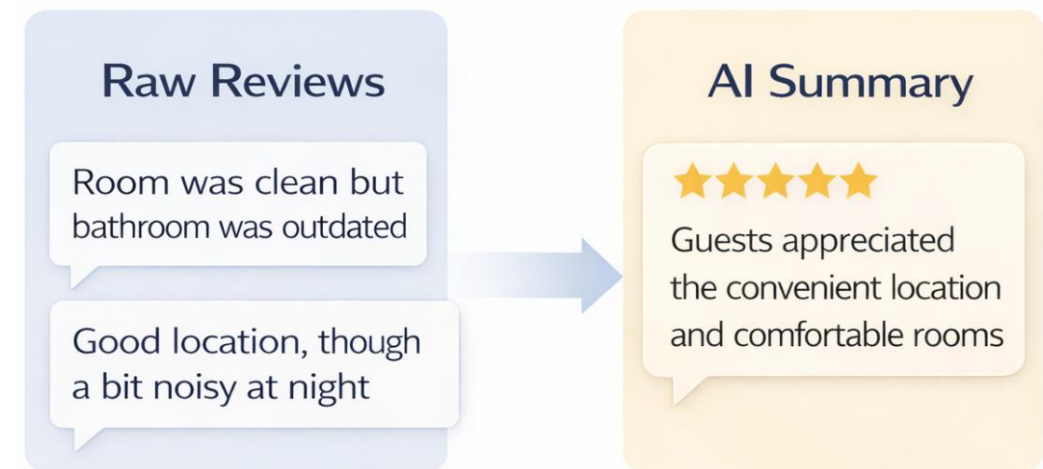


Soft Deception by Design: Algorithmic Positivity Bias in AI-Generated Consumer Review Summaries

Joanna Quan, Wenqi Shen
Virginia Tech

March 19, 2026



AI-generated Summary



Reviews summary

This summary was created by AI, based on recent reviews.

Hyatt Place Dallas Park Central enjoys a convenient location with easy access to major roads and nearby stores, according to many guests. However, reviews reveal mixed experiences with amenities, citing frequent elevator outages and closed facilities like the pool.

Opinions on cleanliness and service are also varied, with some praising the staff and room conditions while others mention housekeeping inconsistencies and noise issues. Overall, the hotel's value and noise levels are common points of dissatisfaction for many travelers.

Motivation Example

Consumer-generated Reviews



darcey113 wrote a review Feb 2023

1 contribution



Bugs and no working elevators

Avoid this place at all costs. Our sink filled with ants and they did not even have an **bug** spray. The elevators don't work and we were on the fourth floor. I can get up and down the steps but people



ksunfire2010 wrote a review Jul 2018

Austin, Texas • 1 contribution • 1 helpful vote



Not what I expected.

Stayed for 4 nights; 2 adults and 3 children. Booked a double room with sleeper sofa non-ADA and got an ADA room. Asked for ground floor and got 2nd. Night two found **bugs**, went to the front desk with the **bug** and changed rooms. Now on the 6th floor. Requested to know where laundry was to wash clothes because of **bugs**, no laundry facility on property. Yet they did offer to shuttle me to a laundry facility so I could do my laundry at 730 at night with 2 pre-teens and a toddler. Went to checkout and was told that the **bug** I found was not a bed **bug** (rrriiiiiggghhhhtttt), yet when asked



Stephanie H wrote a review Oct 2025

4 contributions



Find another hotel.

Incessant noise from the highway along with loud banging noises inside the hotel made it difficult to sleep. The elevator was extremely slow and jerky. There was a **roach** in my room, and my TV only



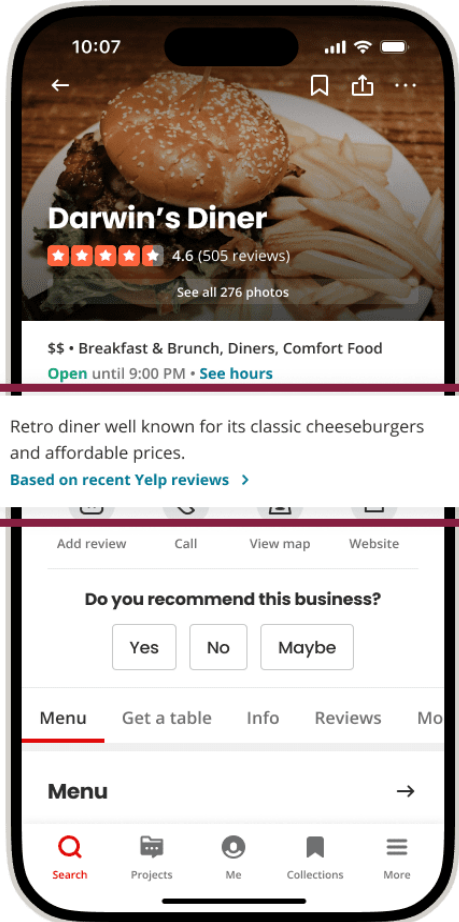
Y6983YAmikea wrote a review Apr 2017

San Antonio, Texas • 1 contribution



Bugs on the Bed

I was woken up to ants crawling on my bed. I had no food or beverages in the room. After calling the front desk I waited 53 minutes for someone to come up. I was then told no other rooms were available and that they would bring **bug** spray. I woke up on the couch waiting, the **bug** spray never made it up.



Customer reviews
 ★★★★★ 4.7 out of 5
 185,057 global ratings

5 star 82%
 4 star 11%
 3 star 3%
 2 star 1%
 1 star 3%

How customer reviews and ratings work

Review this product
 Share your thoughts with other customers
 Write a customer review

Customers say
 Customers praise the Echo Dot for its decent sound quality, reliable performance, and excellent build quality. Customers find it easy to set up and use, and consider it good value for money. Opinions are split regarding connectivity - while some customers report seamless integration with smart home devices, others experience frequent connection issues. The main complaint is that the device frequently stops working unexpectedly.
AI Generated from the text of customer reviews

Select to learn more
 ✓ Sound quality (3.1k) | ✓ Works well (2.2k) | ✓ Quality (2.2k) | ✓ Ease of setup (1.7k) |
 ✓ Value for money (1.1k) | ~ Connectivity (1.1k) | ✗ Playability (860)

Reviews with images

“Up to **98%** of consumers read reviews before making a purchase.”
 - Forbes, 2024



Customer ratings & reviews

4.8 out of 5
 ★★★★★ 801 ratings | 199 reviews
 How item rating is calculated ⓘ

5 stars 88% (705)
 4 stars 6% (51)
 3 stars 2% (17)
 2 stars 1% (5)
 1 star 3% (23)

Reviews summary

Pros

- **Soft material:** Made with a very cuddly material.
- **Affordable price:** Offers great value at this cost.
- **Durable soft fabric:** Made with materials that hold up well to use.
- **Compact size:** Ideal for toddlers to hold and cuddle with.
- **Affordable gift option:** Suitable for small children at a low price.

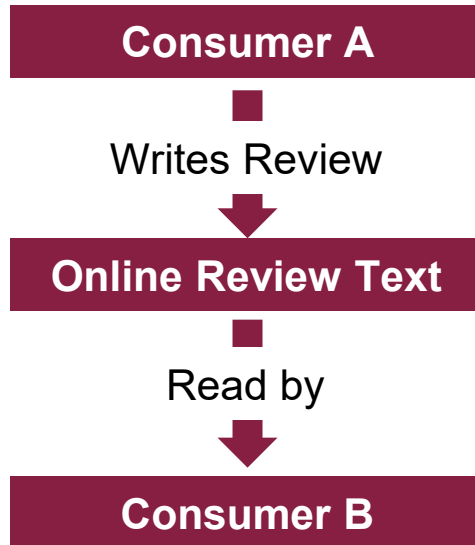
Generated by AI ⓘ [View less](#) ^

“Approximately 3 in 5 (**58%**) U.S. adults said they refer to AI-generated summaries when conducting online searches.”
 - Morning Consult, 2025

Why This Matters NOW

Research Gap – The New Era of Algorithmic Mediation

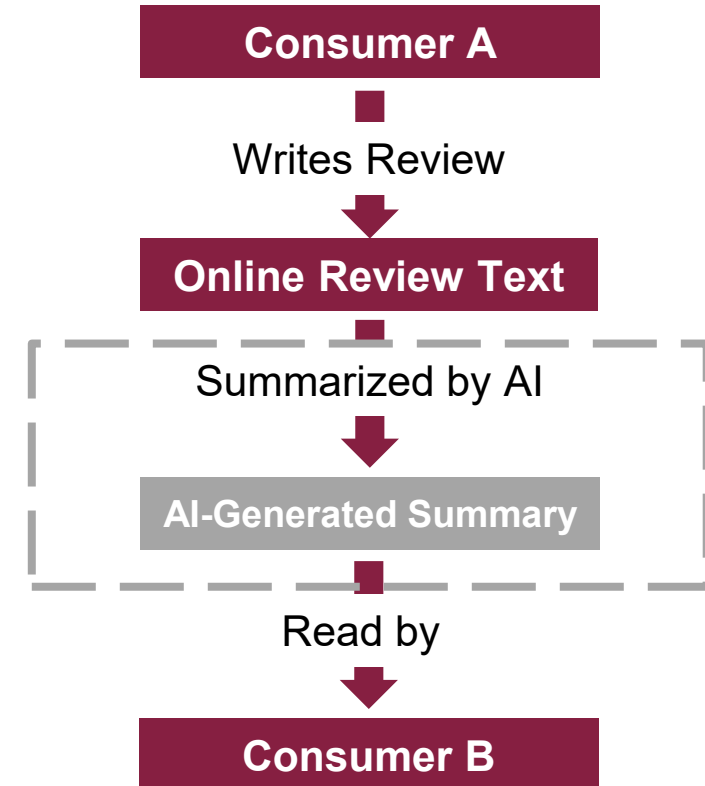
Traditional Online Reviews (Before)



Human Reviewer → Human Consumer

(Chevalier & Mayzlin 2006; Mudambi & Schuff 2010; Shen et al. 2025)

AI-Mediated Review Environment (Today)



Human Reviewer → AI Algorithm → Human Consumer

(Gillespie 2014; Napoli 2014; Seaver 2017; Bucher 2018)

Theoretical Framework – Why AI Might Become Systematically More Positive

Mechanism 1

Positivity Prior

- LLMs trained for fluency, helpfulness, “safety tuning”, user engagement (Bender et al. 2021; Ji et al. 2023)



Mechanism 2

Information Asymmetry

- When review signals are weak (few reviews)
- Model default back to linguistic priors (Akerlof 1970; Chen & Xie 2008)



Systematic Positivity Bias

AI summaries soften negative feedback



Testable Hypotheses

H1: AI summaries more positive than human reviews

H2: Bias strongest when ratings are low and reviews are few

Research Design – Data, Measurement, and Empirical Strategy

TripAdvisor Hotel Dataset

- **905** hotels across **21** major U.S. cities
- Only hotels that had AI summaries
- **96,630** customer reviews
- Reviews from Nov 2024 – Oct 2025
- Average **107** reviews per hotel
- Average hotel rating **4.3 out of 5** stars

AI summaries + customer reviews collected Nov 2025



Sentiment Analysis

VADER

(Hutto & Gilbert 2014)

- Designed for online reviews & social media text
- Sentiment score range: -1 (negative) to +1 (positive)

Sentiment Gap

=

AI Summary Sentiment

-

Average Customer Review Sentiment



Testing the Hypotheses

H1: Overall Bias

- Paired t-test AI summary sentiment vs customer review sentiment

H2: When Bias Is Strongest

- Regression model predicting Sentiment Gap
- Predictors:
 - - Hotel rating
 - - Number of reviews
 - - Price tier
 - - City

Finding 1: AI Summaries Are Systematically More Positive Than Customer Reviews

Average Sentiment

- AI Summaries: **0.94**
- Customer Reviews: **0.71**
- Sentiment Gap: **+0.23**

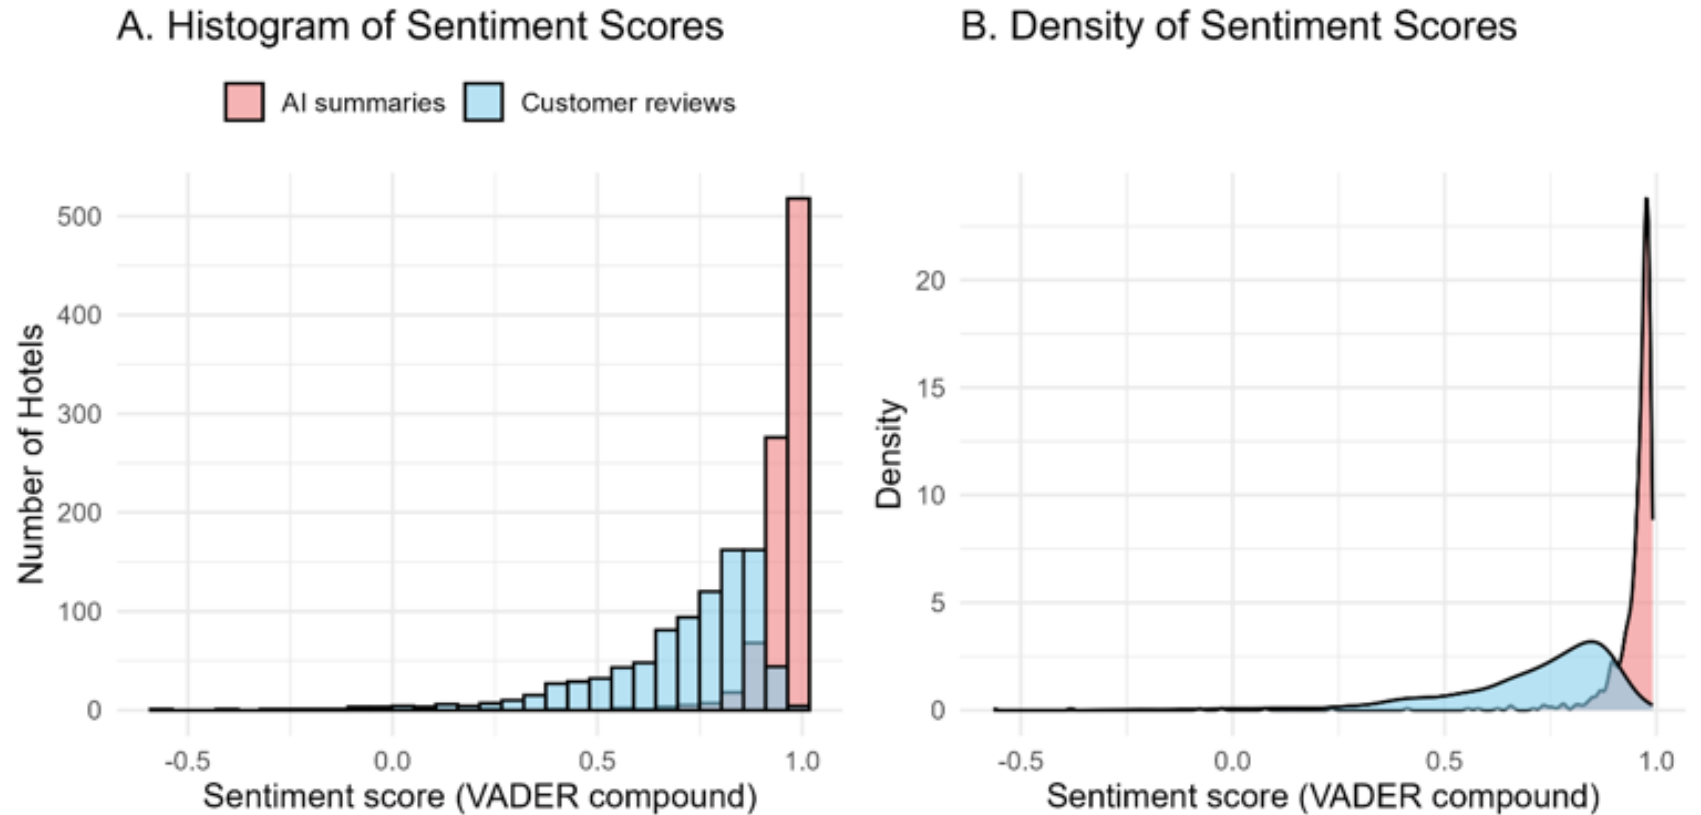
Statistical Tests

- $t(904) = 39.03, p < .001$
- Effect size $d = 1.30$ (very large)

Robustness checks:

- Wilcoxon test $p < .001$
- Recent 6-month sample replicated
- Pattern holds across cities and price tiers

Figure 3. Comparative Distributions of AI and Customer Review Sentiment (N = 905 Hotels)



AI summarization shifts sentiment upward and compresses the variance of sentiment signals. Hotels appear more similar (homogenization effect) and more positive than underlying reviews suggest.

Table 2. Comparison of AI-Generated and Review Sentiment (N = 905 Hotels)

| Measure | M | SD | Min | Max |
|-------------------------------------------------|----------|-----------|------------|------------|
| AI Summary Sentiment (AI_i) | 0.94 | 0.10 | -0.57 | 0.99 |
| Review Sentiment ($Review_i$) | 0.71 | 0.20 | -0.30 | 0.97 |
| Sentiment Gap (Gap_i) | 0.23 | 0.18 | -0.83 | 1.12 |

Table 3. Comparison of AI-Generated Summary and Review Sentiment Across Observation Windows

| Time window | N (hotels) | AI Sentiment M (SD) | Review Sentiment M (SD) | Mean Gap (AI – Review) (SD) | t(df) | p (one-tailed) | d [95% CI] |
|----------------------------------------|-------------------|----------------------------|--------------------------------|------------------------------------|--------------|-----------------------|-------------------|
| 12 months (Nov 2024 – Oct 2025) | 905 | 0.94 (0.10) | 0.71 (0.20) | 0.23 (0.18) | 39.03 (904) | < .001 | 1.30 [1.21, 1.39] |
| 6 months (May – Oct 2025) | 903 | 0.95 (0.10) | 0.69 (0.23) | 0.26 (0.19) | 36.23 (902) | < .001 | 1.21 [1.12, 1.29] |

Finding 2: Bias largest for lower-rated and sparsely reviewed hotels

Predicted Sentiment Gap

- Bias decreases as hotel quality increases
- Inverted-U curve: peaks around 2.5-3.0 stars, then drops toward zero at 5.0 stars

Regression Results

Hotel rating: $\beta = -0.64$, $p < .001$

- Lower rating \rightarrow larger sentiment inflation

Number of reviews: $\beta = -0.24$, $p < .001$

- Few reviews \rightarrow larger bias

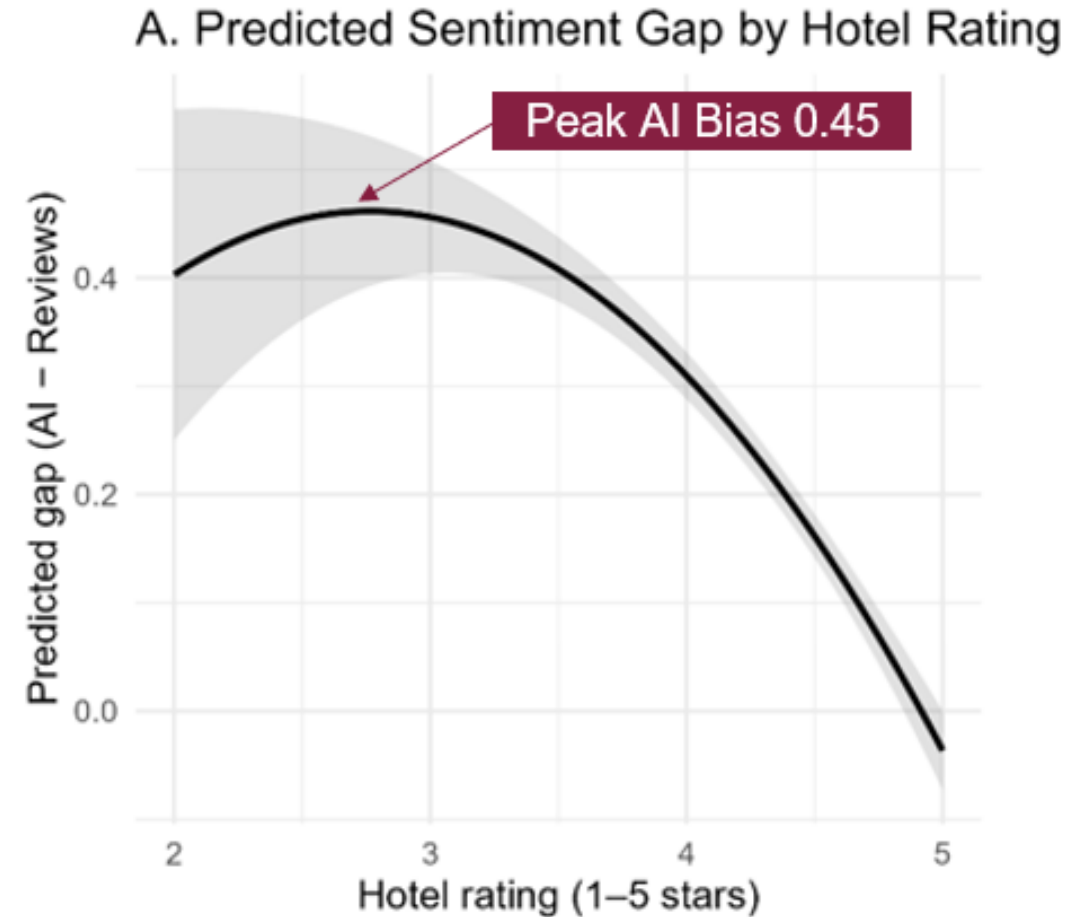
Model fit: $R^2 = 0.39$

- Strong for a field study with real-world data

Robustness checks:

- Quadratic rating term
- Heteroskedasticity-robust SE
- Influential-point diagnostics

AI summaries inflate sentiment most for lower-quality hotels and hotels with few reviews, exactly where consumers need accurate information most.



$$Gap_i = \beta_0 + \beta_1 Rating_i + \beta_2 Rating_i^2 + \beta_3 \log(ReviewVolume_i) + PriceTier_i + City_i + \varepsilon_i$$

Table 8. Nested Model Comparison for Determinants of AI–Review Sentiment Gap

Hierarchical OLS results comparing baseline, linear, and quadratic specifications (HC3-robust).

| | M0 Controls + City FE | M1 + Hotel Rating | M2 + Rating ² (Quadratic) | M2 (sens.) Excl. Outliers |
|--------------------------------------|-----------------------|----------------------|--------------------------------------|---------------------------|
| Hotel rating (centered) | | -0.243*** (0.029) | -0.289*** (0.017) | -0.290*** (0.015) |
| Hotel rating ² (centered) | | | -0.095 (0.056) | -0.055 (0.043) |
| Log ₁₀ (# reviews, 12 m) | -0.056*** (0.008) | -0.037*** (0.007) | -0.036*** (0.007) | -0.036*** (0.007) |
| Num.Obs. | 905 | 905 | 905 | 902 |
| R ² | 0.155 | 0.364 | 0.386 | 0.422 |
| R ² Adj. | 0.132 | 0.346 | 0.368 | 0.405 |
| Std.Errors | HC3 | HC3 | HC3 | HC3 |

* p < 0.05, ** p < 0.01, *** p < 0.001

† p < .10, * p < .05, ** p < .01, *** p < .001. Robust HC3 SEs in parentheses.

Policy Implications for Consumer Protection

Pillar 1 Potential Deception

FTC deception standard:

(FTC, 1983)

- Is likely to mislead
- A reasonable consumer
- Material to decisions

Why AI summaries may qualify:

- Sentiment inflated +0.23 on average. Up to +0.45 for low-quality hotels
- Up to 98% of consumers read reviews before purchasing (Forbes, 2024)
- 58% of consumers read AI search summaries (Morning Consult, 2025)

Mechanism: AI framing softens negative feedback. Similar to dark pattern design risks (FTC, 2022)

Pillar 2 Market Distortion

AI bias advantages lower-quality sellers:

- Bad hotels: up to +0.4 sentiment boost
- Good hotels: +0.2 boost
- → quality differences appear smaller
- AI "fixes" poor reputations instead of rewarding quality

Information asymmetry increases:

- Consumers rely on AI-generated summaries when:
- Reviews are few
- Quality is uncertain

Result: AI amplifies the "lemons problem" (Akerlof, 1970). Disclosure alone may not solve the problem.

AI summarization can unintentionally mislead consumers and distort market competition.

Conclusions

What This Study Demonstrates:

- AI-generated review summaries systematically inflate sentiment
- Bias is strongest where information risk is highest
- AI creates a form of “soft deception”

Why This Matters for Consumer Markets:

- Consumers rely on AI summaries
- Quality signals become distorted
- Information asymmetry increases

Potential FTC Actions:

- **Investigate** AI summarization practices across platforms
- **Apply** deception standards to algorithmic framing
- **Develop** transparency guidelines for AI consumer interfaces
- **Collaborate** with platforms on best practices

Limitations and Future Research:

- Single platform
- Cross-section observational design
- Sentiment measured using VADER



(Image created by ChatGPT)