GeoExplainer: A Visual Analytics Framework for Spatial Modeling Contextualization and Report Generation

Fan Lei, Yuxin Ma, *Senior Member, IEEE*, A. Stewart Fotheringham, Elizabeth A. Mack, Ziqi Li, Mehak Sachdeva, Sarah Bardin, and Ross Maciejewski, *Senior Member, IEEE*



Fig. 1: The model analysis and report authoring interface. The model validation module (A) contains: the model diagnostic panel (A1); a local coefficient list (A2), and; the numerical distribution panel (A3) showing the distribution of the selected coefficient of *pct_age_18_29*. The contextualized exploration module (B) contains: the local parameter explanation view (B1); the map view (B2) where a choropleth map presents the spatial distribution of a selected coefficient or diagnostic value; the external information module, including an external keyphrase list (B3) which retrieves the most important keyphrases from the external context related to the selected range of regions, and; a list of the original external content (B4) that consists of paragraphs of relevant external information related to a selected phrase. The report authoring module (C) supports synthesizing the findings into a report.

Abstract—Geographic regression models of various descriptions are often applied to identify patterns and anomalies in the determinants of spatially distributed observations. These types of analyses focus on answering why questions about underlying spatial phenomena, e.g., why is crime higher in this locale, why do children in one school district outperform those in another, etc.? Answers to these questions require explanations of the model structure, the choice of parameters, and contextualization of the findings with respect to their geographic context. This is particularly true for local forms of regression models which are focused on the role of locational context in determining human behavior. In this paper, we present GeoExplainer, a visual analytics framework designed to support analysts in creating explanative documentation that summarizes and contextualizes their spatial analyses. As analysts create their spatial models, our framework flags potential issues with model parameter selections, utilizes template-based text generation to summarize model outputs, and links with external knowledge repositories to provide annotations that help to explain the model results. As analysts explore the model results, all visualizations and annotations can be captured in an interactive report generation widget. We demonstrate our framework using a case study modeling the determinants of voting in the 2016 US Presidential Election.

Index Terms—Spatial data analysis, local models, multiscale geographically weighted regression, model explanation, visual analytics

- F. Lei, A.S. Fotheringham, M. Sachdeva, S. Bardin, and R. Maciejewski are with Arizona State University. E-mail:
- {flei5,sfotheri,msachde1,sfbardin,rmacieje}@asu.edu.
 Y. Ma is with Southern University of Science and Technology. E-mail:
- mayx@sustech.edu.cn. Y. Ma is the corresponding author.
- E.A. Mack is with Michigan State University. E-mail: emack@msu.edu.
- Z. Li is with Florida State University. E-mail: liziqi1992@gmail.com.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on

1 INTRODUCTION

Geographic models [18, 20, 53] are often used to answer *why* questions regarding underlying spatial phenomena such as why do voter preferences generally exhibit strong spatial dependency. In order to answer such questions, researchers have developed sophisticated spatial modeling techniques and software, such as multiscale geographi-

obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx cally weighted regression (MGWR) [20, 53], to identify the determinants of the spatial patterns of data we observe in both the human and natural environments. However, these models generate output that needs local contextual information to be interpreted properly. Without such information, advanced spatial modeling techniques do not gain their full potential. While a variety of geovisual analytics systems [15, 16, 25, 36, 39, 47] have been developed to explore spatial statistics with the help of interactive maps and narrative annotations, these applications tend to visualize only basic statistical results, and, to our knowledge, none of them explain spatial modeling results.

In this paper, we present GeoExplainer (Figure 1), a visual analytics framework designed to support spatial data modeling, analysis and reporting. Our approach is inspired by research in explainable machine learning [32, 33] and a combination of narrative visualization [62] and storytelling techniques [30]. GeoExplainer provides multiple types of explanation support throughout the spatial analysis pipeline. In the model calibration stage, our framework interprets the functionality of the spatial model configuration and recommends parameter settings. Then, the framework summarizes and explains model outputs by adopting template-based text annotations, linking with external knowledge repositories to provide relevant contextual information. All visualizations, model results, and annotations can be captured in an interactive report authoring widget, enabling analysts to generate documentation that explains their spatial analyses. Our contributions include:

- A novel narrative explanation workflow that supports a dynamic interplay between automatically generated text interpretations and relevant visualizations.
- An automatic spatial cluster detection function that suggests interesting patterns and outliers from the local spatial modeling result.
- The seamless integration of an interactive report authoring tool with the spatial analysis pipeline.
- A web-based visual analytics framework with analytical process recording and sharing functions to facilitate collaboration.

2 RELATED WORK

Our work focuses on facilitating spatial analysis by explaining the outputs of spatial modeling with contextual information through narrative visualization. In this section, we review related work on geographic analysis, narrative visualization, and model explainability.

2.1 Geographic Analysis

Various spatial prediction models [28, 38, 41, 71] and spatial data analysis tools [25, 66] have been developed to support geographic analysis. In this work, we focus on two local spatial models widely used in spatial analysis: Geographically Weighted Regression (GWR) [7], and its recent extension, Multiscale Geographically Weighted Regression (MGWR) [20]. GWR extends the classical linear regression model [31] by capturing spatial heterogeneity with influence spreading over the space in a constant scale. MGWR further improves GWR models where local influences are modeled in different spatial windows.

Support for a variety of spatial models (including GWR and MGWR) is integrated into the most widely used geographic information systems (GIS), e.g. ArcGIS [60], QGIS [23] and GeoDa Web [39], and embedded by libraries (e.g., PySAL [57]) in computational notebook environments such as Jupyter Notebook [53, 54, 58] or R Markdown [44]. However, these systems and notebooks do not support integrated external knowledge sources to help contextualize models. They also require analysts to judiciously choose the proper tools at every stage of the analytical pipeline to build a well-trained model. This process necessitates a solid foundation in geographical concepts and can be time-consuming, even for seasoned analysts. Our goal is to enhance the spatial modeling process through a deliberately designed user-friendly workflow and automatically generated explanatory narratives guided by domain experts to promote the understanding of local contextual information from models. Our work focuses on improving model explainability, enabling interactive report authoring, and supporting contextualization through narrative generation. This is important in local models such as

GWR and MGWR because the main output from such models is a set of local parameter estimates from each process being modeled, and the spatial variation in these estimates needs to be explained in terms of the contextualized environment of each location by providing additional contextual information for users.

2.2 Narratives and Annotations

The design space of narrative visualization [26,62,68] has been widely explored in the visualization community and a diverse set of storytelling and annotation methods have been developed to reveal observations in data and to convey key messages to an audience. Such narrative visualization techniques have been adopted by a variety of geovisual analytics applications [15,47,59] that have integrated story authoring tools with spatial data visualizations. For example, Bespoke Map [6] discusses design implications for map customization tools. NewsViews [21] generates interactive maps with narrative annotations automatically from a given news article. Latif and Beck [36] introduce a bivariate map design that integrates template-based text annotations, and they later extend their work to investigate the interplay of text and visualization in geographic storytelling [37].

Many of these systems utilize annotations to enhance the visual narratives, and Kosara and Mackinlay [30] emphasize the importance of annotations for facilitating storytelling in visualization. Recent annotation work [10, 11, 56, 59] has focused on the semi-automatic creation of presentation-like storytelling visualizations and explored mechanisms for automatically generating annotations by integrating deep-learning feature extraction techniques with a natural language generation process [35, 42]. Annotations in such systems are used to describe salient patterns and facilitate storytelling [9, 12, 14, 64].

In addition to using annotations to explain data features, other works have also explored mechanisms for supplementing narratives from external data corpora. For example, Contextifier [27] and Causeworks [12] leverage relevant external information from news articles or Wikipedia to generate text annotations. Such external information provides audiences with additional context to facilitate the analysis.

While advanced storytelling techniques are being readily adopted to develop narrative reports for geovisualization, the focus has been on supporting simple statistical analysis (e.g., find the maximum, minimum). However, modeling and analysis processes have increased well beyond descriptive statistics, and there is a growing need to expand the narrative generation process to support more sophisticated spatial models, such as MGWR [20,53]. Our work builds upon automatic narrative generations, and contextualization of spatial modeling. We provide templates for text generation for model explainability and develop interactive authoring tools for report generation.

2.3 Model Explainability

Narrative visualization is directly related to the concept of explainability, where the visualization authors seek to couple images and text to explain an underlying data analysis. With respect to model explainability, the visual analytics community has developed a variety of systems to support the interactive explanation of machine learning models (e.g. [4,43,45,46]). Several model-independent approaches, e.g. EnsembleMatrix [67] and RuleMatrix [49], focus on the classifier's input-output behaviors to provide insight into the model classification results. EnsembleMatrix provides a visual summary of the model outputs, RuleMatrix uses a matrix-based visualization to explain classification results, and Prospector [34] explores the relationship between feature values and predictions by using partial dependence diagnostics.

Most closely related to our work are the techniques that explain models from feature-level observation. Mühlbacher and Piringer [50] facilitated feature selection and optimization in regression models by partitioning the feature space into disjoint regions for visualization. Sedlmair et al. [61] proposed an abstract conceptual framework to discuss the visual parameter space analysis problems independent of the application domain. Goodwin et al. [24] extended visual parameter space analysis to the spatial domain, enabling the exploration of correlations between multiple variables that vary geographically at different spatial scales. While these techniques focus on explanations for domain experts, our work is designed to support explanations to experts and support their use of external information to contextualize these relationships and communicate their findings to a general audience. Our choice of text templates, as opposed to large language models, for narrative generation is to support control for reliability and reproducibility. All text generated must have a verifiable source and the generation of text for the analysis should always return the same results to ensure that the resulting analyses are not subject to misinformation.

3 DESIGN OVERVIEW

Our goal is to implement an interactive visual analytical framework with narrative contextualization techniques and report authoring to support the spatial analysis pipeline.

Design study with domain experts. An iterative design process was employed to develop GeoExplainer in collaboration with two domain experts in quantitative geography. The first possesses over thirty years of experience and is renowned in spatial modeling and local statistical analysis. The second expert has over fifteen years of experience in geographic statistical analysis and applied geographic science. Our two-phase design process began with identifying spatial data analysis challenges and establishing initial design tasks through several one-onone semi-structured interviews and group discussions, leading to the development of a conceptual demo. The second phase involved refining these design tasks and improving the framework's functionality and visual design through bi-weekly meetings for a year until all issues were resolved. This led to the identification of five design challenges (C1 - C5) and the formation of a set of analytical tasks (T1 - T5), with C1 demonstrating an architecture-level challenge and C2 to C5 pertain to individual stages of the spatial data analysis pipeline.

C1: Flexible analytical pipeline. Although comprehensive spatial analysis frameworks (e.g., ArcGIS [60], GeoDa [2], PySAL [57]) support sophisticated spatial modeling, they need analysts to select a proper combination among the models and analytical tools to correctly perform the analysis. The analysts have to understand the key concepts of geographical and statistical to perform their own ESDA process [19]. The challenge is making spatial modeling more accessible with a framework integrating a standard ESDA pipeline with dynamic interpretations to support the understanding and tools to enhance the modeling result communication. And this unified framework does not require complex installation, configuration, and coding. Design tasks for this challenge focus on implementing a framework to support a unified spatial analytical pipeline that is accessible to the analysts.

- **T1.1:** Develop an analytical pipeline that unifies ESDA, narrative explanations, and report authoring functionalities.
- **T1.2:** Support coordinated multiple views, each of which can populate information into the report authoring mechanics.
- **T1.3:** Instantiate the framework within modern web browsers to enable broader usage.

C2: Model configuration and validation support. Configuring a modern spatial model, such as MGWR, with properly selected parameters, and validating performance features are fundamental prerequisites for the quality of the model output [55, 63]. However, strong prerequisite knowledge is needed to utilize relevant geostatistical tools and understand their results [53]. The challenge is to develop a framework that can guide analysts to appropriate model configurations and support the model validation process. Design tasks include:

- **T2.1:** Support the interactive selection of dependent variables and model parameters.
- T2.2: Generate interpretable parameter recommendations.
- **T2.3:** Visualize the model diagnostics and parameter estimates associated with each covariate.

C3: Contextualizing model outputs. Once a spatial model is fit to the data, the parameter estimates and diagnostics are often produced in the form of data tables (e.g., comma-separated values files, JSON

files, etc.) [53]. These numerical and statistical records in the table contain heterogeneous local contextual information, which is hard to understand without background information [19, 53]. Most extant tools only use maps to visualize the spatial surfaces that result from local modeling without the ability to bring in further contextual information. The challenge is how to automatically provide relevant contextual information to support explanations in the spatial analysis pipeline. Design tasks include:

- T3.1: Contextualize the model results.
- T3.2: Automatically identify interesting geographical patterns.
- **T3.3:** Explain the local parameter estimates in relation to the identified spatial patterns.
- **T3.4:** Link external contextual information associated with the identified spatial patterns or analyst selections.

C4: Communicating findings. The goal of spatial analysis is to answer questions about spatial phenomena and communicate the findings to a broad audience [9]. However, the multifaceted analysis outcomes from advanced analytical tools are often difficult to interpret for a general audience (C3). The challenge is how to support the analysts in communicating the outputs of their spatial models. Design tasks for this challenge focus on storytelling and report authoring.

- **T4.1:** Highlight, annotate, and organize the analysts' findings in the form of text paragraphs and images.
- T4.2: Transform the findings into reports for a broad audience.

C5: Collaboration and replicability. Analysts often need to save, share, or revert their analytical state during the model development process. The challenge is to develop a asynchronized spatial analysis framework that enables analysts to save, retrieve, and share their configured model, explanations, and mid-state analytics workflow [17]. Design tasks include:

- **T5.1:** Save the analytical state at any point during the pipeline, including the configured model, explored information, and the report.
- T5.2: Retrieve the stored state and continue the analysis.
- T5.3: Share the stored state with other analysts.

4 GEOEXPLAINER

Based on the design overview, we developed a visual analytics framework, GeoExplainer, to support spatial analysis across model creation, validation, contextualization, and reporting.

Framework Pipeline: The GeoExplainer framework has a unified pipeline (Figure 2) that enhances the exploratory spatial data analysis (ESDA) pipeline by adding contextualized interpretations and report synthesis functions (T1.1). The GeoExplainer pipeline is composed of five functional units: model configuration (Sect. 4.1), model validation (Sect. 4.2), narrative generation (Sect. 4.3), contextualized report authoring (Sect. 4.4), and analytical state recording (Sect. 4.5). These functional units are spread across two coordinated multiple view interfaces. Figure 3 illustrates the model configuration interface, and Figure 1 shows the model validation and contextualized report authoring interface. Underlying both of these are narrative generation and analytical state recording tools. The narrative generation tools assist the audience in accessing the information of each module in the ESDA workflow, while the state recording module further supports analysts in saving, loading, and sharing analytical states throughout the pipeline. In each functional unit, we link the dynamic graphics and narrative components with coordinated multiple views (T1.2) to support ESDA [1,2,39].

Web-based, Open-sourced Implementation: Analysts only need a web browser to utilize our framework (**T1.3**) as GeoExplainer is built upon a client-server architecture with a RESTful API between the front-end interface and back-end services. The back-end server is implemented with Python Flask, while the front-end interface uses React [48], Mapbox GL JS [29], and D3.js [5].



Fig. 2: The GeoExplainer pipeline is designed to support the exploratory spatial data analysis (ESDA) workflow with Narrative Generation and Contextualized Report Authoring modules. The analytical state recording functions support asynchronous analysis and collaboration.

The visual design and interactions are robust to various types of spatially referenced data and external knowledge sources for contextualization. Our target audience is spatial data analysts, as our goal is to help them save time and effort in spatial analysis while supporting their improved understanding and communication of the model results. For demonstration purposes, we focus on spatial analyses using (MGWR), and a discussion on the generalization of this work is provided in Sect.6.

4.1 Model Configuration

The model configuration stage acts as an entry point for analysts to inspect data distributions, identify informative features, and find optimal model parameters before a spatial model is trained. Illustrated in Figure 3, the model configuration interface consists of three coordinated views: the **feature view**, the **correlation view**, and the **map view**.

Feature View: The feature view, Figure 3 (A), provides a summary of data distributions and evaluation results for features as well as the model parameter settings. It consists of five panels (A1-A5): the model configuration panel (A1), the dependent variable selection panel (A2), the independent variables selection panel (A3), the feature list (A4), and the detailed feature analysis view (A5).

In the *model configuration panel*, Figure 3 (A1), analysts select the configuration properties for MGWR, an adaptive bisquare spatial kernel is the default setting. For a detailed explanation of MGWR parameters, we refer the reader to the MGWR 2.2 Manual [53].

In the *feature list*, Figure 3 (A4), each row depicts the marginal data distribution of a feature as a histogram. The features can be dragged into the *dependent variable selection panel* (A2) or the *independent variable selection panel* (A3) to configure the model. An analyst starts the exploration of feature relationships by choosing a dependent variable. Then, our framework automatically performs multicollinearity detection using the variance inflation factor(VIF) to estimate how much the variance of a regression coefficient is inflated due to multicollinearity. For each variable in the independent variable selection panel, the corresponding VIF value between the independent and the selected dependent variable is displayed next to the data distribution histogram. A VIF of 1 indicates no correlation between the two variables, while a VIF>10 is a strong indication of a potentially severe multicollinearity effect [22, 70]. To address such severity, green or yellow backgrounds are used to identify VIF values that are smaller or larger than 10.

To inspect the details of a feature, the analyst can click on the "detail button" () at the end of the corresponding row in the feature list (A2-A4) to open the *detailed feature analysis view*. This view provides an annotated histogram with details on normality and skewness tests as well as an autogenerated text annotation illustrating the data distribution. To assist analysts, the text also includes interactive suggestions on appropriate transformations to apply to the data. Analysts can click the Log transformation button corresponding to the suggestions to see the transformed distribution. The Kolmogorov-Smirnov (K-S) test is used to check each feature's normality. When selecting features, the map view updates to show the spatial distribution of the selected variable. **Correlation View:** The correlation view, Figure 3 (B), is designed for analyzing relationships between the variables. The view consists of two interactive panels: a scatterplot view (B1) and a scatterplot matrix (B2). Each view contains automatically generated text annotations to highlight bivariate linear correlations based on Pearson correlation tests, which prompts analysts to consider removing redundant features.

The *scatterplot view* shows the correlation between selected variables, Figure 3 (B1). Analysts can click on the "scatterplot button" ([: . .) on a row in the feature view to render the corresponding scatterplot view.

The *scatterplot matrix* depicts the correlations among multiple selected independent variables, Figure 3 (B2). In addition to the textual suggestions, we annotate the points that show strong linear correlations in **orange** in the scatterplots to help analysts identify potential independent variables. Analysts can click on these scatterplots to open the detailed distribution information, Figure 3 (B3).

Map View: The map view, Figure 3 (C), contains two types of choropleth map designs that link with the operations in the feature view and the correlation view. When clicking on a feature's "detail button" () in the feature view, the univariate choropleth map (C1) is rendered to show its spatial distribution with a sequential blue color scheme in a quantile scale. Clicking on the "scatterplot button" (.) on a row in the feature view will update the map to the bivariate choropleth map mode (C2) to display the spatial distributions of the dependent variable and the selected independent variable as well as their degree of correlation. Our bivariate choropleth design is based on Trumbo's diagonal model [65]. As shown in Figure 3 (C2), the bivariate choropleth map uses a sequential gray color scheme along the main diagonal to show the progression of correlated variables. Combining this with the method from Nusrat et al. [52], we use gradients of blue color (above the diagonal) to depict the regions where the dependent variable values are sufficiently larger than the selected feature, while the gradients of red color (under the diagonal) are used to represent the regions where the dependent variable values are sufficiently smaller.

All views and text are designed to support novices and experts alike. The text generation templates, designed in conjunction with experts in quantitative geography, provide entry level explanations to guide novices to appropriate parameter selections. Underlying all views are visual enhancements, typically colors, that highlight important features within the data. The underlying combination of text and visual representations serves as a novel guidance mechanism for ESDA.

4.2 Model Validation

Once satisfied with model configuration, the analyst can click on the "Train Model" button in Figure 3 (A1). After the spatial model is trained, the model analysis interface (Figure 1) is enabled, which equips a suite of views for model validation (Figure 1(A)), contextualized exploration (Figure 1(B)) and report authoring (Figure 1(C)).

The model validation module (Figure 1 (A)) (**T2.3**) consists of three views: a model diagnostic panel (A1), a list of local coefficient estimates (A2), and a numerical distribution view (A3). The model



Fig. 3: The model configuration interface. In the feature view (A), the *model configuration panel* (A1) provides the configuration properties. The *dependent variable selection panel* (A2) and *independent variable selection panel* (A3) use the features that are dragged from the *feature list* (A4), where each row is associated with a feature in the dataset with the marginal data distribution of the feature depicted as a histogram. The *detailed feature analysis view* (A5) displays the histogram distribution and narrative information suggesting appropriate transformations to apply to the dependent variable. (B) The correlation view (B1) shows the correlation between the dependent variable and a selected independent variable. The *scatterplot matrix* (B2) shows the correlations among multiple selected independent variables. The detailed distribution information with Pearson coefficients and text annotations (B3) is populated by clicking on the scatterplots in the matrix. (C) The map populates a *univariable choropleth map* (C1) to show the spatial distribution of a selected variable. The map can also update a *bivariate choropleth map* (C2) to display the spatial distributions and the degree of correlation between the dependent variable and a selected independent variable.

diagnostic panel contains detailed information of the global and local diagnostic indicators. Analysts can validate the numerical and spatial distributions of either a local diagnostic indicator or a local coefficient estimate by selecting a row in the corresponding list. By clicking on the "numerical distribution button" ([iii]) in the list, the histogram of the selected row can be displayed in the numerical distribution view, where the statistical values of the distribution are depicted in the boxplot under the histogram. Meanwhile, the map view shows the coordinated spatial distribution of the selected estimates, Figure 1 (B2). The choropleth colors map to a given region's coefficient or diagnostic value.

We use the corrected Akaike Information Criterion (AICc), the R^2 , and the adjusted R^2 as our global diagnostic indicators, and local R^2 , Cook's distance, and the standardized residuals as our local indicators to evaluate the overall model performance. Analysts can also explore how the model predicts the relationship between the dependent variable and an independent variable by selecting the corresponding coefficient estimate. The local intercept validates the intrinsic properties of the dependent variable when all the other properties in the model remain constant [19]. To enhance the accessibility of the meaning and context of the above indicators, the framework explains their definitions and the functionalities in the spatial data analysis through pop-over windows. Analysts can click an indicator's name to read the explanation.

Global Diagnostic Indicators:

The corrected *Akaike Information Criterion (AICc)* is an estimator of prediction error [8] defined as $AICc = AIC + \frac{2k(k+1)}{n-k-1}$, where $AIC = 2k - 2logl(\hat{\theta})$, *n* is the sample size and *k* is the number of parameters. *AICc* contains a penalty for model complexity and is a good indicator of model performance when comparing local and global models.

 R^2 and adjusted R^2 are used to evaluate how well the model's dependent variable can be explained by the independent variable(s). They

represent the proportion of the variance for the dependent variable that is explained by the independent variable(s) in a regression model.

Local Diagnostic Indicators:

The *local* R^2 indicates how well the dependent variable is explained by the independent variables. For the choropleth map, we encode the *local* R^2 value for each region using a quantile **blue** color scale with single-hue gradients where darker colors indicate higher values.

Cook's distance (Cook's D) is used to detect influential outliers in the model. Cook's distance D_i of the observation $i \in \{1, ..., n\}$ is defined as the summation of all the changes in the regression model when observation (region) *i* is removed from the model [13]. The distribution of the *Cook's D* is typically a heavy tail distribution, and we apply a quantile single-hue **blue** exponential scale for the choropleth encoding. We only render the regions that are identified as outliers on the map.

The *standardized residuals* help identify if the model is well-specified or if key explanatory variables are missing. The residual values represent the over- and underpredictions for each region for the model. We obtain the residual values by subtracting the fitted dependent values from the predicted dependent values. The residual value in overpredicted/underpredicted areas is greater/less than zero respectively. For the *standardized residuals*, the over-predicted areas in the choropleth map are shaded using a quantile single-hue **blue** color map, and underpredicted areas are shaded using a quantile single-hue **red**.

Local Coefficient List: The local coefficient list contains the estimated local coefficients and local intercepts of the model.

The *local coefficients* provide information about how each independent variable influences the dependent variable in different regions. In a geographical model, every observation (region) has coefficients (β) corresponding to each independent variable. The values of these coefficients represent the strength and the type of relationship that the

🛛 Target	Object—	→ 🖥 Describe Patterns —	ightarrowExplain Results	
Data distr	ibution	{ Normal / skewed } distribution	Recommend a proper model type or data transformation	(A)
Correlatio	on	{ Significant / insignificant } { strong / moderate / weak } { positive / negative } global linear relationship	Independent varieble selection suggestions	for Model Configuration Components
Multicollin detection	nearity	Variables have high multicollinearities	Operation suggestions	
Local	1 Local R ²	Areas with { high / low } local R ² values, based on user-defined threshold	Model performance {decent / poor} regarding to local R ²	
Diagnostic Indicators	2 Cook's D	Areas with { values $\ge 4/N$, N = # of all areas }	Meaning of Cook's D and operation suggestions	(B) Narrative Generation for Model Validation Components
	3 Standard residual	{ Over-predicted / under-predicted } areas	Meaning of the residual with Moran's I results and suggestions	
4 Local inde variable co	pendent [*]	Areas have significant { positive / negative } correlations with dependent variable,based on T-test	The strength and the type of the correlation with statistical context + spatial clustering explanation	

*The pattern description of the local coefficients contains detected geographical clusters.

Fig. 4: The narrative generation pipeline for the model configuration and model validation components.

independent variables have with the dependent variable.

The *intercepts* represent the expected value for the dependent variable if all the independent variables are zero. Given that all the variables are standardized (0,1) prior to calibration by MGWR, the local intercept estimates indicate the value of *y* that would be obtained in each location if each location contained a population with average characteristics.

For the local coefficients and intercept, the framework first performs a t-test to filter out the insignificant coefficients. Then we adopt a quantile single-hue **blue** scale to depict the significant positive relationships (positive coefficient values) between the independent variables and dependent variable, and a quantile single-hue **red** scale is applied to encode the significant negative relationships (negative coefficient values) in the choropleth map. **Gray** is used to indicate insignificance.

4.3 Narrative Generation

Underlying the model configuration and validation functions is our narrative generation pipeline. We build upon previous work from Latif and Beck [36] where text template narratives were instantiated to describe bivariate geographic data. However, these templates focused on maximum, minimum, and correlation discussions. In this work, we expand their concept to adapt to the generation of explanatory text for the outputs from spatial modeling processes (**T3.1 - T3.3**), and to establish a link to external contextual information (**T3.4**).

Template-based Narrative Generation: Based on the collaborations with domain experts and the guidelines learned from the geographic data-driven story design practices [37], we design a novel narrative generation pipeline that provides template-based explanations for the model configuration and validation functions (Figure 4). The templates are generated following a three-step process and mainly consist of two phrases: 1) the pattern description, and; 2) the result explanation and suggestion containing contextual information. Figure A.1 in Appendix A systematically illustrates the narrative generation pipeline and each text template with a corresponding example.

Narratives generation for model configuration components: The pattern description phrase describes the feature analysis results based on specific metrics in different components, while the result explanation phrase generates suggestions on appropriate operations to configure the model or select the parameters (Figure 4 (A)). Two views are dynamically linked with the narrative templates. In the "detailed feature analysis view" (Figure 3 (A5)), a description of the normality test results related to the data distribution of the independent variable is generated. This is accompanied by recommendations for data transformations. The "correlation view" (Figure 3 (B)) describes the correlation test results. This is accompanied by suggestions to prompt analysts to consider removing potentially redundant variables from the list of independent variables in the model.

Narratives generation for model validation components: The model validation components contain a suite of interactive views that are closely

linked with the map, Figure 1 (B). The local parameter explanation view, Figure 1 (B1), interprets model outputs including identified patterns, data distributions, and relevant statistical contexts. The external information view, Figure 1 (B3, B4), retrieves external context related to the selected spatial cluster or area in the map. Just as selections in the model's local diagnostic and coefficient list, Figure 1 (A1, A2), populate the map view (Figure 1 (B2)), they also populate the narrative templates in the local parameter explanation view. Analysts can click each parameter's "show explanation button" (2) to toggle the explanation panels. The local coefficients and three local diagnostic indicators distilled by domain experts that are commonly used in diagnosing local spatial model performance are explained by the template-based narrative generation. The two phrases of the template describe the multivariate spatial distributions and incorporate explanations containing contextual information (Figure 4 (B)). Here, the spatial distribution is based on the detected geographical patterns and extracted range of values. The generated narrative paragraphs are interactively linked with the map view to enhance comprehension.

Narratives Generation for Spatial Distribution: One of the main parts of the text template for describing a selected local parameter is the spatial distribution interpretation. According to the range of values for each local parameter, we classify the parameter data into groups (e.g. the areas that have positive/negative local coefficients). We then populate paragraphs along classified groups to describe the spatial distribution. Each paragraph contains two parts: 1) a user-definable location identifier that indicates the spatial locations of the current data group. Analysts are encouraged to use vernacular geographic language (e.g. downtown Chicago, west coast of the U.S.) [37] to edit the identified location description to facilitate the communication between the geographic entities and target audience, and; 2) a narrative description of the classified value range in this location. We use consistent colors between the map and text to highlight the linkage between spatial narrative and the map view. Hovering the mouse on each paragraph will filter out irrelevant locations on the map.

Local Parameter Explanation View – Local Diagnostics: While our narrative generation process follows the same steps and layouts for populating our templates, various model parameters have unique templates to support their explanation, each of which populates the model explanation view. The templates of the model's local diagnostic indicators are introduced below (**T3.1**):

local R^2 : The framework generates the spatial distribution narratives by describing the number and location of places with high or low local R^2 values, Figure 4 (B.1). Since the validation of local R^2 is highly subjective and depends on the specific task [70], we use a user-definable threshold to denote high and low local R^2 values. We then generate text to explain the meaning and results of different R^2 values.

Cook's D: We identify outliers and describe their spatial distributions using Cook's D values. We also describe the range, threshold, and meaning of the Cook's D, Figure 4 (B.2).

Standardized residual: We describe regions on the map that are overpredicted/under-predicted. The template includes spatial autocorrelation information using Moran's I, Figure 4 (B.3).

Local Parameter Explanation View – Local Coefficients: The local estimates of the coefficients from a local model calibration may contain unique spatial patterns such as outliers and regions with strong spatial autocorrelation. Interpreting these patterns for local coefficients is a key step for understanding local modeling results. In GeoExplainer, underlying the standard narrative generation pipeline in Figure 4 (B.4), we implement an additional set of processes to generate narrative templates to explain the local coefficients and highlight similar patterns (T3.2, T3.3), Figure 5.

Step1. Group the data: We group the description of the spatial distributions of the local coefficients with respect to the dependent variable and the selected independent variable into the areas with positive/negative correlations. We describe all the significant coefficients and intercepts, where significance is measured via a t-test that has been adjusted for dependent multiple tests, Figure 5 (A) - Step 1.



Fig. 5: Workflow for exploring the narrative explanations and external context of the local coefficient of the independent variable *pct_age_18_29* (the percentage of the people in the age between 18 and 29) when exploring the MGWR model results for the 2016 U.S. Presidential Election data. (A) shows the narrative generation pipeline for a selected local coefficient across three steps: 1) measures the significance of the coefficient via a t-test and classifies the significant data into two groups, each has a positive/negative relationship with the dependent variable; 2) detects clusters in each group, and; 3) generates a narrative explanation including detected clusters. The analyst finds there is a significant positive inclination for younger voters to vote Democrat in the four detected clusters. Next they explore the external context for a selected cluster in the state of Ohio (B). They find "Ohio university" is one of the keyphrases extracted from relevant Wikipedia pages (B.1), then they identify Athens county by browsing the original Wikipedia contents (B.2, B.3). Next, they explore keyphrases related to Athens and realize this county is a democratic stronghold.

Step2. Detect clusters: For a selected local coefficient, GeoExplainer automatically detects spatial clusters belonging to each group and generates related interpretations, Figure 5 (A) - Step 2. Before detection, we derive two queen-adjacency graphs from the significant local coefficients data. One graph contains location nodes with positive coefficient values while the other has nodes with negative coefficient values. For each graph, we perform a community detection algorithm to detect spatial clusters. To balance the time efficiency, community modularity, and the stability of the results [3], we use the Leiden algorithm [69] on the server side to identify the spatial communities.

Step3. Populate narrative template: The narrative template is generated in a hierarchical structure with interactions to explain the spatial patterns based on the detected clusters (Figure 5 (A) - Step3). We populate overview paragraphs that describe the spatial distributions with respect to the positive and negative coefficients first. For a selected coefficient, the overview paragraphs also explain how each independent variable influences the dependent variable in the regions with significant coefficients. For the intercept, we describe how the unmeasurable effects (context) of each region affect the value of the dependent variable. Each overview paragraph will include a toggle list of sub-paragraphs that includes additional information about each cluster detected. The list describes the local geographical patterns along clusters through a user-definable location identifier as described in the spatial distribution templates. Analysts have the flexibility to summarize the patterns with vernacular geographic language. All paragraphs describe the spatial pattern overview and local clusters are interactively linked with the map to call attention to interesting patterns. Analysts can click the paragraph to trigger related external information, Figure 5 (B).

External Information View: To help provide further context to the spatial modeling results (**T3.4**), GeoExplainer links to Wikipedia, enabling quick access to information about local geography, politics, etc. Given that distinct datasets exhibit varying spatial resolutions, we align the Wikipedia information to the dataset's resolution (e.g., county level or community areas for a city). The external information view contains an external keyphrases list (Figure 1 (B3)) of the most important keyphrases for the relevant Wikipedia information. Each keyphrase can populate a list of original Wikipedia content (Figure 1 (B4)) consisting of paragraphs relevant to a particular phrase. Figure 5 (B.1–B.4) illustrates the external context exploration workflow.

External Keyphrases List: The external keyphrases list displays the top N (N = 20 by default, can be defined by analysts) most important keyphrases of the Wikipedia information that correspond to a selected range of geographical regions. When analysts explore the narrative

explanation of the spatial distributions, they can click the popup "show keyphrases button" (*Wiki*) on each paragraph to populate the relevant keyphrases list. Analysts can also trigger the relevant keyphrases list by selecting a range of areas on the map. We extract keyphrases based on a Python implementation of the TextRank algorithm [51]. Each extracted phrase is categorized into one topic that is the same as the corresponding section's topic from Wikipedia. We use a qualitative color scheme **Section** to present those topics. Each row in the keyphrases list contains a keyphrase entity and the rank of this phrase, both are encoded in their topic color. Analysts can click operation buttons on each row to observe the original Wikipedia contents that contain that phrase and search the phrase through the online search engine to gain more context knowledge. By default, we sort the list in descending order based on the ranks of phrases. Analysts can rearrange or filter the list by topics as well.

Original Wikipedia Paragraph List: The original Wikipedia paragraph list displays all the paragraphs that contain the selected keyphrase in the corresponding Wikipedia information. This list can be trigged by clicking "original paragraphs button" on each keyphrase row in the external keyphrases list. We highlight the keyphrase with a yellow background in each paragraph. When hovering the mouse over a paragraph, the related map regions are also highlighted.

4.4 Contextualized Report Authoring

Throughout the model configuration and validation process, analysts explore model outcomes from a variety of views. To help proceed from data analysis to result communication, GeoExplainer provides a report authoring module that enables the recording of findings and workflow operations, synthesizes explored knowledge, and communicates the results to wider audiences. (Figure 1 (C)) (T4.1, T4.2). The module has an expandable report-authoring panel which serves as a collection bin for capturing all the outputs from GeoExplainer. Every text narrative and visualization can be included in this panel. The panel also supports the addition of free-form narratives including adding a new paragraph, editing a paragraph, and deleting a paragraph. The contents in the report-authoring panel are arranged as a sorted list. The order of paragraphs and pictures in the panel can be rearranged by clicking the arrow buttons in the popover toolbox $(\land \lor \Box)$ associated with each item in the list. The content styles in the narrative paragraphs can be defined by using inline HTML tags with CSS code. For any map added to the report, instead of using the original color scheme, analysts can pick a new color scheme from a list of color schemes. After authoring a report, the framework can export the narrative report to a PDF file automatically by clicking the "export as PDF button" (D Export PDF).

4.5 Analytical State Recording

GeoExplainer offers assistance to analysts in recording (**T5.1**), recovering (**T5.2**), and sharing (**T5.3**) their analytical outcomes within the exploration pipeline. At any stage of the pipeline, analysts can save progress into a JSON file with the "save state" button. The saved file can be reloaded or shared to reproduce the saved state. Clicking the "load state" button retrieves the current state from the saved breakpoint. In addition to the analytical results, training results of the spatial models are automatically recorded once the training and analysis processes.

5 CASE STUDY AND EXPERT INTERVIEW

In this section, we present a case study and compile results from six expert interviews to demonstrate and evaluate our framework.

5.1 2016 U.S. Presidential Election Analysis

In this case study, we invited two domain experts (DE1 and DE2) to analyze the 2016 U.S. presidential election dataset [19]. The dataset contains 2,813 counties where non-contiguous states and counties with less than 5,000 inhabitants are removed. DE1 is responsible for configuring and training an MGWR model on the election dataset, and DE2 utilizes the explanation and report authoring components.

Model Configuration: Since the objective of this study is to explore how socioeconomic factors influence voter preferences, DE1 selects *pct_democrat* (percentage of people who vote for the Democratic Party) as the dependent variable (**T2.1**) from the Original Feature List (Figure 3 (A4)) after he loads the data. He then browses the populated univariate choropleth map (Figure 3 (C1)) and follows the operation recommendation to normalize the dependent variable by clicking the "log transformation" button (**T2.2**), Figure 3 (A5).

Next, DE1 needs to choose independent variables for the model (T2.1). By dragging and dropping interested features to the Independent Variable List (Figure 3 (A3)), DE1 explores the correlation between each independent variable candidate and the dependent variable. He observes a strong negative global relationship between pct_gop (percentage of people who vote for the Republican Party) and pct_democrat in the Correlation View, Figure 3 (B1). Next, DE1 explores the bivariate map and notes that urban and coastal areas are more likely to vote for the Democratic Party, indicating that there may be underlying latent factors that determine voter preferences, Figure 3 (C2). DE1 notices high multicollinearity among many socioeconomic factors according to the VIF scores and scatterplot matrix, Figure 3 (B2, B3), after discussion with DE2, DE1 compares the trained model's global R^2 and AICc values (**T2.3**) by trying the different combinations of the socioeconomic factors. Based on the exploration and discussion, DE1 chooses the following independent variables: [sex_ratio, pct_black, pct_hisp (percentage of Hispanic population), pct_bach (% of people with a bachelor's degree), income, pct_65_over, pct_age_18_29, gini, pct_manuf (manufacture), log_pop_den (log of population density, a measure of 'urbanness'), pct_3rd_party (votes for a third-party candidate), turn_out (voter turnout), pct_FB (foreign-born), pct_insured (people with health insurance)]. DE1 then trains the model and validates that the model's local performance is reasonable from the generated explanation of local residuals (T3.1). DE1 saves the analytical state (T5.1) and shares the state with DE2 (T5.3).

Exploring the Model Coefficients: After loading the analytical state file (**T5.2**), DE2 explores the local coefficient list and notices an unusual spatial distribution of the coefficient $pct_age_18_29$ (the percentage of the people in the age between 18 and 29) in the map, Figure 1 (B2). DE2 reads the narrative explanations in the Parameter Explanation View (**T3.1**) and realizes that the vast majority of county-specific local coefficients of $pct_age_18_29$ are insignificant, except for counties in Southern California and the counties throughout much of the Northwest. The choropleth map shows younger voters (aged eighteen to twenty-nine) have a significant inclination to vote Republican in these regions. DE2 also notices that the framework has detected two small clusters in Indiana and Ohio (**T3.2**), where there is a significant positive inclination for younger voters to vote Democrat (**T3.3**). DE2 then explores

relevant context information populated in the External Keyphrases List (**T3.4**). While DE2 is familiar with the region in Indiana, quickly identifying it as the home of Indiana University, the region in Ohio is ambiguous. DE2 observed "ohio university" as one of the high-ranking keyphrases of this cluster, and realized Athens County, Ohio in the cluster is a Democratic stronghold and home to Ohio University in the corresponding Wikipedia Paragraph List. Figure 5 illustrates the workflow of exploring the narrative explanation and external context. Here DE2 noted the seamless ability of GeoExplainer to pull contextual information about unfamiliar regions to help hypothesis development.

After recording these findings and relevant maps in the report authoring panel (**T4.1**), DE2 explores the local estimates of the intercept. In MGWR, the local estimates of the intercept indicate the impact of geographical context on the dependent variable while holding everything else in the model constant [19].

In this case, the local intercepts indicate the intrinsic support for the Democratic Party or the Republican Party if all counties had exactly the same mix of population. DE2 explores the spatial distribution of the local intercepts on the map (**T2.3**) and the Parameter Explanation View (**T3.1**). Here,



the spatial distribution and the relevant narrative interpretation suggest that the counties throughout New England, the upper Midwest, and down the Pacific Coast intrinsically tend to vote Democrat, and counties throughout the South, except for Florida, are intrinsically more likely to vote Republican. This spatial distribution pattern is particularly informative to the expert because it represents the unseen, but important, influence of geographical context on behavior. It also represents how each county would vote if it contained an average population composition. That is, if all counties across the US had exactly the same population composition, the distribution of voting behavior would not be constant but as shown on the map. The expert adds this map and the corresponding interpretation to the report panel (**T4.1**).

Performance: Visual elements, including charts, maps, and narrative generation, load and display at a sub-second rate. Spatial cluster detection and keyphrase extraction take roughly 2 seconds, respectively. The MGWR training process took a duration of around 186 seconds. The scalability of the framework is discussed in Sect.6.

5.2 Expert Interviews

GeoExplainer is designed for spatial data analysts who use geographic models to analyze spatial-related phenomena. To further assess our framework, we conducted a group interview with six participants (E0 to E5) via video-conferences. Three participants are geographers who have expertise in spatial data analysis for at least 5 years. The other three participants are Ph.D. students in geographic analysis. All participants have experience using spatial regression models. We recorded the participants' audio and computer screen for subsequent analysis with their consent. We first introduced the background and the analytical tasks of our work, followed by a detailed explanation of the analytical workflow and the interface with a case study. Then, the experts were encouraged to explore the dataset and develop a report using GeoExplainer. During the interview, participants were asked to share their screens with us; they were encouraged to ask questions when necessary. The interview lasted approximately 90 minutes. At the end of the interview, we collected feedback and distributed two questionnaires (a usability evaluation and a functionality questionnaire) to each of the domain experts. Both questionnaires include a series of 7-point Likert-scale questions and are provided in the supplementary material. Figure 6 shows the aggregate results of the surveys.

Combining the questionnaire scores and the feedback from experts, our *unified analytical pipeline* received a high average score (6.23/7) with the respect to its perceived usefulness. The experts commented that the integration of the exploratory spatial data analysis (ESDA) process with the narrative interpretations (**T1.1**) can improve the efficiency of their spatial analytical tasks and save them a lot of time compared to the traditional tools. E1 noted that "*I am happy that the framework has already established a good analytical pipeline, now I can review my*



Fig. 6: Average ratings from six participants regarding the usability and usefulness of the main functional modules and interactions of GeoExplainer. Modules are grouped based on the analytical tasks.

analysis results more efficiently without bothering to import different packages to build my own workflow." The average score (6.76/7) of the perceived ease of use from the questionnaire further indicates that the dynamic narrative visualization design in the framework is able to promote usability. E1 appreciated the flexibility of the interactions and the intuitive visual design of the multiple modules' layout. E2 also had a positive comment with regard to the interactive visualization widgets commenting that "Having used other drag-and-drop type visualization software such as Tableau and ArcGIS, I found the interaction of Geo-Explainer to be really straightforward and easy to use even without the tutorial video." E3 further added that they were impressed by the single dashboard design for the model validation and reporting features. E4 and E5 both like the smooth transitions between different views.

Another favorite feature of GeoExplainer was the narrative explanations and their interplay with the map. All narrative explanation views and inline annotation modules received high average ratings regarding their usefulness and usability (Figure 6 (T3)). The participants agreed that the narrative explanations generated were deemed satisfactory and effective for aiding their analytical tasks. E1 showed a preference for concise interpretation paragraphs that interactively filter the relevant map layer while examining coefficients and local diagnostics (T3.1, T3.3). He mentioned that these explanations facilitated a rapid comprehension of local spatial distributions and guided his focus towards relevant areas on the map. E4 and E5 both appreciated the hierarchical interpretations of spatial clusters for local coefficients (T3.2). They noted that the local context information for each spatial cluster was readily accessible and incorporated into the report. They also appreciated the Wikipedia keyphrases lookup feature with supplemental context information for the selected geographic cluster (T3.4). All student participants found the explanations of the statistical indicators to be highly valuable, as they facilitated the evaluation and understanding of the model's performance and saved considerable time and effort.

The *report authoring and the analytical state record functions* both received high ratings for their usefulness and utility. E2 mentioned: "*I can add whatever worth noticing to the report when exploring [...] I also like the flexibility to arrange my findings and edit them into a report.*" E1 and E3 appreciated they can synthesize their findings into a communicable report without using any 3rd party tools (**T4.2**). E5 especially mentioned analytical process recording function is very useful when she needs to gradually analyze a model's results or share the idea with other collaborators (**T5**). E5 said: "*Now I can save my analytics state, do other jobs, and back to the analysis days later without the need to retrain the model. [...] I also invited my colleagues to do their analysis based on my configured model on GeoExplainer ..."*

Limitations: The domain experts also identified some limitations of our current framework. E0 noted that the map layer loading time needs to be further optimized, and E0, E1, and E2 expected the map

visualization in our framework to support more user-defined operations, such as the selection of color and filtering rules. E2 suggested that the framework's explanatory features could be further expanded to include an explanation of the traditional ordinary least squares (OLS) model's result. E4 suggested that we should work to provide more content automation from the external search function to the model results.

6 DISCUSSION AND CONCLUSIONS

In this paper, we propose a visual analytics framework, GeoExplainer, for supporting spatial data modeling, contextualization, and reporting. We introduce a suite of interactive views that are supported by a template-based narrative generator to provide explanatory text that can directly be utilized in report generation. As analysts create their spatial models, our framework automatically highlights potential issues in model parameter selections, utilizes template-based text generation to summarize model outputs in detail, and links with external knowledge repositories, such as Wikipedia, to provide interactive annotations.

The innovations of this system are in the use of text generation for multivariate spatial models that require locational information to support answering "why" questions about spatial phenomena. While the applied visualizations utilize well-known components, the integration of automatic analysis and human-in-the-loop for exploring and contextualizing model results goes well beyond current systems that support general descriptive statistics (mean, median, correlation). Contributions include methods to automatically detect and explain spatial patterns, all while enabling analysts to capture and share their workflow. The proposed text templates go beyond simple descriptive statistics (mean, mode, correlation) to focus on multivariate patterns and spatial phenomena. As models become more complex, such advances help bridge the gap for explainable artificial intelligence.

Generalizability: Our framework is designed support different spatial regression models that shares the same input and output protocol. Owing to the modularized design, new models can be swapped in to replace MGWR adopted in the case study. The asynchronous model training process can further enable the integration of models with varying computational complexities. For enhancing explanations of diverse spatial models, our narrative generation function, informed by domain experts, selectively presents and elucidates three commonly-used diagnostic indicators (local R^2 , Cook's distance, and standard residual) for local spatial model validation. The input of our framework can be an arbitrary spatial dataset where every observation in the dataset has features corresponding to a geographic location. With respect to the data format, GeoJSON is used in the current prototype, but additional geographic file types can also be supported, such as ESRI shapefiles.

Scalability: Advanced spatial models, e.g., MGWR, may be resourceintensive as the data scale expands [40]. GeoExplainer adapts to varying data scales with the aid of an asynchronized state recording functionality, as described in Sect.4.5. To optimize the loading duration of the views for large datasets, data to be rendered in the visual interface can be processed concurrently with the model training process.

Target Audience: As mentioned in Sect.4, our target audience comprises spatial data analysts requiring enhanced comprehension and communication of model results. Some of the features in our framework like interactive explanations of statistical indicators can also be used by novice audiences with limited spatial analysis knowledge. However, the framework needs further improvement and evaluation of the functionality for novices.

Limitations and Future Work: While template-based narrative explanations facilitate quality control and maintain consistency, these templates are task-oriented and necessitate fine-tuning by domain experts. Possible extensions grounded in large language models may enhance flexibility and adaptability of the narrative generation process. In terms of extending the framework's applicability to a general audience, there are current plans to deploy GeoExplainer in large classroom settings to further assess the report authoring module for junior analysts. We also expect to improve the contextualization support by incorporating knowledge graph-based extraction techniques.

SUPPLEMENTAL MATERIALS

The supplemental materials include (1) a PDF example report generated by the case study, (2) a JSON file as an example state-saving file of the case study in Section 5.1, (3) example expert interview questions in a PDF file, (4) an Excel file containing the evaluation results for creating Figure 6, (5) an appendix, and (6) a demo video. All code is available at https://github.com/VADERASU/geoexplainer.

ACKNOWLEDGMENTS

The authors wish to thank Qiang Xu for the engineering support in the framework development. This work was supported in part by the National Science Foundation under Grant 1639227, the U.S. Department of Homeland Security under Award 2017-ST-061-QA0001, the National Natural Science Foundation of China (No. 62202217), Guangdong Basic and Applied Basic Research Foundation (No. 2023A1515012889), and Guangdong Talent Program (No. 2021QN02X794).

REFERENCES

- L. Anselin. Interactive techniques and exploratory spatial data analysis. Regional Research Institute Working Papers, 200:253–266, 1996. 3
- [2] L. Anselin, I. Syabri, and Y. Kho. GeoDa: An introduction to spatial data analysis. In *Handbook of Applied Spatial Analysis*, pp. 73–89. Springer, 2010. doi: 10.1007/978-3-642-03647-7_5 3
- [3] S. Behera, H. Lu, and A. Nayak. Performance evaluation of community detection algorithms based on relationship strength measurement. In *International Conference on Ubiquitous and Future Networks*, pp. 448– 453. IEEE, 2018. doi: 10.1109/ICUFN.2018.8437001 7
- [4] E. Bertini and D. Lalanne. Surveying the complementary role of automatic data analysis and visualization in knowledge discovery. In *Proceedings of the Special Interest Group on Knowledge Discovery in Data Conference Workshops*, pp. 12–20, 2009. doi: 10.1145/1562849.1562851 2
- [5] M. Bostock, V. Ogievetsky, and J. Heer. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011. doi: 10.1109/TVCG.2011.185 3
- [6] A. M. Brock, B. Hecht, B. Signer, and J. Schöning. Bespoke map customization behavior and its implications for the design of multimedia cartographic tools. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia*, pp. 1–11, 2017. doi: 10.1145/3152832. 3152833 2
- [7] C. Brunsdon, A. S. Fotheringham, and M. E. Charlton. Geographically weighted regression: A method for exploring spatial nonstationarity. *Geographical Analysis*, 28(4):281–298, 1996. doi: 10.1111/j.1538-4632.1996 .tb00936.x 2
- [8] K. P. Burnham and D. R. Anderson. Multimodel inference: Understanding AIC and BIC in model selection. *Sociological Methods & Research*, 33(2):261–304, 2004. doi: 10.1177/0049124104268644 5
- [9] S. Chen, J. Li, G. Andrienko, N. Andrienko, Y. Wang, P. H. Nguyen, and C. Turkay. Supporting story synthesis: Bridging the gap between visual analytics and storytelling. *IEEE Transactions on Visualization and Computer Graphics*, 26(7):2499–2516, 2020. doi: 10.1109/TVCG.2018. 2889054 2, 3
- [10] Y. Chen, S. Barlowe, and J. Yang. Click2Annotate: Automated insight externalization with rich semantics. In *IEEE Symposium on Visual Analytics Science and Technology*, pp. 155–162, 2010. doi: 10.1109/VAST.2010. 5652885 2
- [11] Y. Chen, J. Yang, S. Barlowe, and D. H. Jeong. Touch2Annotate: Generating better annotations with less human effort on multi-touch interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 3703–3708, 2010. doi: 10.1145/1753846.1754042 2
- [12] A. Choudhry, M. Sharma, P. Chundury, T. Kapler, D. W. Gray, N. Ramakrishnan, and N. Elmqvist. Once upon a time in visualization: Understanding the use of textual narratives for causality. *IEEE Transactions on Visualization and Computer Graphics*, 2020. doi: 10.1109/TVCG.2020.3030358
- [13] R. D. Cook. Detection of influential observation in linear regression. *Technometrics*, 42(1):65–68, 2000. doi: 10.1080/00401706.2000.10485981
- [14] Z. Cui, S. K. Badam, M. A. Yalçin, and N. Elmqvist. DataSite: Proactive visual data exploration with computation of insight-based recom-

mendations. Information Visualization, 18(2):251–267, 2019. doi: 10. 1177/1473871618806555 2

- [15] R. Eccles, T. Kapler, R. Harper, and W. Wright. Stories in GeoTime. Information Visualization, 7(1):3–17, 2008. doi: 10.1057/palgrave.ivs. 9500173 2
- [16] N. Ferreira, J. Poco, H. T. Vo, J. Freire, and C. T. Silva. Visual exploration of big spatio-temporal urban data: A study of New York City taxi trips. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2149– 2158, 2013. doi: 10.1109/TVCG.2013.226 2
- [17] A. Fotheringham and P. Rogerson. The SAGE Handbook of Spatial Analysis. SAGE Publications, 2008. 3
- [18] A. S. Fotheringham, C. Brunsdon, and M. Charlton. Geographically Weighted Regression: The Analysis of Spatially Varying Relationships. John Wiley & Sons, 2003. 1
- [19] A. S. Fotheringham, Z. Li, and L. J. Wolf. Scale, context, and heterogeneity: A spatial analytical perspective on the 2016 U.S. presidential election. *Annals of the American Association of Geographers*, 0(0):1–20, 2021. doi: 10.1080/24694452.2020.1835459 3, 5, 8
- [20] A. S. Fotheringham, W. Yang, and W. Kang. Multiscale geographically weighted regression (MGWR). *Annals of the American Association of Geographers*, 107(6):1247–1265, 2017. doi: 10.1080/24694452.2017. 1352480 1, 2
- [21] T. Gao, J. Hullman, E. Adar, B. Hecht, and N. Diakopoulos. NewsViews: An automated pipeline for creating custom geovisualizations for news. *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 3005–3014, 2014. doi: 10.1145/2556288.2557228 2
- [22] J. Gareth, W. Daniela, H. Trevor, and T. Robert. An Introduction to Statistical Learning with Applications in R. Spinger, 2013. 4
- [23] J. Gil, T. Varoudis, K. Karimi, and A. Penn. The space syntax toolkit: Integrating depthmapX and exploratory spatial analysis workflows in QGIS. In *International Space Syntax Symposium*, vol. 10, 2015. 2
- [24] S. Goodwin, J. Dykes, A. Slingsby, and C. Turkay. Visualizing multiple variables across scale and geography. *IEEE Transactions on Visualization* and Computer Graphics, 22:599–608, 1 2016. doi: 10.1109/TVCG.2015. 2467199 2
- [25] D. Guo, J. Chen, A. M. MacEachren, and K. Liao. A visualization system for space-time and multivariate patterns (VIS-STAMP). *IEEE Transactions* on Visualization and Computer Graphics, 12(6):1461–1474, 2006. doi: 10 .1109/TVCG.2006.84 2
- [26] J. Hullman and N. Diakopoulos. Visualization rhetoric: Framing effects in narrative visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2231–2240, 2011. doi: 10.1109/TVCG.2011.255 2
- [27] J. Hullman, N. Diakopoulos, and E. Adar. Contextifier: Automatic generation of annotated stock visualizations. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 2707–2716, 2013. doi: 10.1145/2470654.2481374 2
- [28] M. Kanevski. Machine Learning for Spatial Environmental Data: Theory, Applications, and Software. EPFL press, 2009. 2
- [29] B. Kastanakis. Mapbox Cookbook. Packt Publishing Ltd, 2016. 3
- [30] R. Kosara and J. Mackinlay. Storytelling: The next step for visualization. Computer, 46(5):44–50, 2013. doi: 10.1109/MC.2013.36 2
- [31] W. Krämer and C. Donninger. Spatial autocorrelation among errors and the relative efficiency of OLS in the linear regression model. *Journal of the American Statistical Association*, 82(398):577–579, 1987. doi: 10. 1080/01621459.1987.10478468 2
- [32] J. Krause, A. Dasgupta, J. Swartz, Y. Aphinyanaphongs, and E. Bertini. A workflow for visual diagnostics of binary classifiers using instancelevel explanations. In *IEEE Conference on Visual Analytics Science and Technology*, pp. 162–172. IEEE, 2017. doi: 10.1109/VAST.2017.8585720
- [33] J. Krause, A. Perer, and E. Bertini. INFUSE: Interactive feature selection for predictive modeling of high dimensional data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1614–1623, 2014. doi: 10. 1109/TVCG.2014.2346482 2
- [34] J. Krause, A. Perer, and K. Ng. Interacting with predictions: Visual inspection of black-box machine learning models. In *Proceedings of the* ACM Conference on Human Factors in Computing Systems, pp. 5686– 5697, 2016. doi: 10.1145/2858036.2858529 2
- [35] C. Lai, Z. Lin, R. Jiang, Y. Han, C. Liu, and X. Yuan. Automatic annotation synchronizing with textual description for visualization. *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 1–13, 2020. doi: 10.1145/3313831.3376443 2
- [36] S. Latif and F. Beck. Interactive map reports summarizing bivariate

geographic data. Visual Informatics, 3(1):27–37, 2019. doi: 10.1016/j. visinf.2019.03.004 2, 6

- [37] S. Latif, S. Chen, and F. Beck. A deeper understanding of visualization-text interplay in geographic data-driven stories. *Computer Graphics Forum*, 40(3):311–322, 2021. doi: 10.1111/cgf.14309 2, 6
- [38] J. Li, A. D. Heap, A. Potter, and J. J. Daniell. Application of machine learning methods to spatial interpolation of environmental variables. *Envi*ronmental Modelling & Software, 26(12):1647–1659, 2011. doi: 10.1016/j .envsoft.2011.07.004 2
- [39] X. Li, L. Anselin, and J. Koschinsky. GeoDa Web: Enhancing web-based mapping with spatial analytics. *Proceedings of the ACM International Symposium on Advances in Geographic Information Systems*, 03-06-Nove:1–4, 2015. doi: 10.1145/2820783.2820792 2, 3
- [40] Z. Li and A. S. Fotheringham. Computational improvements to multi-scale geographically weighted regression. *International Journal of Geographical Information Science*, 34(7):1378–1397, 2020. doi: 10.1080/13658816. 2020.1720692 9
- [41] M. Ließ, B. Glaser, and B. Huwe. Uncertainty in the spatial prediction of soil texture: Comparison of regression tree and random forest models. *Geoderma*, 170:70–79, 2012. doi: 10.1016/j.geoderma.2011.10.010 2
- [42] C. Liu, L. Xie, Y. Han, A. Wei, and X. Yuan. AutoCaption: An approach to generate natural language description from visualization automatically. *IEEE Pacific Symposium on Visualization*, 2020-June:191–195, 2020. doi: 10.1109/PacificVis48177.2020.1043 2
- [43] S. Liu, X. Wang, M. Liu, and J. Zhu. Towards better analysis of machine learning models: A visual analytics perspective. *Visual Informatics*, 1(1):48–56, 2017. doi: 10.1016/j.visinf.2017.01.006 2
- [44] R. Lovelace, J. Nowosad, and J. Muenchow. *Geocomputation with R*. Chapman and Hall/CRC, 2019. 2
- [45] J. Lu, W. Chen, Y. Ma, J. Ke, Z. Li, F. Zhang, and R. Maciejewski. Recent progress and trends in predictive visual analytics. *Frontiers of Computer Science*, 11(2):192–207, 2017. doi: 10.1007/s11704-016-6028-y 2
- [46] Y. Lu, R. Garcia, B. Hansen, M. Gleicher, and R. Maciejewski. The state-of-the-art in predictive visual analytics. *Computer Graphics Forum*, 36(3):539–562, 2017. doi: 10.1111/cgf.13210 2
- [47] P. Lundblad and M. Jern. Geovisual analytics and storytelling using HTML5. Proceedings of the International Conference on Information Visualisation, pp. 263–271, 2013. doi: 10.1109/IV.2013.35 2
- [48] Meta Platforms, Inc. React official page, 2022. 3
- [49] Y. Ming, H. Qu, and E. Bertini. RuleMatrix: Visualizing and understanding classifiers with rules. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):342–352, 2018. doi: 10.1109/TVCG.2018.2864812
- [50] T. Mühlbacher and H. Piringer. A partition-based framework for building and validating regression models. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):1962–1971, 2013. doi: 10.1109/TVCG.2013. 125 2
- [51] P. Nathan. PyTextRank, a python implementation of textrank for phrase extraction and summarization of text documents, 2016. doi: 10.5281/ zenodo.4637885 7
- [52] S. Nusrat, M. J. Alam, C. Scheidegger, and S. Kobourov. Cartogram visualization for bivariate geo-statistical data. *IEEE Transactions on Visualization and Computer Graphics*, 24(10):2675–2688, 2017. doi: 10. 1109/TVCG.2017.2765330 4
- [53] T. M. Oshan, Z. Li, W. Kang, L. J. Wolf, and A. S. Fotheringham. MGWR: A python implementation of multiscale geographically weighted regression for investigating process spatial heterogeneity and scale. *ISPRS International Journal of Geo-Information*, 8(6):269, 2019. doi: 10.3390/ ijgi8060269 1, 2, 3, 4
- [54] D. Ozturk, A. Chaudhary, P. Votava, and C. Kotfila. GeoNotebook: Browser based interactive analysis and visualization workflow for very large climate and geospatial datasets. In AGU Fall Meeting Abstracts, vol. 2016, pp. IN53A–1876, 2016. 2
- [55] R. K. Pace and O. W. Gilley. Using the spatial configuration of the data to improve estimation. *The Journal of Real Estate Finance and Economics*, 14(3):333–340, 1997. doi: 10.1023/A:1007762613901 3
- [56] D. Ren, M. Brehmer, B. Lee, T. Höllerer, and E. K. Choe. ChartAccent: Annotation for data-driven storytelling. In *IEEE Pacific Symposium on Visualization*, pp. 230–239, 2017. doi: 10.1109/PACIFICVIS.2017.8031599 2
- [57] S. J. Rey and L. Anselin. PySAL: A python library of spatial analytical methods. In *Handbook of Applied Spatial Analysis*, pp. 175–193. Springer, 2010. doi: 10.1007/978-3-642-03647-7_11 2, 3
- [58] S. J. Rey, L. Anselin, P. Amaral, D. Arribas-Bel, R. X. Cortes, J. D.

Gaboardi, W. Kang, E. Knaap, Z. Li, S. Lumnitz, et al. The PySAL ecosystem: Philosophy and implementation. *Geographical Analysis*, 2021. doi: 10.1111/gean.12276 2

- [59] A. Satyanarayan and J. Heer. Authoring narrative visualizations with ellipsis. *Computer Graphics Forum*, 33(3):361–370, 2014. doi: 10.1111/ cgf.12392 2
- [60] L. M. Scott and M. V. Janikas. Spatial statistics in ArcGIS. In Handbook of Applied Spatial Analysis, pp. 27–41. Springer, 2010. doi: 10.1007/978 -3-642-03647-7_2 2, 3
- [61] M. Sedlmair, C. Heinzl, S. Bruckner, H. Piringer, and T. Möller. Visual parameter space analysis: A conceptual framework. *IEEE Transactions* on Visualization and Computer Graphics, 20:2161–2170, 12 2014. doi: 10.1109/TVCG.2014.2346321 2
- [62] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148, 2010. doi: 10.1109/TVCG.2010.179 2
- [63] Z. Shabrina, B. Buyuklieva, and M. K. M. Ng. Short-term rental platform in the urban tourism context: A geographically weighted regression (GWR) and a multiscale GWR (MGWR) approaches. *Geographical Analysis*, 53(4):686–707, 2021. doi: 10.1111/gean.12259 3
- [64] 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*, 25(1):672–681, 2019. doi: 10.1109/TVCG.2018.2865145 2
- [65] G. Strode, J. D. Morgan, B. Thornton, V. Mesev, E. Rau, S. Shortes, and N. Johnson. Operationalizing Trumbo's principles of bivariate choropleth map design. *Cartographic Perspectives*, 94:5–24, 2019. doi: 10.14714/ CP94.1538 4
- [66] M. Takatsuka and M. Gahegan. GeoVISTA Studio: A codeless visual programming environment for geoscientific data analysis and visualization. *Computers and Geosciences*, 28(10):1131–1144, 2002. doi: 10.1016/ S0098-3004(02)00031-6 2
- [67] J. Talbot, B. Lee, A. Kapoor, and D. S. Tan. EnsembleMatrix: Interactive visualization to support machine learning with multiple classifiers. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 1283–1292, 2009. doi: 10.1145/1518701.1518895 2
- [68] C. Tong, R. Roberts, R. Borgo, S. Walton, R. S. Laramee, K. Wegba, A. Lu, Y. Wang, H. Qu, Q. Luo, and X. Ma. Storytelling and visualization: An extended survey. *Information*, 9(3):1–42, 2018. doi: 10.3390/info9030065 2
- [69] V. A. Traag, L. Waltman, and N. J. Van Eck. From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports*, 9(1):1–12, 2019. doi: 10.1038/s41598-019-41695-z 7
- [70] E. Vittinghoff, D. V. Glidden, S. C. Shiboski, and C. E. McCulloch. Regression Methods in Biostatistics: Linear, Logistic, Survival, and Repeated Measures Models. Springer, 2006. 4, 6
- [71] A.-X. Zhu, G. Lu, J. Liu, C.-Z. Qin, and C. Zhou. Spatial prediction based on third law of geography. *Annals of GIS*, 24(4):225–240, 2018. doi: 10. 1080/19475683.2018.1534890 2

S
ш
Ē
<
1
<u> </u>
Ξ
<u></u>
ш
-
-
~
0
=
2
~
ш.
z
ш
(5
-
ш
>
-
5
2
ш
œ
◄
ブ
<u> </u>
-
4

13	Target Obje	ect → 💷	Describe the Patterns $ ightarrow$	©Explain the Results	Template Framework ²	Narrative Explanation Example ⁴
(∀)	Data distributi	tion {	Normal / skewed } istribution	Recommend a proper model type or data transformation	[Data is in { normal / skewed } distribution.] [Suggest { Gaussian model / Log or square root transformation }.]	Data is in normal distribution. Can use Gaussian model type .
Narrative Generation for Model Configuration Components	Correlation		Significant / insignificant } strong / moderate / weak } positive / negative } lobal linear relationship	Independent varieble selection suggestions	[{ Significant / insignificant } { strong / moderate / weak.} { potitive / negative } global linear relationship between { variable A } and { variable B }.] [Independent variable selection suggestions.]	Significant strong negative global linear relationship between price_pp and income_pc . May consider select income_pc as the independent variable.
	Multicollineari detection	rity /	'ariables have igh multicollinearities	Operation suggestions.	[{ Variables } have high multicollinearities.] [Suggest checking the detail correlation.]	income, pct_bach have high multicollinearities. Check the detail correlations by click the plot.
	Por	ocal R ² u	vreas with { high / low } ocal R² values, based on ser-defined threshold	Model performance {decent / poor} regarding to local R ²	[{ # of areas } in { place } ³ , where the call R values are greater than { threshold }] [Indicating a { decent / poor } prediction of the model in these areas.]	22 community areas in Chicago city , where the local R ² values are greater than 0.7 . Indicating a decent prediction of the model in these areas.
	Co	ok's D	vreas with { values≥4/N, I = # of all areas }	Meaning of Cook's D and operation suggestions	[{ # of areas } in { place } ³ , where the Cook's distances are above the { threshold }.] [These places are considered as influential outliers. Please check these areas for validity.]	5 community areas in Chicago city, where the Cook's distance are above the threshold 0.15. These places are considered as influential outliers. Please check these areas for validity.
(8) Narrative Generation for Model/Valdation Components	Dagoostic Dagoostic Indicators Stat resi	u u sidual	Over-predicted / nder-predicted } areas	Meaning of the residual with Moran's I results and suggestions	[The spatial distribution of the model residuals wariable { dependent variable }] wariable { dependent variable }] [{ # of areas } in { place }, where the predicted auleas are { higher / lower } than expected.] [The expected value of Moran's I if the distribution were random is { value } and the observed value is { value } mod the observed value is { value } with a p-value { value } in \$5% confidence interval.] [There is { clear / no clear } pattern of over or under predictions in the study area.] [The model { may miss key explanatory variables / works good }.]	The spatial distribution of the model residuals between the predicted and observed dependent variable price_pp . 26 community areas in Chicago city , where the predicted values where the predicted values are lower than expected. The expected value of Moran's I if the distribution were random is -0.01 and the observed value is -0.04 with a p-value of 0.72 in 95% confidence intervals. There is no clear pattern of over or under predictions in the study area. The model works good without missing any key explanatory variables.
	Local indepen variable coeffi	ndent	Areas have significant positive / negative } correlations with dependent variable,based on T-test	The strength and the type of the correlation with statistical context	[The spatial distribution of the coefficient of the independent variable { independent variable }. Based on the T-test results, the colored areas in the map indicate the significant relationships between { independent variable } and { dependent variable }.] [[# of areas] in (place), where the { value in percent (%) of the study area.] The significant local coefficients range from { lower bound value } to { upper bound value }, with a mean value of { value }]	The spatial distribution of the coefficient of the independent variable rev_rating. Based on the Trest results, the colored areas in the map indicate the significant relationships between rev_rating and price_pp. 39 community areas in Chicago city, where the positive relationship is significant in 51 % of the study area. The significant local coefficients range from 0.23 to 0.33, with a mean value equal to 0.28.

^TThe pattern description of the local coefficients includes detected geographical clusters, please refer to Fig. 1 (B1) in the paper for a detailed example includes detected clusters. ²Templates in blue background describe the patterns; templates in green background are supportive explanations. ³The interactive user-definable location identifier that indicates the spatial distributions. Analysts can edit the location distribution narratives by themselves with the vernacular geographic language (e.g. downtown Chicago). ⁴For more narrative explanation cases and related interactions, please refer to Fig.1 (B), Fig.2 (A5), Fig.2 (B), and Fig. 5 (A).

Fig. A.1: The narrative generation pipeline and text templates with examples.