Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

AI-Powered Narrative Generator for Data Storytelling in Business Dashboards

[1] Nayan Gowda D
UG Scholar
Department of Computer Science and Engineering
Presidency University,
Bengaluru, India

[2] Harshith H
UG Scholar
Department of Computer Science and Engineering
Presidency University,
Bengaluru, India

[3] Durgesh Nishad
UG Scholar
Department of Computer Science and Engineering
Presidency University,
Bengaluru, India

[4] Abhishek D
UG Scholar
Department of Computer Science and Engineering
Presidency University, Bengaluru, India

[5] Sameera P Simha
Assistant Professor
Department of Computer Science and Engineering
Presidency University, Bengaluru, India

***Abstract*— Business intelligence dashboards provide a more intuitive way to explore data, including with charts, key performance indicators, statistical results, and machine-learning models. Nevertheless, users must still make their own interpretations of these results when trying to determine what the data really means. In this paper, an AI-based narrative generator is introduced which adds a natural-language explanation layer to a business intelligence dashboard. The proposed system involves taking a dataset, applying data profiling and data preprocessing, followed by visualization, statistical analysis, and predictive modeling. The analytical results are gathered together and sent to a module of narrative generation. The system doesn't ask the language model for statistical values but gives it values that are already computed and uses the language model primarily to explain the values in a readable form. The system has a local open-source language model, Hugging Face Inference API, and a default rule-based fallback. A deterministic insight engine is also employed to discover conditions like excessive variation, skewed distributions, high correlation and missing-value clumping. The system discussed in this paper proposes to use analytical computation with natural-language explanations to make business dashboards understandable and to keep the numerical analysis out of the explanations.**

Keywords— Business Intelligence, Data Storytelling, Large Language Models, Natural Language Generation, Data Visualization, Machine Learning, Automated Insights.

I. INTRODUCTION

Data is becoming more critical to the business organizations to understand their business and make the decision. What you can learn from the massive amounts of data is valuable information, including trends, relationships, performance and interesting patterns, but it is not always easy to get this information from the data. Business intelligence systems are designed to solve this issue by creating dashboards with charts, tables, key performance indicators, statistical measures and predictive models. While these dashboards give users a visual overview of data, users must read through various components and make their own interpretations.Data storytelling allows for this process to be more understandable through combining the analytical findings with explanations written in natural language. A data storytelling system can tell the story of what happened, whether there are variables that seemed important, or if the pattern observed might need to be further investigated instead of just a graph or a number. Automated data storytelling and generating explanations based on structured information using large language models have recently been examined [1, 2].In this paper, the project is described that introduces an LLM Narrative Insights Engine into a business intelligence workflow. The system is built on an existing analytical pipeline, which include data ingestion, data profiling, ETL processing, visualization, statistical analysis, predictive modelling, and report generation. The results generated by these analytical components are used to create a proposed narrative layer that is converted to human-readable explanations.One of the key issues in the design of the system is that the language model should not be used for the numerical calculations. For instance, statistical tests, model metrics, correlations and other numerical values are computed before the analytical modules. These verified results are then put together into a structured context and provided to the narrative engine. This enables the language model to concentrate on conveying the results, instead of generating statistical figures on its own.There are also several ways to create the story line in the system. If offline processing is desired, there is an open-source model that can be locally hosted. However, if you have internet connectivity and a suitable API key, another option is available with the Hugging Face Inference API. A simple insight message can also be generated using a deterministic rule-based system if neither of the two options listed above is available.The idea of this system is then to provide a link between the results of the analysis and the human understanding. The narrative generator does not take the place of the existing dashboard elements, but is rather an extra component that clarifies information that has been generated by these elements.

II. RELATED WORK

Business intelligence (BI) dashboards are visual representations that consolidate key performance indicators (KPIs), charts, and data analytics. Users may have to compare several charts and numerical indicators, however, before they can discover some meaningful results or become aware of the meaning of their business. Data storytelling is a solution to this problem, because it is a mix of data analysis, visualisation and text to explain the data. This is the approach taken in the proposed project where natural language explanations, grounded

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

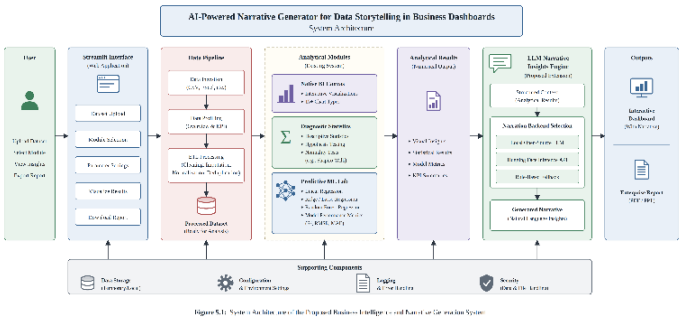
International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

in the dashboard visualization, are created.There has been recent research investigating application of large language models (LLMs) to enhance automated storytelling in data. Zhang, Li and Gao presented a context-aware and multi-dimensional data-storytelling framework, namely MDSF, which relies on large language models [1]. This work illustrates the need to reflect on context and multiple dimensions of data when creating narratives. The proposed system is similar, using structured analysis outputs to a language-generation component. What is more, its main task is to incorporate the capability of generating a story into an already established workflow for business intelligence, including data preprocessing, statistical analyses, visualization, predictive modelling and report generation.Islam et al. introduced DataNarrative, an approach that examines the possibility of automating and creating data-driven stories with the help of visualizations and textual explanations [2]. The focus on linking visual evidence to description to create coherent data stories illustrates the importance of this approach. The proposed system has outputs of computed metrics like KPI summaries, hypothesis-test results, feature-importance values, and predictive-model metrics. These outputs include facts to give context for creating explanations, while the dashboard remains with the same charts and numbers.Studies of the development of a dashboard with the help of AI are also of relevance. Niu et al. introduced Dashboard2Code, which investigates the use of multimodal models to reconstruct interactive dashboards [3]. While dashboard reconstruction is not narrative generation, it does show the application of AI techniques in business-dashboard workflows. An alternative problem is considered, namely the generation of natural language explanations from the results that are created by the existing modules in the dashboard, instead of generating the code of the dashboard.Additional to this, there is another type of basis for communicating analytical findings, called explained artificial intelligence (XAI). Recently, Lundberg and Lee [4] have introduced a framework for understanding model predictions by feature-attribution values, known as SHAP. Ribeiro et al. proposed LIME, that approximates the behaviour of a model locally to explain individual predictions [5]. Selvaraju et al. proposed the Grad-CAM method that is used to obtain visual explanation for decisions made by deep neural network [6]. These techniques can be useful to understand the model's behavior, but they might not be sufficient to be directly interpretable by the model user if they do not understand the model's feature or explanation. The proposed system can use the computed feature importance information to provide to the narrative engine, as a way to avoid asking the LLM to compute the underlying values, and instead render analytical evidence in a more accessible language.The project also expands upon existing analytical techniques that produce the quantitative data that informs the created narratives. A foundation for ensemble-learning models for predictive analysis is provided by Breiman's Random Forest method [7]. Staniak and Biecek presented an overview of the automated exploratory data analysis and associated tools, focusing on the process of systematic data investigation [8]. Shapiro-Wilk test developed by Shapiro and Wilk [9] is a statistical test to see if a sample is normally distributed. The narrative engine is used to convey their computed output, but these techniques are used to support the analytical stage of the proposed system.The included work exemplifies advances in automated data storytelling, automated LLM text generation, dashboard creation and construction, statistical analysis, and explainable AI. But these areas are not necessarily connected to a single flow of work. The proposed project aims to tackle this integration challenge by integrating dataset profiling, ETL processing and interactive visualization, statistical analysis, predictive modelling and report generation with a dedicated layer for creating narratives. It has a locally-hosted open-source model, a hosted inference API, and a deterministic rule-based fallback. Importantly, the analytical modules compute the numerical results before the generation of narratives. These results are used in the selected language-generation backend to generate natural-language explanations. This separation aims to enhance the traceability and to decrease the number of unsupported numerical statements, but the narratives which are generated from the analytical outputs should be checked for consistency.

III. PROPOSED SYSTEM

This proposed system is an enhancement to an existing business intelligence system built using Streamlit, which includes an LLM Narrative Insights Engine. This system integrates the dataset processing, visualization, statistical analysis, predictive modelling, narrative generation and report publishing processes into a unified workflow. Its main objective is to take uploaded datasets and turn them into new knowledge and insights, and present them in natural language explanations.The first step of the workflow is uploading the dataset, followed by dataset profiling, and the ultimate step is the ETL processing of the dataset to prepare it for analysis. The data is then passed on to the visualization, statistics and predictive machine-learning modules. Such modules generate outputs like KPI summaries, statistical test results, feature-importance information, model evaluation metrics, etc. The output is gathered and passed to the narrative engine that will yield explanations based on the chosen language generation backend. The resulting narratives can be shown on the dashboard or output to PDF reports via Enterprise Publisher.

Fig. 1. System architecture of the proposed business intelligence and narrative generation system.



An analytical computation/natural language generation distinction is maintained in the architecture. The analytical modules perform numerical calculations and the proposed narrative engine takes the numerical result of the calculations and turns them into something readable as an explanation. It also enables interactive visualization using 15 chart types, statistical methods like Shapiro–Wilk normality test and predictive models such as Linear Regression, Ridge Regression, Lasso Regression and Random Forest Regression. The narrative generation can be embedded without overwriting the existing analytical elements, using this modular design. The proposed solution is also flexible, with multiple backends for generating a narrative, enabling the system to run on varying hardware and network conditions. It is designed to combine explanations that are generated with dashboard visualizations and PDF reports to make analytical results easier to interpret and communicate with users..

IV. METHODOLOGY

The methodology is a series of steps that converts the raw data into processed data sets, analytical results and natural language narratives. The user uploads a data set to the web app using the Streamlit interface. The system evaluates the structure of the information, the type of data, the missing data and the basic statistical features of the data to get an overview of the information present in the system.

The ETL processing stage involves performing various operations on the data, including data cleaning, handling missing values, imputation, removal of duplicate records, and normalization if applicable, to

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

prepare the data for further analysis. The processed data goes further on to the visualization, statistical analysis and predictive-modelling modules. The visualization component enables the user to explore distributions, trends, comparisons and relationships, and the statistical component runs appropriate tests and provides numerical results, including test statistics and p values.

The predictive-modelling part uses supported machine-learning algorithms, and makes predictions and evaluation metrics as suitable to the selected task. The results of the various activities in the various modules are then gathered in an organised manner. These outputs can be a summary of KPIs, statistics, feature-importance values, or metrics of the predictive-model.

Finally, the structured analytical information is provided to the LLM Narrative Insights Engine. The selected backend creates a human readable description from the given context. The resultant story is fed back to the application and can be displayed with the appropriate visualizations or added to a PDF report. This workflow retains the numerical analysis in the Analytical modules, and relies upon the language-generation component mainly for communicating the results.

Fig. 2. Data flow of the proposed business intelligence and narrative generation system.

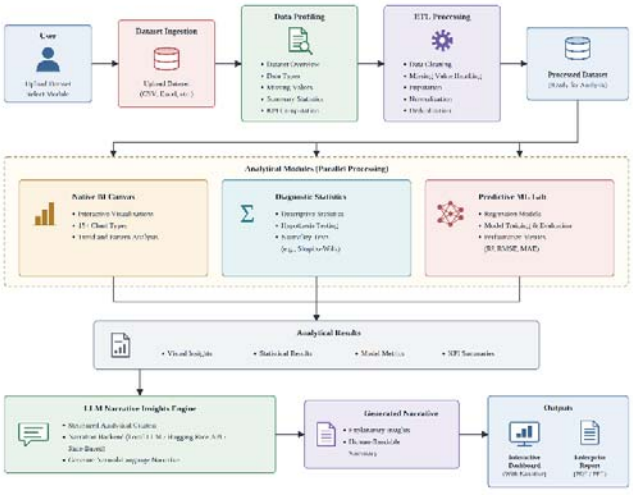


Figure 5.2: Data Flow of the Proposed Business Intelligence and Narrative Generation System

Data flow diagram shows the process flow of the proposed system, which starts from uploading the data sets and ends with the production of analytical narratives and reports. The dataset uploaded gets profiled and processed for ETL.

It is then fed into the data cleaning, missing-value handling, imputation, and normalization blocks before reaching the analytical blocks. The data prepared is processed in the Native BI Canvas, Diagnostic Statistics, and Predictive ML Lab and generates visual insights, statistical results, model metrics, and KPI summaries. The output of these is then passed to the LLM Narrative Insights Engine that produces human-readable explanations using the chosen narration backend. Lastly, the created narrative is displayed on the interactive dashboard or incorporated into the enterprise report, making it easier to understand the analytical results.

V. LLM NARRATIVE INSIGHTS ENGINE

The LLM Narrative Insights Engine is the main enhancement proposed in this project. Most template-based methods produce explanations based on pre-defined sentence structures resulting in common sentences in different data sets and analytical contexts. The proposed engine tries to offer more flexible explanations, basing on the context of the computed analytic results.

The data they deliver to the engine are structured and include summaries of KPIs, results of hypothesis testing, feature importance and metrics for the predictive models. This information is passed to the chosen language-generation back-end that generates a natural-language description of the findings. The language model does not compute the statistics values that underlie it; these values are provided by the analytical modules prior to the generation of the narratives.

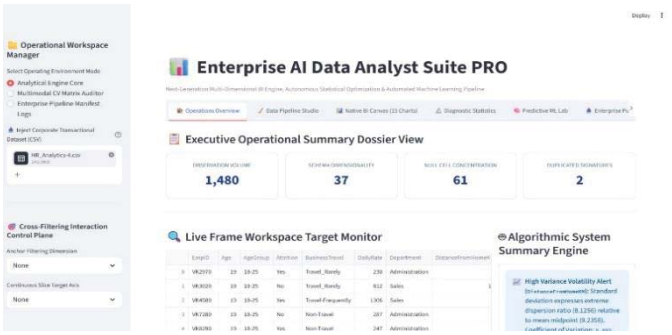
The system has three narration back-ends for versatility in operating conditions. Depending on the computing resources available, the local open-source language model can be used for the generation of the narrative, which is run on the local machine. In the case of internet connectivity and the required credentials, Hugging Face Inference API offers a hosted inference service as an alternative. The rule-based fallback is more rigid in its wording than a language model, but it can be used to provide deterministic explanations when the local model or hosted service isn't available.

Once created, the narrative is sent back to the dashboard and then added to reports created by the Enterprise Publisher. This integration allows users to access charts, numerical results, explanations in the same reporting process. While output that is computed as context can minimize unsupported numerical claims, output generated narratives should still be verified for consistency with the original outputs.

VI. RESULTS AND DISCUSSION

The AI-Powered Narrative Generator is designed to combine data profiling, visualization, statistical analysis, predictive modelling and natural-language explanation in a single BI platform. The interface that is developed provides analytical information in the form of KPI cards, charts, metric evaluation of the models, and summaries produced by AI. These features allow users to analyze the data and make sense of the results beyond a single graph or data point..

Fig. 3. Main Dashboard Showing Dataset Profiling and Operational Summary.

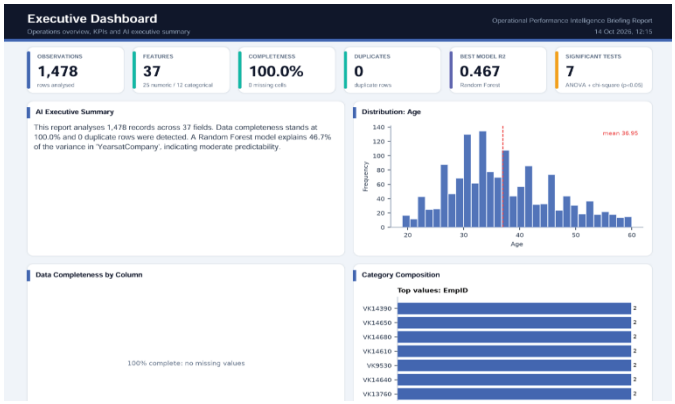


It presents important data including the number of observations, features available, how complete the data is, how many duplicate records there are, and the performance of the predictive models. The dashboard also features visualizations of age distribution and category composition, and an executive summary. This can give an overview of the data set before looking at more detailed analytical data. The dashboard is a visual analysis with explanations in text. In the case of Sthe results from the regression model selected. When this information is displayed with charts, users can see the numerical results and their interpretation at the same time in one user interface. The completeness and duplicate indicators also give users an idea of how good the data being analysed is.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

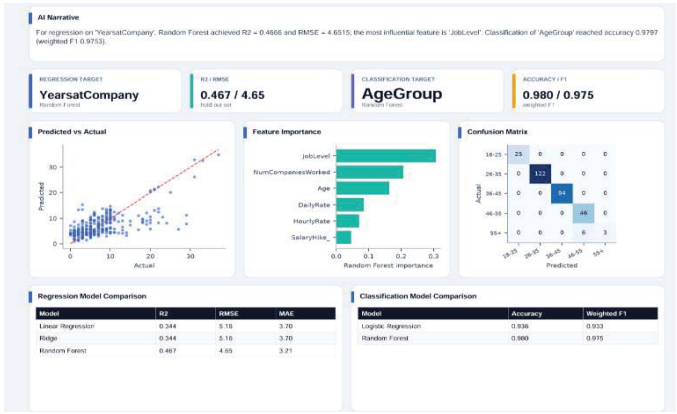
International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Fig. 4. Executive Dashboard Presenting Key Performance Indicators and Data Distributions



The text provides a summary of the selected regression model's R^2 and RMSE values, and concludes that JobLevel is the most important feature included in the analysis displayed. The same view displays results of regression and classification such as model comparison tables, feature importance, predicted versus actual values, and a confusion matrix. These outputs are examples of formatting analytical outputs with a natural language description.The regression comparison shown for the Random Forest model yields a value for R^2 of ~0.467 and RMSE of ~4.65. The accuracy score for the classification comparison is about 0.980 and weighted F1 score is about 0.975 for the Random Forest classifier. The following metrics represent different prediction tasks, and must thus be evaluated independently. The classification results should not be interpreted as a proxy for similar performance on the regression task, in particular.

Fig. 5. AI-Generated Narrative and Predictive Model Evaluation Results.



The comparison of the actual and predicted data, feature importance graph, confusion matrix, and comparison of models have been provided in Fig. 5 for closer look at the predictive analysis. The feature importance plot will present the users with the features that are most important, and the plot of predicted vs actual values will aid the users in studying the correlation between the predicted and actual values and input variables to the displayed Random Forest model contribute the most to the model. Confusion matrix shows the confusion for every class in the classification. These visualizations, when combined, provide users with further evidence to aid in understanding the model results.

VII. CONCLUSION

This project proposes a new method for creating explanations in natural language from business intelligence dashboards using artificial intelligence. Data Profiling, ETL Processing, Interactive Visualization, Statistical Analysis, Predictive Modelling and Report Generation are all integrated into a single system. The LLM Narrative Insights Engine takes the next step by generating user-friendly explanations from the analytical data computed by the LLM:The dashboard and report screenshots show the use of key performance indicators, visualizations, predictive metrics and generated narratives all together. The system also has multiple narration back-ends such as a language model locally hosted, a hosted inference API, and a rule-based fallback. This allows flexibility in computing and connectivity scenarios. One of the main design principles is that analytical computation is separated from the generation of narratives. The numerical result is generated by statistical and machine-learning modules while the language-generation component explains the given information. This process will help to enhance the accessibility and traceability of the results of analysis. The outputs demonstrated are indicative of the proposed workflow but additional systematic testing should be conducted to ensure that the generated narratives are factually coherent, efficient and user-friendly.

VIII. FUTURE SCOPE

There are a number of modifications and extensions that can be made to the proposed system in order to enhance the quality, reliability, and usability of automated data storytelling. A key way forward is to compare the narrative generation engine in a broader spectrum of datasets and analysis scenarios. For example, this testing might involve evaluating the accuracy of the explanations generated and checking for consistency as the data set and/or analysis conditions vary.Another improvement is a new narrative generation that is interactive. Users might want to ask further questions about a chart, want a more simple explanation, or look for the "why" behind a certain trend. This would enable the dashboard to be more appealing to more users with varying analytical skill levels.

This system could be further improved by better factual verification. The generated statements may be reviewed against original metrics and statistical outputs, prior to presentation or incorporation into a report. This would help to disambiguate unsupported claims, and make the narratives more confident. Future work might include a comparison of the local language model, hosted inference API, and rule-based fallback based on response time, model size, narrative clarity, and factual consistency of the results. These comparisons would then enable to decide the most appropriate backend for various deployment scenarios. Last but not least, the system could be expanded with more sources of data, more sophisticated analytical techniques, and domain-specific narrative forms like finance, healthcare, sales, and human-resource analytics. Users studies also could be carried out to determine if generated narratives add value to visualizations for interpreting dashboards better.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

REFERENCES

[1] C. Zhang, C. Li, and S. Gao, "MDSF: Context-Aware Multi-Dimensional Data Storytelling Framework Based on Large Language Model," *arXiv preprint arXiv:2501.01014*, 2025.

[2] M. S. Islam, M. T. R. Laskar, M. R. Parvez, E. Hoque, and S. Joty, "DataNarrative: Automated Data-Driven Storytelling with Visualizations and Texts," *arXiv preprint arXiv:2408.05346*, 2024.

[3] T. Niu, Z. Han, Q. Chen, S. Zhou, B. Shan, H. Fang, Q. Zhu, and W. Che, "Dashboard2Code: Evaluating Multimodal Models on Reconstructing Interactive Dashboards," *arXiv preprint arXiv:2607.04727*, 2026.

[4] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[5] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why Should I Trust You?': Explaining the Predictions of Any Classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.

[6] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2017, pp. 618–626.

[7] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[8] M. Staniak and P. Biecek, "The Landscape of R Packages for Automated Exploratory Data Analysis," *arXiv preprint arXiv:1904.02101*, 2019.

[9] S. S. Shapiro and M. B. Wilk, "An Analysis of Variance Test for Normality (Complete Samples)," *Biometrika*, vol. 52, no. 3/4, pp. 591–611, 1965.

[10] P. Vassiliadis, P. Marcel, F. El Outa, V. Peralta, and D. Gkitsakis, "A Conceptual Model for Data Storytelling Highlights in Business Intelligence Environments," arXiv preprint arXiv:2403.00981, 2024.

[11] R. Puduppully, L. Dong, and M. Lapata, "Data-to-Text Generation with Content Selection and Planning," in *Proc. AAAI Conf. Artificial Intelligence*, vol. 33, no. 1, pp. 6908–6915, 2019.

[12] E. Reiter and R. Dale, "Building Applied Natural Language Generation Systems," *Natural Language Engineering*, vol. 3, no. 1, pp. 57–87, 1997.

[13] N. H. Riche, C. Hurter, N. Diakopoulos, and S. Carpendale, *Data-Driven Storytelling*. Boca Raton, FL, USA: CRC Press, 2018.

[14] D. Ren, M. Brehmer, B. Lee, E. K. Choe, and H. Höllerer, "ChartAccent: Annotation for Data-Driven Storytelling," in *Proc. IEEE Pacific Visualization Symposium (PacificVis)*, 2017, pp. 230–239.

[15] A. Lundgard and A. Satyanarayan, "Accessible Visualization via Natural Language Descriptions: A Four-Level Model of Semantic Content," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 1073–1083, 2022.

[16] C. Dong, Y. Li, H. Gong, M. Chen, J. Li, Y. Shen, and M. Yang, "A Survey of Natural Language Generation," *ACM Computing Surveys*, 2023.

[17] S. Obaid ul Islam, I. Škrjanec, O. Dušek, and V. Demberg, "Tackling Hallucinations in Neural Chart Summarization," in *Proc. 16th International Conference on Natural Language Generation (INLG)*, 2023, pp. 414–423.

[18] A. Mishra, A. Laha, K. Sankaranarayanan, P. Jain, and S. Krishnan, "Storytelling from Structured Data and Knowledge Graphs: An NLG Perspective," in *Proc. 57th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, 2019.

[19] A. Srinivasan, S. M. Drucker, A. Endert, and J. Stasko, "Augmenting Visualizations with Interactive Data Facts to Facilitate Interpretation and Communication," IEEE Transactions on Visualization and Computer Graphics, vol. 25, no. 1, pp. 672–681, 2019.

[20] E. Hoque and M. S. Islam, "Natural Language Generation for Visualizations: State of the Art, Challenges and Future Directions," Computer Graphics Forum, vol. 44, no. 1, Art. no. e15266, 2025, doi: 10.1111/cgf.15266.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)