

UC Riverside

UC Riverside Electronic Theses and Dissertations

Title

A-GWR: Fast and Accurate Geospatial Inference via Augmented Geographically Weighted Regression

Permalink

<https://escholarship.org/uc/item/2s85s181>

Author

Zare Shahneh, Mohammad Reza

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
RIVERSIDE

A-GWR: Fast and Accurate Geospatial Inference via Augmented Geographically
Weighted Regression

A Thesis submitted in partial satisfaction
of the requirements for the degree of

Master of Science

in

Computer Science

by

Mohammad Reza Zare Shahneh

June 2021

Thesis Committee:

Amr Magdy, Chairperson
Samet Oymak
Ahmed Eldawy

Copyright by
Mohammad Reza Zare Shahneh
2021

The Thesis of Mohammad Reza Zare Shahneh is approved:

Committee Chairperson

University of California, Riverside

Acknowledgments

First and foremost, I would like to thank my research supervisors, Prof. Magdy and Prof. Oymak. Without their support and assistance in every step throughout the process, this work would have never been accomplished. I would like to thank you both very much for your support and understanding.

I would like to pay my special regards to my committee member, professor Ahmed Eldawy whose work was, and is, an inspiration to me. My partner Saba – I simply couldn't have done this without you, thank you for making me believe in myself.

Most importantly, I must express my very profound gratitude to my parents and my brother for providing me with unfailing support and continuous encouragement throughout my years of study.

ABSTRACT OF THE Thesis

A-GWR: Fast and Accurate Geospatial Inference via Augmented Geographically Weighted Regression

by

Mohammad Reza Zare Shahneh

Master of Science, Graduate Program in Computer Science
University of California, Riverside, June 2021
Amr Magdy, Chairperson

Geographically Weighted Regression (GWR) is a seminal technique with rich applications in geospatial data analysis. However, it has critical drawbacks in the age of big data in terms of expressiveness, i.e., predictive power, and scalability. This work proposes Augmented GWR (A-GWR) that alleviates these drawbacks. A-GWR adapts a novel technique, Stateless-MGWR or S-MGWR, that enriches the predictive power by allowing distinct bandwidths for individual features of the training data. S-MGWR uses a customized black-box optimization approach for discovering optimal bandwidths in a fast and efficient way. In addition, A-GWR modularly combines S-MGWR or GWR with versatile models such as random forest models. Moreover, A-GWR enables scalability by operating on flexible partitions of the data that can adapt to the computational budget. Our extensive evaluations on various real and synthetic datasets demonstrate the scalability and accuracy benefits of the proposed techniques over state-of-the-art competitors.

Contents

List of Figures	vii
List of Tables	viii
1 Introduction	1
2 Related Work	6
3 Preliminaries	9
3.1 GWR	9
3.2 MGWR	10
4 Stateless-MGWR	12
4.1 Spatial Weighting Scheme	17
4.2 Bandwidth Search Strategies	18
4.2.1 Implementation Details	19
4.3 Prediction Procedure	21
4.4 Accelerating Weight Computation	22
5 Augmented GWR (A-GWR)	25
5.1 Data Splitting	26
5.2 Model Integration	27
6 Computational Complexity	30
6.1 Time Complexity	30
6.2 Space Complexity	31
7 Experimental Evaluation	32
7.1 Model Tuning	36
7.2 Model Performance	38
8 Conclusions	42
Bibliography	43

List of Figures

4.1	S-MGWR Training Phase Algorithm	13
4.2	Compute Error Algorithm	15
4.3	Adaptive bandwidth illustration.	16
4.4	Timeline of error reduction for different bandwidth search methods.	22
4.5	Predict Algorithm	23
5.1	Splitting a sample data into eight sections using grid and k-means methods.	27
5.2	Ensemble model integration order (general-purpose model first).	28
5.3	Pipeline model integration order (spatial model first).	29
7.1	Illustration of the covariate functions. β_0 to β_4 are displayed from left to right, respectively	34
7.2	Accuracy and runtime over synthData1 and kingHousePrices datasets	37

List of Tables

7.1	Error value of different bandwidth search methods	36
7.2	Errors and bandwidths computed by each search method	38
7.3	Performance of S-MGWR, GWR, and MGWR. S-MGWR* is the S-MGWR model with bandwidth computed by MGWR.	39
7.4	Evaluating A-GWR as a framework.	40
7.5	Performance of A-GWR, MGWR, GWR, and Random Forest over different datasets	41

Chapter 1

Introduction

Regression is one of the most traditional and popular learning models in excessive number of applications. In spatial applications, regression models of different types are used for various use cases [4, 36, 34, 2, 27, 26, 40, 3, 10, 13, 35]. Geographically Weighted Regression (GWR) is one of the most popular spatially-varying coefficient (SVC) models [11]. It has been widely adopted in many areas such as climate science [2], health and health care [27], criminology [40], transportation analysis [3], house price modeling [10], data analysis and visualization [13, 35]. GWR and its variations specifically are popular and have been used recently for COVID-19 pandemic data in analyzing the spatio-temporal effects of its driving factors to limits its spread worldwide [26]. SVC models are powerful alternatives to most traditional methods of modeling data when dealing with spatial data. This is because in traditional methods, such as ordinary least squares, the role of location is often neglected or is not properly captured. Despite its efficacy, GWR has two major limitations: expressiveness and runtime scalability. These limitations restricts utilizing GWR for mod-

ern geospatial applications that involve diverse features and thousands to millions of data points. We outline both limitations as follows:

(1) **Expressiveness:** GWR uses a single number, known as bandwidth, to control the contribution of neighboring data around each location. This bandwidth plays the role of the spatial influence parameter, so a point in a location L is only affected by other points that are within L 's bandwidth but not beyond. GWR uses one bandwidth for all features, which enforces same spatial impact for different features and reduces the expressive power of the model. In reality, different features potentially have effects at different spatial scales. For example, a feature that represents *wind speed* might be spatially correlated at a state-level while another feature that expresses *temperature* is spatially correlated at a county-level. Multiple works [25], [41], [19] attempted to overcome this issue with providing a vector of bandwidths replacing the single bandwidth value, assigning a control knob to each of the features. One of the recent notable advancements is Multiscale GWR (MGWR) [12]. MGWR assigns different bandwidths for different features, thus, enabling each parameter surface to operate on a different locality. MGWR uses the backfitting algorithm for calibration involving many uni-variate GWR calibrations in each step. These calibrations results in considerably higher run time than regular GWR. In addition, due to the nature of backfitting, it is not possible to produce the same results only using the learned bandwidth; the complete history of bandwidths throughout the process is required.

Besides the limited expressiveness of using a single bandwidth, GWR also ignores nonlinear feature interactions beyond location. To overcome this, it is important to augment GWR with other powerful supervised learning models, such as random forests and neural

networks. These models can effectively capture nonlinear feature interactions and has shown superiority in various modern applications. While there have been attempts to combine GWR with other learning methods [21], previous works are often optimized for a single problem and do not provide a general method of using other learning models with GWR.

(2) Scalability is arguably a core bottleneck of the GWR, as training a GWR model has a quadratic time complexity in terms of the training data size. Such quadratic growth implies that scalability beyond a few tens of thousands examples becomes highly challenging. This situation is exacerbated with MGWR-type approaches which require many iterations of GWR subproblems to optimize bandwidths. For modern datasets, it is critical to develop resource-aware algorithms which can adapt to the computational constraints and utilize modern computational resources. Even though there have been some works on reducing the run time of GWR[24] and MGWR [23], it is still a major runtime bottleneck to handle large datasets.

Contributions: Towards addressing these challenges, this paper proposes Augmented Geographically Weighted Regression (A-GWR), a novel and scalable framework for applying regression models to spatial data. A-GWR provides scalability for large spatial datasets and benefits from the advantages of modern supervised learning techniques to boost both accuracy and runtime efficiency. A-GWR goes two step beyond the existing GWR-related literature. First, it innovates a stateless version of MGWR [12], called Stateless-MGWR (S-MGWR). S-MGWR eliminates the need for the full history of bandwidth values to optimize bandwidth of different features. This allows S-MGWR to use various out-of-the-box parameter optimization methods, and enables faster, more flexible, and more accurate parameter

tuning. Second, A-GWR augments S-MGWR with a general-purpose model, e.g., random forest model, to bring the benefits of popular supervised learning techniques into spatial regression models. This enables spatial regression models to capture non-linear interaction among different features, in addition to seamlessly integrates with modern machine learning (ML) frameworks.

The flexibility of A-GWR framework allows performing regression on smaller chunks of data to process large datasets with limited computational budgets. Therefore, A-GWR significantly enhances the scalability and expressiveness of GWR via the following novelties:

- We propose S-MGWR, a new adaptation of the MGWR approach to perform geographically weighted regression where each feature has its own spatial scale. S-MGWR is able to adjust each set of bandwidths independently rather than the sequential nature of MGWR, making it an ideal model to be compatible with black-box optimization methods.
- We design a customized black-box optimization algorithm to accelerate the discovery of optimal set of feature bandwidths for S-MGWR.
- We propose a framework for merging spatial regression algorithms such as GWR and S-MGWR with fast ML models via ensemble learning to improve accuracy.
- We propose a divide-and-conquer technique that splits the data and operates on its flexible partitions so as to adapt to the computational budget.
- We provide extensive experimental evaluation that shows superiority of our techniques

in both accuracy and runtime with up to 14.4 times faster runtime.

The remainder of the paper is organized as follows: In Chapter 3, we review the basics of the GWR and MGWR and their drawbacks in more detail. Chapter 4 introduces S-MGWR and Chapter 5 introduces our ensemble framework with general ML models. The experiments and performance analysis are presented in Chapter 7. In the last section, we conclude and discuss potential future directions.

Chapter 2

Related Work

This section outlines our related work that is mainly the literature of geographically weighted regression (GWR) [11] and its variations. In GWR, regression coefficients vary across space to capture the spatial non-stationarity aspect of the data. However, GWR has limitations in both computational scalability (quadratic complexity) and expressiveness (using same spatial scale for all features). To overcome these limitations, several efforts have been made. We classify this literature into three main categories as follows.

Variable spatial scale models. To improve the GWR expressiveness over datasets with varying spatial scales for different features, different models have been proposed in [20, 25, 41, 19, 12]. All of these methods use a vector of bandwidths to express a different spatial scale for each feature instead of a single bandwidth value. To find that vector, they use an iterative method that adjusts the state of coefficients until a convergence.

Non-linear models. Combining GWR weighting schemes with Neural Networks have been used to capture the non-linear relation between data points [14, 7]. Along with predicting

dependant variables, GWR is also used for visualization of the data. The GWR is a popular visualization tool to explore patterns of a data set and acquire valuable information, such as locality of features, for further data analysis [42]. Given that Neural Networks, in general, cannot be properly visualized, Combining GWR with Neural Networks affects the visualization aspect of GWR. Another approach to capture the non-linearity is to use non-linear kernels on data and extract features and then use a GWR on extracted features [21]. One downside of this approach is that it requires finding a good kernel which can be challenging.

Parallel and grid models. To improve the runtime of GWR, one way is to parallelize the process [24, 23]. Since parallelization does not reduce the time complexity, even parallelized versions of GWR are bounded to comparably small data sets. Another approach is to develop grid infrastructures. In this approach, the data points are divided to different sections and different GWR model is trained over these sections [9, 16].

Distinguished from all existing work, A-GWR is the first model that simultaneously enables using variable spatial scales, capturing non-linear feature interactions with customizable general-purpose models, and using both parallelization and grid-based modeling to support large datasets. First, our proposed Stateless-MGWR (S-MGWR) model trains a stateless model to find a vector of bandwidths. This enables better computational scalability for variable spatial scales models, while maintaining the high model expressiveness. Second, our A-GWR multi-stage framework enable combining GWR variants with any general-purpose learning models in a way that effectively visualizes the locality while capturing the non-linear relations between features. The ability of A-GWR to adapt to various spatial and general models makes it fully customizable for use with a problem-specific

kernel or a completely general method. Third, A-GWR framework enables using both parallelization and grid-based modeling. It uses grids to handle large data by splitting data into multiple partitions, each works within a spatial locality. Grid system in A-GWR is general and not limited to equal-sized grids, but uses various grid split criteria including variable cells and k-means clusters. In addition, the bandwidth search method of Stateless-MGWR, which is the bottleneck of the model, is also parallelizable. Having all these features, A-GWR outperforms all existing models as a generic model that support various applications and use cases with high computational scalability and model expressiveness.

Chapter 3

Preliminaries

This section introduces the preliminaries of GWR [11] and MGWR [12].

3.1 GWR

In GWR [11], the parameters can vary by location. Assume that each observation i is in the location (u_i, v_i) . Then the GWR model is

$$y_i = \sum_{j=0}^k \beta_j(u_i, v_i) x_{i,j} + \epsilon_i.$$

$\beta_j(u_i, v_i)$ can be computed using the following equation

$$\hat{\beta}(u_i, v_i) = [X^T W_i X]^{-1} X^T W_i y.$$

W_i in this equation is a diagonal matrix of weights assigned to each data point based on its distance to the i -th observation. The main idea is to assign more weight to data points

closer to i -th observation as close data points tend to share more similarities. Matrix W_i is calculated based on a specified kernel function. This kernel function can be adjusted using an input variable. This variable, known as bandwidth, is determined using a number of trials. In each trial, a bandwidth is selected and then the kernel function uses the bandwidth to construct the weights. GWR then uses these weights to solve the equation. The bandwidth that produces the least error is selected as the trained bandwidth value and is used during the prediction. The problem with GWR is that it bounds the locality of all features to a single bandwidth. That is, all the features are assumed to have the same locality which limits the flexibility of GWR in describing real-world data sets.

3.2 MGWR

MGWR [12] tries to address the the single locality issue by introducing a bandwidth for each feature. The MGWR is formulated under Generalized Additive Model (GAM) where:

$$y = \sum_{j=0}^k f_j + \epsilon \quad (3.1)$$

Where f_j is a smooth function estimated with covariate-specific bandwidths and the ϵ is the i.i.d error. MGWR uses the backfitting algorithm to determine the f_j . To compute the functions f_i , the backfitting algorithm interpolates the function in data points' locations. For the initial state, all the coefficients are assigned with an initial value. Then, MGWR at the first iteration, fixes all the functions $f_{j \neq 1}$ and considers them as constants, then, it solves a uni-variate GWR for f_1 and finds the optimum bandwidth for it. Using this optimal bandwidth, the MGWR updates its state by the new coefficients computed

for function f_1 interpolation. For the second iteration, all the values of updated state, except for the interpolated values of second function f_2 , are considered as constants. This procedure is then repeated for all the different features. After adjusting all the features, MGWR starts to readjust f_i with the updated state. The process continues until a stopping criteria is reached.

MGWR allows each feature to have its own spatial scale, represented by a separate bandwidth. Due to the flexibility of locality for each feature, MGWR fits the training data well and computes a good estimate over the coefficients of training data which makes it a good tool for data analysis. However, the backfitting algorithm introduces a number of problems. The first problem in MGWR is its running time. It solves a GWR for each feature until convergence. The number of iterations for convergence can be large and therefore the run time of the program can be very high. After computing a set of bandwidths, MGWR cannot produce the same result using only that set of bandwidths. During the training, each feature is adjusted using the updated state. This means that for creating the same results, a complete history of all the intermediate bandwidths is required. Due to this problem, MGWR is bound to use iterative methods and the bandwidth search method cannot be parallelized.

Chapter 4

Stateless-MGWR

In this section, we propose Stateless-MGWR (S-MGWR), a stateless spatial model that assigns different localities to different features. S-MGWR is able to fit a set of bandwidths directly, hence, any search algorithm, independent of being parallel or linear, can be used to find the locality of each feature. S-MGWR can fit a model for a given vector of bandwidths, using this ability, a search algorithm can take repeated steps. In each step, a vector of bandwidths is generated and based on the validation error over the trained model, the search algorithm can decide its next step. This is in contrast to MGWR where the bandwidth search happens one feature at the time and the trained model depends both on the bandwidth and the previous state of the model.

Algorithm 4.1 gives the pseudo-code for the training phase of S-MGWR. S-MGWR initializes the model values with the results of a GWR model fitted to the data. These values are stored in *initials* and are not re-computed again. Afterwards, starting from the *initials* values, it proceeds to adjust the model based on the bandwidths. To adjust the model, S-

Figure 4.1: S-MGWR Training Phase Algorithm

Algorithm 1 (S-MGWR Training Phase)

Input: L , Training Data D_t , Validation Data D_v

Output: A vector of k bandwidths

1. $initials = \mathbf{GWR}(D_t \cup D_v)$
2. $bw = initials.bandwidth$
3. *//compute_error is detailed in Algorithm 4.2*
4. $error = \mathbf{compute_error}(3, bw, D_t, D_v)$
5. **While** number of iterations $< L$
 6. $next_bw = \mathbf{generate_bandwidth}(bw, error)$
 7. $next_error = \mathbf{compute_error}(3, next_bw, D_t, D_v)$
 8. **if** $next_error < error$
 9. $error = next_error$
 10. $bw = next_bw$
11. **return** bw

MGWR adjusts features based on the vector of bandwidths one at a time. Starting from the first feature, rest of the coefficients are fixed as constants and updates the dependent values accordingly. A weighted least squares computes the adjusted coefficients of the feature and the algorithm proceeds to the next feature. For each set of bandwidths generated by the bandwidth search algorithm, using *generate_bandwidth* subroutine, S-MGWR adjusts the model based on the bandwidths and computes a validation error based on the adjusted model. The set of bandwidths with minimum validation error is then selected as the trained bandwidths for the S-MGWR model. This process is repeated until a user-defined number of iterations are performed. In each cycle of error computation, we set number of iterations to three as this is enough experimentally to adjust coefficients for the best set of bandwidths in this cycle. However, stopping criteria of error computation could also depend on a certain

error difference threshold or a combination of them.

Algorithm 4.2 details *compute_error* subroutine. The subroutine takes two datasets of points, training and validation data, each point consists of a triple $(X, y, \text{coordinates})$, where X is the features, y is the dependant variable, and *coordinates* are the location coordinates. Line 2 initializes the coefficients to the initial coefficients values of GWR as computed in Algorithm 4.1. Then, for a user-defined number of iterations, it adjusts coefficients of all features one by one using each observation in the input training data. In particular, for each regression point i , *get_weights* subroutine computes the weights of other data point based on its distance to point i . In the next lines, $X_{:,j}$ is the j -th column of X , $X_{:,j}(\text{coefficients}_{:,j})$ refers to all the values in X (coefficients) except for the values in column j , and the $*$ operator is the sum of the rows of the Hadamard product:

$$c = A * B \Rightarrow c_i = \sum_j A_{i,j} B_{i,j}. \quad (4.1)$$

The subroutine *WLS* returns the value of weighted least squares for the features vector and the dependent vector using the weights w_i . The predict subroutine, that is called in Algorithm 4.2 line 2, is the same subroutine that is used to predict new observation. We detail the prediction procedure in Section 4.3.

S-MGWR and MGWR share the core idea of using distinct bandwidths for different features to enable multiple spatial scales. Also, both techniques adjust the model one feature at a time using a variation of weighted least squares. However, there are significant differences between the two models. MGWR uses back-fitting to adjust the coefficients

Figure 4.2: Compute Error Algorithm

Algorithm 2 (compute_error)

Input: T_c , bandwidth vector bw , Training Data D_t , Validation Data D_v

Output: validation error

//each $d \in D_t$ or D_v consists of $(X, y, \text{coordinates})$

1. //initials computed in Algorithm 4.1
2. $\text{coefficients} = \text{initials.coefficients}$
3. **while** number of iterations $\neq T_c$
4. **foreach** feature j
5. **foreach** observation i in D_t
6. $w_i = \text{get_weights}(\text{coordinates}_i, bw_j)$
7. $x' = X_{:,j}$
8. $y' = y - (X_{\bar{j}} * \text{coefficient}_{\bar{j}}^T)$
9. $\text{coefficients}_{i,j} = \text{WLS}(x', y', w_i)$
10. $\hat{y} = \text{predict}(bw, \text{coefficients}, D_v)$
11. **return** residual error of \hat{y}

which means for every iteration and every feature, MGWR searches for the best bandwidth for that specific feature and then updates its state based on the trained coefficients calculated from the bandwidth. For example, based on the initial values, MGWR finds b_1 as the best bandwidth describing the first feature. Then, based on this it will update all the other features. When MGWR tries to readjust the first feature, it will use the updated coefficients which were computed based on b_1 . Choosing a different bandwidth in first iteration can greatly impact the second iteration and the final result. This means, to produce the same result as training MGWR needs the complete history of the bandwidths to adjust the model. Generating the complete history where the length of the sequence is not known extends the search space by orders of magnitude making it very difficult to use a black-box optimization algorithm. This problem does not exist in S-MGWR where




Figure 4.3: Adaptive bandwidth illustration.

the bandwidths adjust the model independently from the values of past bandwidths. In addition to the flexibility to use various optimization algorithms, it also enables S-MGWR to search multiple bandwidths simultaneously in parallel. Such parallelization facilitates running more iterations to find the set of bandwidth that best describes the data. So, it boosts both runtime scalability and accuracy of the model without compromising the model expressiveness.

The remainder of this section detail different part of the S-MGWR technique. Section 4.1 explains the concepts of the weight kernel used to compute the weights of data points in Algorithm 4.2. Section 4.2 discusses different strategies to explore the bandwidth space to find the optimal set of bandwidths. Section 4.3 explains the prediction procedure for the validation as well as the test data. To further improve the runtime performance of S-MGWR, we use a Least Recently Used (LRU) cache for the computed weights that is described in Section 4.4.

4.1 Spatial Weighting Scheme

The weight function can be any function that assigns weights to the data points based on their distance to the regression point. The following equation is an example of computing weight of point k based on regression point i and bandwidth of feature j using a Gaussian weight function.

$$w_i[k] = \exp\left[-\frac{1}{2}\left(\frac{d_{i,k}}{b_j}\right)^2\right]. \quad (4.2)$$

In Equation 4.2, $d_{i,k}$ is the spatial distance between data point i and regression point k . b_j is computed using the bandwidth parameter related to the feature j . The bandwidth value for each feature helps indicate how local or global the impact of that feature is. We use adaptive bandwidths in S-MGWR. For each bandwidth b for feature j and each regression point i , the value b_j (Equation 4.2) is computed based on the distance of b nearest neighbors of data point i . This allows the weighting scheme to adjust the weight kernel based on how dense or sparse the neighborhood of the regression point is. Figure 4.3 is an illustration of fixed and adaptive kernels. In Figure 4.3, the dark blue points represent the data points and the red points represent the regression points. The left plane shows a fixed kernel where for each regression point the same kernel is applied. Either the locality is sparse or dense with data points, the same influence range is used. The plane on the right represents the adaptive kernel where different kernel functions are applied for points in dense or sparse areas. Areas that are sparse and have less data points use larger influence ranges to capture enough nearby points. On the contrary, areas that are dense with data points use smaller influence ranges as they will be enough to capture nearby points. In both cases, the weight for distant points would be very small. Therefore, in order to improve the speed, we introduced a cutoff

distance to our algorithm. For each b_j and cutoff number c , any point with a distance more than $c \times b_j$ from the observation point j is not considered and has a weight of zero. The example kernel of Equation 4.2 will then change to:

$$w_i[k] = \begin{cases} \exp[-\frac{1}{2}(\frac{d_{i,k}}{b_j})^2] & d_{i,k} \leq b_j \times c \\ 0 & otherwise \end{cases} \quad (4.3)$$

Cutoff threshold reduces the search space by allowing the algorithm to omit the data points that are far enough from the regression point.

4.2 Bandwidth Search Strategies

The bandwidth should reflect the rate that regression weights decay around a certain location. Using a sensible weighting function, if the bandwidth is small, weights decrease quickly as the distance between a data point and the regression point increases. On the other hand, using big bandwidths should create a smoother surface and the weights would decline slowly as the distance increases. S-MGWR uses adaptive bandwidths. That means S-MGWR uses the number of neighbors as bandwidths and computes the locality based on the distance of its neighbors. Adaptive bandwidth gives S-MGWR the ability to adjust the locality based on the sparsity of the neighborhood around an observation.

The bandwidth search space is extremely large. If n is the number of observations and m is the number of features, then the bandwidth space size is n^m . Thus, a smart algorithm is necessary to explore the huge search space effectively. We have implemented

Hyperband [22], Simulated Annealing, Bayesian Optimization[37], simultaneous perturbation stochastic approximation (SPSA) [38], and Hill Climbing to search for the best set of bandwidths. We also created a combination of these algorithms to benefit from the strength of each algorithm. Our approach uses a combination of SPSA, Bayesian Optimization, and hill climbing methods. This combination is founded on studying the effect of each method on error improvement over time and its speed of convergence to the best error. SPSA method is used to explore the space as broadly as possible in a short time frame. After that, the Bayesian Optimization is performed with a slight modification that the random states generated are in a close neighborhood of the previous state. This part helps to explore the space around the SPSA’s result. Finally, to find a local optima, our algorithm performs the Hill Climbing search.

4.2.1 Implementation Details

In this section, the bandwidth search methods are explained in more detail. \diamond

Hyperband: Hyperband [22] is an extension of successive halving algorithms [18] where a number of different bandwidths are generated and then evaluated using a fraction of training data. Then the top performing combinations are selected and evaluated using a bigger fraction of training data. This methods takes three parameters, *starting_configs*, *min_budget*, and *eta*. The *starting_configs* determines the number of different combinations to start from. The *min_budget* determines the minimum number of points for evaluating each set of bandwidths. The *eta* is the eta value explained in [22].

\diamond **Simulated Annealing:** Simulated annealing is a well-known probabilistic opti-

mization algorithm that simulates temperature-like heating and cooling moves in the search space. It takes parameters T , α , $steps$, $updates$, μ , and σ . T and α determine the temperature. The neighboring states are generated based on a normal distribution $N(\mu, \sigma)$. For each state, the number $steps$ is reduced by one. If the new state is better than the best state found so far, the number $updates$ is reduced by one. The algorithm returns the best value when $steps$ or $updates$ reach to zero.

◇ **Hill Climbing:** Starting from an arbitrary solution, at each iteration of the Hill Climbing method only one of the features' bandwidth is changed. If the new combination has a better validation result, it is replaced as the answer so far. This method is very similar to the *Simulated Annealing* method. The only difference is in the way of generating neighboring states.

◇ **Bayesian Optimization:** It is a black-box optimization algorithm. It creates a probabilistic model (called surrogate model) of the objective function; with each new point, the surrogate model is updated. To pick a new point for evaluating, an acquisition function is used. At each step, the next point is determined by the acquisition function and after evaluating that point the model is updated. We use the BayesianOptimization library for python [28] and pass the *random_count* and *iter_count* to it. When the *is_local* variable is true, only values in the *locality_range* of best answer so far is considered.

◇ **SPSA:** SPSA is a stochastic approximation algorithm that is capable of finding the global optimum. At each step, it approximates the gradient by measuring the objective function at only two points. Then, it moves in the direction that minimizes the error. The fact that SPSA is independent of the dimension of the features makes it a great choice for

our application where data have more than one feature and computing the gradient for all these takes a lot of time. SPSA takes three parameters, *alpha_decay*, *gamma_decay*, and *steps*.

◇ **Combined Search Strategy:** Our approach uses a combination of all the methods discussed. Hyperband is used to explore the space as broadly as possible in a short time frame. After that, the simulated annealing is performed with a slight modification that the random states generated are in a close neighborhood of the previous state. This part helps to explore the space around the Hyperband’s result. Finally, to find a local optima, our algorithm performs the Hill Climbing search. Combining different bandwidth search method is founded by analyzing the error improvement of each method over time. Figure 4.4 shows the error reduction for different methods over 100 seconds of time. The x-axis represents a timeline, from $t = 0$ to $t = 100$, and the y-axis represents the error value. The figure shows GWR as the fixed upper bound (at error 1.15×10^{-2}) and MGWR as the fixed lower bound (at error 0.18×10^{-2}), and every method starts with error close to GWR and converges to an error close to MGWR. It is noticeable that different methods converge to the best error (MGWR error) at different speeds. This guides the *combined search strategy* to use the faster convergence methods first.

4.3 Prediction Procedure

To predict the dependant value for a new sample, similar to the train process, the algorithm computes the coefficients of this sample. Algorithm 4.5 shows the prediction procedure. For each feature of each new sample, the weight matrix is generated based on

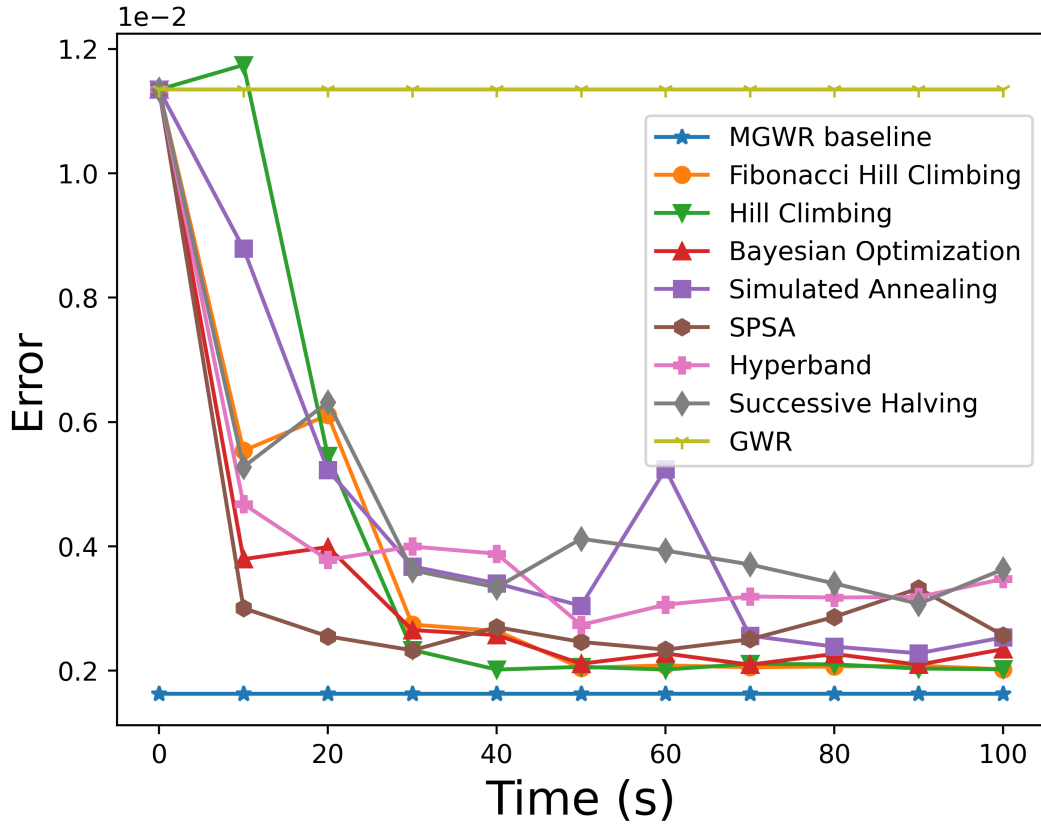


Figure 4.4: Timeline of error reduction for different bandwidth search methods.

the location of the new sample and the training data. After computing the weight matrix, a weighted least squares will compute the coefficient which is used to compute the predicted dependant value of the sample.

4.4 Accelerating Weight Computation

For each bandwidth value, a weight vector with size n is computed for each observation point i where n is the number of training data points. Thus, for each bandwidth, we can store all the weights of all observations on a $n \times n$ matrix. Since S-MGWR uses adaptive bandwidth, the bandwidth can only be a positive integer number smaller than n .

Figure 4.5: Predict Algorithm

Algorithm 3 (predict Procedure)

Input: bandwidth vector bw , trained_coefficients, dataset D

Output: prediction values

1. $m = |D|$
2. $k =$ number of features of $d \in D$
3. $X = m \times k$ matrix of feature values of D
4. $coords = m \times 2$ matrix of location coordinates of D
5. $coefficients = [0]_{m \times k}$
6. **foreach** observation i in D
 7. **foreach** feature j
 8. $w_i = \mathbf{get_weights}(coords_i, bw_j)$
 9. $x' = X_{:,j}$
 10. $y' = y - (X_{\bar{j}} * \mathit{trained_coefficient}_{\bar{j}}^T)$
 11. $coefficients = \mathit{WLS}(x', y', w_i)$
12. $predictions = \mathit{coefficient} * X$
13. **return** predictions

Therefore, all the possible weights would require $\mathcal{O}(n^3)$ space. Since most generated bandwidths have at least one value in common with previous bandwidths, storing the values of previously computed weights can be helpful to reduce the run time of S-MGWR. In our implementation, we are using a cache with *least recently used* (LRU) replacement policy. If the cache stores c (a constant number) bandwidths at a time, then the memory usage will be $\mathcal{O}(c \times n^2) \rightarrow \mathcal{O}(n^2)$. The impact of cache is noticeable, especially when using the Hill climbing algorithm for the bandwidth search. In Hill climbing, only one of the bandwidths changes each time. Hence, the rest of the weights can be read from the cache. We evaluate the caching impact in our experimental evaluation.

Chapter 5

Augmented GWR (A-GWR)

S-MGWR provides a multi-bandwidth spatial regression model that does not need a full history of bandwidth values for computations. Hence, it has eliminated the dependency problem among the consecutive iterations of adjusting the model parameters. To make use of such significant advancement in improving the accuracy and reducing the runtime, we propose augmented GWR (A-GWR) model. A-GWR divides data into different sections. Then, a combination of S-MGWR and a general-purpose supervised learning model is used for each section. S-MGWR, as a spatial model, captures the spatial relationships of the data and the general model captures nonlinear feature interactions which are not necessarily dependent on the location. To combine both models, we use two different methods, namely, *pipeline* and *ensemble*, that apply models in a different order. A great advantage of this framework is that it can be generalized to other spatial regression models as well. This gives flexibility for spatial applications to choose the best models according to their needs and data characteristics. This is also applicable for the general-purpose

supervised model, that can be a simple least-squares algorithm or a complex neural network. However, for specificity, and without loss of generality, we discuss A-GWR in the rest of this section using **S-MGWR** as a spatial model and **Random Forest** as a general-purpose learning model. The rest of the section discusses data splits, model integration, and computational complexity.

5.1 Data Splitting

In order to divide the data, we use two methods that can be applied to different applications.

- **Grid Splitting:** Dividing data into grid cells is a popular method which is also used in [9, 16]. In the grid splitting, the data is divided into grid cells where the number of rows and columns are defined by the user. The grid either divides the space into equal-height cells, by dividing the space to equal height rows and then dividing the rows to equal width sections, as in Figure 5.1a. Another grid option is to divide the space into varying size cell where each of them contains the same number of training data points, as in Figure 5.1b. For a new point, its corresponding grid cell is located. Then, the result of the model associated with that grid cell is reported.

- **Cluster Splitting:** Using randomly generated k-means centers, we perform k-means clustering for training data points. Each cluster represents a section of data containing all the training data points in that cluster. For a new point, after finding its designated cluster, the result of the model associated with that cluster is reported. An example of this method




Figure 5.1: Splitting a sample data into eight sections using grid and k-means methods.

is shown in Figure 5.1c.

Using k-means and grid splitting methods focus more on preserving the spatial neighborhood, which is more in accordance with Tobler’s first law of geography, i.e., “everything is related to everything else, but near things are more related than distant things” [39].

Our empirical evaluation has confirmed the advantage of these splitting methods.

5.2 Model Integration

This section discusses integrating the two models employed by A-GWR: S-MGWR and Random Forest. We apply these to different data sections using two different orders. The first order, *ensemble*, employs the Random Forest model first, while the second order, *pipeline*, employs S-MGWR first. We outline each of them below.

- **Ensemble (general-purpose model first):** In this method, for each section i , first, we train a Random Forest (RF) model on all the training data except for the data in section i . Then, using the trained model, we predict dependent values for the data in section i . Since this data was not in the training process, the model does not overfit the data in section i .

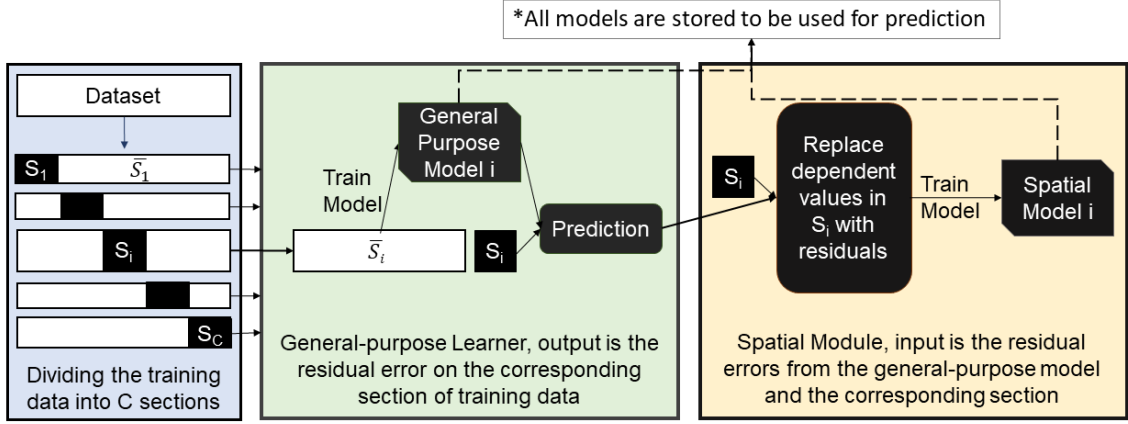


Figure 5.2: Ensemble model integration order (general-purpose model first).

Using the true dependent values and the predicted ones, we compute the residual errors. In the last step, an S-MGWR model is fitted to the data and the residuals. This model captures the spatial relation between data that the RF model fails to capture. In addition, due to running the spatial model on a fraction of data, the runtime is improved.

Figure 5.2 demonstrates the *ensemble* process on C sections of data. For Section i , first, a random forest model is trained over all the data in the data set except for data in section i to create model i . Then, the data of Section i is passed to Model i to compute the residual errors. These residuals are then passed to the spatial S-MGWR model to capture the spatial relativity of the data in each section. For predicting a new sample, if the data splitting method is k-means or grid data, then the new sample is homed into a corresponding data section. Afterward, the data is passed to the corresponding S-MGWR and RF models of that section. The result is the sum of the two outputs. Another option is to run through all the models and then combine the result based on the distance of each section to the point. This can be measured by the distance of the grid center or cluster center of that section to the point.

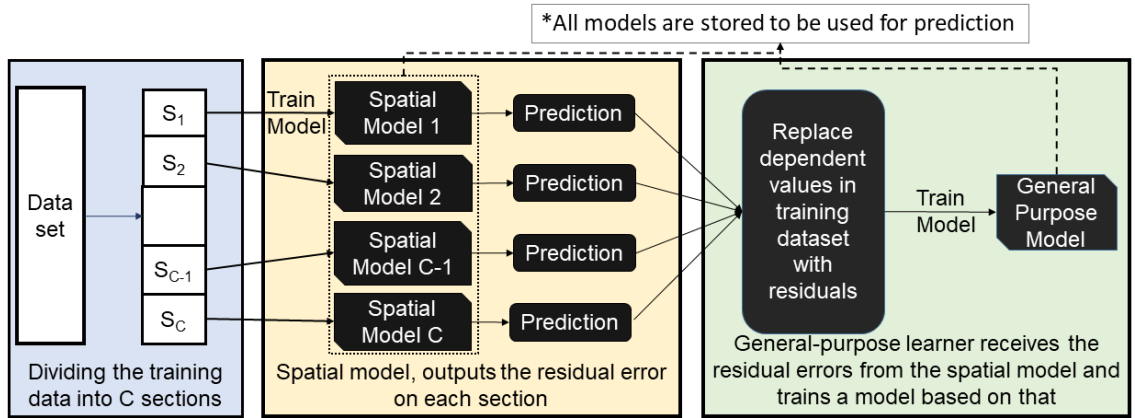


Figure 5.3: Pipeline model integration order (spatial model first).

- Pipeline (spatial model first):** In this method, we feed the training data to the spatial S-MGWR model first. Then, the residual data computed from the spatial model is passed to the RF model. The prediction is the same as the ensemble method. After determining the section, the sum of results from both spatial and fast models is reported as the result. Figure 5.3 shows the pipeline method on C data sections. For each section, a spatial model is trained and then the residual errors are passed to the RF model. The spatial model might overfit the data preventing the RF model to completely capture the global relations between data. However, when using models such as S-MGWR where the base is using less complex methods such as least squares, this is not the issue and the pipeline method provides great improvements on accuracy.

Chapter 6

Computational Complexity

6.1 Time Complexity

Assuming a training dataset of N observations and k features, in Algorithm 4.1, we perform L iterations, each iteration performs bandwidth for the k features (in $\mathcal{O}(k)$) and residual error computation, which is more expensive than the bandwidth search as detailed in Algorithm 4.2. So, the overall complexity is L multiplied by the cost of error computation. The error computation involve T iterations, each goes over each feature and each training observation to get weights and compute coefficients. According to Algorithm 4.2, this performs total of $\mathcal{O}(TkN)$ iterations, each computes weights in $\mathcal{O}(N)$ (Line 2), which gives error computation complexity of $\mathcal{O}(TkN^2)$. So, the overall complexity of the training phase is $\mathcal{O}(LTkN^2)$.

A-GWR breaks data into different sections, each section contains a constant number of observations C , so it operates on $\frac{N}{C}$ sections. At each section, the training cost is $\mathcal{O}(LTkC^2)$. For all sections, the overall complexity is $\mathcal{O}(LTkC^2) \times \frac{N}{C} = \mathcal{O}(LTkCN)$.

Having L , T , and C as constants, then the overall complexity is $\mathcal{O}(kN)$.

This complexity analysis does not consider parallel processing. However, our techniques are parallelizable as different sections are independent and can be trained in parallel. Also, the bandwidth search process is parallelizable, thanks to the stateless nature of S-MGWR, so it is possible to independently validate each set of bandwidths.

6.2 Space Complexity

For each S-MGWR model, A cache size with a capacity of C different bandwidth can be used, in that case the space required by the S-MGWR is $C \times N \times N$ to hold the weight values and $N \times k$ to hold the coefficients. For the ensemble learning, if we have S sections, we need to store S random forest model and S S-MGWR model. The overall space of A-GWR is

$$\mathcal{O}(S \times \text{Random_forest_space}(\frac{N}{S}) + \frac{N}{S} \times S(C \times S^2))$$

. For pipeline model only one random forest is used:

$$\mathcal{O}(\text{Random_forest_space}(N) + \frac{N}{S} \times S(\times C \times S^2))$$

Chapter 7

Experimental Evaluation

We have implemented A-GWR using Python programming language and we use PySAL library [33] modules to run MGWR and GWR techniques with Gaussian kernel and CV criterion for the search method. For Random Forest (RF) learning model, we use scikit-learn Python library with 60 as the number of estimators and default values for the rest of model parameters. All techniques run on an Intel(R) Xenon Silver 4214 machine with CPU @ 2.20GHz and 32 GB of RAM. We evaluate the performance using total of eight different datasets, six real datasets and two synthetic datasets. Unless mentioned otherwise, by default each data section contains 1000 thousands data points. Our performance measures include training runtime and prediction accuracy. Prediction accuracy are measured on the test data using $[1 - (R^2)]$ error measure that is defined as follows:

$$1 - R^2 = \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}$$

Evaluation datasets. We use eight datasets, six real datasets and two synthetic datasets.

Four out of the six real datasets are obtained from PySAL library [33], and these datasets are the standard ones that are used to evaluate various variations of GWR and MGWR techniques [29, 30, 17] . Below is the description of the different datasets.

- **PySAL datasets:** We use four of the datasets provided in the PySAL library, originally hosted by the Spatial Analysis Research Center of the Arizona State University that provided GWR and MGWR software. The datasets represent the socio-economic variables for counties of Georgia, Tokyo Mortality data, the clearwater dataset, and Prenzlauer Berg neighborhood AirBnB data [8]. The dataset sizes are 159, 262, 239, and 2203 respectively. Due to small size of these standard datasets, we solicited the other two large datasets that are described below.

- **New York City Airbnb Open Data:** This dataset, denoted as *NYCAirBnb*, is Airbnb data for the New York city [5]. After removing rows with missing data, we use the room type, minimum nights that a guest can stay, number of reviews, reviews per month, amount of listing per host, and the availability of the unit along with the location to predict the listing price. The price in this dataset is skewed and therefore we use the logarithm of the price. The dataset has 38782 listings after clean up.

- **House Sales in King County, USA:** This dataset, denoted as *kingHousePrices*, has the house sale prices for King County [15]. After removing instances with missing values, we select number of bedrooms, number of bathrooms, square footage of the apartments interior living space, number of floors, house condition and grade, the year built, and the location as our features to predict the price of the house. In total, the dataset has 18708 house sales.




Figure 7.1: Illustration of the covariate functions. β_0 to β_4 are displayed from left to right, respectively

- **Synthetic Dataset 1:** We create a synthetic dataset, denoted as *synthData1*, of 1296 points that is very close to what is described in MGWR [12]. We define five different surfaces for covariates and randomly generate the values for features and compute the dependent values based on them. The synthetic data set is defined by the pair (l, b) where l is the size of the grid and $1 \leq b \leq 5$ is the number of surface functions to use. We define the dependant values using the following equation:

$$y_i = \epsilon_i + \beta_0(u_i, v_i) + \sum_{j=1}^{b-1} \beta_j(u_i, v_i)x_{i,j} \quad (7.1)$$

Where u_i and v_i are the horizontal and vertical location of each point in the grid, $0 \leq u, v < l$. The covariate surface functions (β_0 to β_4) are defined in Equation 7.2 trough

Equation 7.6. Figure 7.1 is an illustration of the surface functions.

$$\beta_0 = 3 \tag{7.2}$$

$$\beta_1 = 1 + \frac{1}{12}(u + v) \tag{7.3}$$

$$\beta_2 = 1 + \frac{1}{l^2} \left[\frac{l^2}{4} - \left(\frac{l}{4} - \frac{u}{2} \right)^2 \right] \left[\frac{l^2}{4} - \left(\frac{l}{4} - \frac{v}{2} \right)^2 \right] \tag{7.4}$$

$$\beta_3 = 5 + \frac{1}{4} \lfloor \frac{u}{5} \rfloor \tag{7.5}$$

$$\beta_4 = l - \frac{1}{500} \left[\frac{l^2}{4} - \left(\frac{l}{4} - \frac{u}{2} \right)^2 \right] \left[\frac{l^2}{4} - \left(\frac{l}{4} - \frac{v}{2} \right)^2 \right] \tag{7.6}$$

For all the points in an $l \times l$ grid, we generate x_1 to x_b randomly from a normal distribution $N(0,1)$ and we combine that with the error term ϵ_i which is randomly generated from a normal distribution $N(0,0.5)$.

• **Synthetic Dataset 2:** We created a second synthetic dataset, denoted as *synthData2*, of 1600 points that is similar in structure to the first synthetic data set. However, one feature is developed in a way that cannot be computed with a simple regression over neighbor data points. The rationale of this dataset is to test different techniques on capturing nonlinear interactions among features (as we will evaluate in Table 7.5 later in the section). The dependant value for this data set is computed as:

$$y_i = \epsilon_i + \beta_0(u_i, v_i) + \sum_{j=1}^3 \beta_j(u_i, v_i)x_{i,j} + 15x_{i,j}^2\beta_4(u_i, v_i).$$

Method	$1 - R^2 (\times 10^{-3})$
Hyperband	3.475
SPSA	2.570
Simulated Annealing	2.535
Bayesian Optimization	2.347
Hill Climbing	2.023
Fibonacci Hill Climbing	2.022
Combined Search	1.955

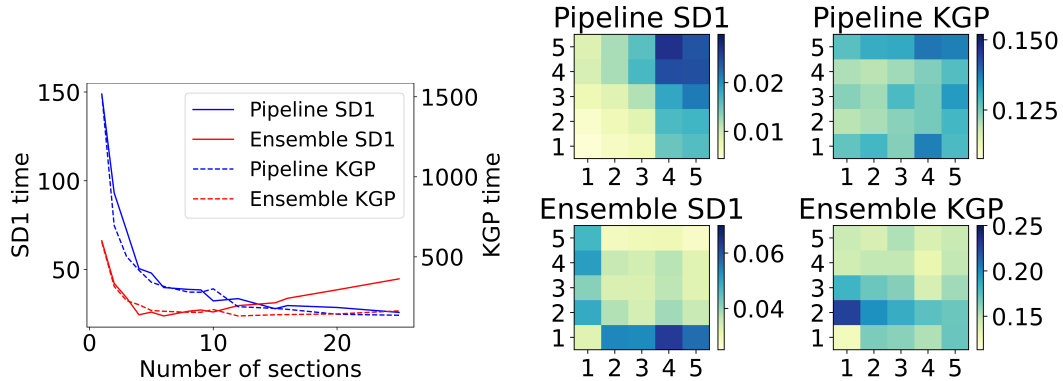
Table 7.1: Error value of different bandwidth search methods

7.1 Model Tuning

We performed a series of experiments to compare different bandwidth search methods and tune the data section size and show their impact on accuracy and run time.

Bandwidth search method. To choose the best strategy to search the bandwidth space, we applied each method over the *synthData1* dataset for the same amount of time and compared the results. Due to the random nature of the algorithms, we performed each experiment ten times and reported the average error value in Table 7.1. The Fibonacci versions are similar to the ordinary versions but the bandwidths are forced to be Fibonacci values. Using Fibonacci numbers limits the searching space and increases the cache hits. The combined search strategy uses Hyperband, Simulated Annealing, and Hill Climbing enabling it to search globally and locally for the optimum set of bandwidths. As shown in Table 7.1, the combined search strategy gives the lowest error compared to other strategies. This shows the superiority of our proposed combined search strategy. Therefore, this strategy is used as the default strategy for bandwidth search in other experiments.

Data section size. In A-GWR, one of the most important factors that affect accuracy and runtime is how many data sections the training data is divided to, and hence



(a) Runtime based on different number of sections (b) R^2 error for different section sizes

Figure 7.2: Accuracy and runtime over synthData1 and kingHousePrices datasets

the size of each data section. There is a trade-off between runtime and accuracy. If there are too many sections, each section contains a few data points and cannot necessarily capture the relations among the data points which reduces the accuracy. If each section has too many points, the runtime will be slow. Figure 7.2 shows the result of different division sizes over two different datasets, *synthData1* and *kingHousePrices*. We use GWR as the spatial module and Random Forest (RF) as the general learning module. Figure 7.2b represents the spatial grid with error value in each grid cell is color-coded according to the associated legend. The figure shows that by dividing the data to more sections, the runtime of the *pipeline* integration method will drop (Figure 7.2a) and error rate will generally increase (Figure 7.2b). The *pipeline* integration method relies on the spatial module as the first model and when the section size is too small, the spatial module cannot capture the relation between data points. For the *ensemble* integration method, the runtime increases after a certain points (10 data sections in Figure 7.2a), this is due to the fact that the increase rate of runtime of Random Forest model is higher than the decrease rate in the run

time of GWR. As the number of sections increases, the training data fed to the Random Forest model increases. Since Random Forest is the first module of the *ensemble* method, it captures most of the relation between data points. Therefore, having more sections and consequently more data, increases the accuracy.

Evaluating cache impact. To evaluate the impact of the LRU cache, we compare the runtime of A-GWR over the two synthetic datasets. Runtimes that are listed in Table 7.2 shows around 30% improvement in time when using cache to store weights. For large datasets, this percentage is considerable and contributes towards faster training phases of A-GWR models.

	Time (s)	
Cache	synthData1	synthData2
Enabled	1133.2258	764.5523
Disabled	1612.3177	1070.3291

Table 7.2: Errors and bandwidths computed by each search method

7.2 Model Performance

We perform a series of experiments to evaluate the performance of A-GWR, showing the separate impact of S-MGWR as well as the impact of combining it with the other parts of A-GWR framework.

Evaluating S-MGWR performance. To compare the performance of S-MGWR with MGWR and GWR, we use the *synthData1* dataset. We have the true values of coefficients for this dataset and we can compute how good each method is computing the coefficient. For a generated dataset we split the data to train and test sections and train

Model	Time (s)	Coeff. RMSE	$1 - R^2 (\times 10^{-3})$
GWR	1.18	3651.77	9.5
MGWR	42.88	467.75	1.2
S-MGWR	28.69	742.42	1.7
S-MGWR*	-	468.20	1.2

Table 7.3: Performance of S-MGWR, GWR, and MGWR. S-MGWR* is the S-MGWR model with bandwidth computed by MGWR.

the S-MGWR, MGWR, and GWR over the train section. We then report the $1 - R^2$ error over the test section and the root mean square error (RMSE) of the coefficients of training points. Since the dataset is generated randomly, we generated ten different datasets and report in Table 7.3 the average values of runtime, $1 - R^2$ error, and the coefficients RMSE over the ten simulations.

Table 7.3 shows that S-MGWR fills the gap between GWR and MGWR. It fits the data more accurately than GWR due to its flexibility and it has a better runtime than MGWR due to the ability to use advanced black-box optimization algorithms. Also, running S-MGWR with the bandwidths computed by MGWR (denoted as S-MGWR*) shows a small difference in coefficients RMSE while maintaining the best $1 - R^2$ error value. The table does not report a runtime for S-MGWR* as it has no training phase and uses the bandwidths computed by MGWR directly. S-MGWR* shows that given a good bandwidth selection method, being stateless does not affect the accuracy. This shows even though S-MGWR does not have the access to MGWR bandwidth history, its performance is comparable to MGWR. This make it a perfect candidate to be used as part of high-level models such as A-GWR or parallelized models.

Evaluating A-GWR framework. This experiment shows the impact of the

general idea of the A-GWR as a framework that could adapt different spatial models and learning models. We apply A-GWR using both MGWR and GWR as the spatial model and Random Forest (RF) as the general learning model on four datasets, *synthData1*, *synthData2*, *kingHousePrices*, and *NYCAirBnb*. For each dataset, we divide the data into sections so that each section contains at most 1000 data points, except the last section that could have less points. The synthetic datasets sections are split with a 36×36 grid and divide it to two columns with equal number of points. The *kingHousePrices* dataset is split using a grid to four rows and five columns with equal number of points. Finally, the *NYCAirBnb* dataset is split using a grid to five rows and six columns with equal number of points. The result of A-GWR w\MGWR and A-GWR w\GWR along with the ordinary versions are shown in Table 7.4.

	synthData1		synthData2		kingHousePrices		NYCAirBnb	
	Time (s)	1 - R^2	Time (s)	1 - R^2	Time (s)	1 - R^2	Time (s)	1 - R^2
GWR	1.5558	0.0121	5.0826	0.9101	512.0763	0.1441	2064.8780	0.4284
A-GWR w\GWR	4.3137	0.0031	9.7359	0.1962	138.563	0.1288	345.7216	0.3932
MGWR	35.7843	0.0016	58.3185	0.8868	80501.9860	0.1378	103713.3542	0.4378
A-GWR w\MGWR	41.0223	0.0016	63.8316	0.1830	5574.1644	0.1143	21460.6177	0.4125

Table 7.4: Evaluating A-GWR as a framework.

Both A-GWR versions of GWR and MGWR outperform the ordinary versions in accuracy in every dataset. In addition, the A-GWR models are considerably faster than the ordinary versions on real datasets. The A-GWR w\MGWR performs up to 14.4 times faster than MGWR and A-GWR w\GWR performs up to 5.9 times faster than GWR.

Finally, we compare A-GWR with S-MGWR as spatial module and Random Forest (RF) as the general learning model with MGWR, GWR, and Random Forest. Table 7.5 shows the error rate of these methods over different datasets where A-GWR outperforms

all other techniques for most of the datasets. It worth noting that in order to have a fair comparison with Random Forest, we have adjusted its hyper-parameters using a random search algorithm to search for the best combination of number of estimators, max features, max depth, min samples split, min samples leaf, and bootstrap setting. This shows the superiority of A-GWR with the best setting of other models.

	pGeorgia	pTokyo	pClrwatr	pBerlin	synthData1	synthData2	kingHousePrices	NYCAirBnb
	$1 - R^2$	$1 - R^2$	$1 - R^2$	$1 - R^2$	$1 - R^2$	$1 - R^2$	$1 - R^2$	$1 - R^2$
A-GWR	0.5962	0.0151	0.8902	0.7868	0.0014	0.2005	0.119	0.4134
MGWR	0.6862	0.029	0.7772	0.7283	0.0014	0.8868	0.1378	0.4378
GWR	0.645	0.0227	0.8365	0.7158	0.0124	0.9101	0.1441	0.4284
Random Forest	0.6195	0.0338	0.8912	0.664	0.0541	0.2685	0.1296	0.3836

Table 7.5: Performance of A-GWR, MGWR, GWR, and Random Forest over different datasets

Chapter 8

Conclusions

This paper has introduced Augmented Geographically Weighted Regression (A-GWR), a spatial regression framework that generalizes MGWR with stateless training and augmented general-purpose learning models. A-GWR introduces Stateless-MGWR (S-MGWR) that is a flexible spatially-varying coefficient (SVC) model with distinct bandwidths for individual features to adjust spatial scale based on the feature impact. S-MGWR is combined with state of the art black-box optimization techniques to find the optimum set of bandwidths efficiently. A-GWR framework achieves higher expressiveness by augmenting powerful machine learning models with spatial models to capture both spatial and non-spatial aspects of the data. In addition, A-GWR achieves scalability by intelligently dividing data to smaller sections while preserving the spatial relations between points. Extensive empirical evaluation on eight real and synthetic datasets has shown that A-GWR is a fast, flexible, and scalable algorithm to model spatial data. It constantly achieves better accuracy while performing up to 14.4 times faster than existing models.

Bibliography

- [1] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(Feb):281–305, 2012.
- [2] Stuart Brown, Vincent Versace, L. Laurensen, Jonathon Fawcett, and Scott Salzman. Assessment of spatiotemporal varying relationships between rainfall, land cover and surface water area using geographically weighted regression. *Environmental Modeling & Assessment - ENVIRON MODEL ASSESS*, 17, 06 2012.
- [3] Osvaldo Daniel Cardozo, Juan Carlos García-Palomares, and Javier Gutiérrez. Application of geographically weighted regression to the direct forecasting of transit ridership at station-level. *Applied Geography*, 34:548 – 558, 2012.
- [4] Henry Crosby, Paul Davis, Theodoros Damoulas, and Stephen A. Jarvis. A spatio-temporal, gaussian process regression, real-estate price predictor. In *ACM SIGSPATIAL*, 2016.
- [5] Dgomonov. New york city airbnb open data, 2019.
- [6] Yadolah Dodge. *Weighted Least-Squares Method*, pages 566–569. Springer New York, 2008.
- [7] Zhenhong Du, Zhongyi Wang, Sensen Wu, Feng Zhang, and Renyi Liu. Geographically neural network weighted regression for the accurate estimation of spatial non-stationarity. *International Journal of Geographical Information Science*, 34(7):1353–1377, 2020.
- [8] Taylor M. Oshan et. al. Pysal datasets, 2021.
- [9] T. Feuillet, H. Commenges, M. Menai, P. Salze, C. Perchoux, R. Reuillon, E. Kesse-Guyot, C. Enaux, J.-A. Nazare, S. Herberg, C. Simon, H. Charreire, and J.M. Oppert. A massive geographically weighted regression model of walking-environment relationships. *Journal of Transport Geography*, 68:118–129, 2018.
- [10] A. Fotheringham, Ricardo Crespo, and Jing Yao. Geographical and temporal weighted regression (gtwr). *Geographical Analysis*, 47, 04 2015.

- [11] Alexander Fotheringham, Chris Brunsdon, and Martin Charlton. Geographically weighted regression: The analysis of spatially varying relationships. *John Wiley & Sons*, 13, 01 2002.
- [12] Alexander Fotheringham, Wenbai Yang, and Wei Kang. Multiscale geographically weighted regression (mgwr). *Annals of the American Association of Geographers*, 107:1247–1265, 08 2017.
- [13] Shuhui Gong, John Cartlidge, Yang Yue, Guoping Qiu, Qingquan Li, and Jingyu Xin. Geographical huff model calibration using taxi trajectory data. In *Proceedings of the 10th ACM SIGSPATIAL Workshop on Computational Transportation Science, IWCTS'17*, page 30–35, New York, NY, USA, 2017. Association for Computing Machinery.
- [14] Julian Hagenauer and Marco Helbich. A geographically weighted artificial neural network. *International Journal of Geographical Information Science*, 0(0):1–21, 2021.
- [15] harlfoxem. House sales in king county, usa, 2016.
- [16] Richard Harris, Alexander Singleton, Daniel Grose, Chris Brunsdon, and Paul Longley. Grid-enabling geographically weighted regression: A case study of participation in higher education in england. *T. GIS*, 14:43–61, 02 2010.
- [17] M. Irfan, A. Koj, H. Thomas, and M. Sedighi. Geographical general regression neural network (ggrnn) tool for geographically weighted regression analysis. In *The Eighth International Conference on Advanced Geographic Information Systems, Applications, and Services, Venice, GEOProcessing 2016*, 2016.
- [18] Kevin G. Jamieson and Ameet Talwalkar. Non-stochastic best arm identification and hyperparameter optimization. *CoRR*, abs/1502.07943, 2015.
- [19] Yin-Yee Leong and Jack Yue. A modification to geographically weighted regression. *International Journal of Health Geographics*, 16, 12 2017.
- [20] Yin-Yee Leong and Jack C. Yue. A modification to geographically weighted regression. *International Journal of Health Geographics*, 16(1):11, Mar 2017.
- [21] Lianfa Li. Geographically weighted machine learning and downscaling for high-resolution spatiotemporal estimations of wind speed. *Remote Sensing*, 11(11), 2019.
- [22] Lisha Li, Kevin G. Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Efficient hyperparameter optimization and infinitely many armed bandits. *CoRR*, abs/1603.06560, 2016.
- [23] Ziqi Li and A. Stewart Fotheringham. Computational improvements to multi-scale geographically weighted regression. *International Journal of Geographical Information Science*, 34(7):1378–1397, 2020.

- [24] Ziqi Li, Alexander Fotheringham, Wenwen Li, and Taylor Oshan. Fast geographically weighted regression (fastgwr): A scalable algorithm to investigate spatial process heterogeneity in millions of observations. *International Journal of Geographical Information Science*, 10 2018.
- [25] Binbin Lu, Chris Brunsdon, Martin Charlton, and Paul Harris. Geographically weighted regression with parameter-specific distance metrics. *International Journal of Geographical Information Science*, 31(5):982–998, 2017.
- [26] Arabinda Maiti, Qi Zhang, Srikanta Sannigrahi, Suvamoy Pramanik, Suman Chakraborti, Artemi Cerda, and Francesco Pilla. Exploring spatiotemporal effects of the driving factors on covid-19 incidences in the contiguous united states. *Sustainable Cities and Society*, 68:102784, 2021.
- [27] Mansour Ndiath, Badara Cissé, Jean Ndiaye, Jules Gomis, Ousmane Bathiery, Anta Dia, Oumar Gaye, and Babacar Faye. Application of geographically-weighted regression analysis to assess risk factors for malaria hotspots in keur soce health and demographic surveillance site. *Malaria journal*, 14:463, 11 2015.
- [28] Fernando Nogueira. Bayesian Optimization: Open source constrained global optimization tool for Python, 2014–.
- [29] Taylor M. Oshan, Ziqi Li, Wei Kang, Levi J. Wolf, and A. Stewart 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), 2019.
- [30] Taylor M. Oshan, Ziqi Li, Wei Kang, Levi John Wolf, and Alexander Stewart Fotheringham. mgwr: A Python implementation of multiscale geographically weighted regression for investigating process spatial heterogeneity and scale. OSF Preprints bphw9, Center for Open Science, October 2018.
- [31] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [32] Alexei Pozdnoukhov and Christian Kaiser. Scalable local regression for spatial analytics. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2011.
- [33] Sergio J. Rey and Luc Anselin. PySAL: A Python Library of Spatial Analytical Methods. *The Review of Regional Studies*, 37(1):5–27, 2007.
- [34] Ibrahim Sabek, Mashaal Musleh, and Mohamed F. Mokbel. Turboreg: a framework for scaling up spatial logistic regression models. In *ACM SIGSPATIAL*, 2018.

- [35] Tayyab Ikram Shah and Scott Bell. Exploring the intra-urban variations in the relationship among geographic accessibility to phc services and socio-demographic factors. In *Proceedings of the Second ACM SIGSPATIAL International Workshop on the Use of GIS in Public Health*, HealthGIS '13, page 68–76, New York, NY, USA, 2013. Association for Computing Machinery.
- [36] Masamichi Shimosaka, Takeshi Tsukiji, Hideyuki Wada, and Kota Tsubouchi. Predictive population behavior analysis from multiple contexts with multilinear poisson regression. In *ACM SIGSPATIAL*, 2018.
- [37] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'12, page 2951–2959, Red Hook, NY, USA, 2012. Curran Associates Inc.
- [38] J.C. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3):332–341, 1992.
- [39] W. R. Tobler. A computer movie simulating urban growth in the detroit region. *Economic Geography*, 46:234–240, 1970.
- [40] Austin Troy, J. Morgan Grove, and Jarlath O'Neil-Dunne. The relationship between tree canopy and crime rates across an urban–rural gradient in the greater baltimore region. *Landscape and Urban Planning*, 106(3):262 – 270, 2012.
- [41] Wenbai Yang, Alexander Fotheringham, and Paul Harris. An extension of geographically weighted regression with flexible bandwidths. In *the 20th Annual GIS Research UK (GISRUK) conference*, 01 2012.
- [42] Bin Zou, Qiang Pu, Muhammad Bilal, Qihao Weng, Liang Zhai, and Janet E. Nichol. High-resolution satellite mapping of fine particulates based on geographically weighted regression. *IEEE Geoscience and Remote Sensing Letters*, 13(4):495–499, 2016.