

Can Nuanced Language Lead to More Actionable Insights? Exploring the Role of Generative AI in Analytical Narrative Structure

VIDYA SETLUR, Tableau Research, USA

LARRY BIRNBAUM, Salesforce, Northwestern University, USA



Fig. 1. Left: Relevant language describing visual features can be useful for authoring summaries to help with readers' takeaways and provide context to what is occurring at the feature. Right: A GPT-4-generated summary that incorporates nuanced language to draw a reader's attention to specific portions of the text that are important to the overall takeaway of the intended message.

Relevant language describing trends in data can be useful for generating summaries to help with readers' takeaways. However, the language employed in these often template-generated summaries tends to be simple, ranging from describing simple statistical information (e.g., extrema and trends) without additional context and richer language to provide actionable insights. Recent advances in Large Language Models (LLMs) have shown promising capabilities in capturing subtle nuances in language when describing information. This workshop paper specifically explores how LLMs can provide more actionable insights when describing trends by focusing on three dimensions of analytical narrative structure: *semantic*, *rhetorical*, and *pragmatic*. Building on prior research that examines visual and linguistic signatures for univariate line charts, we examine how LLMs can further leverage the *semantic dimension* of analytical narratives using quantified semantics to describe shapes in trends as people intuitively view them. These semantic descriptions help convey insights in a way that leads to a *pragmatic* outcome, i.e., a call to action, persuasion, warning vs. alert, and situational awareness. Finally, we identify *rhetorical* implications for how well these generated narratives align with the perceived shape of the data, thereby empowering users to make informed decisions and take meaningful actions based on these data insights.

CCS Concepts: • **Human-centered computing** → **Visualization systems and tools**.

Additional Key Words and Phrases: semantic trends, pragmatics, rhetoric, data analysis.

1 INTRODUCTION

Relevant language describing visual features in charts can be useful for authoring summaries about the data to help with readers' takeaways and provide some context to what is happening in a data trend. However, the language employed in the recommended captions and annotations tends to be simple, ranging from specifying the domain, axes, and encodings to providing statistical information (e.g., extrema) describing specific marks in the chart.

The use of more nuanced language is crucial in data analysis to convey the subtleties and complexities inherent to patterns in the data. This precision in language facilitates a deeper understanding and leads to more effective

Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). Copyright remains with the author(s).

informed decision-making [16, 19]. Hedging is a common communicative strategy used in language to emphasize and bring a reader’s attention to specific portions of the text that are important to the overall takeaway of the intended message [14, 15]. The proliferation of Generative AI along with Large Language Models (LLMs) amplifies this capability by analyzing vast amounts of data and uncovering patterns, trends, and correlations that might be missed in manual analysis [8]. Furthermore, LLMs can simulate various scenarios and predict outcomes, enabling analysts to explore a wider range of possibilities and make more strategic decisions. In this paper, we explore how LLMs such as GPT-4 [17] could be employed to further enrich analytical narratives to elicit effective takeaways and action by the intended reader. We specifically explore the enrichment of analytical narratives across semantic, rhetorical, and pragmatic dimensions to further improve the clarity and impact of data-driven storytelling in the analytical decision-making process.

2 RELATED WORK

Natural language has been used to help facilitate both language interaction and generation to help people reason with data [1, 2, 9]. Systems like Eviza [20] and Analyza [6] incorporate contextual inferencing capabilities to enhance user interactions. Other systems, such as Evizeon [13] and Orko [22], explore pragmatics within analytical conversations by leveraging contextual understanding. Flowsense [23] enables NL-based interactions in dataflow systems. However, existing NLI systems primarily focus on analytical inquiry and do not adequately address the semantic, pragmatic, and rhetorical dimensions required for describing insights in trends, which our work aims to address using GPT models.

iGraph [7] supports querying trends but offers limited support for interpreting their semantic features, particularly in quantifiable properties like slope features. Hoque et al. conducted a comprehensive survey of chart question-answering systems, identifying opportunities for more sophisticated models leveraging language and semantics to support data exploration [12]. Olio, a hybrid search system, combines semantic Q&A search with document-based exploratory search but lacks robust support for querying specific semantics related to trends [21]. Our work extends beyond existing trend querying systems by leveraging GPT models to support the exploration of trends with a comprehensive understanding of their semantic, pragmatic, and rhetorical dimensions, thereby enhancing insights derived from visual data analysis.

3 ANALYTICAL NARRATIVE STRUCTURE FOR DESCRIBING DATA TRENDS

Analytical narrative structure encompasses the organization and presentation of findings to effectively communicate the story within the data [3]. The narrative should carefully incorporate descriptive language that accurately reflects the nuances of the observed trends, ensuring precision and clarity in conveying the message. We now discuss research opportunities that LLMs afford for the integration of semantic precision, pragmatic relevance, and rhetorical persuasion in this analytical narrative structure. We believe such a framework could provide a basis for effectively describing insights in trend data, driving understanding, decision-making, and action.

3.1 Semantic Dimension

Specific word choices play a crucial role in how the meaning or *semantics* of trend shapes and patterns are perceived and understood. The language used to describe data trends significantly influences the audience’s perception and interpretation, making the precision and subtlety of language an essential aspect. Accurate descriptions of trend shapes, such as “spike,” “drop,” “plateau” or “fluctuation,” are not merely descriptive but carry connotations that impact the audience’s understanding. These terms can alter the perceived significance or implications of a trend, necessitating the use of precise language that aligns with the data’s narrative. This premise aligns with the recent work of Bromley and Setlur [4], who explored the importance of labeling semantic visual features in data analysis to facilitate a more

nuanced understanding of trends. Their approach identifies subtle semantic differences in data trends and integrates the the collected dataset with LLMs for textual summary generation. In addition, contextualizing data within a *specific* domain is also helpful in making insights more meaningful. For example, a flat trend in the financial sector may be called ‘constant,’ while a more appropriate label could be ‘unchanged’ for weather forecasting. Incorporating LLMs to generate these bespoke semantics could further augment the semantic dimension of analytical narrative structures for describing trends. These models generate nuanced descriptions, identify contextually relevant linguistic patterns, and suggest alternative phrasings based on the domain that better capture the subtleties of the data.

3.2 Rhetorical Dimension

The rhetorical dimension focuses on conveying the semantics of these trend descriptors to the reader such that the analytical narrative can convey information that leads to either a call to action, persuasion, alert, or providing situational awareness. Language that provides effective rhetoric typically involves the nuanced usage of hedge words and connectives as the atomic units of communication [10]. For instance, the choice between terms like “fell sharply” versus “crashed” is not just about word choice; it conveys different degrees of severity. Similarly, “stagnant” vs. “stable” also convey different expectations, with the former generally having a negative connotation, suggesting a lack of growth or change. In contrast, “stable” has a more neutral or positive connotation, indicating consistency and reliability.

LLMs could be appropriately prompted to emphasize certain aspects of data trends that align with a specific call to action. This process of appropriate narrative construction involves choosing words and constructing sentences that convey the appropriate level of urgency and importance. For data trends specific to a particular domain (e.g., financial reporting, public health data), GPT can employ language that includes explanations, contextual information, and possibly comparisons with past data or predictions for the future. Gricean principles, also known as cooperative principles, are a set of guidelines in the field of pragmatics for effective communication [11]. These principles include four Maxims of Quantity, Quality, Relation, and Manner to support these cooperative properties. The principles have been shown to be appropriate scaffolds within a GPT prompt to help generate an effective rhetoric. For example, in the generated insight, “Revenue fell; however, drilling down further, the revenue from apparel decreased significantly,” supports the Maxims of Manner (i.e., the narrative should be informative as is required) and Quality (i.e., the narrative should provide truthful and adequate evidence) by first stating the overall situation (i.e., revenue fell) and then adding a specific detail (i.e., the apparel revenue’s significant decrease), which is directly relevant to the overall context (Maxim of Relation). The use of connectives, such as, ‘however’ and phrasal transitions, such as ‘drilling down further,’ avoids ambiguity and adds clarity to the narrative (Maxim of Manner).

3.3 Analytical Connectives

In addition to supporting cooperative language construct, we can draw upon several types of analytical connectives for generating rhetoric, especially in the context of data analysis and presentation [5]. By explicating adding directives for the inclusion of analytical connectives in the prompt input to GPT, we can generate various forms of insight narratives that further enhance the rhetorical dimension, described as follows:

Temporal connectives involve connecting ideas or data points based on their temporal relationships. For example, temporal connectives are indicated in **bold** in this GPT-generated summary, “The calm pattern in the stock market over the **last month** with **minimal volatility**, reflects investor confidence and a lack of disruptive market news.”

Part-whole relationships focus on how individual data trends or elements contribute to or are a subset of a larger dataset. Phrases such as ‘a component of’ or ‘an aspect of the broader trend’ are useful in this context. For example, part-whole connectives are indicated in **bold** in this GPT-generated summary, “The increase in overall profits this quarter was largely driven by a surge in online sales, which constituted **60% of the total revenue.**”

Comparison connectives are used to draw parallels or highlight differences between different datasets or elements within a dataset by employing statements of comparison (e.g., ‘compared to,’ ‘unlike,’ ‘similar to’) in the GPT-generated summary indicated in **bold**, “While the company’s revenue has increased, competitors have seen **similar growth**, implying that this trend is industry-wide rather than company-specific.”

Roll-up or drill-down connectives are used to describe trend data in terms of aggregation or disaggregation. A roll-up involves summarizing or aggregating data to a higher level (e.g., from daily sales to monthly sales), while a drill-down involves breaking data into more detailed components (e.g., breaking down annual revenue into product-wise revenue). Phrases like ‘summing up,’ ‘in aggregate,’ or ‘breaking down into’ can be used to generate a summary such as, “**Drilling down** into the total revenue, we observe that the North American market contributes 50% of it, with a significant portion coming from online sales,” indicated in **bold**.

Normalization strategies are used to describe how data can be adjusted or normalized to account for seasonal variations, cyclical trends, or different time scales. For instance, the summary with normalization phrases indicated in **bold** - “When **normalized year over year**, the company’s revenue shows a consistent **seasonal pattern**, with spikes during the holiday season and dips in the second quarter.”

3.4 Pragmatic Dimension

The pragmatic dimension in describing analytical narratives for data trends focuses on the practical implications and actions that arise from the analysis. In other words, this dimension encompasses various aspects such as decision support, predictive analysis for future planning, risk management, effective resource allocation, policy and strategy formulation, performance improvement, and change management [18]. We posit that LLMs can aid in supporting the pragmatic dimension by reassessing the analysis approach, evaluating existing methodologies, and suggesting improvements. For instance, in a scenario where a company observes a stagnation in market growth, the LLM can analyze existing market data and propose a shift in analysis towards emerging trends or unexplored customer segments. Additionally, when existing data is insufficient for informed decision-making, the LLM’s role becomes crucial in identifying gaps and suggesting areas for additional data collection. For example, if a retail business notices a decline in sales, the LLM might suggest gathering data on recent market shifts, competitor strategies, or changing consumer preferences to gain a comprehensive view.

4 TAKEAWAYS AND RESEARCH IMPLICATIONS

In summary, leveraging LLMs to provide semantic, rhetorical, and pragmatic dimensions for describing analytical narratives presents a promising avenue for future research. Specifically, the potential of GPT to generate hedge words and analytical connectives, supplied as a curated word bank in the prompt, opens up possibilities for creating more nuanced and contextually relevant narratives. These narratives, enriched with appropriate linguistic cues, can be effective for memorability and takeaways, thus eliciting a call to action. Investigating the impact of these GPT-generated narratives on decision-making could provide valuable insights into their practical applicability and effectiveness.

REFERENCES

- [1] 2024. IBM Watson Analytics. <http://www.ibm.com/analytics/watson-analytics>.
- [2] 2024. ThoughtSpot. <http://www.thoughtspot.com>.
- [3] N. Allen, J. Templon, P. McNally, L. Birnbaum, and K. Hammond. 2010. Statsmonkey: Automatic Generation of Narrative from Data. In *Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium*. Computational Models of Narrative: Papers from the 2010 AAAI Fall Symposium ; Conference date: 01-01-2010.
- [4] Dennis Bromley and Vidya Setlur. 2023. What Is the Difference Between a Mountain and a Molehill? Quantifying Semantic Labeling of Visual Features in Line Charts. *IEEE Transactions on Visualization and Computer Graphics* (2023).
- [5] Stuart Card, Jock Mackinlay, and Ben Shneiderman. 1999. *Readings in Information Visualization: Using Vision To Think*. Morgan Kaufmann.
- [6] Kedar Dhamdhare, Kevin S. McCurley, Ralfi Nahmias, Mukund Sundararajan, and Qiqi Yan. 2017. Analyza: Exploring Data with Conversation. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces (IUI 2017)*. 493–504.
- [7] Leo Ferres, Avi Parush, Zhihong Li, Yandu Oppacher, and Gitte Lindgaard. 2006. Representing and Querying Line Graphs in Natural Language: The iGraph System. In *Smart Graphics, 6th International Symposium, SG 2006, Vancouver, Canada, July 23-25, 2006, Proceedings (Lecture Notes in Computer Science, Vol. 4073)*, Andreas Butz, Brian D. Fisher, Antonio Krüger, and Patrick Olivier (Eds.). Springer, 248–253. https://doi.org/10.1007/11795018_25
- [8] Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Christian Petersen, Alexis Chevalier, and Julius Berner. 2023. Mathematical Capabilities of ChatGPT. arXiv:2301.13867 [cs.LG]
- [9] Tong Gao, Mira Dontcheva, Eytan Adar, Zhicheng Liu, and Karrie G. Karahalios. 2015. DataTone: Managing Ambiguity in Natural Language Interfaces for Data Visualization. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software Technology (UIST 2015)*. ACM, New York, NY, USA, 489–500.
- [10] G. Goodnight. 2014. Rhetoric, Communication, and Information. *PORO: An Interdisciplinary Journal of Rhetorical Analysis and Invention* DOI: 10.13008/2151-2957.1185 (01 2014), Article 4. <https://doi.org/10.13008/2151-2957.1185>
- [11] H. P. Grice. 1975. Logic and Conversation. In *The Logic of Grammar*, Donald Davidson and Gilbert Harman (Eds.). 64–75.
- [12] E. Hoque, P. Kavehzaheh, and A. Masry. 2022. Chart Question Answering: State of the Art and Future Directions. *Computer Graphics Forum* 41, 3 (2022), 555–572. <https://doi.org/10.1111/cgf.14573>
- [13] Enamul Hoque, Vidya Setlur, Melanie Tory, and Isaac Dykeman. 2017. Applying Pragmatics Principles for Interaction with Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 309–318.
- [14] K. Hyland. 1998. *Hedging in Scientific Research Articles*. John Benjamins Publishing Company. <https://books.google.com/books?id=VP-DQgAACAAJ>
- [15] G. Lakoff. 1973. *Hedges: A Study in Meaning Criteria and the Logic of Fuzzy Concepts*. University of Chicago Press. <https://books.google.com/books?id=H8XbygAACAAJ>
- [16] Bongshin Lee, Nathalie Henry Riche, Petra Isenberg, and Sheelagh Carpendale. 2015. More Than Telling a Story: Transforming Data into Visually Shared Stories. *IEEE Computer Graphics and Applications* 35, 5 (2015), 84–90.
- [17] OpenAI. 2023. GPT-4 Technical Report. arXiv (2023), 2303–08774.
- [18] Klaus P. Schneider and Anne Barron (Eds.). 2014. *Pragmatics of Discourse*. De Gruyter Mouton, Berlin, München, Boston. <https://doi.org/doi:10.1515/9783110214406>
- [19] Edward Segel and Jeffrey Heer. 2010. Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1139–1148.
- [20] Vidya Setlur, Sarah E. Battersby, Melanie Tory, Rich Gossweiler, and Angel X. Chang. 2016. Eviza: A Natural Language Interface for Visual Analysis. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST 2016). ACM, New York, NY, USA, 365–377.
- [21] Vidya Setlur, Andriy Kanyuka, and Arjun Srinivasan. 2023. Olio: A Semantic Search Interface for Data Repositories. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software Technology* (San Francisco, California) (UIST 2023). ACM, New York, NY, USA. <https://doi.org/10.1145/3586183.3606806>
- [22] Arjun Srinivasan and John Stasko. 2018. Orko: Facilitating Multimodal Interaction for Visual Exploration and Analysis of Networks. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 511–521.
- [23] B. Yu and C. T. Silva. 2020. FlowSense: A Natural Language Interface for Visual Data Exploration within a Dataflow System. *IEEE Transactions on Visualization and Computer Graphics* 26, 01 (jan 2020), 1–11. <https://doi.org/10.1109/TVCG.2019.2934668>