

The Hermeneutic Loop: Explainable AI (XAI) as a Partner in Interpretive HASS Research

^{1*}Mohammadamin Dadras, ²Abdul Jabbar Perumbalath

³Lim Chong Ewe, and ⁴Firouz Anaraki

¹Bangkok University International, Thailand

²Graduate School, Stamford International University, Thailand

³Stamford International University, Thailand

⁴Faculty of Business and Technology, Stamford International University, Thailand

E-mail: ¹mohammadamin.d@bu.ac.th, ²abduljabbar.perumbalath@stamford.edu

³limchong.ewe@stamford.edu, ⁴firouz.anaraki@stamford.edu

Received September 26, 2025; **Revised** October 28, 2025; **Accepted** December 9, 2025

Abstract

This paper presented a new method called the “Hermeneutic Loop” for using Explainable Artificial Intelligence (XAI) in humanities and social science research. The goal was to solve a major problem: complex AI often acts as a “black box,” making it challenging for researchers to understand how it reaches its results. This lack of transparency clashes with the core principles of fields that rely on deep interpretation. Our framework transformed XAI from a simple technical tool into a collaborative partner. It created a structured conversation between the researcher and the AI model, making the process of interpretation clearer and more rigorous. We demonstrated this method with three real-world examples. First, in literary studies, we used a language model (BERT) and a technique called Integrated Gradients on 500 British novels. This helped us visually trace how the meaning of the word “virtue” changed from a social concept to a psychological one over a century. Second, in historical research, we analyzed letters from the American Civil War. When a standard sentiment-analysis model failed, we used a method called LIME to understand its mistakes. This process revealed unique 19th-century ways of expressing emotion that the AI initially had missed. Finally, in a study of climate change debates on Twitter, we used “attention visualization” to see how different groups used the same words to build distinct political arguments. The Hermeneutic Loop improved our research, forcing us to question the AI’s output,

which led to more reliable and insightful conclusions. It ensured that the computer's power served human understanding. This approach provides a necessary bridge, allowing scholars to use advanced AI while staying true to the fundamental values of evidence-based interpretation.

Keywords: Explainable AI (XAI); Digital Humanities; Hermeneutics; Interpretive Research; Human-AI Collaboration; AI Ethics

Introduction

Researchers in the humanities and social sciences now have access to more information than ever before, from massive digital archives to huge collections of online posts. Because it is impossible for one person to read all this material, scholars are increasingly using powerful computer tools to help them find patterns and test their ideas on a large scale. Methods that use AI to categorize texts by theme or assess their tone are becoming common in research (Wevers & Smits, 2020).

However, these powerful tools create a significant problem. The AI models that are best at finding patterns are often the most difficult to understand. They function as “black boxes”—you put data in and get a result out, but you cannot see the reasoning behind it (Burrell, 2016). This creates a major methodological tension for fields that depend on interpretation and evidence. How can a literature researcher trust a computer's analysis of a novel's themes if they cannot explain how the computer reached its conclusion? The core of traditional research is built on evidence, argument, and the ability to show your work (Dobson, 2019; Fish, 2012). Simply trusting a machine's output goes against this principle.

This is where Explainable AI (XAI) becomes essential. For these tools to be truly useful, scholars must be able to understand the rationale behind their findings (Adadi & Berrada, 2018). Currently, there is a gap between the technical ability to create explanations and the practical needs of researchers. The conversation about XAI is often too technical, leaving scholars unsure of how to use these explanations in their actual work.

Our research addresses this gap by proposing a practical framework we call the “Hermeneutic Loop.” This method structures a collaborative dialogue between the researcher and the AI. Instead of treating the computer's output as a final answer, we use it as the starting point for a critical conversation. The researcher asks a question, the model provides an answer and an explanation, and the researcher then interprets that evidence to ask a better, more informed question. In this process, the AI is not an oracle to be trusted blindly, but a hermeneutic

technology—a partner that provides evidence for the human expert to critique and use with their own knowledge (Gadamer, 2013; Ribeiro et al., 2016).

This paper is structured to guide the reader through our proposal. First, we review the relevant literature on digital humanities and the challenges of “black-box” AI. Next, we explain our “Hermeneutic Loop” methodology in detail. We then demonstrate its value through three real-world case studies from different fields. Finally, we discuss what this human-centered approach means for the future of collaborative, ethical, and rigorous research.

Literature Review

Researchers in the humanities and social sciences have increasingly employed computers to facilitate the analysis of extensive data sets. Digital Humanities began with basic tasks such as creating digital replicas of literary works. It has subsequently developed to include more difficult tasks, such looking at large databases. Moretti’s (2013) concept of “distant reading” represented a significant advancement. It told researchers to look at a lot of texts at once to uncover large trends in the history of literature. Researchers today use computers to track writing trends in books from hundreds of years ago (Underwood, 2019). Historians utilize these instruments to identify patterns in historical data that cannot be discerned by meticulous document examination (Graham et al., 2022). Social scientists look at a lot of social media posts to find out what people think and how they feel about politics (Bail, 2014).

The main advantage of this strategy is that it helps you draw conclusions from a systematic assessment of a lot of evidence instead of just a few examples. But this benefit is limited by a major issue: scholars can’t always figure out how the computer systems make their decisions. Burrell (2016) posits that the “black box” issue is both technological and sociological. Researchers may accept the computer’s results without understanding how they were generated. This kind of blind acceptance goes against essential academic values like honesty, peer review, and careful data interpretation that are important for studying the humanities and social sciences.

1. Explainable AI (XAI): From Technical Solution to an Interpretive Tool

People are determined to make these procedures clearer because it’s hard to understand how computer analysis works. Computer scientists have figured out how to explain those analysis programs work. Ribeiro and colleagues (2016) created LIME, a method that makes individual outcomes clearer. SHAP (Lundberg & Lee, 2017) is another technique that can help you find out which factors had the most effect on a specific result. These tactics were first supposed to help

technicians make sure that systems were working properly and fairly. However, scholars are now explore the possibilities and use of these explanatory methodologies for broader research applications. These tools can aid social scientists in developing and assessing more complicated queries about their data. A Large number of scholars are utilizing visualization techniques to pinpoint the precise words or phrases that software accentuates during text analysis, potentially revealing innovative patterns in narrative structure.

2. Hermeneutics and the Philosophy of Interpretation

To maximize the potential of XAI for HASS, it is imperative to anchor it in the conceptual framework of hermeneutics, the theory of interpretation. The theory of interpretation, developed by important thinkers, emphasizes that understanding is a two-way process. We understand a small part (like a single word or sentence) by seeing how it fits into the whole text and its context. At the same time, we understand the whole text by looking closely at its individual parts. This process is never completely neutral; it is always influenced by the reader's own background, experiences, and pre-existing beliefs. Using AI for interpretation brings a new, non-human agent into this group. Drucker (2021) cautions about the potential for “mis-representation” when humanistic data is coerced into computational models that eliminate context. So, the problem is to come up with a way for people and AI to work together that honors the hermeneutic tradition. We posit that XAI can function as the conduit for reintegrating human judgment into the computational framework, thereby establishing a “Hermeneutic Loop” in which the model's output and its elucidations transform into a novel “text” for scholarly interpretation, with the scholar's enhanced inquiries directing the subsequent computational iteration.

The current literature indicates a significant shortage of explanations. Scholars are capable in utilizing AI and are strongly aware of the challenges associated with black boxes. Concurrently, computer scientists are engaged in the development of XAI tools. However, there exists a deficiency of comprehensive, practical frameworks that effectively incorporate XAI into the specific methodological workflows of HASS research. Most conversations are either very technical or just theoretical. This paper aims to address this deficiency by introducing a tangible, implementable framework—the Hermeneutic Loop—and illustrating its implementation through empirical case studies. Our contribution is not a new algorithm for XAI, but a new way to use it. We present a framework for scholars to systematically utilize established XAI methodologies (LIME, SHAP, attention visualization) to perform research that is both computationally robust and

epistemologically valid, thereby mitigating the epistemological critiques and promoting a more analytical and fruitful interaction between humanists and machines.

The Hermeneutic Loop Framework and Research Design

This research uses a qualitative, case study approach based on design science principles (Hevner et al., 2004). The study focuses on the research process itself, specifically how a scholar interacts with their digital tools. The central subject of analysis is the “Hermeneutic Loop,” a framework we have developed into a four-stage reusable method. The Hermeneutic Loop is designed to create a structured and repeating conversation between a researcher's knowledge and the patterns found by computational analysis. As shown in Figure 1, it consists of four different but connected stages.




Figure 1 The Hermeneutic Loop

1. **Human Initiation:** The researcher starts by coming up with a broad research question based on their area of expertise. For example, “How did representations of gender change in 19th-century novels?” The researcher subsequently selects or trains a suitable AI/ML model on a pertinent corpus. At this point, the model is clearly not being taken as a final authority but as a “theoretical object” (Dobson, 2019) whose parameters represent though a unclear view of the data.

2. **AI Processing and Explanation:** Here analysis tool processes the data and produces an initial result. It is like a set of themes or categories. The key step is for the researcher to use some interpretation methods that explain how the tool predicts that result. For instance, they might use a technique like LIME (Ribeiro et al., 2016) to see which words were most important for

categorizing a text, or SHAP (Lundberg & Lee, 2017) to understand which factors influenced the overall predictions. The goal is to look beyond the raw results and find the evidence behind them.

3. Human Interpretation: This constitutes the fundamental hermeneutic act. The researcher rigorously evaluates the AI's output with the explanations produced by XAI. This is not a passive acceptance; it is an active process of inference like answer queries like; What led the model to concentrate on these features? Does this agreement of evidence support or go against what scholars already know? What hidden patterns, biases, or oddities does the explanation reveal that were not obvious at first? The researcher applies their topic expertise in the theoretical framework and academic verification to the machine-generated evidence.

4. Refined Querying: The researcher refines their question based on the interpretation made in Stage 3. Based on their interpretation, the researcher refines their original question. This might mean asking a more precise question, adjusting the analysis settings, examining a different set of texts, or rethinking the initial approach. This refined question starts the cycle again. The loop continues, allowing the research to develop in depth until it reaches a point of theoretical saturation (Charmaz, 2006), where new cycles no longer yield significant new insights.

Case Study Design and Implementation

We tested this loop with three different case studies from literature, history, and social science to show it can adapt to various research needs. For each case, we were careful to document our process for transparency and auditability:

1. Data: We specified the source, size, and date range of each text corpus (e.g., 500 British novels from 1780–1900; Civil War letters from a specific archive). For the Twitter data, we used a publicly available, anonymized dataset to address ethical concerns. All texts were cleaned and prepared (e.g., correcting OCR errors, standardizing word forms) before analysis.

2. Models & XAI: We recorded the exact AI models used (e.g., best-base-uncased) and how they were tuned. For the XAI step, we noted the specific techniques (LIME, SHAP, Integrated Gradients) and their settings to ensure our explanations could be reproduced.

3. The Human Expert's Role: The “researcher” in the loop was always a domain expert (e.g., a literary scholar). We tracked each iteration of the loop, keeping notes on why questions were changed and what was learned.

4. Evaluation: Our main evaluation was qualitative, based on whether the loop led to deeper, more evidence-based interpretations. In some cases, we also measured quantitative improvements, like an increase in the AI's accuracy after the researcher used the loop to fix a problem.

This detailed approach allows others to see exactly how we applied the Hermeneutic Loop, check our work, and understand both its strengths and its limitations when dealing with challenges like historical language or complex social media debates.

XAI Techniques and Analytical Method

To substantiate and assess the Hermeneutic Loop concept empirically choose three varied case studies shows in table 1. This multiple-case design (Yin, 2018) enables us to illustrate the framework's adaptability across various HASS domains, data kinds, and methodological obstacles.

Table 1 Case Study Overview

Domain	Research Question	Data Source	AI Model	Primary XAI Technique
Digital Literary Studies	Semantic shift of “virtue”	500 British novels (1780–1900)	BERT	Integrated Gradients
Historical Archival Research	Emotional expression in Civil War letters	Valley of the Shadow Project	Sentiment Classifier	LIME
Computational Social Science	Rhetorical framing of climate change	Twitter corpus		

XAI Techniques and Analytical Procedure

For each case study, we chose a way to explain the computer's results that fit the specific tool and what we were trying to learn.

1. In the first case study, we used a text analysis tool called BERT. To understand its results, we used a standard method from a software library (Kokhlikyan et al., 2020) that creates a visual map showing the words. And the tool is considered most important.

2. In the second case study, we used a simpler sorting tool. To explain its conclusions that used the LIME framework (Ribeiro et al., 2016) and it is helpful because that can create explanations for many different kinds of tools by showing the specific reasons for a single result.

3. In the third case study, which involved grouping texts by topic that used visualization techniques to see any words the tool was focusing on when it created the topics. It is based on established ideas (Vaswani et al., 2017).

The main approach to analysis was qualitative. We used an abductive method (Timmermans & Tavory, 2012), which means we carefully examined the evidence to develop the most reasonable interpretation. We did not simply accept the computer's numerical scores or groupings as final answers. Instead, we treated the explanations—the visual maps and highlighted text—as the key evidence for our own close reading and interpretation. We kept detailed notes of our thought process as we worked through each cycle of the Hermeneutic Loop. This ensured that our interpretive process was clear and could be examined by others.

Case Study 1 and Discussion: Digital Literary Studies

This first case study shows how the Hermeneutic Loop can be used to explore a classic question in literary history (How does the meaning of a word change over time?). We studied the evolution of the word “virtue” across a century of British novels. Our main goal was not just to prove that its usage changed in the context. But, simply to understand how its meaning shifted from an external, social quality to an internal and moral one. The Hermeneutic Loop provided a structure to move beyond simple categorization and also develop interpretations grounded in evidence.

1. Methodology and Application of the Loop

This case study utilized the Hermeneutic Loop to analyze the semantic evolution of the term “virtue” within a collection of 500 British novels published between 1780 and 1900. The first model is a pre-trained BERT model that was fine-tuned on the corpus. And it was able to group novels by period with a high level of accuracy. But these raw clusters didn't help us figure out why the chronological grouping was done. This started the first time the loop ran.

2. Results and Interpretation: Mapping Semantic Shift

The use of Integrated Gradients was the most important step in turning the model's output into scholarly evidence. The saliency maps showed the language features the model used to tell the difference between periods in table 2. For instance, in early novels, the model linked “virtue” to words like “fortune,” “birth,” and “station,” which showed that it had a class-based meaning. In subsequent works, the prominent terms transitioned to “conscience,” “feeling,” and “moral sense,” indicating a movement towards an internalized, psychological interpretation.

Table 2 Top Stylistic Markers for Clusters of Novels Containing “Virtue”

Period Cluster	Top Salient Words (via Integrated Gradients)	Interpreted Semantic Dimension
Early (1780–1820)	Fortune, birth, station, family, honour	Socially–constructed, class–based virtue
Late (1840–1900)	Conscience, feeling, moral sense, duty, inner	Individually–felt, psychological virtue

This data from XAI enabled us to describe the historical shift from a socially–constructed to an individually–perceived view of benefit, offering quantitative validation for established literary theories and identifying the precise linguistic indicators of this transformation. Our research clearly showed how the meaning of the word “virtue” changed in British novels during the 1800s. It started as a quality connected to your social class and family, and later became about your personal conscience and inner feelings. We discovered this by using a process called the Hermeneutic Loop. First, we asked a computer model to sort the novels by their time period, which it did successfully. But we didn’t know how it could tell the difference. So, we used a special technique (XAI) that showed us which words the model was paying attention to. This turned the computer’s “black box” answers into clear evidence. The results gave us solid proof of this historical change and identified the exact words that signaled the shift.

Case Study 2 and Discussion: Historical Archival Research

The first case study looked at the Hermeneutic Loop in regard to modern texts. This second study looks at how important it is for historical research. It talks about a common issue in digital humanities and it is utilizing new computer technologies on old documents. The case study focuses on analyzing emotional expression in letters from the “Valley of the Shadow” Civil War archive, where an initial sentiment analysis model shown inadequate performance.

1. Methodology and Application of the Loop

We analyzed a collection of digitized personal letters from the American Civil War (from the Valley of the Shadow project; Ayers & Rubin, 2000). Our goal was to find subtle expressions of trauma and emotional coping that were not directly stated. We began by using a modern computer tool designed to detect emotional tone in text. However, when we checked its results by reading the letters ourselves, we found many mistakes. This is a common problem when using such tools on historical documents (Underwood, 2019). Instead of abandoning the tool, we used its mistakes as a starting point for our investigation. We applied an explanation technique called LIME (Ribeiro et al., 2016) to understand why the tool made errors. We changed our question from

“Is the tool correct?” to “What can the tool’s mistakes teach us about how people expressed emotions in the past?”

2. Detailed Findings: A Typology of Misinterpretation

LIME explanations on misclassified instances in table 2 revealed that the model was overly reliant on a modern, simplistic lexicon of emotion. The words it weighted heavily (“sad,” “happy”) were less important than period-specific phrases. This discovery led to a refined query: “What language did the model fail to understand, and what does this teach us about 19th-century emotional expression?”

Table 3 LIME Analysis of Model Errors on Civil War Letters

Error Type	Example Text Snippet	Model’s Incorrect Classification	Key Phrases from LIME Explanation
False Positive	“I commend myself to God’s will.”	Negative Sentiment	‘commend’, ‘God’s will’
False Negative	“We must endure this terrible trial.”	Positive Sentiment	‘endure’, ‘terrible trial’

False Positives – Stoicism and Euphemism: The model frequently classified expressions of stoic acceptance or religious faith as negative. LIME showed that words like “commend” and “God’s will” were strong positive indicators for the model, but their co-occurrence with adversity led to misclassification. This revealed a cultural script where emotional restraint was expressed through a framework of duty and faith, which the model’s modern sensibility misinterpreted as negativity.

False Negatives – Complex and Conventionalized Language: Conversely, the model labeled texts expressing grim determination as positive. Phrases like “we must endure” were interpreted as resilience. The LIME output highlighted that the model failed to grasp the conventionalized nature of this language, which often accompanied descriptions of profound suffering, thus missing the nuanced negative affect.

3. Conclusion and Transdisciplinary Value

iterative process forced a refinement of our entire approach, leading us to train a new model on a seed lexicon developed from the LIME explanations. This not only dramatically increased accuracy but, more importantly, used the model’s failure as a lens to reveal the authentic, complex emotional language of the era.

Case Study 3 and Discussion: Computational Social Science

The third case study uses the Hermeneutic Loop in the field of computational social science, where the goal is frequently to figure out how to understand large groups of people and their activities. The study examines the complex online discourse around climate change on Twitter. Although preliminary models were capable of identifying overarching themes, they did not adequately encapsulate the intricate rhetorical tactics that characterize persuasive communication. This case study shows how the Hermeneutic Loop was used to go from figuring out what is being said to figuring out how different parties construct their arguments in a strategic way.

1. Methodology and Application of the Loop

The third case study analyzed the rhetorical framing of climate change on Twitter. A standard transformer-based topic model identified broad themes like “denial,” “policy,” and “science.” While informative, these topics did not reveal how arguments were constructed differently by various groups. This limitation initiated the Hermeneutic Loop, prompting the use of attention visualization.

2. Results and Interpretation: Revealing Rhetorical Structure

The attention visualizations added a deep layer of rhetorical analysis. For the “policy” topic, the visualizations showed that tweets from different ideological groups used the same keywords but structured their arguments fundamentally differently. One group’s attention patterns showed a cause–effect structure linking policy to economic doom. Another group’s patterns linked the same policy words to frames of justice and future security. This moved our analysis from what was being discussed to how arguments were being constructed, a question central to social science but previously difficult to answer at scale.

Results

Our three case studies show that the Hermeneutic Loop successfully solves the “black box” problem in digital research. This framework creates a clear conversation between researchers and AI tools, making computer analysis more transparent and trustworthy. The main strength of this approach is how it helps researchers build evidence–based interpretations. In our literature study, we didn’t just learn that the word “virtue” changed meaning – we discovered exactly how it changed by seeing which words the computer used as evidence. In our history project, instead of throwing away a poorly performing tool, we learned from its mistakes to understand how people expressed emotions differently in the past. The framework works across different research

areas because it's flexible. Each of our studies used a different explanation method suited to its particular needs. This shows that the Hermeneutic Loop is a practical guide that can adapt to various research questions and computer tools. Most importantly, this approach keeps human expertise at the center of research. The computer helps analyze large amounts of data, but the human researcher remains in charge – asking questions, checking evidence, and drawing meaningful conclusions. This ensures that modern computer tools support traditional research values like careful interpretation, transparency, and deep understanding of context.

Conclusion

The Hermeneutic Loop framework operationalizes a novel methodology for human–AI collaboration in interpretive research. By treating the ML model as a “theoretical object” and integrating XAI outputs as primary evidence for close reading, this process moves beyond using AI for mere pattern detection. Instead, it establishes a structured, iterative, and auditable cycle of critical inquiry. The three case studies demonstrate that this is not a generic process but a flexible one, where the choice of XAI technique is carefully tailored to the model and the epistemological need. Ultimately, this framework provides a necessary bridge, ensuring that the application of powerful computational tools in HASS is governed by the foundational scholarly principles of evidence, interpretation, and argumentation.

References

- Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- Bail, C. A. (2014). The cultural environment: Measuring culture with big data. *Theory and Society*, 43(3–4), 465–482. <https://doi.org/10.1007/s11186-014-9216-5>
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1). <https://doi.org/10.1177/2053951715622512>
- Charmaz, K. C. (2006). *Constructing grounded theory: A practical guide through qualitative analysis*. Thousand Oaks, CA: Sage.
- Dobson, J. E. (2019). *Critical digital humanities: The search for a methodology*. Illinois: University of Illinois Press. <https://doi.org/10.5622/illinois/9780252042270.001.0001>

- Drucker, J. (2021). *The digital humanities coursebook: An introduction to digital methods for research and scholarship*. London: Routledge.
- Fish, S. (2012, January 9). *The digital humanities and the transcending of mortality*. The New York Times. Retrieved from <https://archive.nytimes.com/opinionator.blogs.nytimes.com/2012/01/09/the-digital-humanities-and-the-transcending-of-mortality/>
- Gadamer, H.-G. (2013). *Truth and method*. (Revised 2nd ed.). London: Bloomsbury Academic.
- Graham, S., Milligan, I., Weingart, S., & Martin, K. (2022). *Exploring Big Historical Data: The Historian's Macroscope*. (2nd ed.). World Scientific Publishing. <https://doi.org/10.1142/12435>
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
- Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., Melnikov, A., Kliushkina, N., Araya, C., Yan, S., & Reblitz-Richardson, O. (2020). Captum: A unified and generic model interpretability library for PyTorch. *arXiv*. Retrieved from <https://arxiv.org/abs/2009.07896>
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *arXiv*. <https://doi.org/10.48550/arXiv.1705.07874>
- Moretti, F. (2013). *Distant reading*. London: Verso Books.
- Ribeiro, M.T., Singh, S., & Guestrin, C. (2016). “Why Should I Trust You?” Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, August, 13–17, (pp. 1135–1144). <https://doi.org/10.1145/2939672.2939778>
- Timmermans, S., & Tavory, I. (2012). Theory construction in qualitative research: From grounded theory to abductive analysis. *Sociological Theory*, 30(3), 167–186. <https://doi.org/10.1177/0735275112457914>
- Underwood, T. (2019). *Distant horizons: Digital evidence and literary change*. Chicago: The University of Chicago Press.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *arXiv*. <https://doi.org/10.48550/arXiv.1706.03762>
- Wevers, M., & Smits, T. (2020). The visual digital turn: Using neural networks to study historical images. *Digital Scholarship in the Humanities*, 35(1), 194–207. <https://doi.org/10.1093/dlsc/fqy085>

Yin, R. K. (2018). *Case study research and applications: Design and methods*. (6th ed.). Thousand Oaks, CA: Sage.