



Nonrigid Registration Using Gaussian Processes and Local Likelihood Estimation

Ashton Wiens¹ · William Kleiber¹ ·
Douglas Nychka² · Katherine R. Barnhart^{3,4,5}

Received: 11 June 2020 / Accepted: 26 December 2020
© International Association for Mathematical Geosciences 2021

Abstract Surface registration, the task of aligning several multidimensional point sets, is a necessary task in many scientific fields. In this work, a novel statistical approach is developed to solve the problem of nonrigid registration. While the application of an affine transformation results in rigid registration, using a general nonlinear function to achieve nonrigid registration is necessary when the point sets require deformations that change over space. The use of a local likelihood-based approach using windowed Gaussian processes provides a flexible way to accurately estimate the nonrigid deformation. This strategy also makes registration of massive data sets feasible by splitting the data into many subsets. The estimation results yield spatially-varying local rigid registration parameters. Gaussian process surface models are then fit to the parameter fields, allowing prediction of the transformation parameters at unestimated locations, specifically at observation locations in the unregistered data set. Applying these transformations results in a global, nonrigid registration. A penalty on the transformation parameters is included in the likelihood objective function. Combined with smoothing of the local estimates from the surface models, the nonrigid registration model can prevent the problem of overfitting. The efficacy of the nonrigid registration method is tested in two simulation studies, varying the number of windows and

✉ Ashton Wiens
ashton.wiens@colorado.edu

¹ Department of Applied Mathematics, University of Colorado, Boulder, CO, USA

² Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, CO, USA

³ Cooperative Institute for Research in Environmental Sciences, University of Colorado, Boulder, CO, USA

⁴ Department of Geological Sciences, University of Colorado, Boulder, CO, USA

⁵ US Geological Survey, Landslide Hazards Program, 1711 Illinois St., Golden, CO 80401, USA

number of points, as well as the type of deformation. The nonrigid method is applied to a pair of massive remote sensing elevation data sets exhibiting complex geological terrain, with improved accuracy and uncertainty quantification in a cross validation study versus two rigid registration methods.

Keywords Point cloud registration · Point cloud · Structure from motion · Digital elevation model · Geomorphology · Uncertainty quantification

1 Introduction

Aligning or registering multidimensional point sets is a common problem arising in remote sensing, image and video processing, computer graphics, and computer vision. The goal in registration is to estimate a transformation which best aligns the moving data set with the target data set based on a measure of closeness. Rigid registration refers to the estimation and application of a single affine transformation to the coordinate system of the moving data set, whereas nonrigid registration involves a transformation function more general than an affine transformation.

Classical approaches to registration rely on point or surface matching, or some variant thereof. Iterative Closest Point (ICP) (Arthur and Vassilvitskii 2009; Besl and McKay 1992; Ezra et al. 2008) and kernel correspondence methods (Tsin and Kanade 2004) are common and have been developed in both the rigid and nonrigid settings. These algorithms iteratively assign correspondences between points (or surfaces) and then estimate a transformation which minimizes some distance function between the corresponding objects. Representing the point sets as mixtures of Gaussian densities is another strategy for solving the registration problem which reduces the computational complexity of the problem by avoiding discrete point matching, for example in Jian and Vemuri (2005) and Zhou et al. (2018). The review articles by Maiseli et al. (2017) and Tam et al. (2013) detail many popular registration methods. Tam et al. (2013) discriminate among the algorithms based on their optimization constraints and also how they assign correspondences; the majority of existing registration algorithms rely on a closest point criterion in point matching.

Thin plate splines (TPS) and Gaussian processes (GP) play a prominent role in the parameterization of nonrigid deformations (Zhu et al. 2009). Arad and Reisfeld (1995) discuss two-dimensional radial basis function transformations, composed of an affine part and a radial function part. Bookstein (1989) provides a similar mathematical development, however in the context of the physics of TPS as solutions to the biharmonic equation. Gilleland et al. (2010) combine these ideas with climatological forecast verification, and Chui and Rangarajan (2003) combine the TPS parameterization with the robust point matching framework of Gold et al. (1998). Nonrigid registration methods are often developed with an application to medical imaging, and most do not provide uncertainty quantification. Risholm et al. (2010) propose a Bayesian framework, producing a posterior distribution of the deformation parameters. This is an attractive feature, but the computational burden of a Bayesian analysis of a massive sensing data set like the one analyzed in this work would be restrictive. Furthermore, many algorithms are pixel-based and would need to be adapted to the irregular coordinates

in the remote sensing application of this paper or else the irregularly-spaced data summarized to an image. The authors are not aware of nonrigid registration algorithms providing uncertainty quantification capable of analyzing massive data sets.

Motivated by the task of jointly registering and producing a surface estimate from two geologic data sets, Wiens et al. (2020a) adapt methods from spatial statistics to produce a novel registration method based on maximum likelihood estimation. The algorithm embeds a rigid transformation within the GP regression maximum likelihood framework. The statistical nature of this method has the advantage of providing uncertainty measures about the estimated registration parameters. Although this method is not as fast as some traditional methods, results indicate this method can potentially provide a more accurate registration with uncertainty quantification, a feature not available in more algorithmic approaches.

In this work, a novel nonrigid registration algorithm is developed based on a local likelihood estimation approach, expanding on the rigid registration model from Wiens et al. (2020a). The rigid registration likelihood procedure is applied using local windows of data, which gives insight into how these parameters change over the domain. Moreover, a surface model is fit to the estimated parameters, which allows the non-rigid transformation to be predicted at arbitrary locations. What is new in this work is the introduction of a global nonrigid deformation alignment model paired with an estimation approach that is scalable to massive data sets.

Both the rigid and nonrigid likelihood-based algorithms are fundamentally different from point matching and correspondence algorithms, as the points from both data sets are treated jointly as a realization from a stochastic process, so there is no point matching involved. On the other hand, the local likelihood estimation framework does share characteristics with some existing nonrigid registration methods. The methods that can handle data sets containing millions of points, including the nonrigid method developed here, often rely on downsampling the point clouds to reduce the computational burden (Huang et al. 2008).

Practical registration methods must be capable of handling noise and outliers in the data sets. In point matching algorithms, it is common to prune correspondences which seem inconsistent. The likelihood-based methods are able to account for noise in the response variable within the statistical model, provided the parameters are correctly specified. Note that spatial models have been developed with the ability to account for location error (Cervone and Pillai 2015; Cressie and Kornak 2003; Fanshawe and Diggle 2011; Fronterre et al. 2018); however, the errors in location are assumed to follow a distribution to be inferred, rather than viewing the problem as an inconsistency in the coordinate systems of the data sets. Finally, it is common to minimize a loss or energy function as the criterion used to estimate the deformation. In the method developed here, transformation parameters are estimated by minimization of a penalized likelihood function. The likelihood function follows from the standard GP framework, and the inclusion of penalties on the transformation parameters can prevent the intrinsic issue of overfitting when using nonrigid registration methods.

Local likelihood estimation for spatial data is not a new idea (Haas 1990a; Kuusela and Stein 2018; Nychka et al. 2018; Wiens et al. 2020b). However, its application to nonrigid registration provides a completely new approach to this problem. Using local likelihood estimation for nonrigid registration has several advantages. It is amenable to

the task of nonrigid registration because simpler local rigid registrations are estimated and pieced together. Furthermore, it gives flexibility in the estimation of the nonrigid transformation, being data driven and not constrained to the a priori choice of a nonlinear deformation function. Finally, local likelihood estimation also permits the use of parallel computing, significantly accelerating the computation of the registration. Parallelism, combined with subsampling, can allow massive data sets to be accurately registered at tractable computational cost.

In this work, the proposed nonrigid registration method is applied to a geomorphic data set collected via remote sensing. The pair of elevation point sets could be processed into a digital elevation model (DEM), but first require registration. The data were collected by piecing together several hundred photographs taken with a drone using a Structure from Motion (SfM) algorithm (Fonstad et al. 2013; James and Robson 2012; Westoby et al. 2012). With this form of processing, it is hypothesized that a nonrigid deformation may be necessary to align the geologic features seen in the data set.

In Sect. 2, the statistical model for nonrigid registration is presented. In Sect. 3, the estimation algorithm is outlined in the rigid and nonrigid cases, providing implementation guidelines. In Sect. 4, the efficacy of the nonrigid registration method is investigated in two simulation studies, testing how the type and severity of deformation affect estimation accuracy. How accuracy is affected by the number of windows used and the number of subsamples taken in each data window is also investigated. In Sect. 5, the nonrigid registration method is applied to a set of remote sensing scans with over sixty million points each, comparing to two rigid registration methods in a cross validation study.

2 Statistical Model

The nonrigid approach developed in this work is a generalization of the rigid registration approach of Wiens et al. (2020a) to a nonlinear setting. To give visual context to the registration problem, see the top left panel of Fig. 3 which illustrates two misaligned spatial data sets; one of these data sets is shown in its entirety in Fig. 2. In the following, the general registration model which encompasses both the rigid and nonrigid registration cases is presented.

Let $Y_1(\mathbf{s})$ be the data from the fixed point cloud and $Y_2(\mathbf{s})$ be the moving point cloud to be registered, where $\mathbf{s} \in \mathbb{R}^2$ contains the x - and y -positional coordinates and $Y_i(\cdot)$ denotes the measured z -coordinate. The statistical process model for these data is

$$\begin{aligned} Y_1(\mathbf{s}) &= Z(\mathbf{s}) + \varepsilon_1(\mathbf{s}) \\ Y_2(\mathbf{s}) &= \mu_z(\mathbf{s}) + Z(T(\mathbf{s})) + \varepsilon_2(T(\mathbf{s})), \end{aligned} \quad (1)$$

where T is a transformation applied to the moving data set. Both datasets represent the same underlying continuous surface, $Z(\cdot)$, which is modeled as a mean zero GP. Specifically, a process model with a nonstationary Matérn covariance function as in Paciorek and Schervish (2006) is adopted, with a smoothness fixed at unity. The measurement error or microscale variation processes $\varepsilon_i(\cdot)$ for $i = 1, 2$ are assumed to be independent mean zero Gaussian white noise with variances $\tau_i^2(\mathbf{s})$.

The model in Eq. 1 is quite general in allowing for a spatially-varying fixed mean function, $\mu_z(\mathbf{s})$, white noise variance for each process, $\tau_i^2(\mathbf{s})$, and spatial registration function $T(\mathbf{s})$. Some restrictions are required on the classes of functions entertained. In particular, it is proposed that these functions fall into a class of continuous and differentiable functions defined by realizations of a GP, specified below.

The full transformation is contained in the parameter functions $\mu_z : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. μ_z represents the translation in the z -coordinate, and T consists of a translation vector function $\mathbf{r} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and a 2×2 orthogonal rotation matrix function $R : \mathbb{R}^2 \rightarrow \mathbb{M}^{2 \times 2}$ parameterized by an angle $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ as follows

$$\begin{aligned}
 T(\mathbf{s}) &= R(\mathbf{s})\mathbf{s} + \mathbf{r}(\mathbf{s}) \\
 &= \begin{bmatrix} \cos \phi(\mathbf{s}) & \sin \phi(\mathbf{s}) \\ -\sin \phi(\mathbf{s}) & \cos \phi(\mathbf{s}) \end{bmatrix} \begin{bmatrix} s_x \\ s_y \end{bmatrix} + \begin{bmatrix} r_x(\mathbf{s}) \\ r_y(\mathbf{s}) \end{bmatrix}, \tag{2}
 \end{aligned}$$

where $\mathbf{s} = [s_x, s_y]^T$ and $\mathbf{r}(\cdot) = [r_x(\cdot), r_y(\cdot)]^T$.

Our specification is a nonrigid generalization of the model in Wiens et al. (2020a) in which μ_z , \mathbf{r} and ϕ are assumed to be constant over the domain. From a practical standpoint, the mean function $\mu_z(\mathbf{s})$ is included in the model to allow a spatially-varying translation in the z -coordinate, as $R(\mathbf{s})$ allows in the x - and y -coordinates and $\phi(\mathbf{s})$ in the rotation angle. Furthermore, allowing $\tau_i(\mathbf{s})$ to vary locally adapts to potential measurement error behavior in the drone in locations of differing topography, for each data set. Overall, this parameterization produces a flexible nonrigid registration function. The estimation of these spatially-varying functions is done using a local likelihood approach in this work, detailed in Sect. 3; this model is also amenable to hierarchical specification and Bayesian estimation.

3 Estimation

In the following, the estimation strategy is presented for both the rigid and nonrigid cases. The full nonrigid registration estimation algorithm is presented, followed by implementation details describing the options and tuning parameters available in the method.

In either the proposed nonrigid registration approach, or the rigid registration approach of Wiens et al. (2020a) (hereafter referred to as GP embedded rigid registration), the data under investigation are two point clouds whose spatial indices can be distinct. Let $\mathcal{U} = \{(\mathbf{s}_{1,j}, Y_1(\mathbf{s}_{1,j}))\}_{j=1}^{n_u}$ and $\mathcal{V} = \{(\mathbf{s}_{2,j}, Y_2(\mathbf{s}_{2,j}))\}_{j=1}^{n_v}$ be points sets whose elements are in \mathbb{R}^3 . \mathcal{U} will be called the target or fixed frame and \mathcal{V} the moving frame; the choice of Y_1 being the target frame and Y_2 being the moving frame is inconsequential as the model in Eq. 1 can be rewritten to accommodate the opposite setup with an analogous parameterization. Let $\mathbf{Y}_1 = (Y_1(\mathbf{s}_{1,1}), \dots, Y_1(\mathbf{s}_{1,n_u}))$, $\mathbf{Y}_2 = (Y_2(\mathbf{s}_{2,1}), \dots, Y_2(\mathbf{s}_{2,n_v}))$, and $\mathbf{Y} = (\mathbf{Y}_1^T, \mathbf{Y}_2^T)^T$. Analogously define the continuous GP components as $\mathbf{Z}_1, \mathbf{Z}_2$ and \mathbf{Z} .

3.1 Rigid Registration Estimation

In the rigid registration algorithm, global estimation is performed to infer all mean and covariance parameters, including transformation parameters, using all available data \mathbf{Y} simultaneously. Concatenating all parameters to be estimated into a vector $\boldsymbol{\theta}$, our estimator is $\hat{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta}} f(\mathbf{Y}, \boldsymbol{\theta})$, where

$$f(\mathbf{Y}, \boldsymbol{\theta}) = \mathcal{L}(\mathbf{Y}, \boldsymbol{\theta}) + \mathcal{P}(\boldsymbol{\theta}). \quad (3)$$

Here, $\mathcal{L}(\mathbf{Y}, \boldsymbol{\theta})$ is a negative log likelihood, while $\mathcal{P}(\boldsymbol{\theta})$ is an optional penalty term that serves to regularize estimates of $\boldsymbol{\theta}$.

Assuming \mathbf{Y} is a realization from a stationary GP with Matérn covariance, \mathcal{L} is a multivariate normal negative log likelihood with mean \mathbf{m} and covariance matrix $\Sigma + \tau^2 I$, where I is an identity matrix of dimension $n_u + n_v$. For reference, the Matérn covariance function is parameterized in this work as $\operatorname{Cov}(Z(\mathbf{s}), Z(\mathbf{s}')) = \sigma^2 \cdot \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{d}{a}\right)^\nu K_\nu\left(\frac{d}{a}\right)$, where $d = \|\mathbf{s} - \mathbf{s}'\|$ is the Euclidean distance between points, σ^2 is the marginal variance, $a > 0$ a spatial range parameter and $\nu > 0$ the smoothness. Γ is the gamma function, and K_ν is the modified Bessel function of the second kind of order ν . Given the model in Eq. 1, the mean of \mathbf{Y} is $\mathbf{m} = (\mathbf{0}_{n_u}^\top, \mu \mathbf{1}_{n_v}^\top)^\top$, where $\mathbf{0}_{n_u}$ and $\mathbf{1}_{n_v}$ are zero and unit vectors of lengths n_u and n_v , respectively. The covariance matrix for \mathbf{Y} is

$$\Sigma + \tau^2 I = \begin{bmatrix} \operatorname{Cov}(\mathbf{Z}_1, \mathbf{Z}_1) & \operatorname{Cov}(\mathbf{Z}_1, \mathbf{Z}_2) \\ \operatorname{Cov}(\mathbf{Z}_2, \mathbf{Z}_1) & \operatorname{Cov}(\mathbf{Z}_2, \mathbf{Z}_2) \end{bmatrix} + \tau^2 I.$$

One key insight is that only the cross covariances $\operatorname{Cov}(\mathbf{Z}_1, \mathbf{Z}_2)$, $\operatorname{Cov}(\mathbf{Z}_2, \mathbf{Z}_1)$ and the covariance matrix of the second data set $\operatorname{Cov}(\mathbf{Z}_2, \mathbf{Z}_2)$ involve the rigid transformation function T , see Wiens et al. (2020a) for details.

The motivation for including a penalty component arises because, during initial experimentation using the method, the optimization sometimes allows for the translation and rotation parameters to take on very large values, effectively making the supports of \mathbf{Y}_1 and \mathbf{Y}_2 disjoint, while it is expected that their optimal values for alignment are smaller. Thus, a penalty term of the form

$$\mathcal{P}(\boldsymbol{\theta}) = 0.5\lambda(r_x^2 + r_y^2 + \mu_z^2) + \log(I_0(\kappa)) - \kappa \cos \phi \quad (4)$$

is used, where I_0 is the modified Bessel function of order 0, λ is the tuning parameter for the penalty on the translation parameters and κ is the tuning parameter for the penalty on the rotation parameter ϕ . In the Bayesian context, these penalties are equivalent to placing Gaussian and von Mises distribution priors on the translation and rotation parameters, respectively. The penalty terms keep the magnitude of the registration parameters small, and additional box constraints can be placed on these parameters in the optimization.

Minimizing the objective function $f(\mathbf{Y}, \boldsymbol{\theta})$ results in estimates for the transformation parameters $\boldsymbol{\theta}$, which include r_x , r_y , μ_z , and ϕ , as well as τ and the covariance parameters specifying Σ (e.g., the process variance and correlation range).

3.2 Nonrigid Registration Estimation

Estimation of the nonrigid registration model in Eq. 1 can be envisioned in many ways: one possibility would be to take a Bayesian approach, putting, say, GP priors on the parameter functions, followed by sampling their posterior distributions. Another approach would be to explicitly parameterize these surfaces as a linear combination of spatial basis functions. However, neither of these approaches is feasible for our setup with tens of millions of data points per point cloud. Some recent advances in spatial statistics have made modeling massive data sets feasible (Datta et al. 2016; Katzfuss 2017), and one could envision embedding the registration part of the statistical model into one of the former modeling frameworks. Instead, in this work a local likelihood approach to estimation is introduced that allows for computation on very high dimensional data sets.

To briefly outline the estimation approach, the spatial domain is divided using a moving window, the parameters are locally estimated using the rigid registration estimation strategy given in Sect. 3.1 within that window using a subset of the data, and finally local estimates are aggregated using a spatial surface model. More concretely, at any given moving window center, the data are subsampled and then used to fit the rigid registration model given in Sect. 3.1, essentially assuming the registration parameters are locally constant, and the process covariance function is locally stationary. The same isotropic Matérn covariance model is used as in the rigid setting during local estimation and prediction. These local estimates are then used as samples of the parameters θ , to which independent spatial surface models are fit. The full approach is described in Algorithm 1.

Algorithm 1 Gaussian Process Nonrigid Registration

- 1: Set up grid $\mathcal{G} = \{\mathbf{g}_k\}_{k=1}^{n_g}$ to locate the centers of overlapping windows $\{\mathcal{B}_k\}_{k=1}^{n_g}$, and partition the data using these windows yielding $\mathcal{D}^U = \{\mathcal{D}_k^U\}_{k=1}^{n_g}$, where each \mathcal{D}_k^U contains the observations in \mathcal{U} that lie within \mathcal{B}_k . Define \mathcal{D}^V and \mathcal{D}_k^V similarly.
 - 2: For the data set \mathcal{U} , subsample N points from each window \mathcal{D}_k^U , resulting in $\tilde{\mathcal{D}}_k^U$ (and similarly for \mathcal{V}).
 - 3: Perform GP embedded rigid registration using data $\tilde{\mathcal{D}}_k^U \cup \tilde{\mathcal{D}}_k^V$ for each window \mathcal{B}_k , yielding local estimates for $\hat{\theta}_k \supseteq \{\hat{r}_{x,k}, \hat{r}_{y,k}, \hat{\mu}_{z,k}, \hat{\phi}_k, \hat{v}_k, \Sigma_k(\hat{\eta}_k)\}$ for $k = 1, \dots, n_g$, where $\hat{\eta}_k$ are the covariance parameters defining Σ_k . Assign the estimated parameters to the centers of the estimation windows \mathbf{g}_k for $k = 1, \dots, n_g$.
 - 4: For each univariate parameter in the vector $\hat{\theta}_k$, independently fit a predictive spatial surface model using the estimated spatially varying transformation parameters $\hat{\theta}_k$ as observations.
 - 5: Using these fitted models, predict the transformation parameters at all of the observation locations in the moving data set $\{\mathbf{s}_{2,j}\}_{j=1}^{n_v}$. Denote the predictions $\{\theta^*(\mathbf{s}_{2,j})\}_{j=1}^{n_v}$, which is a set of vectors of parameters, each vector of length n_v . In particular, predict the transformation parameters $r_x^*(\mathbf{s}_{2,j}), r_y^*(\mathbf{s}_{2,j}), \mu_z^*(\mathbf{s}_{2,j})$, and $\phi^*(\mathbf{s}_{2,j})$, (thus implying $\mathbf{r}^*(\mathbf{s}_{2,j})$ and $R^*(\mathbf{s}_{2,j})$), for $j = 1, \dots, n_v$.
 - 6: Apply the predicted transformation parameters to the moving data set, resulting in a set $\tilde{\mathcal{V}} = \{(\mathbf{s}_{2,j}^*, Y_{2,j}^*)\}_{j=1}^{n_v}$, $\mathbf{s}_{2,j}^* = R^*(\mathbf{s}_{2,j})\mathbf{s}_{2,j} + \mathbf{r}^*(\mathbf{s}_{2,j})$ and $Y_{2,j}^* = Y_{2,j} + \mu_z^*(\mathbf{s}_{2,j})$, which is registered to the coordinate system of \mathcal{U} .
-

The computational cost of the algorithm can be broken down by step. Setting up the grid is trivial. Partitioning the data set into local windows can be done efficiently using a k-d tree, but this step overall is negligible to the total cost of the algorithm also, as is subsampling points in each window. The main computational cost of the algorithm is in performing the n_g local estimation tasks, which are $\mathcal{O}(N^3)$ each under the standard GP maximum likelihood estimation framework, where N is the number of subsamples in each window. Fitting the GP spatial surface models to the locally estimated parameters is $\mathcal{O}(n_g^3)$, but n_g is typically small ($\leq 1,000$) so this task is not computationally expensive. Finally, predicting the transformation parameters at the observation locations of the moving data set is $\mathcal{O}(n_v n_g^3)$ for each parameter.

As long as enough data are used in each local estimation, one would expect the spatially-varying parameters to be identifiable. This can be seen by observing that local estimation reduces to the rigid registration model of Wiens et al. (2020a). This sample size requirement is easily met when analyzing data collected via remote sensing, as is done in this work.

3.3 Implementation Details

Algorithm 1 will now be discussed in more detail, explaining the modeling choices involved in each step. Some choices track this paper's applications, but might be modified for other registration problems. Five of the most important parameters that need to be chosen are the number of estimation locations n_g , the size of windows, and the overlap among windows; the number of subsamples N used in each window for the local estimation; the penalty type and magnitude of the tuning parameters (e.g., λ and κ), and optimization parameter constraints; the covariance function for the GP $Z(\cdot)$ in which the rigid transformation is embedded, defining Σ_k ; and finally the predictive spatial surface model to fit to the locally estimated parameters.

Many windows can be used to assure the signal is captured in the spatially-varying rigid registration parameters. However, using too many windows increases the computational burden. The authors are unaware of theoretical results concerning the selection of local window size, though numerical studies regarding the accuracy of Matérn parameter estimation using a local likelihood approach have been conducted under the assumption of available spatial replicates (Wiens et al. 2020b). Most importantly, the window size in the local estimation procedure must be chosen to contain features present in both data sets and so allow the local rigid registration GP model to identify and estimate the rigid registration. A simulation study is conducted in the following section to investigate this issue. Overall, one should choose the number of windows and overlap so that the signal of the deformation is observed in the raw estimates before a spatial model is fit. In this work, local windows overlap significantly (50–75% overlap between adjacent windows) so that the estimated parameters result in smoothly varying spatial fields.

For spatially dense point clouds, the number of points in each window may still be too large to make fitting a GP quickly feasible. Therefore, one can subsample points from both data sets in each window to decrease the computational burden and still end up with accurate results. Some windows that include near-vertical terrain (e.g., a

cliff in observations of terrain) will contain many more points overall than the average window. In this case, a lower sampling proportion of points can be taken for the likelihood estimation, or more windows can be used in this region (i.e., windows are more closely spaced). This adaptive strategy is not explored further in this work.

Some device must be used in order to force overlap between the two subsets of data within each window; otherwise, there is no incentive for the algorithm to intersect the two data sets during alignment. Tight constraints on the translation parameters are used in Wiens et al. (2020a) based on empirical observation of the magnitude of the transformation that was needed. In this work, the use of a penalty term on the transformation parameters is utilized. The penalty incentivizes the two subsets of data in each window to overlap without tight box constraints on the transformation parameters. If the penalty is too large, the transformation parameters will be estimated near zero. One can gain insight by carrying out the local likelihood estimation procedure without the penalty parameter. The magnitudes of the estimated transformation parameters combined with the value of the likelihood function at a minimum with the penalty parameter set at 0 can inform how to choose the size of the penalty parameter in subsequent application of the method, this time including a penalty.

In each local rigid registration, the data sets are modeled using a GP with Matérn covariance, following Wiens et al. (2020a), due to its flexibility. However, the rigid registration could be embedded in another covariance model with ease if the data were assumed to follow a specific distribution.

Finally, to piece together the local rigid registration parameters into a surface, either a GP or TPS is used. Both choices provide flexible spatial models allowing for prediction of the parameters at unobserved locations. The nonrigid method produces stable local estimates that vary smoothly over space, if one includes some thought in determining subsample size and penalties on parameters. Either model is able to separate the signal and noise components in these estimates and predict accurately at unestimated locations. TPS are used in the simulation studies. On the other hand, in the data analysis a GP with nonstationary covariance structure is used, provided in the LATTICEKRIG package (Nychka et al. 2016) in R (R Core Team 2018). The default settings for this GP model approximate a TPS but scale to large data sets. Additionally, this model is capable of providing uncertainty estimates about predictions and produce conditional simulations, which enable us to calculate forecast scores such as the continuous rank probability score (CRPS).

Note that local estimation of positive parameters (e.g. τ and other covariance parameters) is performed on a log scale, which would be equivalent to modeling these parameters as locally log-normal random variables. Additionally, the spatial surface models used to piece together the local estimates are fitted to the estimates on the log scale, forcing kriging predictions from these models to also be positive after exponentiating them back to their original scale.

4 Simulation Studies

4.1 Types of Deformations

To test the efficacy of the proposed nonlinear registration algorithm, two simulation studies are implemented showing which conditions cause the method to succeed and fail. The first simulation study investigates how the choice of deformation and the parameters to which the deformation is applied affects the estimation accuracy. The second study explores the interplay between the number of estimation locations and the number of subsamples used for each local estimation.

For both studies, the same general setup is used. A spatial data set is simulated using a GP with zero mean and Matérn covariance function with parameters $a = 3$, $\nu = 1$, $\sigma^2 = 2.5$, $\tau^2 = 0.0001$. The observation locations for both data sets are uniformly sampled within a $[0, 6] \times [0, 6]$ square. The data set is split in half, and one half is de-registered with a known deformation. Then, the nonrigid registration algorithm is used to recover the spatially-varying registration parameters. In all cases, 10,000 spatial points are generated, split in half so each data set contains 5000 points. Note that the process used to generate the data set and the process used to generate the deformation are independent.

The first simulation study attempts to answer what kind of deformations the method is capable of estimating. Two different strategies are used to generate the true deformations that de-register the second data set: a simple quadratic field, and a random realization from a Matérn GP. The Matérn deformations are drawn from a GP with zero mean and Matérn covariance with parameters $a = 6$, $\nu = 2$, $\sigma^2 = 0.01$, $\tau^2 = 0.00001$. One Matérn field is drawn for each translation and rotation parameter, and then applied to the moving data set based on the design of the study. The quadratic deformation for the x -coordinate, for example, is defined as follows

$$m_x = \alpha_1 s_x^2 + \alpha_2 s_y^2 + \beta_1 s_x + \beta_2 s_y,$$

where $\alpha_{1,2}$ are drawn from a uniform distribution $U(-0.002, 0.002)$, and $\beta_{1,2}$ are drawn from a uniform distribution $U(-0.04, 0.04)$. The transformation m_x is then applied to the x -coordinate in the moving data set. The same equation governs how the translations in the y - and z -coordinate m_y and m_z are generated, as well as the rotation m_ϕ , with different draws for $\alpha_{1,2}$ and $\beta_{1,2}$ for each parameter.

To investigate how the complexity of the deformation affects estimation accuracy, the Matérn and quadratic deformations are applied to just single parameters as well as groups of parameters. There are five cases: the deformation is applied to the x -coordinate, the z -coordinate, the rotation angle ϕ , the three translation coordinates (x , y , and z), and all four parameters (x , y , z , and rotation angle ϕ).

To estimate the registration parameters, a 4×4 grid of windows with 50% overlap between adjacent windows is used. The 4×4 grid is shown in the first column of Fig. 1. 100 data points are subsampled within each window, and the penalty tuning parameters on the translation parameters and rotation parameter are set at 5 and 100, respectively. Box constraints are set on the translation parameters of $(-1, 1)$ and on the rotation parameter of $(-\pi/4 + 0.1, \pi/4)$ in the optimization. These settings reflect




Fig. 1 The first column shows the estimated transformation parameters, the second column shows the thin plate spline surface predictions fitted to these data, and the third column shows the true deformation parameters which are applied to de-register the moving data set

those used in the data analysis example. Finally, a TPS is fit with default settings using the TPS function in the FIELDS package (Nychka et al. 2017) in R (R Core Team 2018).

In all cases, to quantify the recovery of the nonlinear registration, the estimated registration parameters, denoted with a hat as in \hat{m}_x , are compared to the true registration parameters, for example m_x . Figure 1 shows an example simulation: the true quadratic deformations are shown in the third column and the raw estimates and fitted surfaces are shown in the first and second columns, respectively.

The results of the simulation study are summarized in Table 1. The experiment is executed $n = 30$ times and the average normalized root mean squared error (NRMSE) is computed for each cell in the table. The root mean squared error (RMSE) is normalized by the mean deformation over all coordinates with nonzero deformation. For example, when the deformation is applied only to the x -coordinate, the NRMSE for the x -coordinate is calculated as follows. Let $r_x = (m_x - \hat{m}_x)^2$ be the vector of squared

Table 1 NRMSE of estimating each deformation parameter for varying deformation types, normalized by the mean deformation of all coordinates with nonzero deformation

Deformation type	Applied to	NRMSE _x	NRMSE _y	NRMSE _z	NRMSE _φ
Quadratic	<i>x</i>	0.114	0.069	0.076	0.169
Quadratic	<i>z</i>	0.109	0.102	0.105	0.124
Quadratic	<i>φ</i>	3.956	3.951	0.421	0.711
Quadratic	<i>xyz</i>	0.132	0.119	0.105	0.196
Quadratic	<i>xyzφ</i>	3.620	3.253	0.343	0.670
Matern	<i>x</i>	0.121	0.080	0.075	0.108
Matern	<i>z</i>	0.083	0.080	0.135	0.074
Matern	<i>φ</i>	3.618	3.634	0.206	0.571
Matern	<i>xyz</i>	0.129	0.118	0.131	0.142
Matern	<i>xyzφ</i>	3.150	3.091	0.264	0.479

errors in the *x*-coordinate. Then

$$\text{NRMSE}_x = \sqrt{\bar{r}_x / \bar{m}_x}, \quad (5)$$

where the bar denotes the sample mean of a vector. When the deformation is applied to all coordinates, the NRMSE for the *x*-coordinate is

$$\text{NRMSE}_x = \sqrt{\bar{r}_x / \bar{m}},$$

where the vector $m = (m_x, m_y, m_z, m_\phi)$ is the concatenation of the deformations applied to each data point for all coordinates. In this way, the RMSE is normalized using the mean of all deformations that are actually applied.

These results indicate that deformations of both types are recovered effectively. Interestingly, a quadratic deformation of the type that is applied here could possibly be estimated using a rigid registration with a scaling/stretch parameter. However, this study indicates that this deformation can also be effectively estimated using piecewise rigid transformations.

It appears that the nonlinear registration algorithm is able to well-estimate deformations applied to the *x*-, *y*-, and *z*-coordinates, but has more trouble recovering deformations involving a rotation. However, the rotation angle is only inaccurately estimated if an angular deformation is introduced. If a deformation is applied to the rotation angle, this also introduces error in estimating the *x*-, *y*-, and *z*-coordinates. Another insight from Table 1 is that when the true parameters are zero (i.e., no translational or rotational misalignment), then the estimated parameters are particularly accurate, suggesting that the approach tends to not introduce deformations when they are not needed.

Table 2 NRMSE of estimating the x -coordinate deformation for varying number of windows and number of subsamples, normalized by the mean deformation of the x -coordinate

Grid size	NRMSE _{x}			
	Subsamples			
	25	50	100	200
3×3	0.612	0.430	0.279	0.112
4×4	0.417	0.285	0.146	0.092
5×5	0.341	0.203	0.139	0.074

4.2 Windowing and Subsampling

In the second simulation study, the effect of the number of estimation locations and the number of subsamples of the estimators is investigated.

Grids of sizes 3×3 , 4×4 , and 5×5 are used to locate the centers for overlapping windows. Subsamples are taken within each window of size 25, 50, 100, and 200. These experimental settings kept the computation reasonable by limiting both the number of local estimation tasks and also the total number of observations in each GP estimation.

The moving data set is deformed using a realization from a GP with Matérn covariance, as in the first simulation study. The covariance parameters used to generate the deformation are known and fixed in the estimation. Furthermore, only the x -coordinate is deformed and estimated, meaning there is only one registration parameter to estimate (ϕ and the y - and z -coordinate translation parameters are fixed at 0). The box constraint on the x -coordinate is set at $(-1, 1)$ in the optimization. The overlap between adjacent windows and the regularization tuning parameters are identical to the previous simulation study, as well as the TPS surface fits with default settings.

This experiment is repeated $n = 30$ times to compute average NRMSE, which is shown in Table 2. In this case, the deformation is applied only to the x -coordinate, so the RMSE is normalized as in Eq. 5. The trend is clear from these results: the nonlinear registration algorithm is more accurate when using more windows and more subsamples. However, since the task becomes much more computationally expensive when using more subsamples and more windows, a balance must be found based on how the registration parameters vary over space. If there is not much variance in the registration parameters over space based on exploratory analysis, few windows and subsamples are needed, such as in this example. If the point clouds exhibit highly complex variation in the response variable, more subsamples and windows may be required to obtain accurate results.

5 Drone-Based Point Clouds from the Chalk Cliffs

The proposed nonrigid Gaussian process registration method is applied to a pair of geomorphic elevation data sets with over 100 million points. The data are point sets processed from drone-based remote sensing scans of the Chalk Cliffs debris flow basin in Colorado. The Chalk Cliffs receives debris from a 0.37 km^2 watershed at the base of Mount Princeton in the Arkansas River Valley. The considered data set results from




Fig. 2 One scan showing the full data set. This is the target data set containing 6,775,704 points in total. The units displayed in the scales are meters

two flyovers of the fan during the same day, 10/16/2017, at 40 m flight altitude. No large scale erosion or deposition occurred between flights. Barnhart et al. (2019) and Wiens et al. (2020a) provide details about the geologic motivation. It is hypothesized that due to the SfM algorithm, a nonrigid registration may be necessary for proper alignment.

To streamline the application for this work, an area is identified using similar geographic features in both scenes, and then the data are cropped manually, resulting in areas whose spatial boundaries are close but not identical. In total, there are around 60 million data locations in each cropped scan of the region; one scan is shown in Fig. 2.

The nonrigid registration method developed in this paper is compared with the likelihood-based rigid registration method, as well as with the ICP algorithm. The predictive quality of each method is assessed in a cross validation study, holding out 0.1% of the points from the fixed data set to be used as the testing set, and the remaining observations in the fixed data set and the entire moving data set are used to estimate the registration parameters. ICP is performed in the open source CloudCompare software, version 2.9.1, which is also used to generate Figs. 2 and 3. To match the registration capabilities of the rigid and nonrigid likelihood-based algorithms, the rotation is restricted to the xy -plane when using ICP.

For the nonrigid registration, a 4×12 grid of windows is used with 66% overlap. 250 subsamples from each data set are used in each local estimation, and the regularization tuning parameters are fixed at $\lambda = 5$ and $\kappa = 100$. The box constraints in the optimization are set at $(-1.2, 1.2)$ on the translation parameters and $(-\pi/4 + 0.1, \pi/4)$ on the rotation parameter. The rotation angle constraints are unrestrictive, and the translation parameter constraints encompass the true values, so in this case are also unrestrictive, but help reduce the parameter search space and prevent complete non-




Fig. 3 The top row shows both data sets before and after the nonrigid registration is applied, zoomed in centered on a ravine. The bottom row shows two sample registrations based on conditional simulations of the registration parameters. The colorbar scale is identical to Fig. 2

overlap between the data sets. The locally estimated rigid registration parameters are fit with LATTICEKRIG GP models.

Since the data sets are too large to perform standard kriging, a local kriging procedure is used (Haas 1990a, b, 1995) to predict the response at the locations of the held-out observations for all registration methods. In particular, the 1,000 closest points to the observation location to be predicted are used, and with these points standard kriging is performed using the stationary Matérn covariance function. The covariance parameters used in the kriging prediction are either constant for the two rigid registration methods, or spatially varying as estimated in the nonrigid registration case, detailed below.

The covariance parameters for the rigid registration method are estimated as part of the rigid likelihood-based procedure. Since the full data sets cannot be used, 500 observations from each data set are subsampled and used to fit the rigid registration model. The transformation parameters are applied to the moving data set, and the covariance parameters are used to perform the local kriging. Similarly, for ICP, after the rigid registration parameters are applied to the moving data, the covariance parameters are estimated using 500 observations from both data sets, then used in the local kriging procedure.

For the nonrigid registration, spatially varying covariance parameters are estimated along with the transformation parameters. Similar to the transformation parameters, the estimated covariance parameter fields are smoothed using LATTICEKRIG GP models, and then used to predict the covariance parameters at the locations of the observations in the testing set. The LATTICEKRIG model makes prediction at the large number of locations in the testing set tractable.

Table 3 RMSE and CRPS for the held-out testing set, comparing three registration methods

	ICP	Rigid likelihood	Nonrigid likelihood
RMSE	0.086	0.118	0.080
CRPS	0.044	0.027	0.021

The locally estimated values of τ^2 range between a minimum of 5.5×10^{-7} and a maximum of 0.0513, with the median over all 48 local estimates equaling 0.002 (units are in meters). Similarly, the range for σ^2 is 0.41 to 13.67 with median 1.06, and the range of a is 2.21 to 28.13 with median 4.18. These results are close to those estimated in the rigid registration case (Wiens et al. 2020a), and the values of τ^2 are consistent with the expected measurement error from a data set constructed via SfM (Harwin and Lucieer 2012).

The overall elapsed time during the local estimation step of the algorithm took 1.51 h, and the mean time of local estimation was 113 s averaged over all 48 local windows.

The cross validation results are shown in Table 3. In addition to calculating RMSE using the true and predicted values of the response for the testing set, CRPS is calculated to assess the accuracy of the uncertainty quantification. To calculate the variance of the predictions, a simulative approach is used. For all three methods, CRPS is calculated using 30 simulations.

For the nonrigid method, the locally-estimated rigid registration parameters have uncertainties approximated in the optimization by the numerical Hessian. Using these uncertainties, registration parameters are simulated at the locations of the centers of the moving windows. A LATTICEKRIG model is fit to these simulated parameters, and then a conditional simulation is produced at the observation locations of the moving data set, and then applied. This technique propagates the two levels of uncertainty into the registration, from the local estimation and the predictive surface model fitting. This procedure was repeated 30 times and these empirical simulations of the predictions at the testing set locations were used to calculate the CRPS.

For the rigid likelihood-based method, the rigid registration parameters are simulated using the uncertainties from the numerical approximation to the Hessian, and then applied to the entire moving data set. Then local kriging is performed using simulations of the Matérn covariance parameters at the testing set locations, once again from the approximate distribution given by the numerical Hessian. For ICP, only the covariance parameters are simulated prior to local kriging as just described. Specifically, rigid registration using ICP produces no uncertainty measures, so the parameter estimates are treated as fixed.

Overall, the nonrigid registration algorithm achieves lower RMSE and CRPS than the two rigid registration methods, indicating superior accuracy of the registration and also quantification of the uncertainty. Moreover, visually the two data sets are more accurately aligned as seen in the top row of Fig. 3, which shows a close-up view of both data sets before and after registration using the nonrigid registration method in the left and right panels, respectively. There is still clear misalignment, albeit minor. The fact that the approach is statistical is critical here. While there may be a best

alignment, it is acknowledged that the estimated nonrigid registration is not perfect; however, the uncertainty about the registration can be quantified.

The bottom row of Fig. 3 displays two registrations incorporating uncertainty using the simulative technique used in calculating CRPS. The bottom right panel contrasts a more accurate registration with a less accurate registration in the bottom left panel. Visually this shows how the uncertainty estimates encompass the true nonrigid deformation.

6 Conclusions

The nonrigid registration algorithm developed in this paper provides an extremely flexible method for registering complex data sets. Moreover, the local likelihood estimation within each window can be treated as an independent task, and can thus utilize parallel processing. For these reasons, this method is capable of handling massive data sets.

When using the nonrigid registration method, several modeling choices must be leveraged to obtain accurate results. One must balance computational demands with model flexibility. The choice of the number of windows and the number of subsamples used in each window in the estimation procedure influences the degree to which the registration parameters are allowed to vary over space, and also controls the computational burden of the algorithm. If a highly nonrigid deformation is required, many windows should be used. If the surface is complex or noisy, many subsamples should be used in each window to ensure proper local rigid registration.

Overfitting is intrinsically associated with nonrigid registration methods, and regularization is essential to avoid this problem (Tam et al. 2013). The method developed in this paper has the capability for accurate and flexible nonrigid registration, but it also suffers from the possibility of overfitting. If too many windows are used in conjunction with too small a penalty parameter, the data can be erroneously registered. However, the inclusion of a penalty parameter and the smoothing of the spatial surface models help combat overfitting associated with nonrigid registration methods. While the penalty parameter is advantageous in this regard, it can also influence the estimation of the registration parameters if it is too large.

Experimentation with different values of λ and κ was performed for the applications in this paper, and the values used were chosen and fixed because they worked well at producing realistic estimates of the deformation. The authors recommend using large penalty parameters initially, causing the transformation parameters to be estimated near 0, then reducing the penalty parameters until stable local estimates are obtained. One could use generalized cross-validation or information criteria to choose the penalty parameters, but this route was not pursued due to the computational demands of the nonrigid algorithm proposed in this paper.

The choice of predictive spatial surface model can also affect the final results. Both GP and TPS provide flexible models that allow for effective smoothing of the locally estimated rigid registration fields. One could use a stationary or nonstationary covariance function in the GP model based on the context, or a TPS is viable if uncertainty quantification is not important in the analysis.

The likelihood-based nonrigid registration method is a unique and novel statistical approach to solving this problem which can help further the field of registration. Through two simulation studies, it has been shown that the method can effectively recover deformations of different types, and the effect of using different tuning parameters in the method was explored. Finally, in the data analysis example, more accurate predictions and uncertainties were achieved in a cross validation study using the non-rigid method compared with two rigid registration approaches.

Acknowledgements This work was supported by the National Science Foundation (Awards DMS-1811294 and DMS-1923062). Additionally, Barnhart was supported by NSF Award EAR-1725774. This work utilized resources from the University of Colorado Boulder Research Computing Group, which is supported by the National Science Foundation (Awards ACI-1532235 and ACI-1532236), the University of Colorado Boulder, and Colorado State University (Anderson et al. 2017).

References

- Anderson J, Burns PJ, Milroy D, Ruprecht P, Hauser T, Siegel HJ (2017) Deploying RMACC summit: an HPC resource for the Rocky Mountain region. In: Proceedings of PEARC17, New Orleans, LA, USA
- Arad N, Reisfeld D (1995) Image warping using few anchor points and radial functions. *Eurogr Assoc* 14:35–46
- Arthur D, Vassilvitskii S (2009) Worst-case and smoothed analysis of the ICP algorithm, with an application to the k-means method. *SIAM Journal on Computing* 39(2):766–782
- Barnhart K, Rengers F, Ghent J, Tucker G, Coe J, Kean J, Smith J, Kleiber W, Wiens A (2019) Topographic change detection at Chalk Cliffs, Colorado, USA, using airborne lidar and UAS-based structure-from-motion photogrammetry. In: 7th international conference on debris-flow hazards mitigation, pp 85–92
- Besl PJ, McKay ND (1992) A method for registration of 3-D shapes. *IEEE Trans Pattern Anal Mach Intell* 14(2):239–256
- Bookstein FL (1989) Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans Pattern Anal Mach Intell* 2(6):567–585
- Cervone D, Pillai NS (2015) Gaussian process regression with location errors. [arXiv:1506.08256](https://arxiv.org/abs/1506.08256) statME preprint
- Chui H, Rangarajan A (2003) A new point matching algorithm for non-rigid registration. *Comput Vis Image Underst* 89(2–3):114–141
- Cressie N, Kornak J (2003) Spatial statistics in the presence of location error with an application to remote sensing of the environment. *Stat Sci* 18(4):436–456
- Datta A, Banerjee S, Finley AO, Gelfand AE (2016) Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *J Am Stat Assoc* 111(514):800–812
- Ezra E, Sharir M, Efrat A (2008) On the performance of the ICP algorithm. *Comput Geom* 41(1–2):77–93
- Fanshawe TR, Diggle PJ (2011) Spatial prediction in the presence of positional error. *Environmetrics* 22(2):109–122
- Fonstad MA, Dietrich JT, Courville BC, Jensen JL, Carbonneau PE (2013) Topographic structure from motion: a new development in photogrammetric measurement. *Earth Surf Process Landf* 38(4):421–430
- Fronterre C, Giorgi E, Diggle P (2018) Geostatistical inference in the presence of geomasking: a composite-likelihood approach. *Spat Stat* 28:319–330
- Gilleland E, Chen L, DePersio M, Do G, Eilertson K, Jin Y, Kang EL, Lindgren F, Lindström J, Smith RL, Xia C (2010) Spatial forecast verification: Image warping. NCAR/TN-482+STR
- Gold S, Rangarajan A, Lu CP, Pappu S, Mjolsness E (1998) New algorithms for 2D and 3D point matching: pose estimation and correspondence. *Pattern Recogn* 31(8):1019–1031
- Haas TC (1990a) Kriging and automated variogram modeling within a moving window. *Atmos Environ Part A Gener Top* 24(7):1759–1769
- Haas TC (1990b) Lognormal and moving window methods of estimating acid deposition. *J Am Stat Assoc* 85(412):950–963

- Haas TC (1995) Local prediction of a spatio-temporal process with an application to wet sulfate deposition. *J Am Stat Assoc* 90(432):1189–1199
- Harwin S, Lucieer A (2012) Assessing the accuracy of georeferenced point clouds produced via multi-view stereopsis from unmanned aerial vehicle (UAV) imagery. *Remote Sens* 4(6):1573–1599
- Huang QX, Adams B, Wicke M, Guibas LJ (2008) Non-rigid registration under isometric deformations. *Eurogr Symp Geom Process* 27(5):1449–1457
- James MR, Robson S (2012) Straightforward reconstruction of 3D surfaces and topography with a camera: Accuracy and geoscience application. *Journal of Geophysical Research* 117(F3):
- Jian B, Vemuri BC (2005) A robust algorithm for point set registration using mixture of Gaussians. In: Tenth IEEE international conference on computer vision (ICCV'05), vol 2, pp 1246–1251
- Katzfuss M (2017) A multi-resolution approximation for massive spatial datasets. *J Am Stat Assoc* 112(517):201–214
- Kuusela M, Stein ML (2018) Locally stationary spatio-temporal interpolation of Argo profiling float data. *Proc R Soc A Math Phys Eng Sci* 474(2220):20180400
- Maiseli B, Gu Y, Gao H (2017) Recent developments and trends in point set registration methods. *J Vis Commun Image Represent* 46:95–106
- Nychka D, Furrer R, Paige J, Sain S (2017) *Fields: tools for spatial data*. University Corporation for Atmospheric Research, Boulder, CO, USA
- Nychka D, Hammerling D, Krock M, Wiens A (2018) Modeling and emulation of nonstationary Gaussian fields. *Spat Stat* 28:21–38
- Nychka D, Hammerling D, Sain S, Lenssen N (2016) LatticeKrig: multiresolution kriging based on Markov random fields. University Corporation for Atmospheric Research, Boulder, CO, USA
- Paciorek CJ, Schervish MJ (2006) Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* 17:483–506
- R Core Team (2018) *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna
- Risholm P, Pieper S, Samset E, Wells WM (2010) Summarizing and visualizing uncertainty in non-rigid registration. volume 13. Springer, Berlin, pp 554–561
- Tam GKL, Cheng Z, Lai Y, Langbein FC, Liu Y, Marshall D, Martin RR, Sun X, Rosin PL (2013) Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. *IEEE Trans Vis Comput Gr* 19(7):1199–1217
- Tsin Y, Kanade T (2004) A correlation-based approach to robust point set registration. In: Pajdla T, Matas J (eds) *Computer vision—ECCV 2004*. Springer, Berlin, pp 558–569
- Westoby MJ, Brasington J, Glasser NF, Hambrey MJ, Reynolds JM (2012) ‘Structure-from-Motion’ photogrammetry: a low-cost, effective tool for geoscience applications. *Geomorphology* 179(15):300–314
- Wiens A, Kleiber W, Barnhart KR, Sain D (2020a) Surface estimation for multiple misaligned point sets. *Math Geosci* 52:527–542
- Wiens A, Nychka D, Kleiber W (2020b) Modeling spatial data using local likelihood estimation and a Matérn to spatial autoregressive translation. *Environmetrics* 31:e2652
- Zhou Z, Tu J, Geng C, Hu J, Tong B, Ji J, Dai Y (2018) Accurate and robust non-rigid point set registration using student’s-t mixture model with prior probability modeling. *Sci Rep* 8(1):8742
- Zhu J, Hoi SCH, Lyu MR (2009) Nonrigid shape recovery by Gaussian process regression. In: 2009 IEEE conference on computer vision and pattern recognition, Miami FL, pp 1319–1326