Taveesh Sharma taveesh@uchicago.edu University of Chicago USA Paul Schmitt pschmitt@hawaii.edu University of Hawaii, Manoa / Invisv USA

Nick Feamster feamster@uchicago.edu University of Chicago USA

# ABSTRACT

Despite significant investments in access network infrastructure, universal access to high-quality Internet connectivity remains a challenge. Policymakers often rely on large-scale, crowdsourced measurement datasets to assess the distribution of access network performance across geographic areas. These decisions typically rest on the assumption that Internet performance is uniformly distributed within predefined social boundaries, such as zip codes, census tracts, or community areas. However, this assumption may not be valid for two reasons: (1) crowdsourced measurements often exhibit nonuniform sampling densities within geographic areas; and (2) predefined social boundaries may not align with the actual boundaries of Internet infrastructure.

In this paper, we model Internet performance as a spatial process. We apply and evaluate a series of statistical techniques to: (1) aggregate Internet performance over a geographic region; (2) overlay interpolated maps with various sampling boundary choices; and (3) spatially cluster boundary units to identify areas with similar performance characteristics. We evaluated the effectiveness of these using a 17-monthlong crowdsourced dataset from Ookla Speedtest. We evaluate several leading interpolation methods at varying spatial scales. Further, we examine the similarity between the resulting boundaries for smaller realizations of the dataset. Our findings suggest that our combination of techniques achieves a 56% gain in similarity score over traditional methods that rely on aggregates over raw measurement values for performance summarization. Our work highlights an urgent need for more sophisticated strategies in understanding and addressing Internet access disparities.

## **1** INTRODUCTION

Measuring the performance of Internet access networks is important for understanding the quality of service offered Nicole P. Marwell nmarwell@uchicago.edu University of Chicago USA

to users by ISPs [11] and identifying access gaps in Internet performance in both urban and rural areas [20, 36]. Over the past few decades, there have been significant advancements in measuring access network performance, both in terms of novel measurement infrastructure [2, 19, 39], and analysis techniques [8, 31, 40, 41]. These advancements have primarily been focused towards the performance of a single access link, using metrics such as throughput, latency, jitter, and packet loss. In the present day, the Measurement Lab (M-Lab) [1] and Ookla Speedtest [3] datasets are the most widely used for understanding Internet performance of an access link. Their growing prevalence has also enabled researchers to use these datasets to ask a broader set of questions [7, 16, 24, 29]. A question of significant recent interest, especially for specific social and policy-related questions, is understanding the distribution of Internet performance over a geography [11, 21, 26, 34]. However, there remain significant challenges in understanding how to bridge the gap between point measurements and regional summaries given uneven distributions of crowdsourced measurement datasets.

Francesco Bronzino

francesco.bronzino@ens-lyon.fr

**ENS Lyon** 

France

Both the M-Lab and Ookla datasets produce crowdsourced "point" measurements from a subset of Internet users over a given geography. These measurements are often concentrated irregularly over space, and are recorded when a user decides to run a speed test for multiple reasons. It thus becomes necessary to understand how these small, self-selected samples can be used to make generalizations about Internet performance for the entire resident population. A key challenge towards this goal is to identify the boundaries in which Internet performance can be sampled, and the methods that can be used to summarize these point measurements over a geography (in other words, "over space"). Additionally, the extent of noise associated with a single measurement can be significant due to many factors such as the testing infrastructure, access media, and client's hardware or software platform [22, 28]. It is therefore necessary to also understand the extent and

the spatial granularity up to which these measurements can be de-noised, and aggregated to form conclusions about the performance of the network over a specific geography.

Prior work has approached these questions with a focus on understanding the distribution of Internet performance over conventional boundaries such as zip codes, census tracts, or community areas [17, 20, 27, 31, 33]. These studies have shown that there is significant variation in Internet performance over these boundaries [33], and have used this information to identify areas that may need additional policy interventions. However, these approaches suffer from a few important limitations. First, the use of aggregate measures such as median and inter-quantile range (IQR) [27, 33] on point measurements directly assumes that the data is uniformly distributed across the geographic boundary. In reality, these measurements are often clustered in some portions of the space and dispersed in others [20], addressing which involves careful point-pattern analysis. Second, there is a lack of studies that compare the accuracy of prior modelling techniques [17, 20] in summarizing Internet performance over a geography. Finally, correlating Internet performance with population measures such as median income and population density [17, 27, 31] using existing social boundaries may not be accurate due to imperfect alignment with infrastructure boundaries.

Our work aims to address these concerns by applying and comparing a new combination of statistical techniques to aggregate point measurements over a geography with stable sampling boundaries, that is, boundaries that do not change when subject to variations in the underlying data. We evaluate these techniques on the basis of their ability to predict latency for in-sample measurements in five zip codes in Chicago we have sampled more densely. These techniques allow us to interpolate to synthetic, out-of-sample locations for areas that are otherwise unsampled in crowdsourced datasets. We use this capability to summarize latency within small, polygonal tessellations of varying resolutions, as well as community area boundaries within Chicago. Further, we cluster these smaller units to discover the edges of sampling boundaries. To evaluate the goodness of the resulting clusters, we measure the stability of these boundaries using the Jaccard similarity between samples drawn from the interpolated dataset. We show that these techniques achieve a median pairwise Jaccard similarity score of 99%, which provides a 56% gain over computing raw averages for community area boundaries. Our work makes the following contributions:

- We apply existing spatial interpolation techniques on a crowdsourced measurement dataset to summarize latency over a city. Our findings suggest that the precision of these techniques in interpolating latency improves with an increase in spatial resolution (Section 4).
- We extend these approaches to discover sampling boundaries that are stable across multiple samples drawn from the

same dataset. We show that these boundaries provide a significant stability gain over those obtained after employing raw averages for performance summarization (Section 5).

- We discuss how network operators and regulators can use our combination of techniques to identify areas with similar latency characteristics to plan investment and policy interventions (Section 5).
- We open-source our code and datasets containing proposed sampling boundaries in a major US city. We also release a dataset quantifying our latency estimates over these boundaries for future research.<sup>1</sup>

## 2 BACKGROUND & RELATED WORK

A variety of approaches can be employed towards discovering boundaries for sampling Internet performance in a region. One approach is to directly use network availability data from ISPs and align performance sampling boundaries with accurate coverage maps. The accuracy of these maps has come under scrutiny in recent times [37], which has given rise to the BEAD challenge process [23]. We acknowledge that the construction of accurate coverage maps is an ongoing process, so relying on this approach may be premature at the time of writing. A second approach could be to leverage crowdsourced measurement data from speed test providers like Ookla or M-Lab to statistically infer sampling boundaries. A key challenge to this approach is the under-representation of areas with access gaps, or areas where users are less likely to conduct speed tests. For ensuring reliability, it is thus important to adopt techniques that accurately estimate Internet performance in sparsely sampled areas. In this work, we adopt the second approach. To address the sparsity challenge, we consider the application and evaluation of existing spatial interpolation techniques in the context of Internet measurement data. Finally, we explore the use of a spatial clustering technique to identify areas with similar Internet performance characteristics. We next provide an overview of relevant spatial statistics literature.

## 2.1 Spatial Interpolation

Spatial Interpolation is a common class of techniques used to minimize noise in geographically sampled data and capture the underlying spatial structure. These techniques are often used to estimate a spatial variable of interest in areas which are sparsely sampled. Spatial Interpolation techniques can be broadly classified into two subcategories: deterministic and stochastic. Deterministic techniques assume that the underlying spatial process is static and smooth over space. The outcome of applying these techniques is often a surface that provides an estimate of the target variable over a region. A

<sup>&</sup>lt;sup>1</sup>The artifacts will be made available post peer-review.

few examples of deterministic techniques include Inverse Distance Weighting (IDW) [35], Kriging [12], LOESS [10], and Self-tuning Bandwidth in Kernel Regression (STBKR) [20]. Stochastic techniques, on the other hand, assume that the underlying spatial process is discontinuous and has random trends along the geography. A few examples of stochastic techniques include Gaussian Processes [4], Random Forests [32], and Neural Networks [30]. Exploring a wide range of interpolation techniques is beyond our current scope, so we focus mainly on deterministic techniques in this paper. We chose deterministic techniques because prior work has applied them on crowdsourced speed test data [20, 36, 38]. In Section 4, we complement these studies with an empirical comparison and demonstrate how their accuracy may be impacted by the size of the sampled geography. We particularly use IDW, LOESS and STBKR for comparisons and extension.

We find very few studies with a notable use of spatial interpolation in prior Internet measurement literature. The LOESS technique has been used in some spatial applications in astronomy [10], but not in the context of Internet performance. IDW found its use in Sommers *et al.* [36] to understand the spatial distribution of cellular and Wi-Fi performance in metro areas. They observed a degradation in performance as one moves further away from metro areas. Stuyvesant *et al.* [38] used Kriging to estimate the spatial distribution of broadband speeds in Michigan. Jiang *et al.* [20] proposed the STBKR technique but applied it on cellular speed test data. Our work differs from these studies in that it focuses primarily on latency measurements, and that it conducts a comparison between interpolation methods at varying spatial scales.

## 2.2 Spatial Clustering

Spatial clustering involves the process of grouping similar data points together based on their spatial proximity, or sometimes another attribute of interest. Common spatial clustering algorithms include K-Means, DBSCAN, and Hierarchical Clustering. The output of applying these algorithms on spatial data is a set of clusters, which may or may not be contiguous in space. A specific form of spatial clustering is regionalization, also known as spatially constrained clustering. Clusters formed using regionalization are contiguous in space. Common regionalization algorithms include Automatic Zoning Procedure (AZP) [5], the Max-P algorithm [13], and Spatial 'K'luster Analysis by Tree Edge Removal (SKATER) [6]. In this work, we consider the use of regionalization in identifying areas with similar Internet performance in a citywide geography. Our analysis is limited to regionalization techniques because of our prior assumptions about Internet infrastructure. Internet infrastructure is often laid out in contiguous regions, implying that areas with similar Internet performance are likely to be contiguous as well. We use the SKATER algorithm

to identify such regions because of its graph-based methodology, which resembles the underlying structure of Internet infrastructure.

## 3 METHODS

This section describes our analysis methodology in detail. We first discuss the choice of our dataset and the preprocessing steps applied to it. Then, we describe the overall approach to combining spatial analysis techniques for discovering the proposed sampling boundaries. Finally, we discuss the evaluation metrics used to compare interpolation techniques and the quality of our boundaries.

## 3.1 Analysis Scope

*3.1.1* Choice of dataset. We use the Ookla dataset [25] for our analysis as it is the largest crowdsourced measurement dataset for access network performance in the present day. As opposed to the M-Lab dataset, Ookla provides access to a greater number of geolocations with higher accuracy, which is crucial for spatial analysis. Ookla uses a combination of GPS and IP geolocation to triangulate a user. We found a lack of availability of an accuracy measure for the geolocated measurements, which renders these measurements unreliable for a high-resolution spatial analysis. So, we choose to focus on GPS geolocated measurements for our analysis. We conducted our analyses for the city of Chicago because (1) it provided us the second-largest overall sample size at a city level, and (2) it is a city with a well documented history of sampling bias across its subdivisions [31, 33].

3.1.2 Choice of performance metric. We use latency as our primary metric for spatial analysis because it is often considered a better proxy for end-user quality of experience (QoE) than metrics such as throughput or packet loss [9]. A higher latency can lead to increased buffering times for real-time applications such as video streaming, video conferencing and online gaming. Additionally, latency can allow us to compare performance across different access technologies such as DSL, Fiber and Cable. Throughput, on the other hand, is often impacted by factors such as the speed tier of the user and the number of users on the network at a given time, which may skew our results. Finally, any sampling boundaries discovered using latency can help network operators and regulators identify regions with suboptimal user satisfaction levels from Internet services. This may further help in devising relevant policy interventions for improving Internet performance in these regions.

# 3.2 Data Preprocessing

The Ookla dataset contains a large repository of crowdsourced measurements. These measurements are collected using the Speedtest application, which is available across a wide range

Filtering step	Description	Before count	After count	Percentage reduction
0	Extract Chicago measurements	76,414,601	5,924,004	92.25%
0	Extract Comcast measurements	5,924,004	1,066,213	82.0 %
6	Filter out Web measurements	1,066,213	579,747	45.63%
4	Filter out Ethernet measurements	579,747	541,150	6.66%
6	Filter out GeoIP measurements	541,150	407,638	24.67%
6	Filter out inaccurate geolocations	407,638	375,132	7.97%
Ø	Filter out distant target servers	375,132	361,151	3.73%
8	Filter out suboptimal target servers	361,151	360,126	0.28%

Table 1: A summary of the filtering steps applied to our initial sample of Ookla measurements.

of users, platforms and devices. Given the large scale of this dataset, it is crucial to apply a series of filtering steps to narrow down our focus on a subset of measurements that are likely to provide insights into the geographic variations in latency, without being affected by noise.

In this section, we thus describe a series of filtering steps applied to our initial sample of Ookla measurements. These steps are different from spatial interpolation/de-noising as they are applied to the raw data before performing any spatial analysis. Interpolation, on the other hand, allowed us to capture the geographical trends in the data by aggregating measurements in high resolution spatial bins.

Our initial sample consisted of 76,414,601 measurements from the Ookla dataset, collected over a period of 17 months. • We focus on 7.75% of these measurements that originated from Chicago because this city provided a large sample size, has a rich set of demographics, and there is evidence for considerable sampling bias across its subdivisions. @ Further, we only consider measurements conducted by users subscribing to Comcast because it was observed to be the largest fixed line ISP in the city in terms of number of measurements. Our focus in this paper is only on fixed line ISPs because users are more likely to consume popular high-bandwidth services over fixed line networks. Additionally, we expected fixed line ISPs to offer a superior spatial correlation compared to mobile ISPs, which are more likely to be affected by the design of cellular networks. A single Packet Data Network Gateway (P-GW) is expected to serve large urban areas for cellular networks, which may bias the spatial correlation calculations. 3 Next, we filter out measurements that are conducted over Ookla's web portal, as these measurements did not allow us to discern between the testing device's platform for contextualizing our results. <sup>(4)</sup> We also filter out measurements conducted over Ethernet connections because we found them to have about a 1:14 ratio of coverage compared to WiFi connections. We defer the evaluation of sensitivity of our approach to inclusion of Ethernet measurements as future work. <sup>(5)</sup> Next, we filter out all IP geolocated measurements due to their lack

of reliability for high resolution spatial analysis. <sup>(6)</sup> We filter out measurements that had a higher than 200 meters location error associated with them. We expect most cable nodes to be situated about a few hundred meters away from a serviceable location, so locations that breach this range may not provide us accurate information. @ Next, we filter out measurements conducted against servers that are not located near Chicago. This is done because most Content Delivery Networks (CDNs) tend to deploy their content caches close to end users [18] for achieving low latencies. This choice therefore allows us to get a close estimate of typical latencies experienced by user-facing applications. S Finally, we use a heuristic to filter out measurements destined to servers that are located near Chicago but likely reported higher than typical latencies observed across all measurements. Figure 1 shows the overall as well as per-server distribution of latency for measurements conducted in Chicago. We observe that the distribution for most of the servers was clustered around the overall distribution. However, a few servers reported higher than typical latencies, likely due to misconfigurations or their distant locations. We filter out these servers on the criterion that their median latency was higher than 17 ms, which is equivalent to the  $85^{th}$  percentile of the overall latency distribution.

Applying these filters resulted in a final sample that primarily consisted of latency measurements from iOS and Android users across Comcast users in Chicago. We evaluate the sensitivity of our methods across these user groups in subsequent sections of the paper. As an additional consideration, we analyzed the age of GPS locations in the final sample. Across our 17-month sample, we found the 95<sup>th</sup> percentile location age to be about 43 seconds across all measurements, indicating that most locations were reasonably recent. Table 1 summarizes the measurement counts before and after applying the above filters to our initial sample. Most reduction in the dataset size occurs at the initial filtering steps, which is expected because we are focusing on a specific geography and ISP. We believe that the final sample is representative of the population of users for a major ISP in a large US city, and is thus suitable



Figure 1: A distribution of latency for measurements conducted against servers located near Chicago. The black line represents the overall latency distribution. The blue lines represent latency distributions for individual servers. The dashed line represents our threshold of 17 ms for filtering out suboptimal servers using their median latencies.

for our analysis. We defer the evaluation of sensitivity of our approach to the inclusion of other ISPs and cities as future work. We summarize basic descriptive statistics for the final sample in Table 2.

## 3.3 Analysis Approach

Motivated by the present unreliability of the current FCC coverage maps, our analysis aims to discover statistical sampling boundaries for latency over Chicago. Given the variable sampling density of measurements over the geography, an important first step is to develop a uniform surface model of latency values. Such a model not only allows us to estimate latency at unsampled locations, but also ensures that our conclusions do not suffer from biases arising due to uneven sampling densities. We do so by applying spatial interpolation techniques. Before proceeding with a citywide interpolation, it is necessary to evaluate interpolation performance in predicting latency at in-sample locations, and quantify the dependence of their accuracy on the size of the geography. A larger geography is expected to contain a greater number of measurements. The aggregation of a greater number of measurements, in turn, is prone to greater noise. An understanding of the accuracy-resolution dependence is thus critical to drive the decision of picking the right resolution, without significantly trading off accuracy. Upon interpolating to a citywide map, it is important to determine which spatial units can be used to aggregate latency for a stable clustering, and how they impact our conclusions regarding latency. We thus perform overlays of the interpolated map with different spatial units, and aggregate latency within these units. Finally, we use SKATER regionalization, a spatial clustering algorithm

Descriptive	Value
Measurement Duration	Jan 2022 – Jun 2023
# Measurements	352,890
# iOS Measurements	263,098
# Android Measurements	89,792
# distinct vantage points (VP)	55,526
Median # samples per VP	2
# target servers	38

Table 2: Basic descriptives of the final sample of Ookla measurements.

that preserves contiguity, to identify regions with similar latency characteristics. Figure 2 shows an analysis approach that incorporates these design goals. Next, we describe some key aspects of our analysis approach below.

3.3.1 Interpolation (Sampling Unit Choice). For spatial interpolation, we compare three interpolation techniques under different sampling resolutions. To control for the sampling resolution, we choose to tessellate the geographic map into hexagons of adjustable sizes. This choice was made because hexagons have the highest perimeter-to-area ratio among regular polygons, which allows them to tessellate the map with minimal overlap. The Federal Communication Commission (FCC) commonly uses the H3 tessellation system [15] to map broadband availability across the United States. H3 is a hierarchical geospatial indexing system widely used in realtime applications such as taxi demand forecasting and urban planning. These hexagons can be constructed at 16 different resolutions ranging from 0 to 15, with a higher resolution representing hexagons of smaller edge lengths. The FCC uses a resolution of 8 for their broadband availability maps. We also experiment with 6, 7, 9 and 10 to evaluate the sensitivity of our results to the choice of resolution.

3.3.2 Interpolation (Evaluation Criteria). To evaluate the interpolation accuracy of the three methods, we use a 5-fold cross validation. In each iteration, we leave out 20% of the data as a test set and interpolate using the remaining 80% of the data. We calculate the Root Mean Squared Error (RMSE) between the interpolated and ground-truth latency values at a cell-level. We use RMSE as our evaluation metric because it is sensitive to outliers and is a common choice for spatial interpolation problems [42]. Our evaluation is limited to celllevel estimates because of the assumption that point-level estimates may be significantly affected by noise due to human factors such as individual load patterns or suboptimal access equipment. Aggregation into hexagonal cells thus allows us to capture the underlying spatial patterns in latency for nearby users that may share infrastructure. In Section 4, we analyze



# Figure 2: Overview of our analysis approach. First, we construct an interpolated map of the region. Then, we use this map to perform spatial clustering.

the distribution of RMSE values across different resolutions and interpolation methods.

3.3.3 Regionalization (Data Preparation). To perform regionalization, we first interpolate latency measurements over the complete geography of Chicago. The target points for interpolation are chosen to be grid points spread across the city at a regular spacing of 100 meters. This choice was made to ensure that the interpolated map was continuous and smooth, which allows an unbiased calculation of latency aggregates such as averages and percentiles. Further, we overlay the interpolated map with a hexagonal tessellation of resolution 8 over Chicago. For our clustering metric, we select the mean latency in each hexagonal cell. Mean latency is chosen because it helps smooth out any local fluctuations within the cells that may arise due to user behavior or random noise. Finally, we use SKATER to perform spatial clustering on the cells to identify regions with similar latency characteristics.

3.3.4 Regionalization (Evaluation Criteria). In addition to visual comparisons, we use two empirical criteria for evaluating the quality of the resulting clusters. Our primary criteria is the Jaccard Similarity Index between clusters obtained using smaller samples drawn from the dataset. This choice is reasonable because it allows us to evaluate whether the clusters obtained using our approach were consistent, and whether there are merits to conducting prior interpolation to bring out spatial trends in latency. We make these comparisons across three different scenarios: (1) using raw measurements with existing units, (2) using interpolated measurements with existing spatial units, and (3) using interpolated measurements with the hexagonal tessellation. Our secondary criteria is the within-cluster to between-cluster variance ratio of latency. We used this ratio to determine the optimal SKATER parameters while assessing the homogeneity of the resulting clusters.

## **4** INTERPOLATION

#### **Summary of findings**

- The precision of interpolation techniques improves as the sampling resolution increases. But we lose reliability in our estimates at finer resolutions due to fewer data points and users. We suggest that a resolution of 8 may help achieve a balance between precision and reliability.
- LOESS, a regression based technique, shows invalid predictions (≤ 0 ms) for a fraction of points at fine resolutions due to potential under-fitting. This suggests that the accuracy of this method is dependent on the number of data points available for interpolation.
- STBKR and IDW use a weighted average of nearby measurements to estimate latency. Both these techniques showed lower errors compared to LOESS at higher resolutions (9 and 10), indicating that they may be better suited for more confined geographies with fewer data points.

In this section, we formally introduce the interpolation problem and the chosen techniques – LOESS, STBKR and IDW. Finally, we present an empirical comparison of the three techniques in estimating latencies at unsampled locations in the Ookla dataset.

# 4.1 **Problem Formulation & Techniques**

Assume that we are given *n* observed locations with latency values  $z_i$  at locations  $(x_i, y_i)$ , i = 1, 2, ..., n, and we are interested in estimating the latency *z* at an unmeasured location (x, y). Let Z(x, y) denote the random variable representing the latency value at location (x, y). We are interested in obtaining an estimate  $\hat{Z}(x, y)$  for Z(x, y) at (x, y). We describe three of many techniques to calculate  $\hat{Z}(x, y)$  below. We summarize the three chosen interpolation methods and their parameters in Table 3.

Interpolation Method	Description	Parameters
Inverse Distance Weighting (IDW)	Computes weighted average of nearby measure- ments in proportion to their relative distance from an unsampled location.	p (Impact of distance on weights)
Locally Estimated Scatter- plot Smoothing (LOESS)	Fits local regression lines to de-noise latency across space.	<i>span</i> (Proportion of data points for regression)
Self-tuning Bandwidth in Kernel Regression (STBKR)	Computes weighted averages of nearby measure- ments as estimates for unsampled locations using a Gaussian Kernel that models point densities.	c (Controls bandwidth of the kernel), k (Number of nearest neighbors)

Table 3: A summary of chosen interpolation methods and their parameters.

4.1.1 Inverse Distance Weighting (IDW). IDW assigns weights to each nearby data point based on its distance from an unsampled location. It uses these weights to calculate a linear combination of nearby values as an estimate of the target metric at an unsampled location. The relationship between the similarity of nearby data points and their distance is assumed to be inverse in nature. The IDW estimate  $\hat{Z}(x, y)$  at location (x, y) is given by:

$$\hat{Z}(x,y) = \frac{\sum_{i=1}^{n} \frac{z_i}{d_i^p}}{\sum_{i=1}^{n} \frac{1}{d_i^p}}$$

where  $d_i$  is the Euclidean distance between the target location (x, y) and the  $i^{th}$  data point  $(x_i, y_i)$ , and  $p \ge 1$  is a parameter used to control the influence of nearby points. A higher value of p indicates a greater influence.

4.1.2 Locally Estimated Scatterplot Smoothing (LOESS). LOESS [10] is a non-parametric regression technique that fits a smooth curve to a scatterplot of data points. By fitting a set of local polynomials to the spatial data, it smoothes any discontinuities and effectively captures the underlying spatial patterns. LOESS uses a smoothing parameter  $\alpha$ , commonly known as the span, to control the extent of smoothing. It assigns weights to the nearby data points  $(x_i, y_i)$  depending on their distance from an unsampled location (x, y) using a Tri-cube Kernel. The weights,  $w((x, y), (x_i, y_i))$  are given by:

$$\begin{cases} \left(1 - \left(\frac{\|(x_i, y_i) - (x, y)\|}{h}\right)^3\right)^3 & \text{if } \|(x_i, y_i) - (x, y)\| \le h, \\ 0 & \text{otherwise} \end{cases}$$

The bandwidth of the Kernel *h* is set in such a way that approximately  $\alpha \times n$  neighbors are included in each local regression, where *n* is the total number of data points.  $||(x_i, y_i) - (x, y)||$  denotes the Euclidean distance between a sampled

and an unsampled location. The final estimate  $\hat{Z}(x, y)$  at location (x, y) is given by  $\hat{Z}(x, y) = \hat{\beta}_0 + \hat{\beta}_1 x + \hat{\beta}_2 y$ . The coefficients  $\hat{\beta}(x, y) = \{\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2\}$  are determined by minimizing the weighted sum of squared residuals, akin to traditional regression methods:

$$\hat{\beta}(x,y) = \arg\min_{\beta} \sum_{i=1}^{n} w((x,y), (x_i, y_i))(z_i - \hat{Z}(x_i, y_i))^2$$

It is worth noting that the above formulation uses a linear polynomial. It is possible to use higher-order polynomials to fit the data, though this may lead to overfitting. In our work, we restrict the scope to linear polynomials due to their low complexity and high interpretability.

4.1.3 Self-tuning Bandwidth in Kernel Regression (STBKR). The STBKR technique proposed in Jiang *et al.* [20] uses a Gaussian Kernel regression method to estimate mobile Internet quality. Their approach allows the bandwidth of the Kernel to be tuned automatically, depending on the density of measurements in the local neighborhood of an unsampled location. The STBKR estimate of  $\hat{Z}(x, y)$  at location (x, y) is given by:

$$\hat{Z}(x,y) = \frac{\sum_{i=1}^{n} K_{h(x,y)} \left( \| (x_i, y_i) - (x, y) \| \right) z_i}{\sum_{i=1}^{n} K_{h(x,y)} \left( \| (x_i, y_i) - (x, y) \| \right)}$$

where  $K_{h(x,y)}$  is the Gaussian Kernel function with band-

width *h*, given by  $K(u) = \frac{1}{\sqrt{2\pi}h(x,y)}e^{-\frac{u^2}{2h^2(x,y)}}$ 

The adaptive bandwidth h(x, y) is given by  $h(x, y) = cR_k(x, y)^2$ , where *c* is a parameter to control the bandwidth, and  $R_k(x, y)$ is the average distance between (x, y) and its *k* nearest neighbors. Parameters *c* and *k* are both tuned using cross-validation.

### 4.2 Parameter Choices

We evaluate the effectiveness of the interpolation procedures using their ability to handle out-of-sample data. To this end, we performed a 5-fold cross-validation on the dataset to tune



Figure 3: Distribution of RMSE values for LOESS, STBKR and IDW across different resolutions for Android and iOS users. The range of RMSE values decreases as resolution increases.

the parameters of each method. We hold-out 20% of the dataset in each fold and invoke the interpolation techniques on the remaining 80%. We perform LOESS with *span* values chosen out of 0.1, 0.5, 0.9 and 1. A linear model is chosen for this purpose because of its computational efficiency and lower likelihood of overfitting. For STBKR, we varied the number of nearest neighbors from 5 to 500. We choose a Gaussian kernel for this method because of its wide range of applications and ability to capture the underlying spatial distribution of the data. The *c* parameter for adjusting the kernel bandwidth was varied from  $10^{-5}$  to  $10^3$  on a logarithmic scale. For IDW, we vary the power parameter *p* from 1 to 3 in steps of 1.

## 4.3 Results

4.3.1 Overall Comparisons. Figures 3a and 3b show the overall performance in terms of root mean squared errors (RMSE) for the three methods across different resolutions and user groups. A general trend observed across all methods is that the RMSE decreases as the spatial resolution becomes finer. This suggests that as the granularity of the sampling geography increases, the precision of the interpolation techniques improves. This is likely because a finer resolution allows the models to capture more detailed spatial variations and anomalies that are lost at coarser resolutions. In Figure 3c, we analyze the percentage drop in RMSE range relative to the previous resolution value to check if it follows a consistent trend. At a resolution of 7, all three methods show a similar drop in range, suggesting that the precision of the three techniques is comparable at this resolution. However, at higher resolutions, STBKR and IDW show a greater drop in range compared to LOESS because of the local regressors generalizing to random noise or fewer data points. Interestingly, we observe a local peak for STBKR and IDW at a resolution of 8, indicating a



Figure 4: Distribution of pairwise separation between Android and iOS measurements in resolution-8 hexagons. Android measurements showed a spike close to zero meters, indicating a higher density of measurements in certain areas.

significant drop in range from 7 to 8. A resolution of 8 corresponds to a hexagon edge length of 460 meters - an area that is likely to contain measurement devices with shared infrastructure nodes. As a result, these measurements are expected to yield similar latency estimates. STBKR and IDW use weighted averages of these measurements, leading to more reliable estimates at this resolution, and thus the drop in RMSE range. Higher resolutions, such as 9 or 10 are expected to have fewer data points and may encompass only a few users and devices, leading to higher precision in STBKR and IDW estimates. It is worth noting that these resolutions may include fewer data points and users despite a higher precision, leading to a loss in reliability. We suggest that a resolution of 8 may therefore help achieve a balance between precision and reliability in latency estimates. This is also in agreement with the FCC's default resolution for availability maps.

4.3.2 Comparison over user types. Each technique shows different levels of performance and variability across the platforms. LOESS, in particular, shows higher variability in RMSE values compared to the other algorithms, as evident from the larger interquartile ranges and longer whiskers in the boxplots. This could be indicative of the sensitivity of LOESS to the underlying noise and data distribution, which may vary significantly between Android and iOS datasets. In contrast, STBKR and IDW appeared more stable and consistent across different resolutions and platforms, suggesting that they might be better suited for locations where data consistency and robustness to technique are critical.



Figure 5: Distribution of interpolated and ground truth latency for LOESS in resolution-8 hexagons for iOS users. LOESS produced invalid predictions ( $\leq 0$  ms) for a fraction of points due to potential overfitting.

Interestingly, the RMSE values for iOS are generally higher than those for Android across most algorithms and resolutions. To investigate these differences further, we looked at the distribution of the pairwise separation between measurements in each group in Figure 4. We observe a local peak for Android users close to zero meters, which could be indicative of a higher density of measurements in certain areas. Measurements clustered together are likely to have similar latencies due to their proximity to the infrastructure. In contrast, iOS users show a lower spike compared to Android, which could indicate measurements more evenly distributed across the geography and higher RMSE values compared to Android users.

4.3.3 LOESS Parameter Sensitivity. Figure 5 represents the distribution of LOESS predictions against the ground truth for iOS users. We observe that LOESS produces invalid predictions ( $\leq 0$  ms) for a fraction of points regardless of the choice of *span*. This could be indicative of potential overfitting of the local regressors to the noise in the data. Furthermore, we notice that hexagons with a lower sample size showed greater RMSE values compared to those with a higher sample size. An additional observation from Figure 5 is that the fraction of invalid latency measurements reduces as the *span* increases



(a) Best-case parameters for the hexagonal cells. Each color denotes different range combinations of *c* and *k*.



(b) Best-case RMSE distribution for each parameter range combination.

Figure 6: Distribution of best-case parameters and RMSE for STBKR in resolution-8 hexagons for iOS users. We observed the majority of cells to have low c and low k values. Hexagons with a higher k generally assumed higher RMSE values.

from 0.1 to 1.0. A higher *span* indicates a higher degree of smoothing due to a greater number of data points for local regression. This suggests that the accuracy of LOESS is dependent on the number of data points available for interpolation.

4.3.4 STBKR Parameter Sensitivity. Next, we look at the sensitivity of the STBKR approach towards the choice of c and *k*. In Figure 6a we show the distribution of *c* and *k* over the hexagonal cells in the regions of interest. We categorize the c and k values into different ranges to understand the distribution of the best-case parameters. We term  $c \leq 1$  as low and c > 1 as high. Similarly, we term  $k \le 200$  as low and k > 200as high. We find that the majority (64%) of the hexagons fall into the Low c Low k range, indicating that a lower smoothing parameter and a few neighbors are sufficient to achieve lower interpolation errors across space in reasonably dense regions. In Figure 6b, we show the distribution of the best-case RMSE values for each hexagon cell. We observe that the best-case RMSE values are generally higher for hexagons with a higher k irrespective of the c value. This suggests that a higher number of neighbors might not always lead to better estimates.

4.3.5 *IDW Parameter Sensitivity.* For IDW, we find the RMSE to be less sensitive to the choice of p, with p = 1, p = 2 and p = 3 achieving median RMSE values of 6.42, 6.60 and 6.73 milliseconds respectively at a cell-level. This suggests that the choice of p does not significantly affect the interpolation accuracy for IDW. We also find IDW to be computationally intensive because it calculates pairwise distances between each pair of points without considering a local neighborhood. This may limit the scalability of this method to larger datasets.

## **5 REGIONALIZATION**

#### Summary of findings

- We find the optimal number of clusters to be N = 4 for the hexagon cell overlay with a LOESS *span* of 0.075 and the minimum number of clustering units kept to 50. The within and between-cluster variance for this choice of parameters is found to plateau for N > 4.
- We observe significant visual differences between the clusters obtained for different sampling approaches, namely, raw community areas, interpolated community areas, and interpolated hexagon cells.
- We evaluate the stability of our clusters by measuring the Jaccard similarity between clusterings for multiple smaller samples of the dataset. The median pairwise Jaccard similarity between these clusters is found to be 0.96 when we overlaid the city's map with hexagons. For community area overlay, we find a similarity of 0.99 when we use prior interpolation, and 0.43 when we use raw measurement aggregates.

To construct the sampling boundaries for Chicago, we first interpolate latency measurements across the citywide geography for iOS users. Then, we overlay these measurements with two sampling unit choices, namely, neighborhood boundaries of Chicago and hexagon cells. Finally, we apply SKATER with averages calculated within these boundary units as the clustering metric. The outcome of SKATER results in distinct yet contiguous latency clusters over space.

## 5.1 Problem Formulation & Technique

The challenge of discovering statistical sampling boundaries for latency can be restructured as an unsupervised learning problem. Consider a geographical region  $\Omega$  and a set of its partitions  $\mathcal{H} = \{H_i\}_{i=1}^n$ . Further, consider a set of latency measurements conducted over the region as  $\mathcal{X} = \{x_i\}_{i=1}^n$ . Our goal is to find a set of spatially contiguous clusters  $C = \{C_i\}_{i=1}^N$ such that each cluster  $C_i$  is a subset of  $\mathcal{H}$ , and the latencies within each cluster are drawn from a common distribution. To achieve this, we calculate the mean latency for each partition  $H_i$  as  $\mu_i = \frac{1}{|H_i|} \sum_{x \in H_i} x$  and assign a feature vector  $\mathbf{v}_i = [\mu_i]$  to each partition. To calculate the dissimilarity among the partitions, we consider the use of Euclidean distance,  $d(\mathbf{v}_i, \mathbf{v}_j) = \|\mathbf{v}_i - \mathbf{v}_j\|$ . Using this dissimilarity function, we apply a spatial clustering algorithm to group the partitions into *N* clusters. Finally, we define the existence of a sampling boundary (*B*) between two partitions of the region  $\Omega$  as:

$$B(H_i, H_j) = \begin{cases} \text{True,} & \text{if } H_i \in C_k \text{ and } H_j \in C_l, \text{ with } k \neq l \\ \text{False,} & \text{otherwise} \end{cases}$$

Due to the resemblance between the graph-based methodology of SKATER with Internet infrastructure topology, we use it as our default clustering algorithm. SKATER involves three main steps. First, it constructs a graph where each node represents a spatial unit, e.g, a census tract boundary, hexagonal units, geographic coordinates of Internet users, or a community area. The edges between the nodes denote spatial adjacency, i.e., two nodes are connected if they share a common boundary. In case of points, the edges are constructed using a distance threshold. The weights of these edges are determined using the dissimilarity between the nodes, which is Euclidean distance in our case. In the second step, SKATER constructs a Minimum Spanning Tree (MST) from the graph. An MST is a tree that connects all the nodes in the graph using the minimum possible edge weights. The use of MST in this step ensures a faster runtime, as considering all edges in the graph is infeasible. In the final step, the MST is iteratively pruned by removing edges with the highest weights. This results in a set of connected components, one for each spatially contiguous cluster. The number and size of the clusters can be controlled using two parameters, N and floor. N denotes the number of clusters, and *floor* denotes the minimum number of nodes in each cluster.

### 5.2 Results

5.2.1 Parameter Sensitivity. Upon overlaying the interpolated map with hexagonal units, we evaluate the sensitivity of the resulting SKATER clusters to the choice of *N*, *floor*, and LOESS *span*. We analyze three different values of *floor* – 50, 75 and 95. *N* is selected from the range [2, 12] at steps of 1, while *span* is chosen from the set {0.075, 0.1, 0.25}. We evaluate parameter sensitivity by calculating the ratio of within-cluster variance (*W*) to the between-cluster variance (*B*) of latency. This ratio ( $\rho$ ) is given by:

$$\rho = \frac{\sum_{i=1}^{N} \sum_{j=1}^{n_i} (x_{ij} - \bar{x}_i)^2}{\sum_{i=1}^{N} n_i (\bar{x}_i - \bar{x})^2}$$

where *N* denotes the total number of clusters,  $n_i$  is the number of hexagon cells in the  $i^{th}$  cluster,  $x_{ij}$  is the mean latency of



(a) Sensitivity towards floor at span = 0.075. We obtain the lowest  $\rho$  for floor = 50.



(b) Within and between-cluster variance for span = 0.075, floor = 50 plateaus for N > 4.

Figure 7: Parameter sensitivity analysis for span = 0.075 and N in the range [2, 12]. We observe N = 4 to be optimal for the choice of span = 0.075 and floor = 50.



(a) Resulting clusters from (b) Distribution of mean cell-SKATER for optimal *N* level latency for the clusters.

Figure 8: Resulting clusters for N = 4, span = 0.075, and floor = 50 and the distribution of cell-level average latencies in each cluster.

the  $j^{th}$  hexagon cell in the  $i^{th}$  cluster,  $\bar{x}_i$  is the average latency of the *i*<sup>th</sup> cluster, and  $\bar{x}$  is the overall average latency. The numerator in this equation is a proxy for the compactness of the clusters, while the denominator measures the separation between the clusters. A low  $\rho$  value indicates that the clusters are well-separated and compact. Typically, the optimal number of clusters is chosen at the "elbow" of the  $\rho$  curve, where using a greater N only introduces marginal changes in  $\rho$ . Figure 7 shows the sensitivity of the clusters to the choice of *floor*, *N*, and *span*. Figure 7a shows the  $\rho$  values for different *floor* values at *span* = 0.075. We observe that the  $\rho$  values are lowest for floor = 50, indicating that the clusters are well-separated and compact. Figure 7b shows the within and between-cluster variance for span = 0.075 and floor = 50. We observe that both the within and between-cluster variance change only slightly beyond N = 4, indicating that a greater N would

only show insignificant differences in latency among clusters. Figure 8a shows the resulting clusters for span = 0.075, floor = 50, and N = 4, while Figure 8b shows the distribution of mean cell-level latency in these clusters. The dark blue cluster (cluster-0) shows a multi-modal distribution, which indicates that it contains hexagon cells with a wide range of latencies. The red cluster (cluster-1) shows moderate latency values, while the pink cluster (cluster-2) shows the highest latency values. The sky-blue region towards the west (cluster-3) has the fewest hexagon cells and shows a wider range of interpolated latency means.

The above analysis highlights the key implications of our clustering approach. Using a similar analysis, ISPs and regulators can make informed decisions about where to prioritize infrastructure investments. For example, regions with high latencies, such as cluster-2, may benefit from targeted interventions such as infrastructure upgrades or improved connectivity. This targeted approach can help improve overall network performance and provide more equitable access to quality Internet across the city.

*5.2.2 Visual Comparisons.* We first evaluate whether the boundaries obtained for different sampling approaches were visually consistent. For our choice of sampling approaches, we consider three combinations of sampling unit choices and aggregation methods: (1) community area units with raw measurement aggregates [33] (Community-Raw), (2) community area units with interpolated measurement aggregates (Community-Interpolated), and (3) hexagon cell units with interpolated measurement aggregates (Hex-Interpolated). For the Community-Interpolated and Hex-Interpolated approaches,



Figure 9: Regionalizations of Chicago for three sampling approaches using the same number of clusters (N = 4). Figure 9a shows community area clusters calculated using raw aggregates, Figure 9b shows the community area boundaries after LOESS interpolation, and Figure 9c shows hexagon cell clusters after LOESS interpolation. The choice of aggregation method and sampling unit can significantly affect resulting sampling boundaries, and hence our conclusions about the spatial distribution of latency.

we use LOESS with a span of 0.075 to interpolate at a regular grid spacing of 100 meters (Section 3). Since we perform interpolation over the complete city, we do not observe any invalid predictions for LOESS. This is because of the greater abundance of data for local regression in comparison to individual hexagons. Before applying SKATER, we fixed the number of clusters to N = 4, and adjust *floor* such that we obtained the same set of clusters for each approach for all  $N \geq 4$ . This is done to ensure the reliability of our analysis, as we want to make comparisons across the most consistent clusterings for the three sampling approaches. We use a floor of 18 each for Community-Raw and Community-Interpolated, and 50 for the Community-Interpolated and Hex-Interpolated approaches. A higher floor value is chosen for the Hex-Interpolated approach to account for the increased number of cells compared to community areas.

Figure 9 shows the clusters obtained for the three sampling approaches. We observe that the choice of aggregation method and sampling unit significantly affects the resulting boundaries. The transition from Figure 9a to Figure 9b shows that applying interpolation to raw aggregates can result in a fundamentally different perception of the underlying spatial trends. For instance, the red cluster in Figure 9a extends up to the northern boundary of Chicago, while in Figure 9b, it is confined to the central region. We also notice fewer community areas in this cluster in Figure 9b compared to Figure 9a. Further, the transition from Figure 9b to Figure 9c shows that the choice of sampling unit can also significantly affect the resulting boundaries. The northernmost blue cluster remains consistent across the first two approaches, but it split into distinct regions for the Hex-Interpolated approach. Hex-Interpolated could thus capture fine-grained spatial differences that were



Figure 10: A comparison of pairwise Jaccard similarity between bootstrapped clusters for three different boundary identification approaches. Boundaries constructed post interpolation showed higher median Jaccard similarity among bootstrapped clusters.

not apparent in the other two approaches. These visual differences highlight the need for a careful choice of sampling units for regulatory purposes. Depending on the policy goals, one might choose to use community areas for their convenience towards local governance, or select a regular tessellation such as hexagonal cells for their ability to capture fine-grained spatial differences.

*5.2.3 Stability Evaluation.* We compare the similarity between the clusters obtained for the three approaches by employing the use of Bootstrap resampling technique. Bootstrap resampling involves drawing a number of samples with replacement to estimate the distribution of a given statistic, whose distribution is often unknown. In our use case, we are using this technique to simulate the perturbations across clusters to estimate the distribution of a similarity metric. Continuing with our choice of N = 4, we draw 100 bootstrap samples for the three approaches. A total of 100 bootstraps yields 4950 distinct pairs of regionalizations for each baseline approach. A total of 1000 samples is often considered adequate for most practical use cases [14], so this sample size is reasonably large. These bootstraps are drawn at the level of the sampling unit, i.e., community areas for Community-Raw and Community-Interpolated, and hexagon cells for Hex-Interpolated. In case of the interpolation-based approaches, these bootstraps are drawn from the interpolated grid points, while for the raw aggregates, they are drawn directly from the preprocessed dataset. Then, we apply SKATER to individual bootstraps and compute the Jaccard similarity between each pair of samples. The similarity scores are calculated between the assigned clusters for each pair of bootstraps. As an example, consider the calculation of Jaccard similarity for the  $3^{rd}$  cluster for two samples A and B with N = 4. Now, consider the set of hexagon cells that belong to the  $3^{rd}$  cluster in sample A as  $A_3$ , and in sample B as  $B_3$ . The Jaccard similarity between these two sets is given by  $|A_3 \cap B_3| / |A_3 \cup B_3|$ . We repeat this process for all pairs of samples and all clusters. The distribution of Jaccard similarity scores for the three approaches is shown in Figure 10. We observe that Community-Interpolated shows the highest median Jaccard similarity of 0.99, followed by 0.96 for Hex-Interpolated, and 0.43 for Community-Raw. The higher similarity of Community-Interpolated and Hex-Interpolated indicates that post-interpolation aggregation of latencies at finer spatial scales can lead to more reliable sampling boundaries. Another compelling evidence for the benefits of prior interpolation is the large difference between the Jaccard scores of Community-Interpolated and Community-Raw. We observe that for the same choice of sampling units, similarity between the resulting boundaries can be significantly improved by spatially interpolating raw measurements. In addition to reaffirming the importance of the choice of sampling unit and aggregation method, our findings highlight the benefits of prior spatial interpolation. This further implies that taking into account the spatial distribution of measurements can lead to more reliable sampling boundaries.

## **6** CONCLUSION & FUTURE WORK

This work presents a novel approach for discovering Internet performance sampling boundaries within a city using crowdsourced latency measurements. The findings of this study underscore the importance of spatial analysis in network planning and the benefits of targeted infrastructure investments for equitable Internet access. Our methods treated each point measurement individually for interpolation to aggregate latency within hexagonal cells. Future research should explore the use of novel metrics that allow for more accurate representation of latency patterns. Further, our methods aim to provide a nuanced spatial view of Internet performance using a crowdsourced dataset. Future research should build on this approach by incorporating more comprehensive data and temporal analysis to improve our understanding of Internet performance. Finally, our study considered the use of hexagonal tessellations as a tool for capturing spatial continuity. Future research should explore the accuracy of tessellations of varying shapes and sizes.

## **ETHICS**

We used the *Ookla for good* dataset for our analysis. It is a publically available dataset [25]. In this dataset, the geolocations of Ookla users were truncated upto 4 decimal places, which allows a margin of a few hundred meters. The IP addresses were masked up to the last octet, which ensures anonymity. We did not find any other personally identifiable information in the dataset. Our research, therefore, does not raise any ethical concerns.

## REFERENCES

- 2024. MLab Test Your Speed. https://speed.measurementlab.net/. Accessed: 2022.
- [2] 2024. Netalyzr. http://netalyzr.icsi.berkeley.edu/. Accessed on April 12, 2024.
- [3] 2024. Ookla Speedtest. https://www.speedtest.net/. Accessed: 2024.
- [4] Jeremy Aldworth and Noel Cressie. 1999. Sampling designs and prediction methods for Gaussian spatial processes. In *Multivariate analysis, design of experiments, and survey sampling.* CRC Press, 25–78.
- [5] Luc Anselin. 2018. Spatial Clustering (2). Disponible en (2018).
- [6] Renato M Assunção, Marcos Corrêa Neves, Gilberto Câmara, and Corina da Costa Freitas. 2006. Efficient regionalization techniques for socio-economic geographical units using minimum spanning trees. *International Journal of Geographical Information Science* 20, 7 (2006), 797–811.
- [7] Battle For the Net. 2022. Internet Health Test based on Measurement Lab NDT. https://www.battleforthenet.com/internethealthtest/
- [8] Zachary S Bischof, Fabian E Bustamante, and Nick Feamster. 2017. Characterizing and improving the reliability of broadband internet access. arXiv preprint arXiv:1709.09349 (2017).
- [9] Broadband Internet Technical Advisory Group (BITAG). 2022. Latency Explained. https://www.bitag.org/documents/BITAG\_latency\_ explained.pdf
- [10] M. Cappellari, R. M. McDermid, K. Alatalo, L. Blitz, M. Bois, F. Bournaud, M. Bureau, A. F. Crocker, R. L. Davies, T. A. Davis, P. T. de Zeeuw, P.-A. Duc, E. Emsellem, S. Khochfar, D. Krajnović, H. Kuntschner, R. Morganti, T. Naab, T. Oosterloo, M. Sarzi, N. Scott, P. Serra, A.-M. Weijmans, and L. M. Young. 2013. The ATLAS<sup>3D</sup> project - XX. Mass-size and mass-σ distributions of early-type galaxies: bulge fraction drives kinematics, massto-light ratio, molecular gas fraction and stellar initial mass function. *MNRAS* 432 (2013), 1862–1893. https://doi.org/10.1093/mnras/stt644 arXiv:1208.3523
- [11] David D Clark and Sara Wedeman. 2021. Measurement, Meaning and Purpose: Exploring the M-Lab NDT Dataset. In TPRC49: The 49th Research Conference on Communication, Information and Internet Policy.

- [12] Noel Cressie. 1988. Spatial prediction and ordinary kriging. Mathematical geology 20 (1988), 405–421.
- [13] Juan C Duque, Luc Anselin, and Sergio J Rey. 2012. The max-p-regions problem. *Journal of Regional Science* 52, 3 (2012), 397–419.
- [14] Bradley Efron and Robert J Tibshirani. 1994. An introduction to the bootstrap: CRC press. Ekman, P., & Friesen, WV (1978). Manual for the facial action coding system (1994).
- [15] Uber Engineering. 2024. Introducing H3: Uber's Hexagonal Hierarchical Spatial Index. https://www.uber.com/blog/h3/. Accessed: Date of access.
- [16] Federal Communications Commission. 2022. FTC Takes Action Against Frontier for Lying about Internet Speeds and Ripping Off Customers Who Paid High-Speed Prices for Slow Service. Press Release. https://www.ftc.gov/news-events/news/pressreleases/2022/05/ftctakes-action-against-frontier-lying-about-internet-speeds-rippingcustomers-who-paid-highspeed.
- [17] Peter Cody Fiduccia. 2022. Deconstructing the Digital Divide: The Geography, Demography, and Spatial Dependence of Internet Stability in the US. Cornell University.
- [18] Syed Hasan, Sergey Gorinsky, Constantine Dovrolis, and Ramesh K Sitaraman. 2014. Trade-offs in optimizing the cache deployments of CDNs. In *IEEE INFOCOM 2014-IEEE conference on computer communications*. IEEE, 460–468.
- [19] Cheng Huang, Angela Wang, Jin Li, and Keith W Ross. 2008. Measuring and evaluating large-scale CDNs. In ACM IMC, Vol. 8. 15–29.
- [20] Hanyang Jiang, Henry Shaowu Yuchi, Elizabeth Belding, Ellen Zegura, and Yao Xie. 2023. Mobile Internet Quality Estimation using Self-Tuning Kernel Regression. arXiv preprint arXiv:2311.05641 (2023).
- [21] Hyeongseong Lee, Udit Paul, Arpit Gupta, Elizabeth Belding, and Mengyang Gu. 2023. Analyzing Disparity and Temporal Progression of Internet Quality through Crowdsourced Measurements with Bias-Correction. arXiv preprint arXiv:2310.16136 (2023).
- [22] Kyle MacMillan, Tarun Mangla, James Saxon, Nicole P Marwell, and Nick Feamster. 2023. A comparative analysis of ookla speedtest and measurement labs network diagnostic test (ndt7). Proceedings of the ACM on Measurement and Analysis of Computing Systems 7, 1 (2023), 1–26.
- [23] Jonatas Marques, Alexis Schrubbe, Nicole P Marwell, and Nick Feamster. 2024. Are We Up to the Challenge? An analysis of the FCC Broadband Data Collection Fixed Internet Availability Challenges. arXiv preprint arXiv:2404.04189 (2024).
- [24] New York State Office of the Attorney General. 2020. New York Internet Health Test. https://ag.ny.gov/SpeedTest
- [25] Ookla. 2024. Ookla for Good: Open Data. https://www.ookla.com/ ookla-for-good/open-data Accessed: 2024-05-15.
- [26] Udit Paul, Vinothini Gunasekaran, Jiamo Liu, Tejas N Narechania, Arpit Gupta, and Elizabeth Belding. 2023. Decoding the Divide: Analyzing Disparities in Broadband Plans Offered by Major US ISPs. In Proceedings of the ACM SIGCOMM 2023 Conference. 578–591.
- [27] Udit Paul, Jiamo Liu, Vivek Adarsh, Mengyang Gu, Arpit Gupta, and Elizabeth Belding. 2021. Characterizing performance inequity across us ookla speedtest users. arXiv preprint arXiv:2110.12038 (2021).
- [28] Udit Paul, Jiamo Liu, Mengyang Gu, Arpit Gupta, and Elizabeth Belding. 2022. The importance of contextualization of crowdsourced active speed test measurements. In *Proceedings of the 22nd ACM Internet Mea*surement Conference. 274–289.
- [29] Pennsylvania State University and Measurement Lab. 2019. Broadband Availability and Access in Rural Pennsylvania. https://www.rural.pa. gov/publications/broadband.cfm
- [30] Juan P Rigol, Claire H Jarvis, and Neil Stuart. 2001. Artificial neural networks as a tool for spatial interpolation. *International Journal of Geographical Information Science* 15, 4 (2001), 323–343.

- [31] James Saxon and Dan A Black. 2022. What we can learn from selected, unmatched data: measuring Internet inequality in Chicago. Computers, Environment and Urban Systems 98 (2022), 101874.
- [32] Aleksandar Sekulić, Milan Kilibarda, Gerard BM Heuvelink, Mladen Nikolić, and Branislav Bajat. 2020. Random forest spatial interpolation. *Remote Sensing* 12, 10 (2020), 1687.
- [33] Ranya Sharma, Tarun Mangla, James Saxon, Marc Richardson, Nick Feamster, and Nicole P Marwell. 2022. Benchmarks or Equity? A New Approach to Measuring Internet Performance. A New Approach to Measuring Internet Performance (August 3, 2022) (2022).
- [34] Taveesh Sharma, Jonatas Marques, Nick Feamster, and Nicole P Marwell. 2023. A First Look at the Spatial and Temporal Variability of Internet Performance Data in Hyperlocal Geographies. Available at SSRN 4568668 (2023).
- [35] Donald Shepard. 1968. A two-dimensional interpolation function for irregularly-spaced data. In Proceedings of the 1968 23rd ACM national conference. 517–524.
- [36] Joel Sommers and Paul Barford. 2012. Cell vs. WiFi: On the performance of metro area mobile connections. In *Proceedings of the 2012 internet measurement conference*. 301–314.
- [37] Sean Stokes and Kathleen Slattery Thompson. 2022. With Billions of Dollars of Broadband Funding at Stake, the Timing of the Challenge Process to the FCC's Broadband Map Under Increasing Scrutiny. The National Law Review (2022). https://natlawreview.com/article/billions-dollars-broadbandfunding-stake-timing-challenge-process-to-fcc-s-broadband
- [38] Amy Stuyvesant. 2023. Michigan Broadband Personas-Assessing Why Households Lack Reliable Service Using Survey Responses, Speed Tests, and Location. Speed Tests, and Location (July 31, 2023) (2023).
- [39] Srikanth Sundaresan, Sam Burnett, Nick Feamster, and Walter De Donato. 2014. {BISmark}: A testbed for deploying measurements and applications in broadband access networks. In 2014 USENIX Annual Technical Conference (USENIX ATC 14). 383–394.
- [40] Srikanth Sundaresan, Walter De Donato, Nick Feamster, Renata Teixeira, Sam Crawford, and Antonio Pescapè. 2011. Broadband internet performance: a view from the gateway. ACM SIGCOMM computer communication review 41, 4 (2011), 134–145.
- [41] Srikanth Sundaresan, Xiaohong Deng, Yun Feng, Danny Lee, and Amogh Dhamdhere. 2017. Challenges in inferring internet congestion using throughput measurements. In Proceedings of the 2017 Internet Measurement Conference. 43–56.
- [42] Cort J Willmott and Kenji Matsuura. 2006. On the use of dimensioned measures of error to evaluate the performance of spatial interpolators. *International journal of geographical information science* 20, 1 (2006), 89–102.