Spatio-Temporal Kriging for Texas Weather

STA 5377 Final Report

Carson Slater

May 7, 2025

Abstract

This paper details the application of spatio-temporal kriging to predict weather patterns in central Texas, shifting from an initial focus on residential water usage due to a lack of spatial and temporal correlation in that dataset. The study was motivated by the need to predict the average daily temperature at an outdoor wedding venue on a specific date. The methodology involved exploratory data analysis, spatio-temporal variogram modeling, and the implementation of Ordinary and Universal Kriging techniques. Cross-validation strategies were employed to assess model performance, and the results indicated that Universal Kriging, incorporating a spatial trend, provided more accurate predictions. The final model was used to estimate the average daily temperature at the wedding venue on October 26, 2024. The findings highlight the effectiveness of spatio-temporal kriging for weather data interpolation and suggest avenues for future work, including the incorporation of additional covariates and the exploration of non-stationary models.

1 Introduction

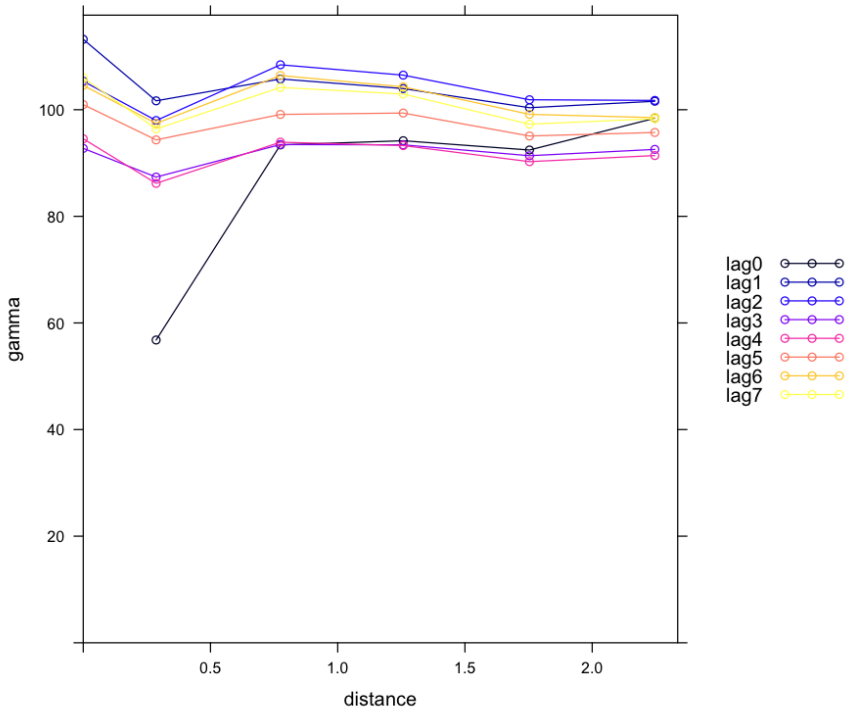

The initial focus of this project was to apply spatio-temporal kriging to model residential home water usage. However, preliminary analysis of the water usage data revealed a lack of significant spatial correlation and temporal autocorrelation, as indicated by the spatio-temporal variogram. This absence of dependence structures rendered spatio-temporal kriging unsuitable for that particular dataset. Consequently, the project's direction shifted towards a more relevant and personally significant application: predicting weather patterns for a specific date and location. This pivot was motivated by the upcoming outdoor wedding of the author and his fiancé, scheduled for October 26, 2025, at an outdoor venue. The desire was to develop a methodology to estimate the average daily temperature at the wedding venue on the corresponding date in the previous year (October 26, 2024) and to establish a framework for future precipitation forecasting. This paper details the process of conducting spatio-temporal variogram analysis, implementing kriging techniques, and generating a temperature prediction for the specified date and location in central Texas.

2 Goals

The primary objectives of this study were threefold:

1. To perform a rudimentary spatio-temporal variogram and trend analysis on weather data collected from central Texas and its surrounding areas. This involved investigating the spatial and temporal dependencies within the temperature data.
2. To identify and implement the most appropriate spatio-temporal kriging method for the given dataset. This necessitated evaluating different kriging techniques and selecting the one that best accounted for the observed spatio-temporal correlation structure.
3. To utilize the chosen kriging method to interpolate the average daily temperature at the wedding venue (Old Bethany Venue) on October 26, 2024. This prediction serves as a practical application of the developed methodology.

3 Literature Review

This study draws upon foundational and contemporary literature in the field of spatio-temporal statistics. Key resources that informed the methodological approach include:

1. Cressie and Wikle's Statistics for Spatio-Temporal Data [@cressie2015statistics] provides a canonical theoretical framework for understanding and modeling spatio-temporal data, covering various aspects from exploratory analysis to advanced modeling techniques like kriging.
2. Cressie, Wikle, and Zammit-Mangion's Spatio-Temporal Statistics with R [@wikle2019spatio] offers practical guidance on implementing spatio-temporal statistical methods using the R programming language.
3. Dhaher et al. [@dhaher2023kriging] compared kriging with a dynamic factor model for forecasting temperatures in Iraq, finding that kriging outperformed the dynamic factor model for spatial prediction.
4. Huang et al. [@huang2024hybrid] proposed a hybrid spatial interpolation method that integrates kriging with machine learning algorithms to improve meteorological data estimation in complex terrains.

These textbooks provided the necessary theoretical background and practical guidance for conducting the spatio-temporal analysis and kriging implemented in this project. The papers provided helpful context on where the methodology is being developed currently.

4 The Data

Originally, the dataset utilized in this study comprises average daily temperature (`tavg`) readings collected from 12 weather stations across Texas. The temporal extent of the data spans from September 1, 2024, to October 31, 2024. This period was chosen to encompass the time leading up to the target prediction date (October 26, 2024). The data were obtained from the National Centers for Environmental Information (NCEI) (National Centers for Environmental Information 2025), a reliable source for meteorological information. These data were later supplemented with records from TexMesonet (Texas Mesonet 2025), accessed on May 10, 2025.

Following the final presentation, feedback was given that the standard error estimates for the kriging results were far too high. A search for more data began, and more data was found. Ultimately, there ended up being 92 weather stations that had daily average temperatures for their respective locations from September 1, 2024 through October 31, 2024. These data were obtained from TexMesonet and accessed on May 10, 2025. Figure 1 shows the locations of these weather stations, in addition to the location of the wedding venue of interest where we are interested in forecasting the temperature. The location of interest, the Old Bethany wedding venue, is described further on their website (Old Bethany Venue 2025).

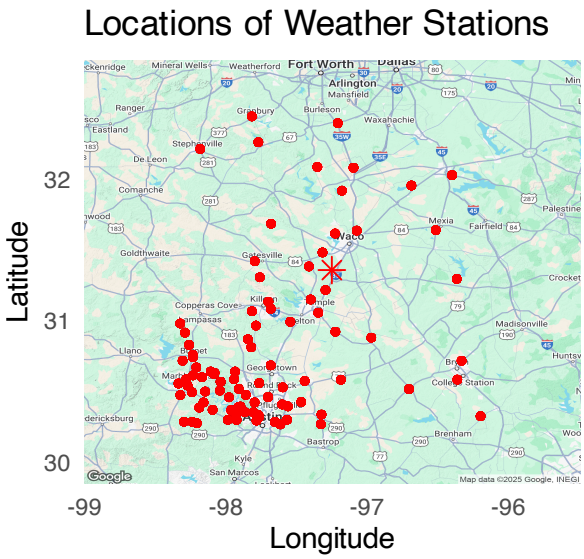

Figure 1: Locations of the 92 weather stations from the Texas Mesonate database.

5 Exploratory Analysis

All modeling was conducted in R (R Core Team 2024) using the gstat (Pebesma 2023) and spacetime (Pebesma and Graeler 2023) packages for kriging and variogram modeling. Machine learning models were implemented using the tidymodels framework (Kuhn and Wickham 2024). Prior to implementing kriging, exploratory data analysis was conducted to understand the characteristics of the temperature data and inform subsequent modeling steps.

5.1 Transforming the Data

Kriging assumes that the underlying spatial process can be modeled as a Gaussian process. To better align the temperature data with this assumption, the original NCEI data (12 stations) required a power transformation was considered using the Box-Cox method. The Box-Cox transformation aims to find an optimal power (*lambda*) to make the data more normally distributed and stabilize the variance. The analysis suggested a transformation with $\lambda \approx 1.434$. Applying this transformation aimed to improve the normality of the data, which is beneficial for kriging.

For the Texas Mesonet data (92 stations), a tranformation was not deemed necessary. When grouping by date, only 12 of the 62 days rejected the null hypothesis for the Shapiro-Wilk normality test at the $\alpha = 0.05$ level after the Benjamini-Hochberg false-discovery rate adjustment. Figure 2 shows histograms of select dates in the training data before October 19, 2024.[1] This is likely because having eighty more weather stations in the sample better informed the sample.

5.2 Empirical Spatial Mean

To visualize the spatial trends in the average daily temperature, the empirical spatial mean was calculated for each weather station over the study period (September-October 2024). Mapping these means revealed the spatial variability of temperature across the region.

5.3 Empirical Temporal Mean

Similarly, the empirical temporal mean of the average daily temperature was calculated across all spatial locations for each day within the study period. Plotting this temporal mean against time allowed for the

[1]Note not all dates are shown for brevity's sake.

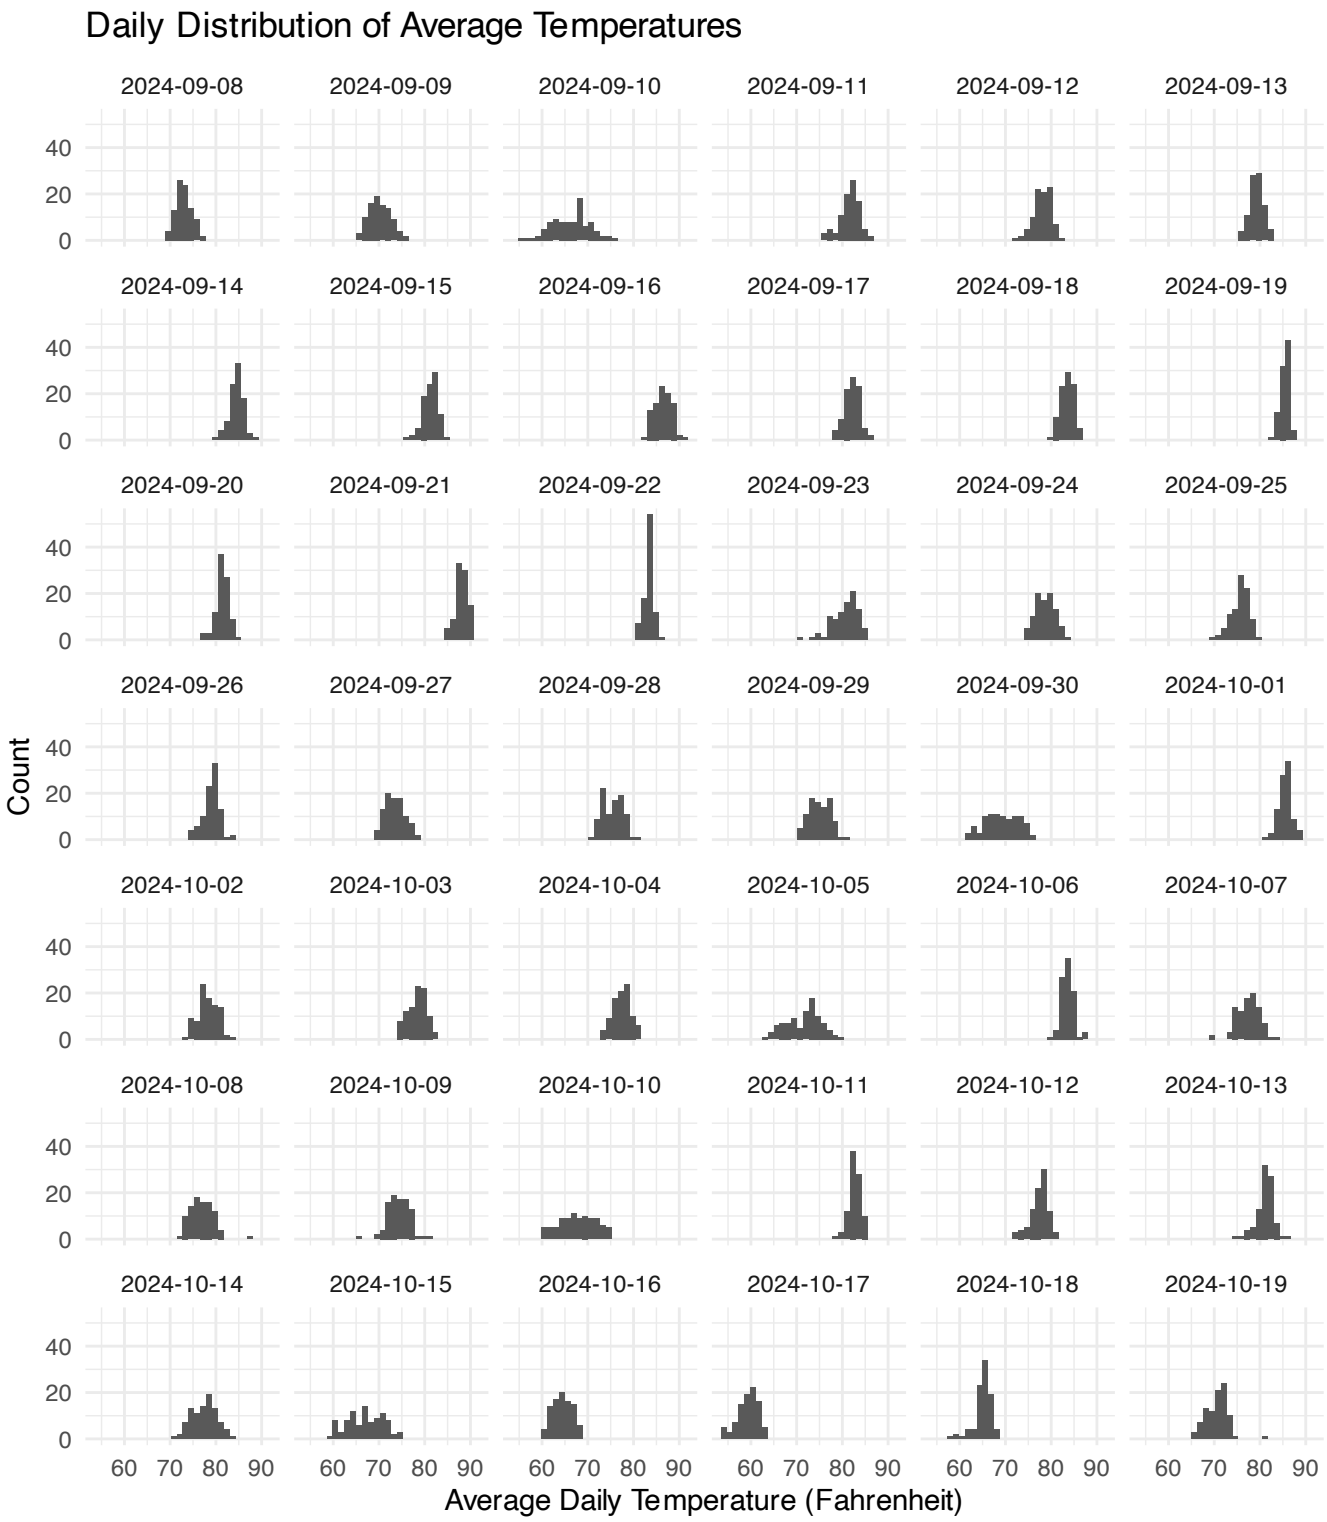

Figure 2: Visualizing the distribution of the transformed data across time, for selected dates ranging from September 7 - Octber 19, 2024.

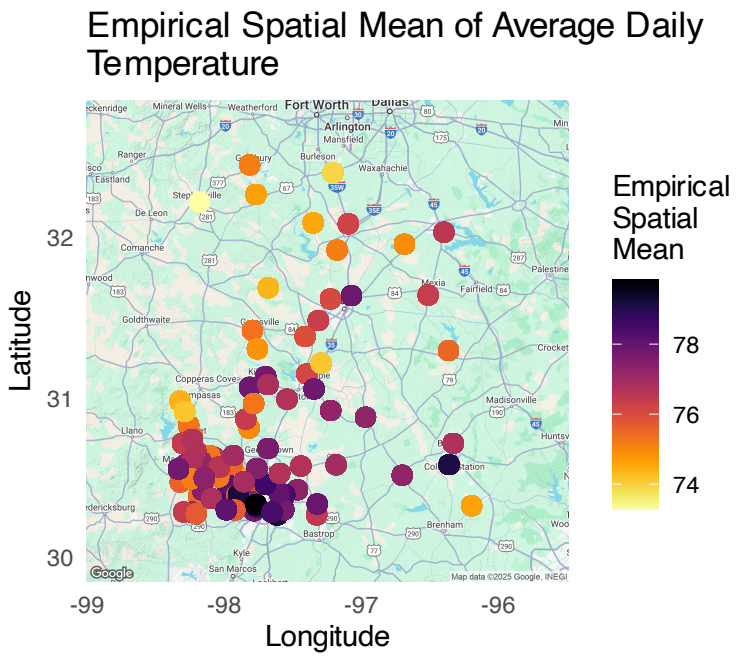

Figure 3: The empirical spatial mean of temperature (Fahrenheit) from Sepember 1, 2024 to October 31, 2024.

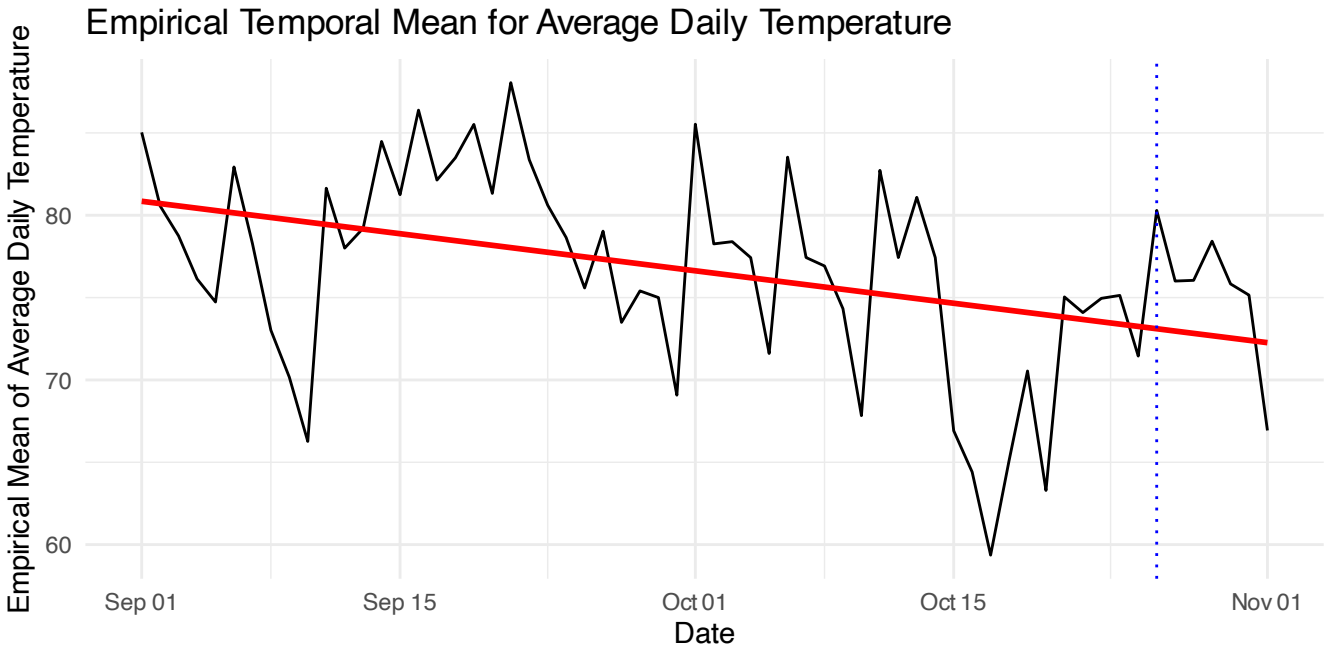

Figure 4: The empirical temporal mean of temperature (Fahrenheit) from Sepember 1, 2024 to October 31, 2024.

examination of temporal trends and the identification of any systematic changes in temperature over the two-month period. A linear trend line was fitted to the temporal mean to identify overall decreasing trend. The date of interest, October 26, 2024, was also marked on this plot for reference.

5.4 Spatial Correlation

To formally assess the presence of spatial autocorrelation in the temperature data (`tavg`) for each day, the global Moran's I test was performed. This test quantifies the degree of spatial clustering or dispersion of the data. For each day, a spatial weights matrix was constructed based on the k-nearest neighbors (with a minimum of 4 neighbors). The Moran's I statistic and its corresponding p-value were calculated. To account for multiple testing across the days, the Benjamini-Hochberg (BH) procedure was used to adjust the p-values. The proportion of days with statistically significant spatial autocorrelation (at an $\alpha = 0.05$ level after adjustment) was then determined. The results indicated that a substantial portion of the days exhibited significant spatial correlation, supporting the use of spatial interpolation techniques like kriging. Overall, 54 of the 62 days (approximated 87%) exhibited significant spatial autocorrelation globally. We have reason to believe that there exists spatial autocorrelation in these data, which was a problem with the inital residential water consumption dataset.

5.5 Anisotropy

Anisotropy, the directional dependence of spatial correlation, was investigated by computing directional spatial variograms for each day in the training dataset. Variograms were calculated in different directions (defined by angles) to assess whether the spatial variability of temperature differed based on the direction. Visualizing these directional variograms across time allowed for the examination of any consistent anisotropic patterns in the spatial correlation structure. Visual assessment yielded the conclusion that this data is anisotropic.

5.6 Temporal Autocorrelation

Temporal autocorrelation, the correlation of a variable with itself across time, was examined by aggregating the transformed average daily temperature by date and calculating the autocorrelation function (ACF) of this daily mean. The ACF plot provides insights into the temporal dependencies present in the data, indicating how strongly temperature values at one time point are correlated with values at previous time points. Figure 5 shows the ACF graphic of the daily mean average temperature across all 62 spatial locations. There appears to be significant autocorrelation at the first and second lags, but not after that.

6 Spatio-Temporal Variogram

To model the spatio-temporal correlation structure of the temperature data, a spatio-temporal variogram analysis was conducted. This involves examining how the semivariance (a measure of dissimilarity) between data points changes with both spatial distance and temporal lag.

6.1 Joint Space-Time Variogram

The analysis focused on a joint space-time variogram, considering both the spatial distance between weather stations (in degrees of latitude/longitude) and the temporal lag between observations (in days, ranging from 0 to 7 days). A spatial bin width of 0.5 degrees was used for aggregating spatial pairs.

6.2 Model Selection

Four common types of spatio-temporal variogram models were considered:

1. **Separable Model**: This model assumes that the spatial and temporal dependencies are independent and can be modeled separately, with the spatio-temporal variogram being the product of a purely

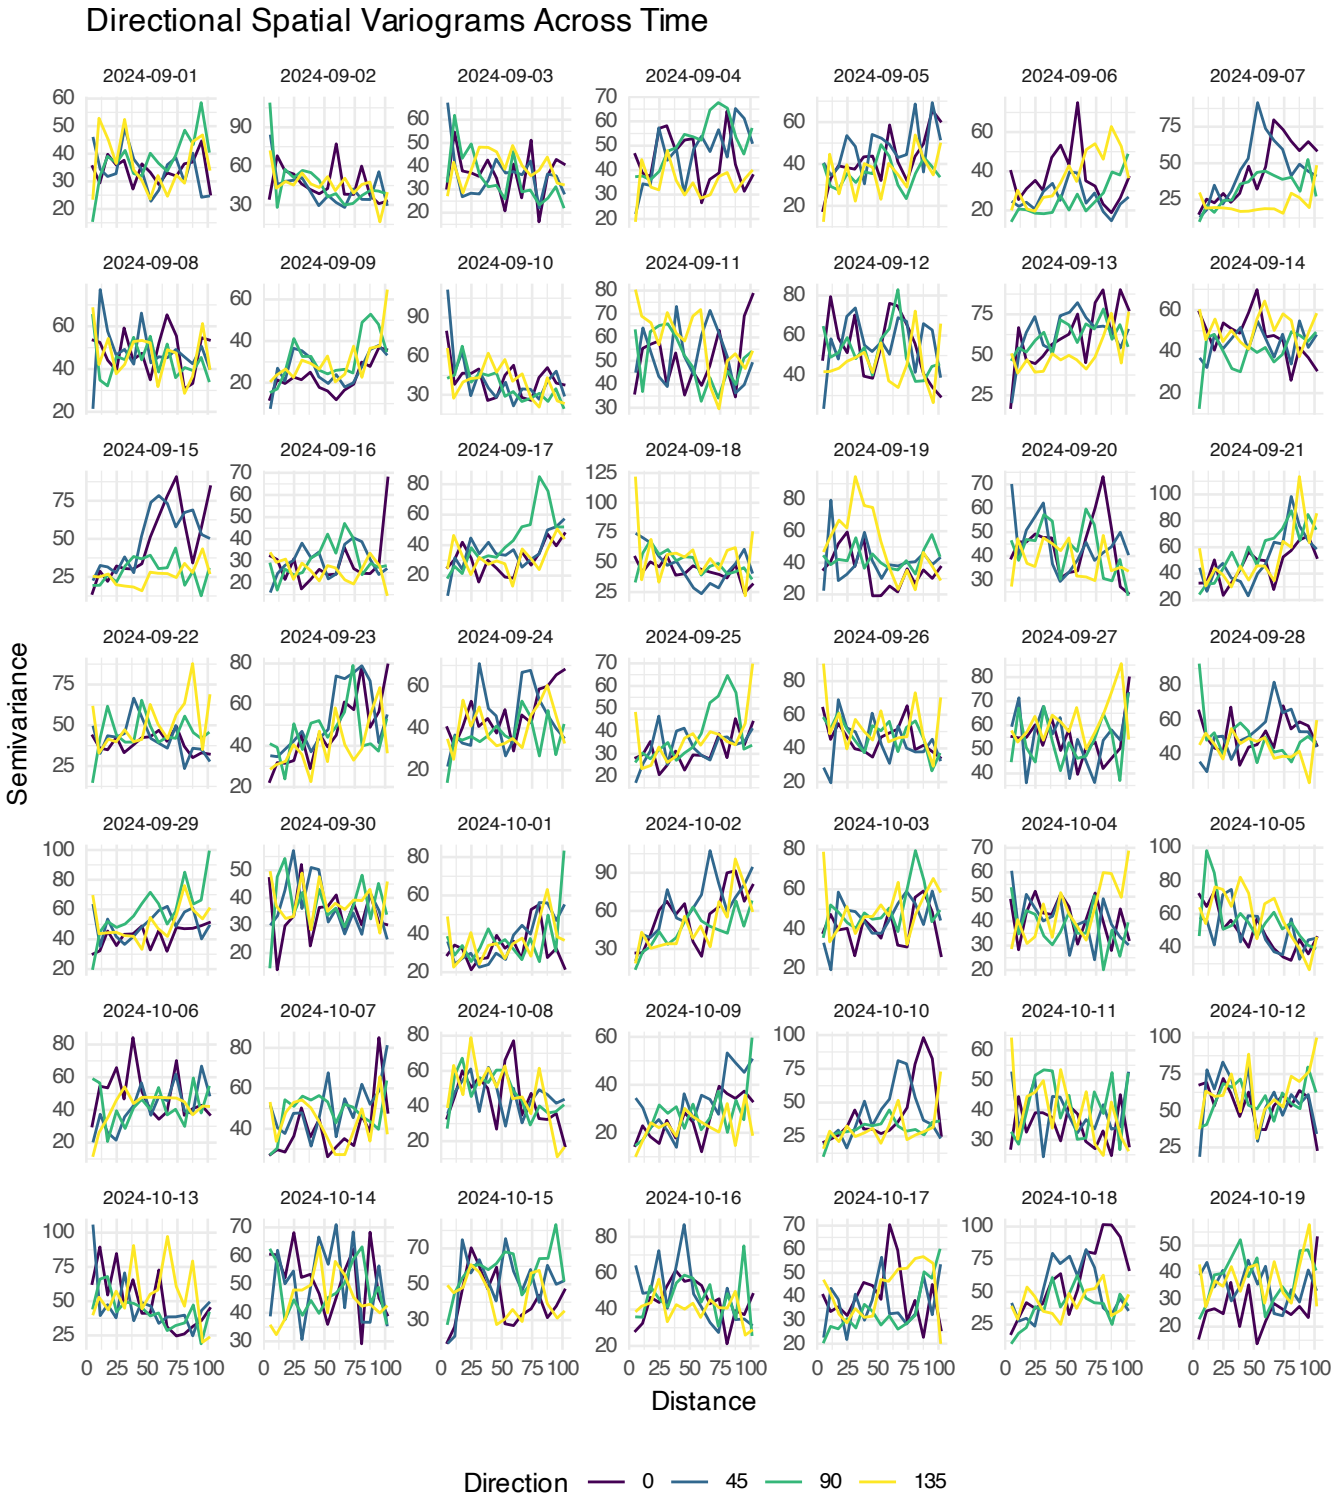

Figure 5: Directional variograms to assess anisotropy.

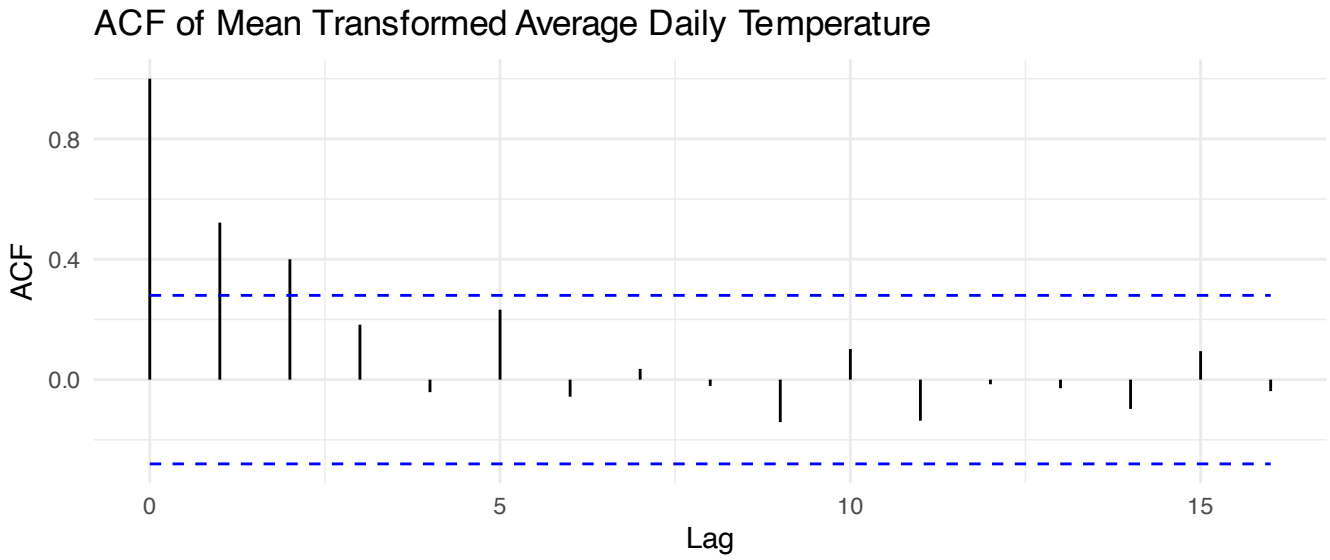

Figure 6: Temporal autocorellation plot for extttavg.

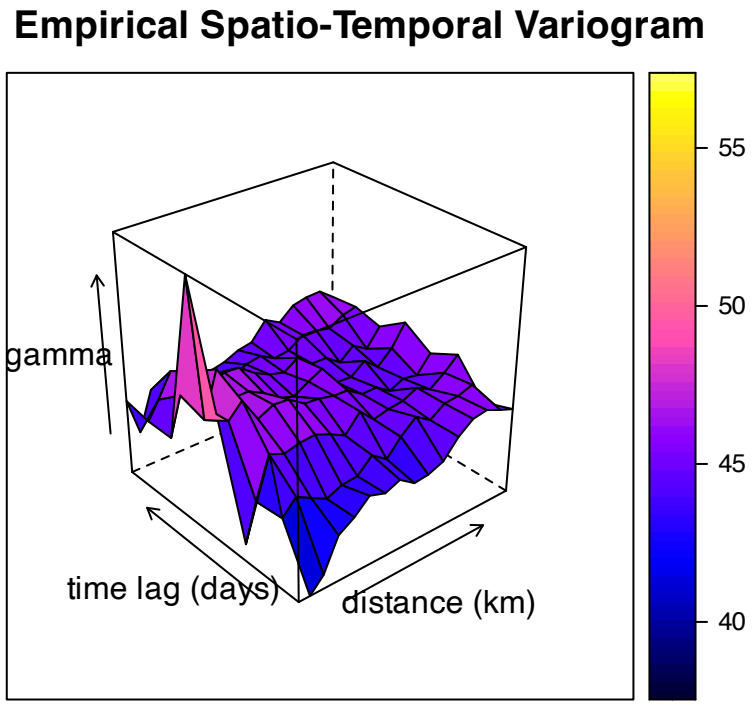

Figure 7: Empirical spatio-temporal variogram for the Texas Mesonet weather data.

spatial and a purely temporal variogram.

$$\gamma_{ST}(h_s, h_t) = \gamma_S(h_s) \times \gamma_T(h_t)$$

2. **Product-Sum Model**: This model allows for interaction between the spatial and temporal components, represented as a weighted sum of the product of the spatial and temporal variograms and the individual spatial and temporal variograms.

$$\gamma_{ST}(h_s, h_t) = k\gamma_S(h_s)\gamma_T(h_t) + \gamma_S(h_s) + \gamma_T(h_t)$$

3. **Metric Model**: This model uses a combined space-time distance metric, where temporal distances are scaled by a factor (α) to make them comparable to spatial distances. The spatio-temporal variogram is then modeled as a function of this combined distance.

$$\gamma_{ST}(h_s, h_t) = \gamma_J \left(\sqrt{h_s^2 + \left(\frac{h_t}{\alpha}\right)^2} \right)$$

4. **Sum-Metric Model**: This model combines the properties of both product-sum and metric models by expressing the spatio-temporal variogram as the sum of a purely spatial component, a purely temporal component, and a joint space-time metric component.

$$\gamma_{ST}(h_s, h_t) = \gamma_S(h_s) + \gamma_T(h_t) + \gamma_{ST} \left(\sqrt{h_s^2 + \left(\frac{h_t}{\alpha}\right)^2} \right).$$

6.3 Variogram Fitting

Sample spatio-temporal variograms were computed from the data, and the three candidate models were fitted to these empirical variograms using a weighted least squares approach. The parameters of each model (e.g., nugget, sill, range for the spatial and temporal components, and the interaction parameter k for the product-sum model, and α for the metric model) were estimated by minimizing the difference between the model and the empirical variogram values.

6.4 Variogram Fitting: Results

All four spatio-temporal variogram models were successfully fitted to the data. The estimated parameters for each model are presented in the table below.

The fitted models reveal interesting characteristics about the spatio-temporal structure of the data. The Separable model maintained non-zero partial sills for both spatial (Exponential, range = 300.0) and temporal (Spherical, range = 4.1) components, with substantial nugget effects in both dimensions. This suggests that while the model captures some structured spatial and temporal variation, there is considerable small-scale variability or measurement error.

The Product-Sum model converged to a solution with zero partial sills for both spatial and temporal components, with the interaction parameter $k = 7.47$ indicating strong spatio-temporal interactions. The high nugget values suggest this model primarily captures the variance through interaction effects rather than through separate spatial and temporal structures.

The Metric model, which uses a joint space-time distance function, showed a moderate partial sill (7.258) with a relatively short effective range (36.0) and a high nugget effect (38.429). The space-time anisotropy parameter (stAni = 47.18) indicates the scaling factor between spatial and temporal distances.

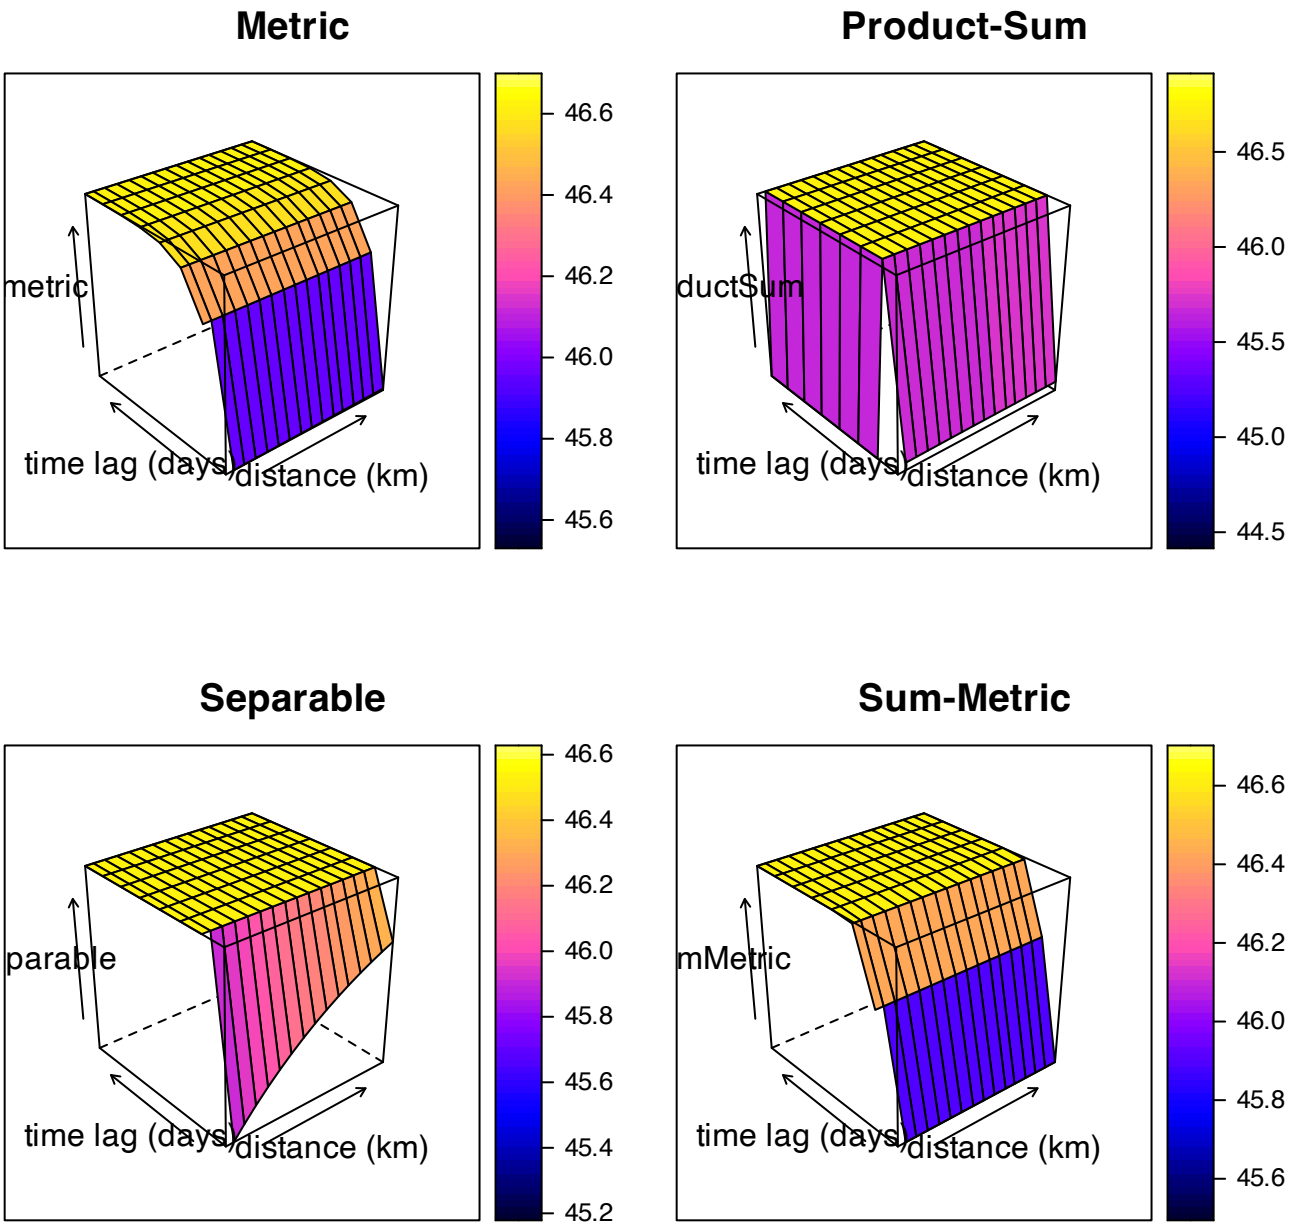

Model	Component	Structure	Psill	Range	Nugget
Separable	Space	Exponential	0.058	300.0	0.942
	Time	Spherical	0.028	4.1	0.972
	Joint parameters	Sill: 45.49			
Product-Sum	Space	Exponential	0.000	250.5	2.147
	Time	Spherical	0.000	38.5	2.553
	Joint parameters	k: 7.47			
Metric	Joint	Exponential	7.258	36.0	38.429
	Joint parameters	stAni: 47.18			
Sum-Metric	Space	Exponential	0.477	250.4	0.000

Model	Component	Structure	Psill	Range	Nugget
	Time	Spherical	0.000	0.7	0.000
	Joint	Exponential	3.661	142.6	41.981
	Joint parameters	stAni: 154.45			

Finally, the Sum-Metric model, which combines aspects of both additive and metric approaches, revealed a moderate spatial correlation structure (Exponential, range = 250.4, psill = 0.477) but negligible pure temporal structure. The joint component showed considerable structured variation (psill = 3.661, range = 142.6) alongside a large nugget effect (41.981). The high space-time anisotropy (stAni = 154.45) suggests that spatial distances contribute much more strongly to the joint distance metric than temporal distances.

Based on these results, the Sum-Metric model appears to provide the most detailed characterization of the spatio-temporal dependence structure, capturing both pure spatial effects and joint space-time interactions, while the Product-Sum model emphasizes interaction effects with minimal pure spatial or temporal components.

7 Kriging and Results

Kriging, a geostatistical interpolation technique, was employed to predict temperature values at unsampled locations and times. Several cross-validation strategies were implemented to assess the performance of different kriging methods.

7.1 Cross-Validation Strategies

5-fold Spatial Cross-Validation: This approach divides the spatial locations into 5 folds randomly. For each fold, the stations within that fold are held out, and the model is trained on the remaining stations to predict the temperature at the held-out locations. This provides a more robust evaluation of spatial prediction accuracy by testing the model on spatially distinct subsets of the data. Table 1 shows the results from the variogram models for 5-fold cross validation. For some unknown reason we were unable to compute results for the separable and product-sum models. It is likely that sparsity in the folds yielded issues with computation.

Model	RMSE
Metric	6.95
Sum-Metric	8.36

Table 1: 5-Fold Cross-Validation Metrics for Universal Kriging.

Fixed Time Cross Validation: In this approach, a spatio-temporal kriging model was trained on data from September 1, 2024 through October 19, 2024, and predictions were made for the data after that date. We performed this using ordinary kriging, universal kriging, and a few machine learning models. The performance metrics (RMSE and MAE) were then calculated to assess the model's ability to forecast future temperatures. For the machine learning models, a random forest, XGBoost, and single-layer neural network was considered. These models had hyperparameters tuned on the training data using 5-fold cross validation, where the best models were selected using RMSE. All models were trained and tested using the `tidymodels` modeling framework. The machine models used space and time as predictors. The table below shows the machine-learning models used:

Model	Preprocessor Steps	Engine	Model Type	Main Arguments
Random Forest	`step_date()`, `step_normalize()`	ranger	Random Forest (regression)	`mtry = 3`, `trees = 500`, `min_n = 8`

Model	Preprocessor Steps	Engine	Model Type	Main Arguments
XGBoost	`step_date()`, `step_normalize()`, `step_dummy()`	xgboost	Boosted Tree (regression)	`mtry = 4`, `trees = 500`, `tree_depth = 4`, `learn_rate = 1.023`
Neural Network	`step_date()`, `step_normalize()`	nnet	Neural Network (regression)	`hidden_units = 7`, `penalty = 1e-04`, `epochs = 200`

Overall, the performance results can be show in the table below for the fixed-time cross validation (i.e. forecasting):

	UK Metric Variogram Model	Neural Network	Random Forest	XGBoost
RMSE	5.790888	8.723968	8.145796	8.037257
MAE	2.011414	8.030071	7.522727	6.951712

Discussion of Model Performance

The performance of Universal Kriging, employing the Metric variogram model, was rigorously compared against several machine learning methodologies, specifically Neural Networks, Random Forests, and XGBoost, within a fixed-time cross-validation framework. Evaluation metrics included both the Root Mean Squared Error (RMSE) and the Mean Absolute Error (MAE). The results clearly indicate that Universal Kriging demonstrated superior accuracy, achieving the lowest RMSE of 5.791 and the lowest MAE of 2.011 across all models. This outcome signifies that Universal Kriging not only minimizes the average magnitude of prediction errors but also reduces the average absolute error in comparison to machine learning techniques. Among the machine learning models, XGBoost exhibited the relatively strongest performance, registering lower RMSE and MAE values than both Neural Networks and Random Forests. [cite: 118] In conclusion, these quantitative findings provide compelling evidence for the enhanced effectiveness of Universal Kriging, particularly when integrated with the Metric variogram model, for temperature prediction within this study, as it consistently outperformed alternative kriging approaches and machine learning algorithms.

Predicting Across a Grid

The accompanying figures visually represent the output of spatio-temporal kriging, providing both predicted temperature values and their associated standard errors across a spatial grid over a sequence of time steps. Figure 7 illustrates the kriging predictions themselves, with a color gradient indicating the estimated temperature in degrees Fahrenheit for each grid cell and time point. These plots reveal a relatively smooth spatial pattern in predicted temperatures across the region, with a slight temporal shift in the overall temperature range over the 12-day period from October 20th to October 31st, 2024. Complementing this, Figure 8 displays the standard errors associated with these predictions. Notably, the standard errors exhibit a consistent spatial pattern and magnitude across all time steps, suggesting a uniform level of uncertainty in the predictions throughout the period. The relatively low standard error values, as indicated by the color scale, imply a high degree of confidence in the kriging estimates. Together, these graphics offer a comprehensive view of both the predicted temperature trends and the reliability of those predictions, highlighting the spatio-temporal dynamics of temperature in the study area.

The accompanying figures visually represent the output of spatio-temporal kriging, providing both predicted temperature values and their associated standard errors across a spatial grid over a sequence of time steps. The first set of plots, labeled 'Predicted Average Daily Temperature,' illustrates the kriging predictions themselves, with a color gradient indicating the estimated temperature in degrees Fahrenheit for each grid cell and time point. These plots reveal a relatively smooth spatial pattern in predicted temperatures across the region, with a slight temporal shift in the overall temperature range over the 12-day period from October 20th to October 31st, 2024. Complementing this, the second set, 'St. Errors for Average Daily Temperature,' displays the standard errors associated with these predictions. Notably, the standard errors exhibit a consistent spatial

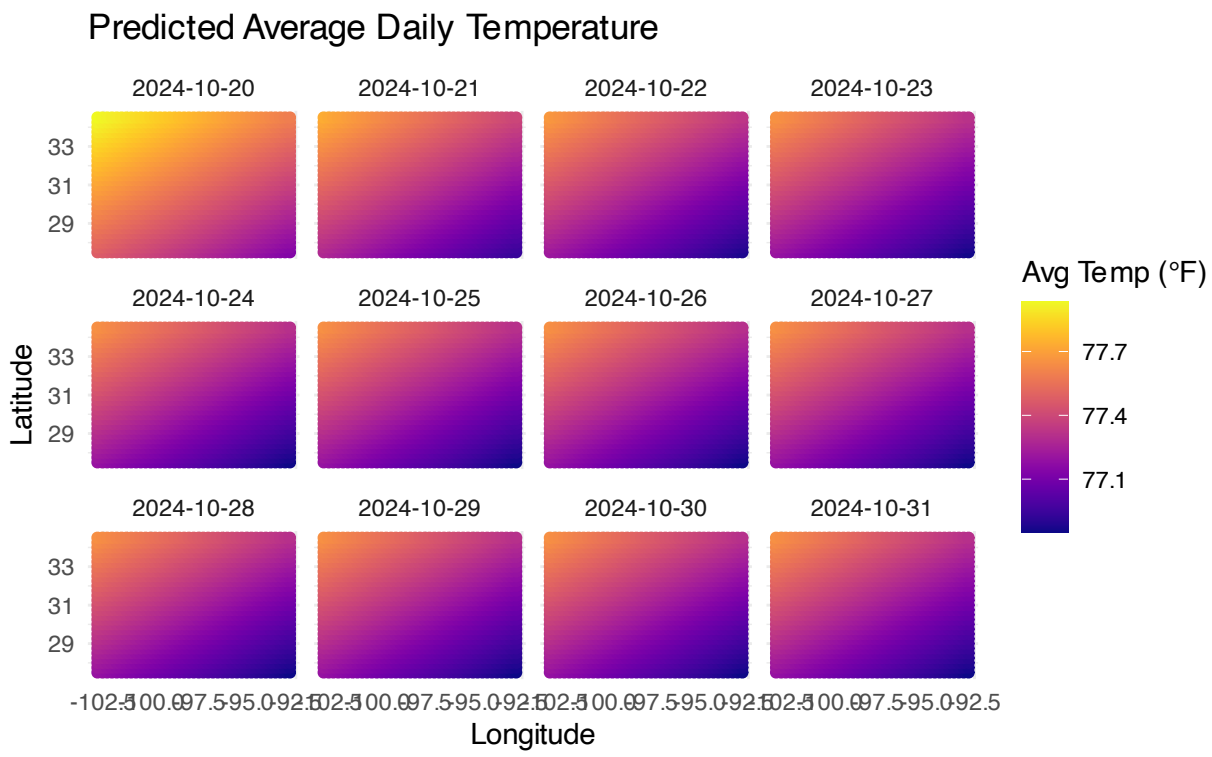

Figure 8: Spatio-temporal predictinos from universal kriging on a grid centered at the wedding venue.

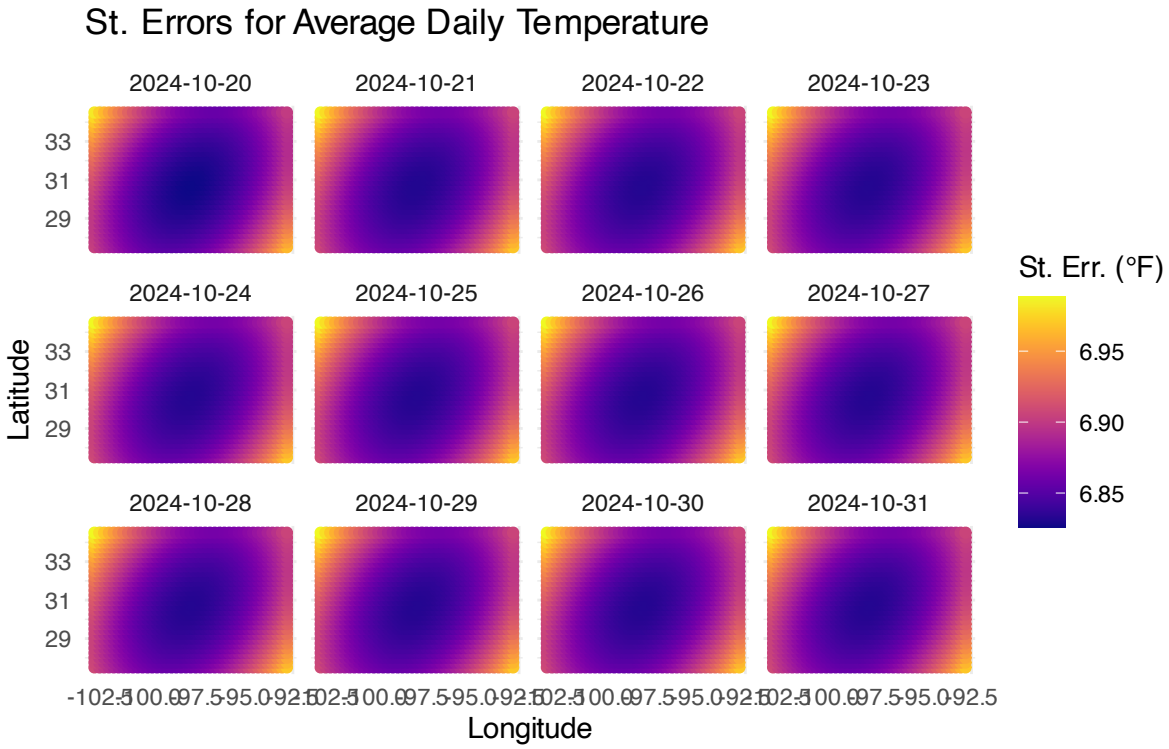

Figure 9: Spatio-temporal predictinos from universal kriging on a grid centered at the wedding venue.

pattern and magnitude across all time steps, suggesting a uniform level of uncertainty in the predictions throughout the period. The relatively low standard error values, as indicated by the color scale, imply a high degree of confidence in the kriging estimates. Together, these graphics offer a comprehensive view of both the predicted temperature trends and the reliability of those predictions, highlighting the spatio-temporal dynamics of temperature in the study area.

7.2 Improvement from the Original Predictions

It is noteworthy that these predictions are substantially better than the original Texas NCEI weather data predictions. For those, the average standard were generally upwards of 60 degrees Fahrenheit (which means the variance was upwards of 3600). These predictions were pragmatically useless since their variance was so high. Increasing the number of spatial locations from 12 to 92 decreased the standard errors 10-fold using the same approach. The authors were very pleased to see these results, albeit there are still many limitations in the project and improvements that can be made.

7.3 Prediction at Old Bethany Wedding Venue

The honest truth is that because we were unable to implement universal kriging with temporal trend, our models indicated small change over time for these dates. Based on our models, we would expect the average daily temperature at Old Bethany wedding venue on October 26 to be approximately 77.4 degrees Fahrenheit.

8 Limitations

A key limitation of this study is the relatively short temporal window of the dataset, which spans only two months (September through October 2024). This restricted timeframe limits the ability of the model to capture broader seasonal trends, interannual variability, or long-range temporal dependencies that may influence temperature patterns. As a result, while the kriging predictions for the fall of 2024 are likely reasonable, the generalizability of these results to future years—particularly for planning purposes related to the wedding date in 2025—is inherently constrained. Furthermore, the inability to implement universal kriging with an explicit temporal trend component further limits the capacity to model gradual temporal changes and trends in temperature that might otherwise enhance the predictive power of the approach.

A second limitation pertains to the computational challenges encountered during model fitting and cross-validation. Specifically, the separable and product-sum variogram models failed to produce stable results under 5-fold spatial cross-validation. This failure may be attributed to data sparsity or model instability within certain spatial partitions, but it nonetheless restricted the scope of model comparison and precluded a full evaluation of their potential performance. In addition, the lack of external validation using data from a different time period (e.g., a different year or season) means that the robustness and generalizability of the model cannot be fully assessed. Although Universal Kriging with the metric variogram performed well relative to machine learning benchmarks, its effectiveness beyond the observed data window remains uncertain.

9 Discussion

This study encountered and surmounted several challenges, ultimately leading to a more robust analysis and insightful results.

The initial focus of the research was to apply spatio-temporal kriging to model residential home water usage. However, the preliminary analysis revealed a lack of significant spatial and temporal correlation in the water usage data, rendering spatio-temporal kriging unsuitable for this purpose. This led to a shift in the study's direction towards predicting weather patterns, specifically average daily temperature, for a specific date and location.

A significant challenge was the high standard errors in the initial kriging results, which made the predictions pragmatically useless. To address this, a search for additional data was conducted, which substantially

increased the number of weather stations from 12 to 92. This increase in spatial data points led to a 10-fold decrease in the standard errors, greatly improving the precision and reliability of the predictions.

Despite these improvements, it is important to acknowledge the limitations of the study. The relatively short temporal window of the dataset (September to October 2024) limits the model's ability to capture broader seasonal trends and long-range temporal dependencies. Additionally, computational challenges were encountered during model fitting and cross-validation, specifically with the separable and product-sum variogram models, which may be attributed to data sparsity or model instability within certain spatial partitions.

Bibliography

Campitelli, Ethan. 2023. *Ggnewscale: Multiple Fill and Colour Scales in Ggplot2*. https://cran.r-project.org/package=ggnewscale.

Cressie, Noel, and Christopher K. Wikle. 2015. *Statistics for Spatio-Temporal Data*. John Wiley & Sons.

Dhaher, Hadi A., Hadi M. Hassan, and Suhaib T. Abbas. 2023. "A Comparative Study of Spatio-Temporal Kriging and Dynamic Factor Models for Temperature Forecasting in Iraq." *Iraqi Journal of Statistical Science* 23 (1): 15–30.

Fox, John, and Sanford Weisberg. 2023. *Car: Companion to Applied Regression*. https://cran.r-project.org/package=car.

Huang, Wei, Li Chen, and Yu Zhang. 2024. "A Hybrid Spatial Interpolation Framework Integrating Machine Learning and Kriging for Meteorological Data Estimation." *International Journal of Climatology* 44 (2): 229–46.

Kahle, David, and Hadley Wickham. 2023. *Ggmap: Spatial Visualization with Ggplot2*. https://cran.r-project.org/package=ggmap.

Kuhn, Max, and Hadley Wickham. 2024. *Tidymodels: Easily Install and Load the 'Tidymodels' Packages*. https://www.tidymodels.org/.

National Centers for Environmental Information. 2025. "NCEI Climate Data Online." https://www.ncei.noaa.gov/cdo-web/search?datasetid=GHCND.

Old Bethany Venue. 2025. "Old Bethany Wedding Venue." https://oldbethanyvenue.com.

Padgham, Mark, Bob Rudis, Robin Lovelace, and Morgan Salmon. 2023. *Osmdata: Import 'OpenStreetMap' Data as Simple Features or Spatial Objects*. https://cran.r-project.org/package=osmdata.

Pebesma, Edzer J. 2023. *Gstat: Spatial and Spatio-Temporal Geostatistical Modelling, Prediction and Simulation*. https://cran.r-project.org/package=gstat.

Pebesma, Edzer J., and Benjamin Graeler. 2023. *Spacetime: Classes and Methods for Spatio-Temporal Data*. https://cran.r-project.org/package=spacetime.

R Core Team. 2024. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.R-project.org/.

Sjoberg, Daniel D., Kyle Whiting, Michael Curry, and James A. Lavery. 2023. *Gtsummary: Presentation-Ready Data Summary and Analytic Result Tables*. https://cran.r-project.org/package=gtsummary.

Texas Mesonet. 2025. "Custom Weather Data Downloads." https://www.texmesonet.org/DataProducts/CustomDownloads.

Wikle, Christopher K., Andrew Zammit-Mangion, and Noel Cressie. 2019. *Spatio-Temporal Statistics with r*. Chapman; Hall/CRC.