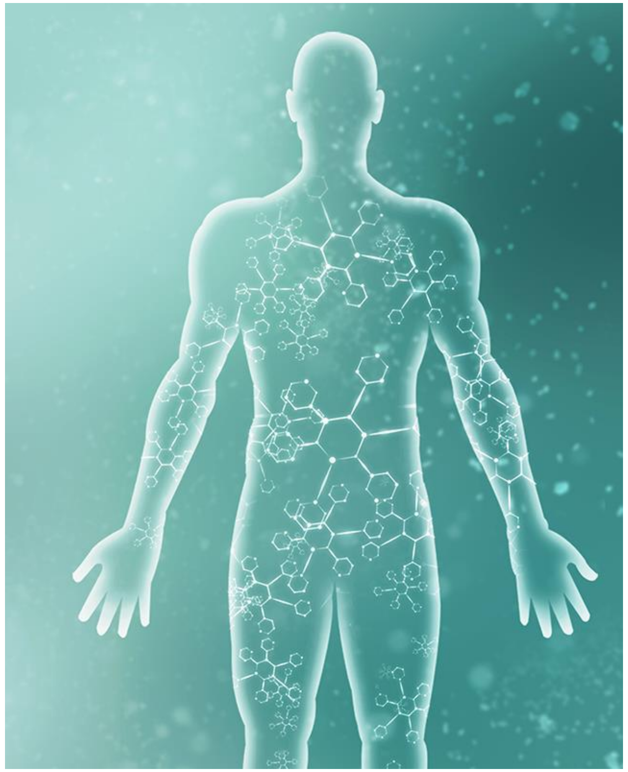

...but that's not always feasible.

So we use computer-aided text analysis (CATA) to help us

DICTION

LIWC

The typical approach

1) Identify articles that include the organization’s name.
2) Read the first paragraph to determine if the article is relevant.*
    1) Discard articles including > X organizations, where X ≈ 1-4
3) Run the entire article through LIWC (or another “bag of words” program that counts the occurrence of words within a corpus); the program classifies text as positive and/or negative.

**if sample size permits*

Bundy et al., 2021; Gamache et al., 2015; Pfarrer et al., 2010; Pollock et al., 2008; Zavyalova et al., 2012

Assessing these models

LIWC

Pros

1) Provides reliable, consistent measurement.
2) Requires no technical knowledge or ability.
3) Is accessible and user-friendly.
4) Has become a standard in applied fields such as management.

Cons

1) Misattributes words to the chosen focal org that refer to other objects (i.e., “type I” errors).
2) Does not attribute words to other orgs in your sample that refer to them (i.e., “type II” errors).
3) Relatedly, we are not capturing within-article variance between coverage across multiple orgs.

Bundy et al., 2021; Gamache et al., 2015; Pfarrer et al., 2010; Pollock et al., 2008; Zavyalova et al., 2012

Recent research may address some of these concerns

Some have created custom sentiment classifiers by:

1) Identifying articles that include the organization's name.
2) Manually classifying a subset of sentences that include the names/keywords.
3) Training a natural language processing algorithm against the manually labeled data.
4) Applying the trained algorithm to the remaining text.

While this process appears to mitigate "type I" errors, it does little to address "type II" errors—and, as we'll see, "type I" errors persist.

Choudhury et al., 2019; Suarez et al., 2021

An example helps illustrate the measurement error

Headline: Following Walmart's lead, Kroger asks customers not to openly carry firearms in stores (Washington Post; Sept 4, 2019)

"**Kroger** is respectfully asking that customers no longer openly carry firearms into our stores, other than authorized law enforcement officers," Jessica Adelman, group vice president of corporate affairs, said in a statement to CNBC on Tuesday. "We are also joining those encouraging our elected leaders to pass laws that will strengthen background checks and remove weapons from those who have been found to pose a risk for violence."

[…]

The announcement came hours after **Walmart**'s. But the Bentonville, Ark.-based retailer went further, saying it would stop selling ammunition for military-style weapons and complete its exit from the handgun business. The company had been under pressure from gun-control advocacy groups, politicians and its own employees since the two store shootings. Roughly 40 white-collar workers in California walked off the job to protest **Walmart**'s gun policies last month, and e-commerce workers in Portland, Ore., and Brooklyn urged the company to stop selling firearms and organized a Change.org petition, which has since garnered more than 140,000 signatures.

An example helps illustrate the measurement error

Headline: Following Walmart's lead, Kroger asks customers not to openly carry firearms in stores (Washington Post; Sept 4, 2019)

"**Kroger** is respectfully asking that customers no longer openly carry firearms into our stores, other than authorized law enforcement officers," Jessica Adelman, group vice president of corporate affairs, said in a statement to CNBC on Tuesday. "We are also joining those encouraging our elected leaders to pass laws that will strengthen background checks and remove weapons from those who have been found to pose a risk for violence."

The problem(s): While the first paragraph would suggest this article is about Kroger, it includes text that refers to another company. Additionally, the companies are referred to by words beyond their names.

The announcement came hours after **Walmart**'s. But the Bentonville, Ark.-based retailer went further, saying it would stop selling ammunition for military-style weapons and complete its exit from the handgun business. The company had been under pressure from gun-control advocacy groups, politicians and its own employees since the two store shootings. Roughly 40 white-collar workers in California walked off the job to protest **Walmart**'s gun policies last month, and e-commerce workers in Portland, Ore., and Brooklyn urged the company to stop selling firearms and organized a Change.org petition, which has since garnered more than 140,000 signatures.

A solution: Coreference Resolution

In linguistics, coreference occurs when two or more expressions in a text refer to the same entity

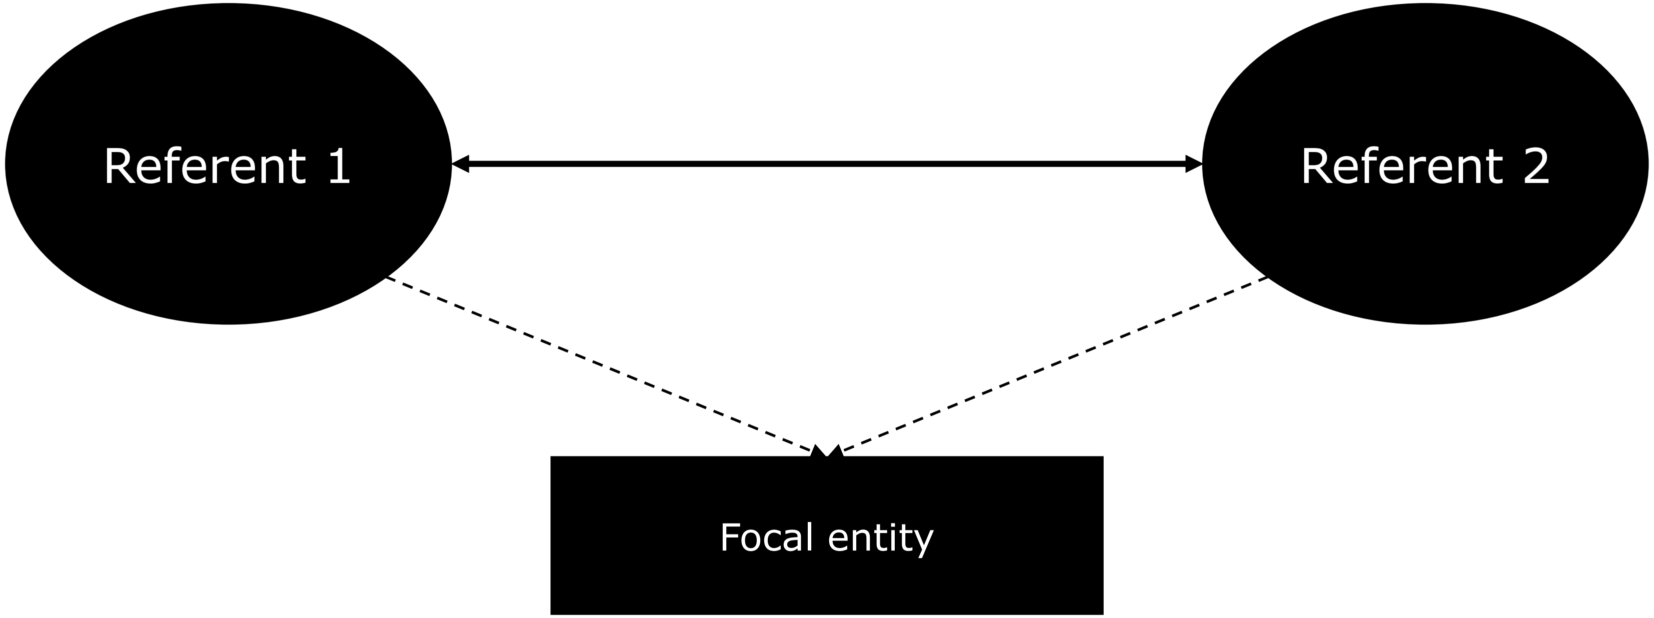

Coreference resolution is actually a series of independent, entity-specific neural networks

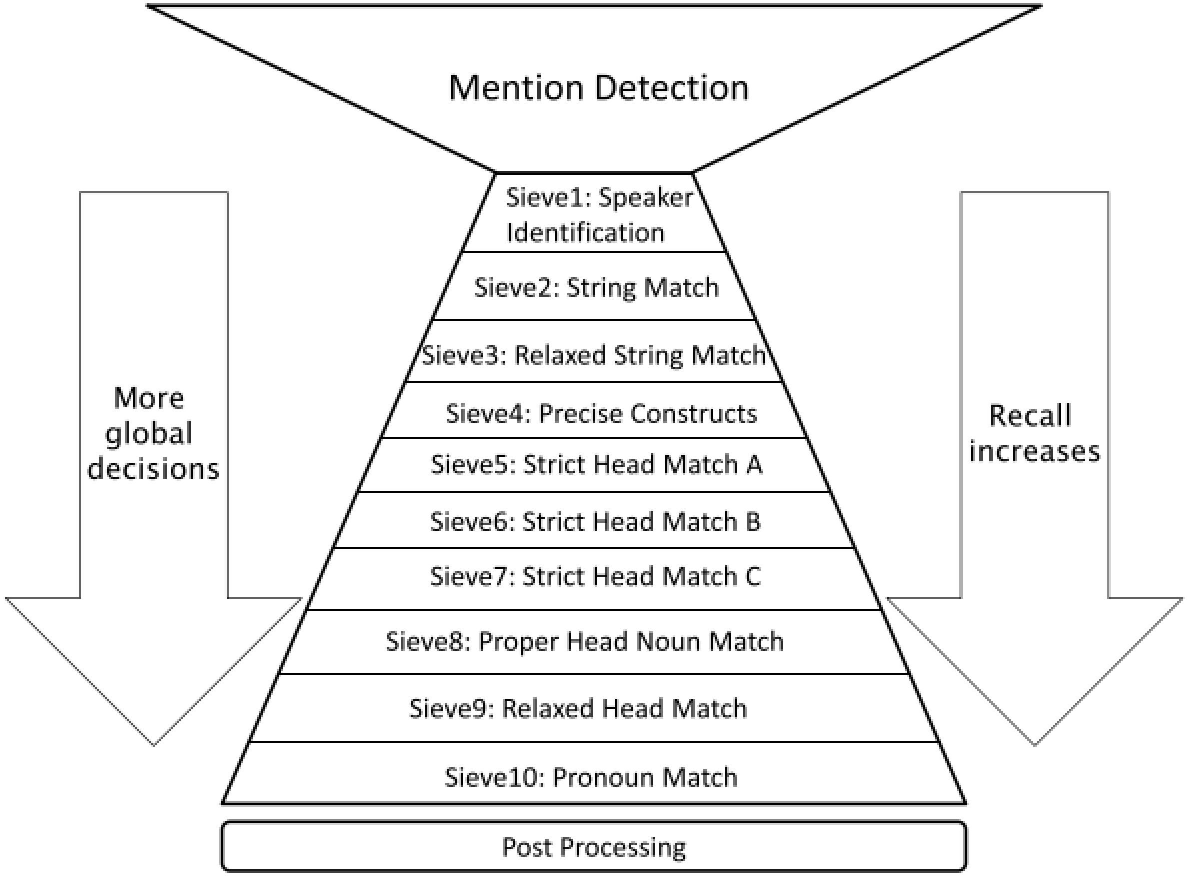

Lee et al. (2013)

A great example from the authors:

Table 1
A sample run-through of our approach, applied to a made-up sentence. In each step we mark in **bold** the affected mentions; superscript and subscript indicate entity id and mention id.

Input:	John is a musician. He played a new song. A girl was listening to the song. "It is my favorite," John said to her.
Mention Detection:	[John]$^{1}_{1}$ is [a musician]$^{2}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. "[It]$^{7}_{7}$ is [[my]$^{9}_{9}$ favorite]$^{8}_{8}$," [John]$^{10}_{10}$ said to [her]$^{11}_{11}$.
Speaker Sieve:	[John]$^{1}_{1}$ is [a musician]$^{2}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. "[It]$^{7}_{7}$ is [**[my]**$^{9}_{9}$ favorite]$^{8}_{8}$," **[John]**$^{9}_{10}$ said to [her]$^{11}_{11}$.
String Match:	**[John]**$^{1}_{1}$ is [a musician]$^{2}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. "[It]$^{7}_{7}$ is [[my]$^{1}_{9}$ favorite]$^{8}_{8}$," **[John]**$^{1}_{10}$ said to [her]$^{11}_{11}$.
Relaxed String Match:	[John]$^{1}_{1}$ is [a musician]$^{2}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. "[It]$^{7}_{7}$ is [[my]$^{1}_{9}$ favorite]$^{8}_{8}$," [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Precise Constructs:	**[John]**$^{1}_{1}$ is **[a musician]**$^{1}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. "**[It]**$^{7}_{7}$ is **[[my]**$^{1}_{9}$ **favorite]**$^{7}_{8}$," [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Strict Head Match A:	[John]$^{1}_{1}$ is [a musician]$^{1}_{2}$. [He]$^{3}_{3}$ played **[a new song]**$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to **[the song]**$^{4}_{6}$. "[It]$^{7}_{7}$ is [[my]$^{1}_{9}$ favorite]$^{7}_{8}$," [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Strict Head Match B,C:	[John]$^{1}_{1}$ is [a musician]$^{1}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{4}_{6}$. "[It]$^{7}_{7}$ is [[my]$^{1}_{9}$ favorite]$^{7}_{8}$," [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Proper Head Noun Match:	[John]$^{1}_{1}$ is [a musician]$^{1}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{4}_{6}$. "[It]$^{7}_{7}$ is [[my]$^{1}_{9}$ favorite]$^{7}_{8}$," [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Relaxed Head Match:	[John]$^{1}_{1}$ is [a musician]$^{1}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{4}_{6}$. "[It]$^{7}_{7}$ is [[my]$^{1}_{9}$ favorite]$^{7}_{8}$," [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Pronoun Match:	**[John]**$^{1}_{1}$ is [a musician]$^{1}_{2}$. **[He]**$^{1}_{3}$ played [a new song]$^{4}_{4}$. **[A girl]**$^{5}_{5}$ was listening to **[the song]**$^{4}_{6}$. "**[It]**$^{4}_{7}$ is [[my]$^{1}_{9}$ favorite]$^{4}_{8}$," [John]$^{1}_{10}$ said to **[her]**$^{5}_{11}$.
Post Processing:	[John]$^{1}_{1}$ is **a musician**. [He]$^{1}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{4}_{6}$. "[It]$^{4}_{7}$ is **[my]**$^{1}_{9}$ **favorite**," [John]$^{1}_{10}$ said to [her]$^{5}_{11}$.
Final Output:	[John]$^{1}_{1}$ is a musician. [He]$^{1}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{4}_{6}$. "[It]$^{4}_{7}$ is [my]$^{1}_{9}$ favorite," [John]$^{1}_{10}$ said to [her]$^{5}_{11}$.

Lee et al. (2013)

Coreference Resolution: A Practical Example

Headline: Following Walmart's lead, Kroger asks customers not to openly carry firearms in stores (Washington Post; Sept 4, 2019)

"**Kroger** is respectfully asking that customers no longer openly carry firearms into our stores, other than authorized law enforcement officers," **Jessica Adelman, group vice president of corporate affairs**, said in a statement to **CNBC** on Tuesday. "We are also joining those encouraging our elected leaders to pass laws that will strengthen background checks and remove weapons from those who have been found to pose a risk for violence."

[…] "A year ago, **Kroger** made the conscious decision to completely exit the firearm and ammunition business when **we** stopped selling them in **our** Fred Meyer stores in the Pacific Northwest," **Adelman** said in the statement to **CNBC**. "**Kroger** has demonstrated with **our** actions that **we** recognize the growing chorus of Americans who are no longer comfortable with the status quo and who are advocating for concrete and common sense gun reforms."

Cincinnati-based **Kroger** is the second-largest grocer in the United States and counts Harris Teeter and Mariano's among **its** many brands. **It** operates nearly 2,800 grocery stores in 35 states, many of which also include on-site pharmacy and jewelry stores.

Coreference Resolution: A Practical Example (continued)

Headline: Following Walmart's lead, Kroger asks customers not to openly carry firearms in stores (Washington Post; Sept 4, 2019)

[…] The announcement came hours after **Walmart's**. But **the Bentonville, Ark.-based retailer** went further, saying **it** would stop selling ammunition for military-style weapons and complete its exit from the handgun business. **The company** had been under pressure from gun-control advocacy groups, politicians and **its** own employees since the two store shootings. Roughly 40 white-collar workers in California walked off the job to protest **Walmart**'s gun policies last month, and e-commerce workers in Portland, Ore., and Brooklyn urged **the company** to stop selling firearms and organized a Change.org petition, which has since garnered more than 140,000 signatures.

"In a complex situation lacking a simple solution, **we** are trying to take constructive steps to reduce the risk that events like these will happen again," **Walmart chief executive Doug McMillon** said in a memo to employees on Tuesday. "The status quo is unacceptable."

The decision was a blow to gun rights advocates, some of whom had been showing up at **Walmart** locations carrying guns on their hips in the hope that **the retailer** would not shift **its** policies.

Why use coreference resolution?

Once properly trained, we can classify sentiment for multiple organizations within one article – mitigating both types of errors.

Beefing it up: By pairing coreference resolution with part-of-speech tagging, we can also identify and remove potentially irrelevant sentences, e.g.:

> The decision was a blow to $\overbrace{\text{gun rights advocates, some}}^{\text{subject}}$ of whom had been showing up at **Walmart** locations carrying guns on their hips in the hope that **the retailer** would not shift **its** policies.

This decision, of course, is dependent on your research question.

Methods of implementation

Stanford NLP

spaCy

Natural Language
Analyses with NLTK

Let’s implement CR on our running example

Headline: Following Walmart’s lead, Kroger asks customers not to openly carry firearms in stores (Washington Post; Sept 4, 2019)

	Score**
LIWC positive	1.41
LIWC negative	1.93
Walmart – positive	57.14
Walmart – negative	28.57
Kroger – positive	71.43
Kroger - negative	14.28

- 22 sentences total* (7 Kroger-specific; 7 Walmart-specific; 569 total words).

- LIWC calculates the percentage of words falling into “positive emotion” and “negative emotion”.

- <u>Devil’s advocate</u>: Most NLP classifiers are trained on higher levels of data than is LIWC (e.g., sentence, paragraph, corpus). So, differences in sentiment scores are expected—not because of CR, but because of the classifier.

*Sentence classifiers can sometimes consider multiple sentences as just one if there is continuity between them (e.g., a quotation by a person that extends to two sentences).

**LIWC scores represent the percentage of words in the article reflecting the corresponding variable; the scores reported for this analysis have thus been comparably manipulated.

Let’s implement CR on our running example

	Score	Z-score*
LIWC positive	1.41	-0.814
LIWC negative	1.93	0.216
LIWC (W) positive	0.84	-1.270
LIWC (W) negative	1.26	-0.374
LIWC (K) positive	1.98	-0.360
LIWC (K) negative	1.58	-0.092

The negative affect appears to stem from parts of the article not specifically related to Kroger or Walmart

Walmart’s score is almost a full standard deviation below Kroger’s

Even when we run the parsed data through LIWC, we see variance in its scores, suggesting that coreference resolution leads to differential evaluations of media coverage within an article.

**Based on a subsample of 226,126 articles published in the third quarter of 2019 across ~30 media outlets*

When does this variance matter?

- When you’re working with text sources that can be attributed to multiple focal entities.
  - In other words, this is not a problem with CEO/firm communications such as letters to shareholders, quarterly earnings calls, tweets, etc.

- When you’re studying behaviors that are more likely to have varied coverage.
  - In my case, sociopolitical acts by firms are more likely to be covered differentially by media outlets (even within one article).

- When you’re studying behaviors that incorporate a breadth of organizations.
  - E.g.: If you’re studying product launches, it is less likely that a given article discussing the launch will feature multiple companies (and, if multiple companies are mentioned, it is unlikely that their mentions are substantive).

Closing thoughts

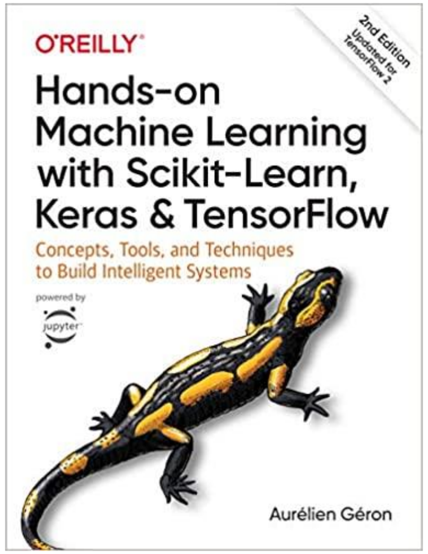

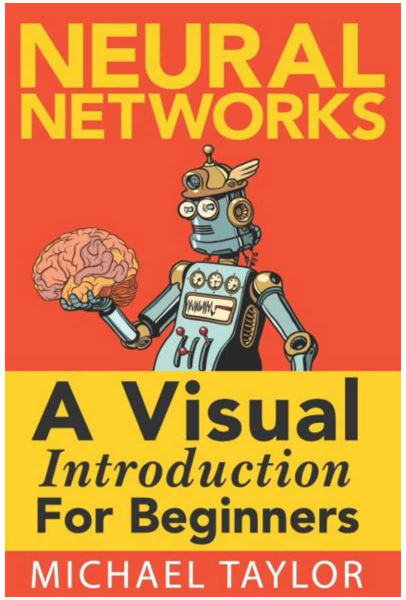

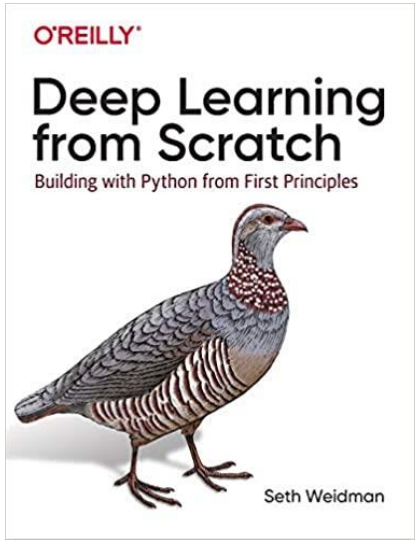

Garbage in, garbage out:
The importance of having quality data, and preprocessing it properly, can never be overstated.

Know what you're doing:
Make sure you know what's going on "under the hood," and be able to explain it (especially to reviewers).

Thank you!

Farhan Iqbal
farhan.iqbal@uga.edu

UNIVERSITY OF GEORGIA

1785

References

Bednar, M. K. 2012. Watchdog or lapdog? A behavioral view of the media as a corporate governance mechanism. *Academy of Management Journal*, *55*(1): 131-150.

Bednar, M. K., Boivie, S., & Prince, N. R. 2013. Burr under the saddle: How media coverage influences strategic change. *Organization Science*, *24*(3): 910-925.

Bitektine, A. 2011. Toward a theory of social judgments of organizations: The case of legitimacy, reputation, and status. *Academy of Management Review*, *36*(1): 151-179.

Bundy, J., Iqbal, F., & Pfarrer, M. D. (2021). Reputations in flux: How a firm defends its multiple reputations in response to different violations. *Strategic Management Journal*, *42*(6): 1109-1138.

Bushee, B. J., & Miller, G. S. 2012. Investor relations, firm visibility, and investor following. *The Accounting Review*, *87*(3): 867-897.

Choudhury, P., Allen, R. T., & Endres, M. G. 2021. Machine learning for pattern discovery in management research. *Strategic Management Journal*, *42*(1): 30-57.

Choudhury, P., Wang, D., Carlson, N. A., & Khanna, T. (2019). Machine learning approaches to facial and text analysis: Discovering CEO oral communication styles. *Strategic Management Journal*, *40* (11): 1705-1732.

Deephouse, D. L. 2000. Media reputation as a strategic resource: An integration of mass communication and resource-based theories. *Journal of Management*, *26*(6): 1091-1112.

Dyck, A., Volchkova, N., & Zingales, L. (2008). The corporate governance role of the media: Evidence from Russia. *The Journal of Finance*, *63* (3): 1093-1135.

Engelberg, J. E., & Parsons, C. A. 2011. The causal impact of media in financial markets. *The Journal of Finance*, *66* (1): 67-97.

Gamache, D. L., McNamara, G., Mannor, M. J., & Johnson, R. E. (2015). Motivated to acquire? The impact of CEO regulatory focus on firm acquisitions. *Academy of Management Journal*, *58* (4): 1261-1282.

Graf-Vlachy, L., Oliver, A. G., Banfield, R., König, A., & Bundy, J. 2020. Media coverage of firms: Background, integration, and directions for future research. *Journal of Management*, *46*(1): 36-69.

Gurun, U. G., & Butler, A. W. 2012. Don't believe the hype: Local media slant, local advertising, and firm value. *The Journal of Finance*, *67* (2): 561-598.

Hubbard, T. D., Pollock, T. G., Pfarrer, M. D., & Rindova, V. P. (2018). Safe bets or hot hands? How status and celebrity influence strategic alliance formations by newly public firms. *Academy of Management Journal*, 61 (5), 1976-1999.

Lee, H., Chang, A., Peirsman, Y., Chambers, N., Surdeanu, M., & Jurafsky, D. (2013). Deterministic coreference resolution based on entity-centric, precision-ranked rules. Computational linguistics, 39(4), 885-916.

Kuypers, J. A. 2002. *Press bias and politics: How the media frame controversial issues*: Greenwood Publishing Group.

Pfarrer, M. D., Pollock, T. G., & Rindova, V. P. (2010). A tale of two assets: The effects of firm reputation and celebrity on earnings surprises and investors' reactions. *Academy of Management Journal*, *53* (5): 1131-1152.

Pollock, T. G., & Rindova, V. P. 2003. Media legitimation effects in the market for initial public offerings. *Academy of Management journal*, *46* (5): 631-642.

Pollock, T. G., Rindova, V. P., & Maggitti, P. G. 2008. Market watch: Information and availability cascades among the media and investors in the US IPO market. *Academy of management journal*, *51* (2): 335-358.

Shani, G., & Westphal, J. D. 2016. Persona non grata? Determinants and consequences of social distancing from journalists who engage in negative coverage of firm leadership. *Academy of Management Journal*, *59* (1): 302-329.

Stephens, M. 2007. *A history of news*: Oxford University Press.

Suàrez, J. L., White, R. E., Parker, S., & Jimenez-Mavillard, A. (2021). Entrepreneurship and the mass media: Evidence from big data. *Academy of Management Discoveries*, (ja).

Wartick, S. L. (1992). The relationship between intense media exposure and change in corporate reputation. *Business & Society*, *31*(1): 33-49.

Westphal, J. D., & Deephouse, D. L. 2011. Avoiding bad press: Interpersonal influence in relations between CEOs and journalists and the consequences for press reporting about firms and their leadership. *Organization Science*, *22* (4): 1061-1086.

Zavyalova, A., Pfarrer, M. D., Reger, R. K., & Shapiro, D. L. (2012). Managing the message: The effects of firm actions and industry spillovers on media coverage following wrongdoing. *Academy of Management Journal*, 55 (5): 1079-1101.