

Refining Coarse-Grained Spatial Data Using Auxiliary Spatial Data Sets with Various Granularities

Yusuke Tanaka,^{1,3} Tomoharu Iwata,² Toshiyuki Tanaka,³ Takeshi Kurashima,¹
Maya Okawa,¹ Hiroyuki Toda¹

¹NTT Service Evolution Laboratories, NTT Corporation

²NTT Communication Science Laboratories, NTT Corporation

³Graduate School of Informatics, Kyoto University

{tanaka.y, iwata.tomoharu, kurashima.takeshi, okawa.maya, toda.hiroyuki}@lab.ntt.co.jp, tt@i.kyoto-u.ac.jp

Abstract

We propose a probabilistic model for refining coarse-grained spatial data by utilizing auxiliary spatial data sets. Existing methods require that the spatial granularities of the auxiliary data sets are the same as the desired granularity of target data. The proposed model can effectively make use of auxiliary data sets with various granularities by hierarchically incorporating Gaussian processes. With the proposed model, a distribution for each auxiliary data set on the continuous space is modeled using a Gaussian process, where the representation of uncertainty considers the levels of granularity. The fine-grained target data are modeled by another Gaussian process that considers both the spatial correlation and the auxiliary data sets with their uncertainty. We integrate the Gaussian process with a spatial aggregation process that transforms the fine-grained target data into the coarse-grained target data, by which we can infer the fine-grained target Gaussian process from the coarse-grained data. Our model is designed such that the inference of model parameters based on the exact marginal likelihood is possible, in which the variables of fine-grained target and auxiliary data are analytically integrated out. Our experiments on real-world spatial data sets demonstrate the effectiveness of the proposed model.

1 Introduction

Many cities around the world are now collecting large amounts of spatial data from a wide range of sources. Governments and other organizations are releasing data on items such as poverty rate, air pollution, traffic flow, energy consumption and crime (Shadbolt et al. 2012; Goldstein and Dyson 2013; Barlacchi et al. 2015). Analyzing such spatial data is of critical importance in improving the life quality of citizens in many fields such as socio-economics (Rupasinghaa and Goetz 2007; Smith, Mashhadi, and Capra 2014), public health (Jerrett et al. 2013), public security (Bogomolov et al. 2014; Wang et al. 2016) and urban planning (Yuan, Zheng, and Xie 2012). For example, knowing the spatial distribution of poverty enables us to optimize allocation of resources for remedial action. Likewise, the spatial distribution of air pollution is useful in creating policies that can control air quality and thus protect human health.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: The distribution of poverty rates at different spatial granularities.

Naturally, information at fine spatial granularity is preferred because it allows us to identify key regions that require intervention to improve city environments efficiently. As an example, Figures 1(a) and 1(b) visualize the distributions of poverty rates in New York City by community district and by borough, respectively; darker hues represent poorer regions. Clearly, to better understand socio-economic problems, Figure 1(a) is better than Figure 1(b). In practice, however, such information is often aggregated into coarse granularities as in Figure 1(b). It is usually thought to be too time-consuming and costly to conduct a *census* over the whole population of a city, and a *sample survey* is conducted instead. Accordingly, the number of samples associated with each fine-grained region may not be large enough to provide a statistically significant estimate of the value associated to this region; the typical response is to aggregate samples over larger regions (Smith, Mashhadi, and Capra 2014).

With the recent increase in data availability, utilizing auxiliary spatial data sets on the same region is an effective way of refining coarse-grained target data (Bogomolov et al. 2014; Park 2013; Smith and Capra 2016; Smith, Mashhadi, and Capra 2014; Wotling et al. 2000). In these works, the regression models are used for estimating the relationships between target data (e.g., poverty rate) and auxiliary data sets (e.g., unemployment rate). These existing methods, however, require that the spatial granularities of all the auxiliary data sets are the same as the desired granularity of target data. This requirement prevents us from making full use of the auxiliary data sets with various geographi-

cal partitions. For example, New York City has released various spatial data sets portioned into boroughs, community districts, zip code, police precincts and so on.

We propose a probabilistic model for refining coarse-grained target data via the effective use of auxiliary data sets with various granularities. An important characteristic is discerning the usefulness of each auxiliary data set which depends on not only the strength of relationship with the target data but also the level of spatial granularity. For example, consider the case of two auxiliary data sets that have the same strength of relationship with the target data, but different granularities. In that case, the finer-grained one is seen as more helpful for refining the coarse-grained target data.

With the proposed model, the fine-grained target data are assumed to follow a Gaussian process (GP) (Rasmussen and Williams 2006) whose mean function is modeled by a linear regression of the auxiliary data sets. This GP-based modeling allows us to consider the spatial correlation in the target data and the auxiliary data sets simultaneously. Since the target data are observed not at fine granularity but at coarse granularity, we model a spatial aggregation process to transform the fine-grained target data into the coarse-grained target data. Furthermore, to handle auxiliary data sets with various granularities, we apply GP regression to each auxiliary data set to derive a predictive distribution defined on the continuous space; this conceptually corresponds to spatial interpolation. A key idea is that it hierarchically incorporates the predictive distributions into the model; that is, it does not use point estimates. This enables us to consider uncertainty in the prediction of auxiliary data sets. The uncertainty is governed by several factors, one of which is sample density, i.e., spatial granularity of the auxiliary data; the finer the granularity is, the lower the uncertainty is. Incorporating the uncertainty leads to effectively learning the usefulness of the auxiliary data with consideration of the levels of spatial granularity; this allows our model to accurately refine the coarse-grained target data. We predict the fine-grained target data via a Bayesian inference procedure. The proposed model is designed such that the estimation of model parameters based on the exact marginal likelihood is possible: By analytically integrating out the variables of fine-grained target and auxiliary data, we can estimate the parameters without explicitly obtaining these variables. We construct the predictive distribution of the fine-grained target data by using the estimated parameters.

2 Related Work

The problem of refining coarse-grained spatial data has been studied in various fields such as socio-economics (Smith and Capra 2016; Smith, Mashhadi, and Capra 2014), agricultural economics (Howitt and Reynaud 2003; Xavier et al. 2016), epidemiology (Sturrock et al. 2014), meteorology (Wilby et al. 2004; Zorita and von Storch 1999) and geographical information system (GIS) (Boucher and Kyriakidis 2006; Goovaerts 2010). This problem is also called *statistical downscaling*, *spatial disaggregation*, and *areal interpolation*. The previous works can be categorized into two cases in terms of target data availability.

In the first case, in which a large amount of coarse- and fine-grained target data are available, we can predict the fine-grained target data by using a mapping function from coarse- to fine-grained data. The mapping function can be learnt by using various machine learning methods including linear regression models (Hessami et al. 2008), neural networks (Cannon 2011; Misra, Sarkar, and Mitra 2017) and support vector machines (Ghosh 2010). Recently, super-resolution techniques based on deep neural networks have been applied for refining coarse-grained spatial data (Vandal et al. 2017; 2018). The super-resolution techniques aim to learn a mapping function from low- to high-resolution images (Dong et al. 2014). The method by (Vandal et al. 2017) is based on the analogy between gridded spatial data and images; values at grid cells are regarded as values at pixels. The large amount of fine-grained data needed for training are, however, not available in many cases (e.g., poverty survey), and often only coarse-grained data are available. These methods are not applicable in such situations.

In the second case, in which only coarse-grained target data are available, many regression-based methods have been proposed that use auxiliary spatial data sets to refine coarse-grained target data (Flaxman, Wang, and Smola 2015; Smith, Mashhadi, and Capra 2014; Wang et al. 2016; Zheng, Liu, and Hsieh 2013; Zheng et al. 2015). Regression models (linear and non-linear) are used for estimating the relationships between target data and auxiliary data sets. A few methods can construct the regression models under the spatial aggregation constraints (Murakami and Tsutsumi 2011; Park 2013). The constraints state that a value associated with a coarse-grained region is a linear average of their constituent values in a fine-grained partition. In order to satisfy the spatial aggregation constraints, the regression residuals at the coarse-grained regions are allocated to the fine-grained regions by using the spatial interpolation method, i.e., kriging (Stein 1999). These methods, however, assume that the auxiliary data sets have spatial granularities equivalent to that of fine-grained target data to be estimated. This assumption makes it difficult to utilize multiple auxiliary data sets with various granularities.

Several regression methods have been developed for estimating relationships between multi-scale spatial data sets (Miller et al. 2015; Diodato et al. 2010; Xu 2017; Xu et al. 2018). These methods predict the target data with the same granularity as that of the training data by utilizing multi-scale auxiliary data sets. They do not, however, consider the spatial aggregation constraint, which is a critical factor in refining the coarse-scale target data.

There have been several hierarchical Bayesian models to predict fine-grained target data using fine-grained auxiliary data sets. Although they introduce a fully Bayesian inference (Taylor, Andrade-Pacheco, and Sturrock 2018; Wilson and Wakefield 2018; Keil et al. 2013) or a variational inference (Law et al. 2018) for model parameters, the uncertainty in the prediction of auxiliary data sets is ignored: They cannot discern the usefulness of each auxiliary data set considering their levels of spatial granularity.

Different from prior works, the proposed model can effectively make use of auxiliary data sets with various granu-

Table 1: Notation.

Symbol	Description
\mathcal{S}	set of indices of auxiliary spatial data sets
s	index of auxiliary spatial data set, $s \in \mathcal{S}$
\mathcal{X}	total region of a city
\mathbf{x}	location point represented by latitude and longitude coordinates, $\mathbf{x} \in \mathcal{X}$
$\mathcal{P}^{\text{coar}}$	coarse-grained partition of \mathcal{X} of target data
i	region in the coarse-grained partition of target data, $i \in \mathcal{P}^{\text{coar}}$
$\mathcal{P}^{\text{fine}}$	fine-grained partition of \mathcal{X} of target data
j	region in the fine-grained partition of target data, $j \in \mathcal{P}^{\text{fine}}$
\mathcal{P}_s	partition of \mathcal{X} of s th auxiliary data set
p	region in the partition of s th auxiliary data set, $p \in \mathcal{P}_s$
a_i	value associated with region i in coarse-grained target data, $a_i \in \mathbb{R}$
z_j	value associated with region j in fine-grained target data, $z_j \in \mathbb{R}$
$y_{s,p}$	value associated with region p in s th auxiliary data set, $y_{s,p} \in \mathbb{R}$

larities by hierarchically incorporating Gaussian processes. This hierarchical modeling allows us to effectively learn the usefulness of each auxiliary data set considering the levels of spatial granularity. Our model also considers the spatial aggregation constraints by integrating the Gaussian processes with a spatial aggregation process to transform the fine-grained target data into the coarse-grained target data.

3 Problem Formulation

In this section, we describe the spatial data this study focuses on, and define our problem of refining coarse-grained spatial data by using, for the same region, auxiliary spatial data sets with various granularities. Assume that we have a target spatial data set with coarse granularity, and we would like to obtain a fine-grained version. Let \mathcal{S} be the collection of indices of auxiliary data sets. The notations used in this paper are listed in Table 1.

Partition: Let \mathcal{X} be a total region of a city, and $\mathbf{x} \in \mathcal{X}$ be a location point represented by its coordinates (e.g., latitude and longitude). Partition \mathcal{P} of \mathcal{X} is a collection of disjoint subsets, called *regions*, of \mathcal{X} , whose union is equal to \mathcal{X} . Let $|\mathcal{P}|$ denote the number of regions in \mathcal{P} . We can consider several partitions of \mathcal{X} as follows. Let $\mathcal{P}^{\text{coar}}$ be the coarse-grained partition, i.e., that of the coarse-grained target data. Let $\mathcal{P}^{\text{fine}}$ be the fine-grained partition, of the desired fine-grained target data. For $s \in \mathcal{S}$, let \mathcal{P}_s be the partition of the s th auxiliary data set.

Spatial data: Let $\mathbf{a} = (a_1, \dots, a_{|\mathcal{P}^{\text{coar}}|})^\top$ be a $|\mathcal{P}^{\text{coar}}|$ -dimensional vector consisting of the coarse-grained target values, where $a_i \in \mathbb{R}$ is the value associated with region $i \in \mathcal{P}^{\text{coar}}$. For $s \in \mathcal{S}$, let $\mathbf{y}_s = (y_{s,1}, \dots, y_{s,|\mathcal{P}_s|})^\top$ be a $|\mathcal{P}_s|$ -dimensional vector consisting of the s th auxiliary data values, where $y_{s,p} \in \mathbb{R}$ is the value associated with region $p \in \mathcal{P}_s$ of the s th auxiliary data set.

Problem: Suppose that we have coarse-grained target data \mathbf{a} whose partition is $\mathcal{P}^{\text{coar}}$, auxiliary data sets with the respective partitions $\{(\mathcal{P}_s, \mathbf{y}_s) \mid s \in \mathcal{S}\}$, and the desired

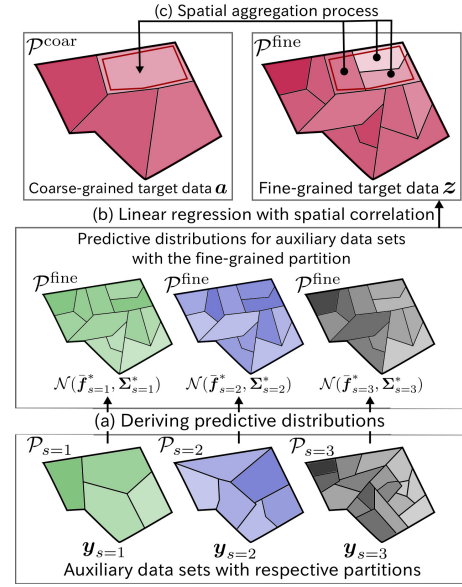


Figure 2: Generative process of coarse-grained target data given three auxiliary data sets.

fine-grained partition $\mathcal{P}^{\text{fine}}$, we wish to estimate a $|\mathcal{P}^{\text{fine}}|$ -dimensional vector $\mathbf{z} = (z_1, \dots, z_{|\mathcal{P}^{\text{fine}}|})^\top$ consisting of the fine-grained target values, where $z_j \in \mathbb{R}$ is the value associated with region $j \in \mathcal{P}^{\text{fine}}$. Here, the values a_i , $y_{s,p}$ and z_j are assumed to be intensive quantities such as ratios; that is, they are independent of the area scale of the respective regions. When the values are extensive quantities such as population, they can be transformed into intensive quantities by dividing them with the areas of regions.

4 Proposed Model

We propose a probabilistic model that allows auxiliary spatial data sets with various granularities to be used in refining coarse-grained spatial data. Our model is based on Gaussian process (GP) (Rasmussen and Williams 2006), which is a flexible non-parametric model for non-linear functions in a continuous domain. We model the generative process for coarse-grained target data \mathbf{a} , given the auxiliary data sets with known partitions $\{(\mathcal{P}_s, \mathbf{y}_s) \mid s \in \mathcal{S}\}$, coarse-grained partition $\mathcal{P}^{\text{coar}}$, and fine-grained partition $\mathcal{P}^{\text{fine}}$. In other words, we model the conditional probability $p(\mathbf{a} \mid \{\mathbf{y}_s\}_{s \in \mathcal{S}})$ instead of the joint probability of \mathbf{a} and $\{\mathbf{y}_s\}_{s \in \mathcal{S}}$. It enables us to adopt two-step inference approach described in Section 5, which is advantageous in the computational cost for learning model parameters.

The generative process (given three auxiliary data sets) is illustrated schematically in Figure 2, where darker hues represent regions with higher values. This process contains the following three steps: (a) Deriving the predictive distribution over continuous space for each auxiliary data set \mathbf{y}_s via GP regression, which corresponds to spatial interpolation; (b) generating the fine-grained target data \mathbf{z} via a GP whose mean function is modeled as the linear regression of the con-

tinuous predictive distributions of the auxiliary data sets; (c) generating the coarse-grained target data \mathbf{a} by spatially aggregating the constituent values in a fine-grained partition.

In our problem, each value is associated with a region in a partition rather than a single location point in \mathcal{X} ; this prevents us from directly applying GP. We thus associate each region in a partition with its centroid, and regard each value as being associated with the centroid of that region. This assumption, while significantly simplifying computations involved, might worsen the fit of the GP to the data set, which however is appropriately taken into account in the following steps as increased uncertainty of the GPs for both the respective auxiliary data sets (described in (5)) and the target data (described in (6)). For $s \in \mathcal{S}$, let $\mathbf{X}_s = (\mathbf{x}_{s,1}, \dots, \mathbf{x}_{s,|\mathcal{P}_s|})$ be the set of the centroids in partition \mathcal{P}_s , where $\mathbf{x}_{s,p}$ is the centroid of region $p \in \mathcal{P}_s$. Similarly, for fine-grained partition $\mathcal{P}^{\text{fine}}$, let $\mathbf{X}^{\text{fine}} = (\mathbf{x}_1, \dots, \mathbf{x}_{|\mathcal{P}^{\text{fine}}|})$ be the set of centroids in $\mathcal{P}^{\text{fine}}$. Thus, our problem is now reformulated as estimating $\mathbf{z} = (z_1, \dots, z_{|\mathcal{P}^{\text{fine}}|})^\top$, where $z_j \in \mathbb{R}$ is a target value at the centroid of region $j \in \mathcal{P}^{\text{fine}}$, as indicated by the auxiliary spatial data sets $\{(\mathbf{X}_s, \mathbf{y}_s) \mid s \in \mathcal{S}\}$.

(a) Deriving predictive distributions of auxiliary spatial data sets: In order to handle auxiliary spatial data sets with various granularities, we use GP regression to derive a posterior Gaussian process for a latent continuous random function on \mathcal{X} ; this conceptually corresponds to spatial interpolation of each auxiliary spatial data set. We then evaluate the predictive distribution on the basis of the posterior Gaussian process. Let $f_s(\mathbf{x})$ be a noise-free latent function for the s th auxiliary data set at location \mathbf{x} . We assume that $f_s(\mathbf{x})$ follows a Gaussian process, $f_s(\mathbf{x}) \sim \mathcal{GP}(0, k_s(\mathbf{x}, \mathbf{x}'))$, with mean zero and a covariance function $k_s(\mathbf{x}, \mathbf{x}')$. Though our model does not depend on any particular choice of the covariance function, for simplicity we consider the well-known covariance function, i.e., squared-exponential kernel, which is widely used for measuring the similarity between function values in spatial coordinates (Rasmussen and Williams 2006). The squared-exponential kernel is defined as

$$k_s(\mathbf{x}, \mathbf{x}') = \alpha_s^2 \exp\left(-\frac{1}{2\gamma_s^2} \|\mathbf{x} - \mathbf{x}'\|^2\right), \quad (1)$$

where γ_s is the scale parameter, α_s^2 is a signal variance that controls the magnitude of the covariance, and $\|\cdot\|$ is the Euclidean norm. We assume that the s th auxiliary data \mathbf{y}_s is generated with an additive Gaussian noise with noise variance σ_s^2 . Defining $f_s^*(\mathbf{x}_p)$ as the prediction of the s th auxiliary data set for the centroid \mathbf{x}_p of the fine-grained partition, the predictive distribution of $\mathbf{f}_s^* = (f_s^*(\mathbf{x}_1), \dots, f_s^*(\mathbf{x}_{|\mathcal{P}^{\text{fine}}|}))^\top$ is as follows:

$$p(\mathbf{f}_s^*) = \mathcal{N}(\mathbf{f}_s^* \mid \bar{\mathbf{f}}_s^*, \Sigma_s^*), \quad (2)$$

where $\bar{\mathbf{f}}_s^* = \mathbf{K}_{s*}^\top (\mathbf{K}_s + \sigma_s^2 \mathbf{I})^{-1} \mathbf{y}_s$ is the predictive means, and $\Sigma_s^* = \mathbf{K}_{s**} - \mathbf{K}_{s*}^\top (\mathbf{K}_s + \sigma_s^2 \mathbf{I})^{-1} \mathbf{K}_{s*}$ is the covariance matrix, whose diagonal elements represent the uncertainties in the prediction at the test points \mathbf{X}^{fine} . Incorporation of the predictive distributions (2) is expected to allow the usefulness of auxiliary data to be effectively learnt as it allows consideration of the uncertainty in the prediction. Details are

given in (7) in Section 5. Here, \mathbf{K}_s is a $|\mathcal{P}_s| \times |\mathcal{P}_s|$ covariance matrix whose entries are covariances between training points \mathbf{X}_s . \mathbf{K}_{s*} is a $|\mathcal{P}_s| \times |\mathcal{P}^{\text{fine}}|$ covariance matrix whose entries are covariances between training points \mathbf{X}_s and test points \mathbf{X}^{fine} . \mathbf{K}_{s**} is a $|\mathcal{P}^{\text{fine}}| \times |\mathcal{P}^{\text{fine}}|$ covariance matrix whose entries are covariances between test points \mathbf{X}^{fine} .

(b) Generative process of fine-grained target data: We model a generative process for the fine-grained target data \mathbf{z} . Let $z(\mathbf{x})$ be a noise-free latent function for the fine-grained target data at location \mathbf{x} . We assume that $z(\mathbf{x})$ follows a Gaussian process, $z(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, with mean function $m(\mathbf{x}) = \sum_{s \in \mathcal{S}} w_s f_s(\mathbf{x}) + w_0$, where $w_s \in \mathbb{R}$ and $w_0 \in \mathbb{R}$ are the regression coefficient of the s th auxiliary data set and the bias parameter, respectively. The covariance function $k(\mathbf{x}, \mathbf{x}')$ is a squared-exponential kernel with the scale parameter γ and signal variance α^2 . Given the predictive values for the auxiliary data sets from (2), the conditional distribution of \mathbf{z} at the centroids \mathbf{X}^{fine} is given by

$$p(\mathbf{z} \mid \mathbf{F}^*) = \mathcal{N}(\mathbf{z} \mid \mathbf{F}^* \mathbf{w}, \mathbf{K}), \quad (3)$$

where $\mathbf{w} = (w_0, \dots, w_{|\mathcal{S}|})^\top$ and \mathbf{K} is a $|\mathcal{P}^{\text{fine}}| \times |\mathcal{P}^{\text{fine}}|$ covariance matrix defined by $k(\mathbf{x}, \mathbf{x}')$. Here, we let $|\mathcal{S}|$ be the number of auxiliary data sets. We define the augmented matrix as the $|\mathcal{P}^{\text{fine}}| \times (|\mathcal{S}| + 1)$ matrix $\mathbf{F}^* = (\mathbf{f}_1^*, \dots, \mathbf{f}_{|\mathcal{S}|}^*, \mathbf{1})$, in which $\mathbf{1}$ is a column vector of 1's. This GP-based modeling enables us to consider the spatial correlation in the target data and the auxiliary data sets simultaneously.

(c) Generative process of coarse-grained target data: We design a spatial aggregation process to transform the fine-grained target data \mathbf{z} into the coarse-grained target data \mathbf{a} , in order to encourage consistency between \mathbf{z} , which is to be estimated, and the available coarse-grained target data \mathbf{a} . In the spatial aggregation process, a value associated with one region in the coarse-grained partition is obtained by aggregating the values in the fine-grained regions contained in the coarse-grained region (see the upper part of Figure 2). Then, \mathbf{a} is generated from the following conditional distribution given \mathbf{z} ,

$$p(\mathbf{a} \mid \mathbf{z}) = \mathcal{N}(\mathbf{a} \mid \mathbf{H}\mathbf{z}, \sigma^2 \mathbf{I}), \quad (4)$$

where σ^2 is the noise variance for the coarse-grained target data, and \mathbf{H} is a $|\mathcal{P}^{\text{coar}}| \times |\mathcal{P}^{\text{fine}}|$ aggregation matrix, whose entries are nonnegative weighting coefficients; the row sum of \mathbf{H} should equal 1. We set the coefficients in accordance with the property of the target data. For example, in cases where target data are incidences of disease, then the (i, j) -entry $H(i, j)$ of \mathbf{H} would be proportional to the population in the intersection of the coarse-grained region i and the fine-grained region j . In the following, for simplicity, we consider a simple aggregation matrix, in which entry $H(i, j)$ is $1/|\mathcal{P}_i^{\text{fine}}|$ if the fine-grained region j is contained in the coarse-grained region i , and zero otherwise. Here, $\mathcal{P}_i^{\text{fine}}$ is a subset of $\mathcal{P}^{\text{fine}}$, all the elements of which are contained in the coarse-grained region $i \in \mathcal{P}^{\text{coar}}$.

5 Inference

Given the coarse-grained target data \mathbf{a} , the auxiliary spatial data sets with centroids $\{(\mathbf{X}_s, \mathbf{y}_s) \mid s \in \mathcal{S}\}$, the centroids

Algorithm 1: Bayesian inference procedure of the fine-grained target data \mathbf{z}

Input : $\mathbf{a}, \{(\mathbf{X}_s, \mathbf{y}_s) \mid s \in \mathcal{S}\}, \mathbf{X}^{\text{fine}}, \mathbf{H}$

Output: Predictive distribution of \mathbf{z}

- 1: Initialize model parameters, $\{\alpha_s \mid s \in \mathcal{S}\}, \{\gamma_s \mid s \in \mathcal{S}\}, \{\sigma_s \mid s \in \mathcal{S}\}, \mathbf{w}, \alpha, \gamma, \sigma$
 - 2: /* first inference step */
 - 3: **for** $s \in \mathcal{S}$ **do**
 - 4: Estimate $\alpha_s, \gamma_s, \sigma_s$ by maximizing the logarithm of (5)
 - 5: **end for**
 - 6: /* second inference step */
 - 7: Estimate $\mathbf{w}, \alpha, \gamma, \sigma$ by maximizing the logarithm of (6)
 - 8: Construct predictive distribution of \mathbf{z} by (8) using the estimated model parameters
-

of fine-grained partition \mathbf{X}^{fine} and the aggregation matrix \mathbf{H} , we aim to predict the fine-grained target data \mathbf{z} via a Bayesian inference procedure. In order to calculate the predictive distribution of \mathbf{z} , we need to estimate the model parameters. The problem of estimating the model parameters can be divided into two steps: 1) Estimate hyperparameters $\alpha_s, \gamma_s, \sigma_s$ for each auxiliary data set and 2) estimate regression coefficient \mathbf{w} and hyperparameters α, γ, σ for the target data. Although one could also opt for estimating all the model parameters simultaneously (i.e., one-step inference), it will increase the computational cost of inference drastically; we adopt the efficient two-step inference as described in the following paragraphs. We finally construct the predictive distribution of \mathbf{z} by using the estimated parameters. Details of the inference procedure are shown in Algorithm 1.

The first inference step: Given the s th auxiliary spatial data set with centroids $(\mathbf{X}_s, \mathbf{y}_s)$, the marginal likelihood of \mathbf{y}_s is given by

$$p(\mathbf{y}_s \mid \alpha_s, \gamma_s, \sigma_s) = \mathcal{N}(\mathbf{y}_s \mid 0, \mathbf{K}_s + \sigma_s^2 \mathbf{I}). \quad (5)$$

The hyperparameters $\alpha_s, \gamma_s, \sigma_s$ are estimated by maximizing the logarithm of (5). We solve the optimization problem through the use of the BFGS method (Liu and Nocedal 1989). By solving the optimization problem for each auxiliary data set independently, we obtain the set of the estimated hyperparameters for all auxiliary data sets. The predictive distribution of \mathbf{f}_s^* corresponding to (2) is obtained using the estimated hyperparameters.

The second inference step: Given the coarse-grained target data \mathbf{a} and the centroids of fine-grained partition \mathbf{X}^{fine} , the marginal likelihood of \mathbf{a} is given by

$$\begin{aligned} p(\mathbf{a} \mid \mathbf{w}, \alpha, \gamma, \sigma) &= \int \int p(\mathbf{a} \mid \mathbf{z}) p(\mathbf{z} \mid \mathbf{F}^*) \prod_{s \in \mathcal{S}} p(\mathbf{f}_s^*) d\mathbf{F}^* d\mathbf{z} \\ &= \int \int \mathcal{N}(\mathbf{a} \mid \mathbf{H}\mathbf{z}, \sigma^2 \mathbf{I}) \mathcal{N}(\mathbf{z} \mid \mathbf{F}^* \mathbf{w}, \mathbf{K}) \\ &\quad \times \prod_{s \in \mathcal{S}} \mathcal{N}(\mathbf{f}_s^* \mid \bar{\mathbf{f}}_s^*, \Sigma_s^*) d\mathbf{F}^* d\mathbf{z} \\ &= \mathcal{N}(\mathbf{a} \mid \mathbf{H}\bar{\mathbf{F}}^* \mathbf{w}, \mathbf{\Lambda}), \end{aligned} \quad (6)$$

where $\bar{\mathbf{F}}^* = (\bar{\mathbf{f}}_1^*, \dots, \bar{\mathbf{f}}_{|\mathcal{S}|}^*, \mathbf{1})$ is a $|\mathcal{P}^{\text{fine}}| \times (|\mathcal{S}| + 1)$ matrix, and we analytically integrate out the latent variables \mathbf{F}^* and \mathbf{z} with the help of the conjugacy of the distributions (2), (3), and (4). $\mathbf{\Lambda}$ is a $|\mathcal{P}^{\text{coar}}| \times |\mathcal{P}^{\text{coar}}|$ covariance matrix represented by $\mathbf{\Lambda} = \sigma^2 \mathbf{I} + \mathbf{H}\mathbf{\Omega}\mathbf{H}^\top$, where $\mathbf{\Omega} = \mathbf{K} + \sum_{s \in \mathcal{S}} w_s^2 \Sigma_s^*$. The (i, i') -entry $\Lambda(i, i')$ of $\mathbf{\Lambda}$ is shown in (7). Here, $\delta_{\bullet, \bullet}$ in (7) represents Kronecker delta; $\delta_{A, B} = 1$ if $A = B$, and $\delta_{A, B} = 0$ otherwise. The residual variance term in (7) represents the residual variance in the regression of $z(\mathbf{x}_j)$. This term contains the uncertainty in the prediction of $f_s(\mathbf{x}_j)$, i.e., $\Sigma_s^*(j, j)$, which is weighted by w_s^2 . The spatial correlation term in (7) represents the strength of spatial correlation between $z(\mathbf{x}_j)$ and $z(\mathbf{x}_{j'})$. This term contains the covariance between $f_s(\mathbf{x}_j)$ and $f_s(\mathbf{x}_{j'})$, i.e., $\Sigma_s^*(j, j')$, which is weighted by w_s^2 . On the basis of the marginal likelihood (6) with this covariance matrix $\mathbf{\Lambda}$, our model can effectively learn the regression coefficient \mathbf{w} while taking into consideration the prediction uncertainties and the spatial correlations from the auxiliary data sets with various granularities, simultaneously. The parameter \mathbf{w} and the hyperparameters α, γ, σ are estimated by maximizing the logarithm of (6). We solve the optimization problem by using the BFGS method (Liu and Nocedal 1989). The derivatives of the logarithm of (6) with respect to $w_s, \alpha, \gamma, \sigma$ are described in Appendix A.

Predictive distribution of fine-grained target data: Using the estimated model parameters, the predictive distribution of the fine-grained target data \mathbf{z} is given by

$$p(\mathbf{z}^*) = \mathcal{N}(\mathbf{z}^* \mid \bar{\mathbf{z}}^*, \text{cov}(\mathbf{z}^*)), \quad (8)$$

where $\bar{\mathbf{z}}^* = \bar{\mathbf{F}}^* \mathbf{w} + \mathbf{\Omega} \mathbf{H}^\top \mathbf{\Lambda}^{-1} (\mathbf{a} - \bar{\mathbf{F}}^* \mathbf{w})$ is the predictive means, and where $\text{cov}(\mathbf{z}^*) = \mathbf{\Omega} - \mathbf{\Omega} \mathbf{H}^\top \mathbf{\Lambda}^{-1} \mathbf{H} \mathbf{\Omega}$ is the covariance matrix. We can obtain the refinement results, i.e., the estimated fine-grained target data, by using the predictive means $\bar{\mathbf{z}}^*$. By analyzing the covariance matrix $\text{cov}(\mathbf{z}^*)$, we can also evaluate the confidence of the refinement results.

$$\begin{aligned} \Lambda(i, i') &= \sigma^2 \delta_{i, i'} + \frac{1}{|\mathcal{P}_i^{\text{fine}}| |\mathcal{P}_{i'}^{\text{fine}}|} \\ &\quad \times \sum_{j \in \mathcal{P}_i^{\text{fine}}} \sum_{j' \in \mathcal{P}_{i'}^{\text{fine}}} \left[\underbrace{\delta_{j, j'} \left(\alpha^2 + \sum_{s \in \mathcal{S}} w_s^2 \Sigma_s^*(j, j') \right)}_{\text{residual variance term}} + (1 - \delta_{j, j'}) \underbrace{\left(\alpha^2 \exp \left(-\frac{1}{2\gamma^2} \|\mathbf{x}_j - \mathbf{x}_{j'}\|^2 \right) + \sum_{s \in \mathcal{S}} w_s^2 \Sigma_s^*(j, j') \right)}_{\text{spatial correlation term}} \right] \end{aligned} \quad (7)$$

Table 2: MAPE L and standard errors for the predictions of the fine-grained target data.

	PM2.5	Ozone	Formaldehyde	Benzene	Elemental carbon	Poverty rate
Proposed model	0.038 ± 0.005**	0.030 ± 0.005*	0.078 ± 0.010**	0.138 ± 0.021*	0.100 ± 0.012*	0.202 ± 0.024**
2-stage SD	0.052 ± 0.007	0.035 ± 0.007	0.101 ± 0.013	0.181 ± 0.032	0.123 ± 0.016	0.228 ± 0.028
LR-based method	0.056 ± 0.007	0.040 ± 0.007	0.108 ± 0.013	0.185 ± 0.031	0.123 ± 0.016	0.234 ± 0.028
GPR	0.072 ± 0.010	0.062 ± 0.011	0.191 ± 0.020	0.267 ± 0.029	0.195 ± 0.019	0.344 ± 0.046

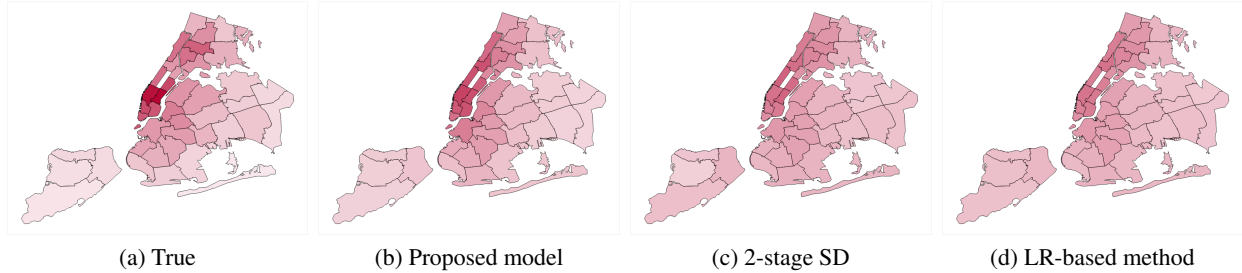


Figure 3: Comparison of the predicted fine-grained target data for PM2.5 data set.

6 Experiments

Data description: We evaluated the proposed model using real-world spatial data sets from NYC Open Data¹. There are 44 data sets that contain a variety of categories such as social indicators, land use, air quality and taxi traffic. Each data set is associated with one of six geographical partitions, i.e., school district (32), UHF42 (42), community district (59), police precinct (77), zip code (186) and taxi zone (249), where each number in parenthesis denotes the number of regions in the corresponding partition. In our experiments, we try to refine the poverty rate data set and the five air pollution data sets (i.e., PM2.5, ozone, formaldehyde, benzene, elemental carbon). The experimental setting is as follows: 1) Given the poverty rate data set with the borough partition ($|\mathcal{P}^{\text{coar}}| = 5$), we would like to refine the data into the community district partition ($|\mathcal{P}^{\text{fine}}| = 59$), and 2) given each air pollution data set with the borough partition ($|\mathcal{P}^{\text{coar}}| = 5$), we aim to refine the data into the UHF42 partition ($|\mathcal{P}^{\text{fine}}| = 42$). Appendix B details the data sets and the experimental settings.

Baselines: The existing methods can be applied to auxiliary data sets with various granularities if pre-processing is applied, i.e., spatial interpolation, so that the granularities of the auxiliary data sets match with that of the fine-grained target data. Accordingly, we first performed spatial interpolation of each auxiliary data set \mathbf{y}_s by using GP regression; we then obtained the predictive values $\bar{\mathbf{f}}_s^*$ at the centroids \mathbf{X}^{fine} of the target fine-grained partition so that the spatial granularities of all auxiliary data sets equaled that of the fine-grained target data. We compared the proposed model with three baselines: GP regression (GPR) (Rasmussen and Williams 2006), Linear regression-based method (LR-based method) (Smith, Mashhadi, and Capra 2014) and Two-stage statistical downscaling method (2-stage SD) (Park 2013). Here, GPR is a simple spatial interpolation, namely, it pre-

Table 3: Top-10 relevant auxiliary data as estimated by our model and 2-stage SD for PM2.5 data set.

	Proposed model		2-stage SD	
	Auxiliary data	w_s	Auxiliary data	w_s
1.	Fire incident (Zip code)	0.173	1-2 fam. bldg (Comm.)	-0.088
2.	Taxi dropoff (Taxi zone)	0.139	Hospital (Comm.)	0.069
3.	311 call (Zip code)	0.135	Public school (Comm.)	0.069
4.	Public telephone (Zip code)	0.114	Lots of vacant (Comm.)	-0.067
5.	Natural gas (Zip code)	0.109	Crime (Police precinct)	0.064
6.	Mean commute (Comm.)	-0.109	Unemployment (Comm.)	0.063
7.	1-2 fam. bldg (Comm.)	-0.089	Pct. served parks (Comm.)	0.062
8.	Pct. served park (Comm.)	0.075	Library (Comm.)	0.061
9.	GHG emission (Zip code)	0.068	Fire incident (Zip code)	0.059
10.	Population (Comm.)	0.062	Park (Comm.)	0.058

dicts the fine-grained target data \mathbf{z} by using only the coarse-grained target data \mathbf{a} . Details of these baselines are given in Appendix C.

Fine-grained target data prediction: We evaluated our model in terms of its performance in predicting fine-grained target data \mathbf{z} . The evaluation metric is the mean absolute percentage error (MAPE) in fine-grained target values: $L = \frac{1}{|\mathcal{P}^{\text{fine}}|} \sum_{j \in \mathcal{P}^{\text{fine}}} \left| \frac{z_j^{\text{true}} - z_j^*}{z_j^{\text{true}}} \right|$, where z_j^{true} is the true value associated with region j in the target fine-grained partition; z_j^* is its predicted value. Table 2 shows the MAPE and the standard error of absolute percentage error for the proposed model, 2-stage SD, LR-based method and GPR. For all data sets, our model performed better than the baselines, and the differences between our model and the baselines are statistically significant (Student's t-test). In Table 2, the single star (*) and the double star (**) indicate significant difference at the levels of $P < 0.05$ and $P < 0.01$, respectively. We found similar results using other evaluation metrics (e.g., MAE, RMSE, RMSPE). These results show that our model well utilized the auxiliary data sets with various granularities to accurately predict the fine-grained target data.

¹<https://opendata.cityofnewyork.us>

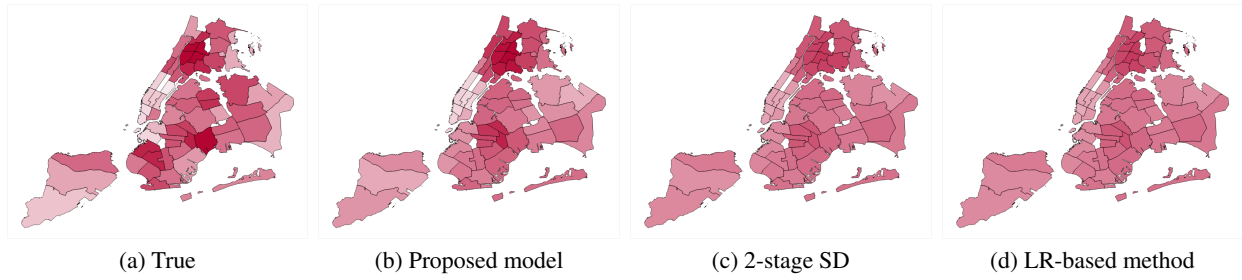


Figure 4: Comparison of the predicted fine-grained target data for poverty rate data set.



Figure 5: Top-2 auxiliary data sets ranked by the proposed model for PM2.5 data set.

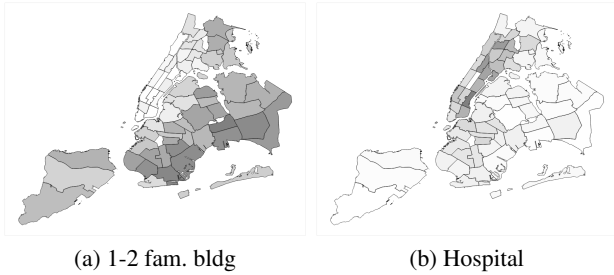


Figure 6: Top-2 auxiliary data sets ranked by the 2-stage SD for PM2.5 data set.

Figures 3 and 4 visualize the predicted fine-grained target data z for the PM2.5 data set and for the poverty rate data set, respectively. We illustrate the true fine-grained data on the left in Figures 3 and 4, and the predictions made by the proposed model, 2-stage SD and LR-based method on the right. Here, the predictive values of each method were normalized to the range $[0, 1]$, and darker hues represent regions with higher values. As shown in these figures, our model refined the coarse-grained data more precisely than the other methods. In particular, in both data sets, our model achieved significant improvement in the north part of the map (i.e., Manhattan). Such visualization results are useful for finding key regions, e.g., the poorest regions of a city.

Evaluation of auxiliary spatial data sets: Table 3 shows the top ten relevant auxiliary data sets as determined by our model and 2-stage SD for the PM2.5 data set. These auxiliary data sets are arranged in descending order of the absolute values of the estimated regression coefficient w , each of which is listed in the “ w_s ” columns of Table 3. By com-

paring the sorted list of the auxiliary data sets created by the proposed model with that yielded by 2-stage SD, we can confirm that the proposed model assigned relatively large regression coefficients to the auxiliary data sets with finer-grained partitions (i.e., Zip code and Taxi zone).

Figures 5 and 6 visualize the top two relevant auxiliary data sets as estimated by our model and 2-stage SD for the PM2.5 data set, respectively. Comparing these visualizations with that of the true target data in Figure 3(a) shows that our model emphasized the most useful auxiliary data sets, i.e., those that are both strongly related with the target data and have fine granularities; 2-stage SD evaluated the usefulness of auxiliary data sets only in terms of the strength of relationships between the target data and the auxiliary data sets in the coarse-grained partition.

Figure 7 shows the relation between the regression coefficient and the uncertainty in the prediction of auxiliary data sets estimated by the proposed model for the PM2.5 data set. In this figure, each auxiliary data set is depicted by a dot whose color indicates its partition. The horizontal axis shows the averages of the variances in the predicted values of each auxiliary data set; for the s th auxiliary data set, the average of variances was calculated by $(1/|\mathcal{P}^{\text{fine}}|) \sum_{j \in \mathcal{P}^{\text{fine}}} \Sigma_s^*(j, j)$, which is the degree of uncertainty in predicting the s th auxiliary data set; the vertical axis shows the absolute values of the estimated coefficients. As shown, the absolute coefficient values estimated by our model were likely to be higher for the auxiliary data sets that had lower degrees of uncertainty. These results indicate that our model can effectively learn the usefulness of each auxiliary data set by considering the uncertainty in the prediction of auxiliary data sets. Consequently, the proposed model can precisely refine the coarse-grained target data by effectively utilizing auxiliary data sets with various granularities.

7 Conclusion

This paper has proposed a probabilistic model for refining coarse-grained spatial data by utilizing auxiliary spatial data sets with various granularities on the same region. Our model can effectively make use of auxiliary data sets with various granularities by hierarchically incorporating Gaussian processes. Our model also has the advantage of allowing the inference of model parameters based on the exact marginal likelihood, in which the variables of fine-grained target and auxiliary data are analytically integrated out. Us-




Figure 7: Relation between the coefficients and the uncertainties for PM2.5 data set.

ing multiple real-world spatial data sets in New York City, we confirmed that our model can predict the fine-grained target data more precisely compared with the baselines.

Our future work is to consider shapes of regions as in the previous study (Rathbun 1998): The assumption of using the centroid of each region allows for GP-based formulations and significantly simplifying computations involved; meanwhile, it might worsen the fit of the GP to the exotic shaped regions (e.g., extremely elongated). Another future work is to incorporate fully Bayesian treatment for model parameters. It can be expected to provide the better results.

A Derivatives of model parameters

The log-marginal likelihood of \mathbf{a} is given by

$$\log p(\mathbf{a} | \mathbf{w}, \alpha, \gamma, \sigma) = -\frac{1}{2}(\mathbf{a} - \mathbf{H}\bar{\mathbf{F}}^*\mathbf{w})^\top \mathbf{\Lambda}^{-1}(\mathbf{a} - \mathbf{H}\bar{\mathbf{F}}^*\mathbf{w}) - \frac{1}{2} \log(\det(\mathbf{\Lambda})) - \frac{|\mathcal{P}^{\text{coar}}|}{2} \log 2\pi. \quad (9)$$

We describe the first derivatives of (9) with respect to w_s , α , γ , σ , which is required for estimating the parameter based on the BFGS method. The derivative of (9) with respect to w_s is given by

$$\frac{\partial}{\partial w_s} \log p(\mathbf{a} | \mathbf{w}, \alpha, \gamma, \sigma) = \frac{\partial(\mathbf{a} - \mathbf{H}\bar{\mathbf{F}}^*\mathbf{w})}{\partial w_s} \mathbf{p} + \frac{1}{2} \text{tr} \left((\mathbf{p}\mathbf{p}^\top - \mathbf{\Lambda}^{-1}) \frac{\partial \mathbf{\Lambda}}{\partial w_s} \right), \quad (10)$$

where $\mathbf{p} = \mathbf{\Lambda}^{-1}(\mathbf{a} - \mathbf{H}\bar{\mathbf{F}}^*\mathbf{w})$ and $\partial \mathbf{\Lambda} / \partial w_s$ is a matrix of elementwise derivatives. The derivative of the element $\Lambda(i, i')$ (7) is obtained by

$$\frac{\partial \Lambda(i, i')}{\partial w_s} = \frac{1}{|\mathcal{P}_i^{\text{fine}}| |\mathcal{P}_{i'}^{\text{fine}}|} \sum_{j \in \mathcal{P}_i^{\text{fine}}} \sum_{j' \in \mathcal{P}_{i'}^{\text{fine}}} 2w_s \Sigma_s^*(j, j'). \quad (11)$$

Denoting $\theta \in \{\alpha, \gamma, \sigma\}$, the derivative of (9) with respect to θ is given by

$$\frac{\partial}{\partial \theta} \log p(\mathbf{a} | \mathbf{w}, \alpha, \gamma, \sigma) = \frac{1}{2} \text{tr} \left((\mathbf{p}\mathbf{p}^\top - \mathbf{\Lambda}^{-1}) \frac{\partial \mathbf{\Lambda}}{\partial \theta} \right). \quad (12)$$

The matrix of elementwise derivatives $\partial \mathbf{\Lambda} / \partial \theta$ is trivial. The derivative of the element $\Lambda(i, i')$ (7) with respect to each hyperparameter is as follows:

$$\begin{aligned} \frac{\partial \Lambda(i, i')}{\partial \alpha} &= \frac{1}{|\mathcal{P}_i^{\text{fine}}| |\mathcal{P}_{i'}^{\text{fine}}|} \sum_{j \in \mathcal{P}_i^{\text{fine}}} \sum_{j' \in \mathcal{P}_{i'}^{\text{fine}}} 2\alpha \exp \left(-\frac{1}{2\gamma^2} \|\mathbf{x}_j - \mathbf{x}_{j'}\|^2 \right), \end{aligned} \quad (13)$$

$$\begin{aligned} \frac{\partial \Lambda(i, i')}{\partial \gamma} &= \frac{1}{|\mathcal{P}_i^{\text{fine}}| |\mathcal{P}_{i'}^{\text{fine}}|} \sum_{j \in \mathcal{P}_i^{\text{fine}}} \sum_{j' \in \mathcal{P}_{i'}^{\text{fine}}} \alpha^2 \left(\frac{1}{\gamma^3} \|\mathbf{x}_j - \mathbf{x}_{j'}\|^2 \right) \\ &\quad \times \exp \left(-\frac{1}{2\gamma^2} \|\mathbf{x}_j - \mathbf{x}_{j'}\|^2 \right), \end{aligned} \quad (14)$$

$$\frac{\partial \Lambda(i, i')}{\partial \sigma} = 2\sigma \delta_{i, i'}. \quad (15)$$

B Description of real-world spatial data sets

We used the real-world spatial data sets from NYC Open Data ². for evaluating the proposed model. The data sets were collected and released for improving the urban environment in New York City, and contain a variety of categories such as social indicators, land use, air quality and taxi traffic. Details of the data sets are listed in Table 4. There are multiple data sets in each category, with the total number of data sets being 44. Each data set is associated with one of six geographical partitions, i.e., school district, UHF42, community district, police precinct, zip code and taxi zone. These partitions have various spatial granularities; the number of regions in each partition is shown in Table 4. These data sets are gathered once a year using the time ranges shown in Table 4; the values of data are divided by the number of observation times. When the values of data are extensive quantities (i.e., proportional to the scale of areas, e.g., population), the values are divided by the areas of respective regions; the resulting values are intensive quantities (i.e., independent of area scale, e.g., population density).

In our experiments, we try to refine the poverty rate data set in the social indicator category and the five air pollution data sets in the air quality category. The poverty rate data set contains the values of poverty rates associated with each region in the community district partition as visualized in Figure 1(a). The air pollution data sets contain the average concentrations of pollutants (i.e., PM2.5, ozone, formaldehyde, benzene, elemental carbon) associated with each region in the UHF42 partition. In order to evaluate the performance in refining coarse-grained data, we used the data that were aggregated into a coarser-grained partition, i.e., borough partition, via spatial averaging, where the borough partition has five regions as illustrated in Figure 1(b). The experimental setting is as follows: 1) Given the poverty rate data set with borough partition ($|\mathcal{P}^{\text{coar}}| = 5$), we would like to refine the data into the community district partition ($|\mathcal{P}^{\text{fine}}| = 59$), and 2) given each air pollution data set with the borough partition ($|\mathcal{P}^{\text{coar}}| = 5$), we aim to refine the data into the UHF42

²<https://opendata.cityofnewyork.us>

Table 4: Spatial data sets.

Category/Name	#data sets	Partition	#regions	Time range	Description
Education	3	School district	32	2010	Class size, ratio of #pupils to #teachers, SAT score
Air quality	8	UHF42	42	2009–2010	Average concentration of pollutants
Social indicator	13	Community district	59	2009–2013	Poverty rate, population, mean commute time, etc.
Land use	11	Community district	59	2009–2013	Area percentage for commercial office, parking, etc.
Crime	1	Police precinct	77	2010–2016	Number of crimes
Incident	2	Zip code	186	2010–2016	#311 calls, #fire incidents
Telecommunication	2	Zip code	186	2016	#public telephones, #free Wi-Fi hotspots
Consumption	2	Zip code	186	2010–2014	GHG emission, natural gas consumption
Taxi traffic	2	Taxi zone	249	2014–2016	#taxi pick-up and drop-off events

partition ($|\mathcal{P}^{\text{fine}}| = 42$). In the setting for the poverty rate data set, we used all data sets other than the target data as auxiliary data sets, so the number of auxiliary data sets $|\mathcal{S}|$ was 43. In the setting for the air pollution data sets, we used all data sets not contained in the air quality category, so $|\mathcal{S}|$ was 36.

C Baselines description

For GPR, we predict the fine-grained target data \mathbf{z} based only on the coarse-grained target data \mathbf{a} . For LR-based method and 2-stage SD, given the coarse-grained target data \mathbf{a} and the predictive values of all auxiliary data sets \mathbf{F}^* , we predict the fine-grained target data \mathbf{z} . Details of these baselines are given below.

Gaussian process regression (GPR): We compared our proposed model with a simple spatial interpolation (i.e., GPR) of the coarse-grained spatial data \mathbf{a} . This baseline assumes that the target data are explained by only the spatial correlation. Given \mathbf{a} and the set of centroids of the coarse-grained partition $\mathcal{P}^{\text{coar}}$, we predicted the fine-grained target data \mathbf{z} by using the predictive distribution. Note that this baseline does not use the auxiliary spatial data sets.

Linear regression-based method (LR-based method): We used a linear regression-based method that has been applied in various studies (Bogomolov et al. 2014; Smith, Mashhadi, and Capra 2014). The linear regression model is used for estimating the relationships between the coarse-grained target data and the auxiliary data sets. The procedure in the training phase is as follows: 1) aggregate all auxiliary data sets into the coarse-grained partition of target data via spatial averaging; 2) estimate the regression coefficients \mathbf{w} of the respective auxiliary data sets by using the coarse-grained target data and the auxiliary data sets aggregated via spatial averaging. In the prediction phase, generate unknown values \mathbf{z} for the target fine-grained partition by applying the estimated relationships to the predictive values of auxiliary data sets \mathbf{F}^* as follows: $\mathbf{z} = \mathbf{F}^* \hat{\mathbf{w}}$, where $\hat{\mathbf{w}}$ is the estimated regression coefficient.

Two-stage statistical downscaling method (2-stage SD): We used the statistical downscaling method proposed in (Park 2013). This method assumes that coarse-grained target data \mathbf{a} can be decomposed into linear regression terms and residual terms. The downscaling procedure is divided into two stages. In the first stage, we obtain the regression coefficients \mathbf{w} in a manner similar to the training phase of

the LR-based method. In the second stage, given the estimated coefficient $\hat{\mathbf{w}}$, the fine-grained target data \mathbf{z} are estimated to be those that satisfy the following relation:

$$\begin{aligned}
 a_i &= \hat{w}_0 + \underbrace{\sum_{s \in \mathcal{S}} \hat{w}_s \left[\frac{1}{|\mathcal{P}_i^{\text{fine}}|} \sum_{j \in \mathcal{P}_i^{\text{fine}}} f_s(\mathbf{x}_j) \right]}_{\text{linear regression term}} + \underbrace{R_i^{\text{coar}}}_{\text{residual term}} \\
 &= \frac{1}{|\mathcal{P}_i^{\text{fine}}|} \sum_{j \in \mathcal{P}_i^{\text{fine}}} \left[\hat{w}_0 + \sum_{s \in \mathcal{S}} \hat{w}_s f_s(\mathbf{x}_j) + R_j^{\text{fine}} \right] \\
 &= \frac{1}{|\mathcal{P}_i^{\text{fine}}|} \sum_{j \in \mathcal{P}_i^{\text{fine}}} z_j.
 \end{aligned} \tag{16}$$

This relation expresses the spatial aggregation constraint, i.e., the assumption that value a_i associated with coarse-grained region i is the linear average of the constituent values in the fine-grained partition. Here, R_i^{coar} and R_j^{fine} are the residuals in the coarse-grained and fine-grained partitions, respectively. To obtain the fine-grained target data \mathbf{z} , the residual value R_j^{fine} in the fine-grained partition must be determined. Since the linear regression terms have already been fixed in the first stage, R_i^{coar} is obtained from (16); the residuals in the fine-grained partition are predicted by applying the spatial interpolation method, i.e., simple kriging (Kyriakidis 2004), to the residuals R_i^{coar} in the coarse-grained partition.

References

- Barlacchi, G.; Nadai, M. D.; Larcher, R.; Casella, A.; Chitic, C.; and G. Torrisi et al., j. 2015. A multi-source dataset of urban life in the city of Milan and the province of Trentino. *Scientific Data* 2.
- Bogomolov, A.; Lepri, B.; Staiano, J.; Oliver, N.; Pianesi, F.; and Pentland, A. 2014. Once upon a crime: Towards crime prediction from demographics and mobile data. In *ICMI*, 427–434.
- Boucher, A., and Kyriakidis, P. C. 2006. Super-resolution land cover mapping with indicator geostatistics. *Remote Sensing of Environment* 104:264–282.
- Cannon, A. J. 2011. Quantile regression neural networks: Implementation in R and application to precipitation downscaling. *Computers & Geosciences* 37(9):1277–1284.
- Diodato, N.; Bellocchi, G.; Bertolin, C.; and Camuffo, D. 2010. Multiscale regression model to infer historical temperatures in a central mediterranean sub-regional area. *Climate of the Past Discussions* 6:2625–2649.

- Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2014. Learning a deep convolutional network for image super-resolution. In *ECCV*, 184–199.
- Flaxman, S. R.; Wang, Y. X.; and Smola, A. J. 2015. Who supported Obama in 2012?: Ecological inference through distribution regression. In *KDD*, 289–298.
- Ghosh, S. 2010. SVM-PGSL coupled approach for statistical downscaling to predict rainfall from GCM output. *Journal of Geophysical Research: Atmospheres* 115(D22).
- Goldstein, B., and Dyson, L. 2013. *Beyond Transparency: Open Data and the Future of Civic Innovation*. Code for America Press.
- Goovaerts, P. 2010. Combining areal and point data in geostatistical interpolation: Applications to soil science and medical geography. *Mathematical Geosciences* 42(5):535–554.
- Hessami, M.; Gachon, P.; Ouarda, T. B.; and St-Hilaire, A. 2008. Automated regression-based statistical downscaling tool. *Environmental Modeling & Software* 23(6):813–834.
- Howitt, R., and Reynaud, A. 2003. Spatial disaggregation of agricultural production data using maximum entropy. *European Review of Agricultural Economics* 30(2):359–387.
- Jerrett, M.; Burnett, R. T.; Beckerman, B. S.; Turner, M. C.; Krewski, D.; Thurston, G.; and et al. 2013. Spatial analysis of air pollution and mortality in California. *American Journal of Respiratory and Critical Care Medicine* 188(5):593–599.
- Keil, P.; Belmaker, J.; Wilson, A. M.; Unitt, P.; and Jetz, W. 2013. Downscaling of species distribution models: a hierarchical approach. *Methods in Ecology and Evolution* 4(1):82–94.
- Kyriakidis, P. C. 2004. A geostatistical framework for area-to-point spatial interpolation. *Geographical Analysis* 36(3):259–289.
- Law, H. C. L.; Sejdinovic, D.; Cameron, E.; Lucas, T. C. D.; Flaxman, S.; Battle, K.; and Fukumizu, K. 2018. Variational learning on aggregate outputs with Gaussian processes. In *NeurIPS*, 6084–6094.
- Liu, D. C., and Nocedal, J. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical programming* 45(1–3):503–528.
- Miller, B. A.; Koszinski, S.; Wehrhan, M.; and Sommer, M. 2015. Impact of multi-scale predictor selection for modeling soil properties. *Geoderma* 239–240:97–106.
- Misra, S.; Sarkar, S.; and Mitra, P. 2017. Statistical downscaling of precipitation using long short-term memory recurrent neural networks. *Theor. Appl. Climatol.*
- Murakami, D., and Tsutsumi, M. 2011. A new areal interpolation technique based on spatial econometrics. *Procedia-Social and Behavioral Sciences* 21:230–239.
- Park, N. W. 2013. Spatial downscaling of TRMM precipitation using geostatistics and fine scale environmental variables. *Advances in Meteorology* 2013.
- Rasmussen, C. E., and Williams, C. K. I. 2006. *Gaussian Processes for Machine Learning*. MIT Press.
- Rathbun, S. L. 1998. Spatial modelling in irregularly shaped regions: Kriging estuaries. *Environmetrics* 9:109–129.
- Rupasinghaa, A., and Goetz, S. J. 2007. Social and political forces as determinants of poverty: A spatial analysis. *The Journal of Socio-Economics* 36(4):650–671.
- Shadbolt, N.; O'Hara, K.; Berners-Lee, T.; Gibbins, N.; Glaser, H.; Wendy, H.; and Schraefel, M. C. 2012. Linked open government data: Lessons from data.gov.uk. *IEEE Intelligent Systems* 27(3):16–24.
- Smith, C. C., and Capra, L. 2016. Beyond the baseline: Establishing the value in mobile phone based poverty estimates. In *WWW*, 425–434.
- Smith, C. C.; Mashhadi, A.; and Capra, L. 2014. Poverty on the cheap: Estimating poverty maps using aggregated mobile communication networks. In *CHI*, 511–520.
- Stein, M. L. 1999. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer.
- Sturrock, H. J. W.; Cohen, J. M.; Keil, P.; Tatem, A. J.; Menach, A. L.; Ntshalintshali, N. E.; Hsiang, M. S.; and Gosling, R. D. 2014. Fine-scale malaria risk mapping from routine aggregated case data. *Malaria Journal* 13:421.
- Taylor, B. M.; Andrade-Pacheco, R.; and Sturrock, H. J. W. 2018. Continuous inference for aggregated point process data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 12347.
- Vandal, T.; Kodra, E.; Ganguly, S.; Michaelis, A.; Nemani, R.; and Ganguly, A. R. 2017. DeepSD: Generating high resolution climate change projections through single image super-resolution. In *KDD*, 1663–1672.
- Vandal, T.; Kodra, E.; Ganguly, S.; Michaelis, A.; Nemani, R.; and Ganguly, A. R. 2018. Generating high resolution climate change projections through single image super-resolution: An abridged version. In *IJCAI*, 5389–5393.
- Wang, H.; Kifer, D.; Graif, C.; and Li, Z. 2016. Crime rate inference with big data. In *KDD*, 635–644.
- Wilby, R. L.; Zorita, S. P.; Timbal, E.; Whetton, B.; and Mearns, L. O. 2004. *Guidelines for Use of Climate Scenarios Developed from Statistical Downscaling Methods*.
- Wilson, K., and Wakefield, J. 2018. Pointless spatial modeling. *Biostatistics*.
- Wotling, G.; Bouvier, C.; Danloux, J.; and Fritsch, J. M. 2000. Regionalization of extreme precipitation distribution using the principal components of the topographical environment. *Journal of Hydrology* 233(1–4):86–101.
- Xavier, A.; Freitas, M. B. C.; Rosário, M. D. S.; and Fragoso, R. 2016. Disaggregating statistical data at the field level: An entropy approach. *Spatial Statistics* 23:91–103.
- Xu, J.; Liu, X.; Wilson, T.; Tan, P. N.; Hatami, P.; and Luo, L. 2018. Muscat: Multi-scale spatio-temporal learning with application to climate modeling. In *IJCAI*, 2912–2918.
- Xu, J. 2017. Multi-task learning and its application to geospatio-temporal data. *ProQuest Dissertations Publishing*.
- Yuan, J.; Zheng, Y.; and Xie, X. 2012. Discovering regions of different functions in a city using human mobility and pois. In *KDD*, 186–194.
- Zheng, Y.; Yi, X.; Li, M.; Li, R.; Shan, Z.; Chang, E.; and Li, T. 2015. Forecasting fine-grained air quality based on big data. In *KDD*, 2267–2276.
- Zheng, Y.; Liu, F.; and Hsieh, H. P. 2013. U-air: When urban air quality inference meets big data. In *KDD*, 1436–1444.
- Zorita, E., and von Storch, H. 1999. The analog method as a simple statistical downscaling technique: Comparison with more complicated methods. *Journal of Climate* 12:2474–2489.