IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 1 March 2024, accepted 9 April 2024, date of publication 12 April 2024, date of current version 22 April 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3388292

RESEARCH ARTICLE

Bayesian Inference of Elevation to Reduce Large Interpolation Errors in 2-d Road Features Draped Over Digital Elevation Models

CRISPIN H. V. COOPER
Department of Computer Science and Informatics, Cardiff University, CF24 4AG Cardiff, U.K.
e-mail: cooperch@cardiff.ac.uk

ABSTRACT The usual approach for adding elevation data to two dimensional (2-d) vector features in a Geographic Information System (GIS) is to infer heights from a Digital Elevation Model (DEM), either through traditional (naïve) interpolation, Kriging, or deep learning. Where the terrain contains steep slopes, however, any of these approaches can generate large errors due to the limited resolution of the DEM, and model error in the DEM concept itself. In the case of road networks, these errors have a severe and nonlinear effect on cycling route planners and transport models, especially those based on open elevation data. This paper introduces a Bayesian maximum likelihood approach to correcting interpolated heights, by combining a DEM with prior expectations of feature gradient. The topological network defined by feature shapes is used as auxiliary information. Correcting the output of naïve interpolation shows reduction of mean errors, and reduced overprediction of elevation change outliers, compared to both naïve interpolation and Kriging.

INDEX TERMS Bayesian, digital elevation model, drape, geographic information system, interpolation, Kriging, network, road.

I. INTRODUCTION

Geographic Information System (GIS) practitioners often need to add elevation (height) data to a 2-dimensional vector dataset. Typically this is conducted with a naïve interpolation [1], also known in software implementations as a "drape" operation, e.g. Grass v.drape [2] or ESRI Interpolate Shape [3]. These tools perform bilinear or bicubic interpolation between adjacent sample points in an underlying digital elevation model (DEM) to assign suitable elevations to vertices on the 2-d dataset, which in most cases will lie in between points of different heights on the DEM.

A bilinear interpolation, for example, can be obtained from the one-dimensional linear interpolation

$$z_x = f(x, x_0, x_1, z_0, z_1) = z_0 + \frac{x - x_0}{x_1 - x_0}(z_1 - z_0) \quad (1)$$

where $f(x, x_0, x_1, z_0, z_1)$ is the linear interpolation function giving z_x the height at coordinate x, from x_0, x_1, z_0, z_1 the coordinates and heights for the nearest known values to either side of x. Given any terrain tile in a 2-d rectilinear DEM, in which any row of measurements is one-dimensional, the interpolated height $z_{x,y}$ can be obtained by first performing one-dimensional interpolations at points (x, y_0) and (x, y_1), for the nearest rows of samples at y_0, y_1: $z_{x,y0} = f(x, x_0, x_1, z_{x0,y0}, z_{x1,y0})$ and $z_{x,y1} = f(x, x_0, x_1, z_{x0,y1}, z_{x1,y1})$ respectively. These outputs are then interpolated again to give $z_{x,y} = f(y, y_0, y_1, z_{x,y0}, z_{x,y1})$. Given more adjacent points, a cubic or higher-order interpolation can be used; alternatively, Kriging [4] or DEM super-resolution [5].

Any of these approaches, however, are prone to large errors in the case where large changes in elevation occur between adjacent sample points in the DEM, as with the DEM alone, we lack information on exactly where the elevation change occurs relative to the draped feature. A stark example is found in the case of a road near to the top of a cliff (Fig. 1, top). Terrain data is inevitably of limited resolution, ranging from 30-50m sampling intervals for open data [6], [7], [8] to 5m for

The associate editor coordinating the review of this manuscript and approving it for publication was Jenny Mahoney.

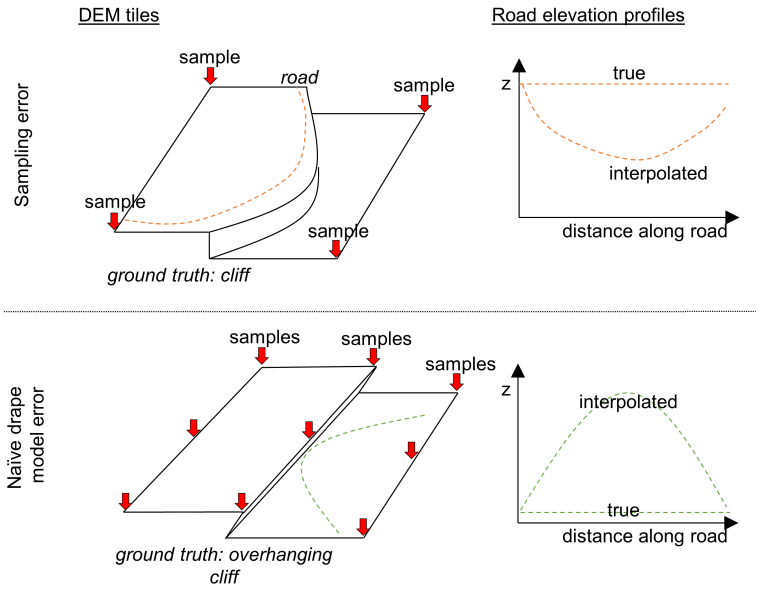

FIGURE 1. Illustration of errors in the naïve drape operation. [TOP] Sampling error causes a road (orange dotted line) to be interpolated partway down the cliff height, even though the road's ground truth is level. In principle this can be resolved with sufficient sample density. [BOTTOM] Model error causes another road (green dotted line) to be interpreted as climbing the full height of the cliff, even though the road's ground truth is level. No DEM can resolve this model error, regardless of resolution.

commonly used commercial data [9]. Even the higher of these resolutions has a horizontal inter-sample distance exceeding the width of some roads. We are unlikely to estimate the correct centreline height of a road adjoining the top, or base of a cliff, based on low resolution DEM data alone - regardless of the technique used - as the elevation change across the DEM tile reflects the presence of a cliff which the road does not cross. This can be classed as a sampling error, because the DEM sample points are not representative of the underlying terrain. Fig. 3(a) shows an example of this error in draped imagery, for a road following the top of an 80 metre cliff.

Furthermore, the DEM abstraction itself assumes only a single height for each (x, y) coordinate, and is thus subject to model mis-specification error in the case of overhangs (Fig. 1, bottom). In these cases, the true terrain has multiple elevations for the same (x, y) coordinates, so no DEM can be interpreted correctly regardless of the sampling interval. The maximum error is limited only by the height of the cliff itself.

This paper presents and tests a method for correcting such errors, in the absence of more accurate survey data or further sources of information. The outputs of the naïve drape process are corrected by incorporating priors on vector feature elevation, into a Bayesian likelihood model (Fig. 2). The priors are quantified by our pre-existing expectations for the distributions of its gradients, and the random and sampling errors in the drape process. As each gradient is affected by two vertex elevations, and most vertex elevations affect multiple gradients, likelihood information must propagate through the topological network of the vector features. The vector topology itself can thus be considered as auxiliary data. Tests are conducted to determine the resulting error reduction in the naïve drape output, and outputs are also compared to Kriging.

The focus is on networks of road centrelines, however, the approach should be similarly applicable to GPS traces of people and vehicles following terrain, whether or not on a road (e.g. those recorded from fitness tracker apps). Other datasets will also exist in which it is possible to quantify priors on the shape of interpolated features and hence use a similar method.

The remainder of this paper is structured as follows. Section I-A elaborates on the current motivation for this study. Section II describes related work. Section III describes the methods used, both for the model itself and for a test against validation data. Section IV shows test results, including the sensitivity to choice of prior. Section V concludes with recommendations on suitable choice of priors for future users of the model.

A. MOTIVATION

The issue of interpolation errors applies both to Triangular Irregular Networks (TINs) derived through survey, and to interpolated rectilinear DEMs derived from TINs, which are more typically available to the end user of the data. Although such problems can be resolved through more detailed surveying techniques, in practice many users are constrained to use off-the-shelf DEM data.

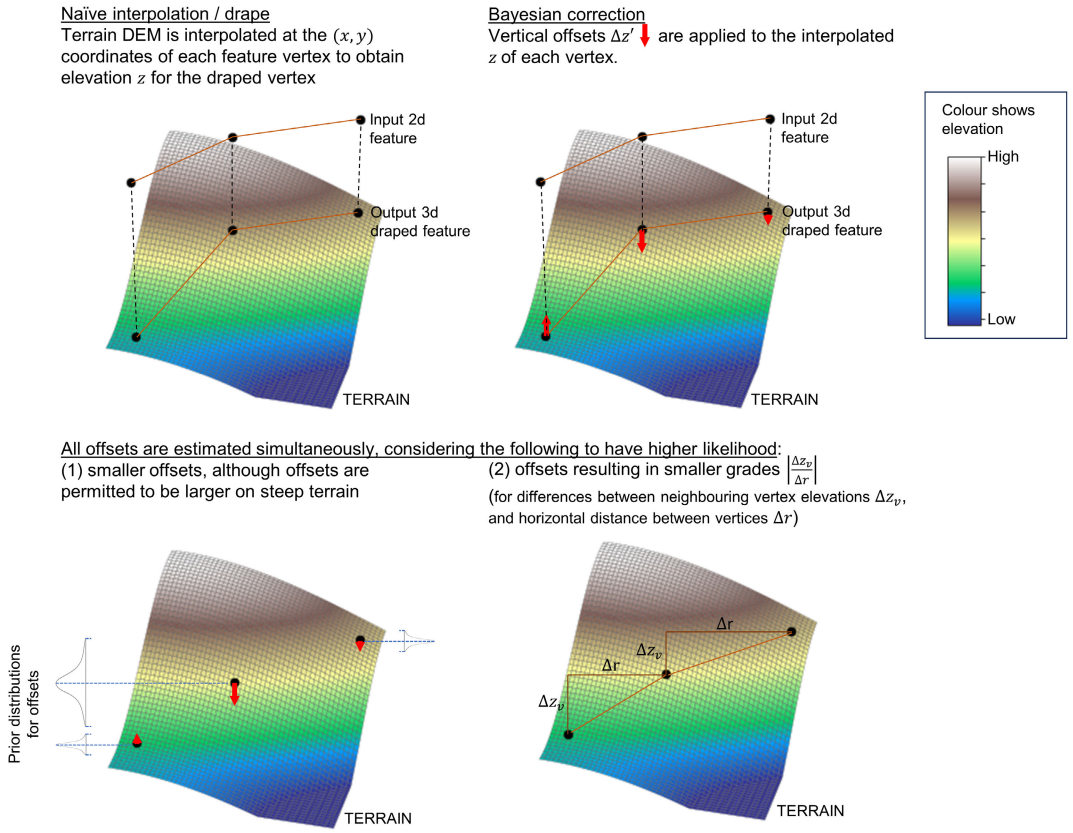

FIGURE 2. **Illustration of the usual naïve drape and proposed Bayesian improvement. Terrain is illustrated as an interpolated surface, though in reality the ground truth would be approximated by a lower resolution DEM.**

In the case of road networks, such errors are a problem of current practical importance in the drive to expand traditional transport models to better simulate 'active travel', i.e. walking [10], [11], [12], [13], [14], [15], [16], [17] and cycling [18], [19], [20], [21]. Both in transport models, and also public route planners such as Google Maps and CycleStreets, routes of active travel users are assumed to minimize some definition of generalized cost. In the case of transport models, some randomization / distribution over similar routes may be applied, but the resulting routes will all be approximations to the minimum cost route nonetheless [22]. As steep roads substantially deter people from walking and cycling otherwise short journeys, the definition of generalized cost typically includes gradient either directly, or in the computation of trip time [20], [21]. However, this approach suffers from multiple sources of nonlinearity. Firstly there is nonlinearity in computing minimum cost routes: if a link which in reality forms part of the minimum cost route is modelled with a cost which substantially exceeds the ground truth, then it will no longer be the minimum cost route in the model. The modelled route will then be wrong, and modelled flows will be assigned elsewhere. Secondly, cyclist speed is nonlinear with respect to gradient; the equilibrium speed of a cyclist uphill shows better (negative) fit to the gradient squared [18]. A similar effect is likely to exist for the generalized cost perceived by the cyclist. These factors combine to make cycling route planners and transport models extremely sensitive to gradient outliers. The link in Fig. 3(c) is a case in point. This is a traffic-free section of a major route in the UK National Cycle Network [23] which caused problems in cycle modelling [24] due to errors in elevation from a naïve drape. The model predicted that cyclists would instead use a busy and steep trunk road. The height inaccuracies remain regardless of whether a 5m resolution or 50m resolution model is used, and the routing inaccuracies persist in spite of motor vehicle traffic also serving as a deterrent to cyclists on the alternative route. If such errors continue to escape detection, inaccurate transport models could lead to real world consequences such as (1) deterring people from travel by bicycle in cases where route planners recommend bad routes; (2) underinvestment in sustainable transport in areas where it is erroneously deemed too inconvenient; (3) overinvestment in infrastructure at inappropriate locations, or (4) errors in quantifying the safety of some junctions if cyclist flows are different to what is expected.

The importance of correctly modelling height change can be seen in Fig. 3(b), the Cefn Coed viaduct near Merthyr Tydfil. The town comprises several steep-sided valleys which deter cycling. This local authority had the lowest levels of cycling to work in the 2011 census of England and Wales [25]. The Wales Active Travel Act mandates identification of

(a)

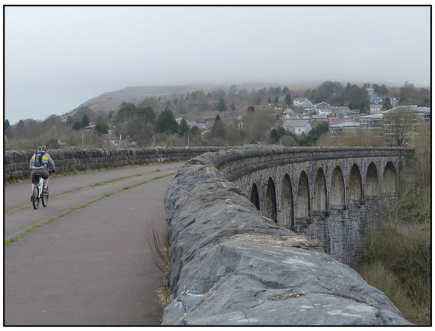

(b)

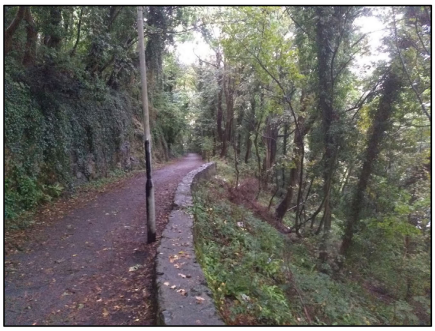

(c)

FIGURE 3. (a) Incorrect draping in 3-d visualisation of a satellite photograph over a DEM in Google Earth - a level road along a clifftop is displayed with major height changes. Although the current paper deals with vectorized roads, not raster images, this image illustrates the cause of error (© Google, 2021). (b) Example of terrain in which accurate modelling of height in sustainable transport is crucial (© Robin Drayton, licensed under CC-BY-SA-2.0). (c) Level link traversing steep slope, with high potential for incorrect elevation estimates. This is part of the National Cycle Network and causes problems in existing transport models.

suitable cycling routes on an ongoing basis [26], a process which is informed in part by modelling. Routes thus identified are prioritized for future investment.

II. RELATED WORK

A. ALTERNATIVE METHODS FOR SMOOTHING LINEAR DATA

The slopes R package [27] aims to resolve elevation errors in transport modelling by use of weighted mean of segment gradients. It only interpolates polyline vertices and does not consider what happens in between them. This in turn is based on a similar method applied to rivers [28]. Although the choice to ignore elevation between vertices may serve to smooth interpolated elevations somewhat, the approach risks over-smoothing, and also assumes that the interpolated height at each vertex is accurate in the same way as the drape tools already discussed [2], [3].

The usual approach to removing spurious variance in input data to any model (geospatial or otherwise) is to apply some form of outlier correction or smoothing. If we encode a prior for the likely gradient between successive vertices of a linear feature, then the sequence of true elevations along the feature could be considered a Markov chain: each elevation depends both on the true elevation of the previous vertex, and the likelihood of the gradient implied by any change in elevation between successive vertices (a physical model). For a single polyline feature such as the GPS trace of a cyclist, elevation data can therefore be smoothed using a Bayesian filter [29]. This considers elevations interpolated from the terrain model as incoming measurements subject to error, and computes the maximum a posteriori likelihood for the elevation of each successive vertex. It thus corrects incoming measurements based on our prior beliefs on what is likely based on previous estimates and our physical model.

In the general case of a geographical feature, however, we are not dealing with an ordered sequence of vertices but a topological model in which each vertex can have multiple neighbours, each of which influence our beliefs about the true elevation of the vertex. This requires a different estimation procedure. A further complication is that the expected elevation error in a naïve drape of a road centreline over terrain, will vary depending on the steepness of the terrain. If we are to distinguish between genuine versus spurious elevation changes, the smoothing process must therefore be informed directly by the terrain model rather than postprocessing a naïve drape. Fig. 4 illustrates how different ground truths can lead to identical draped features - one of which exhibits spurious height change due to the drape process, while the other genuinely reflects the ground truth. These identical draped features must be treated differently to

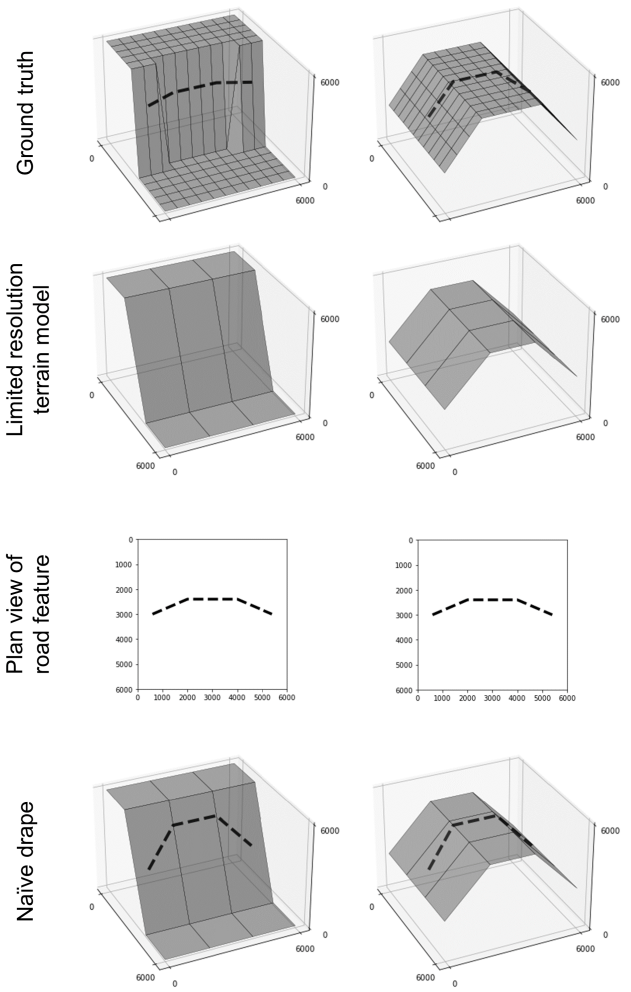

FIGURE 4. Why smoothing the draped feature is insufficient: two different ground truths give rise to two identical 3-d drapes of 2-d road polylines, so a smoothing algorithm applied to the draped feature would not be able to distinguish these situations. [LEFT] A level road follows a contour of a steep slope. The limited resolution DEM does not capture the precise slope contours, so the varying y coordinate of the 2-d feature means that a naïvely interpolated elevation profile will appear to climb and descend the low resolution DEM tile. [RIGHT] A road climbs and descends a ridge. In this case the naïve interpolation of the 2-d vertices on the DEM gives rise to an elevation profile which accurately reflects the ground truth. All axis measurements in metres.

correct the errors, so we cannot simply smooth the draped features.

B. ALTERNATIVE METHODS FOR INTERPOLATING DIGITAL ELEVATION MODELS

Naïve interpolation is not the only option for inferring a height from a DEM: Kriging, and DEM super-resolution using deep learning, may also be used.

Kriging [4] estimates the value of a variable at an unmeasured location, based on Maximum A Posteriori (MAP) interpolation of nearby measurements. Many variants exist with similar underlying principles. A variogram is first estimated which describes spatial dependency in the data as a function of expected difference in the output variable Z at two different points p_1, p_2: $E(|Z(p_1) - Z(p_2)|^2)$. Several alternative variogram models are available, typically a function of the distance between p_1 and p_2. Estimators are then derived for the unmeasured value of $Z(p)$ at any point p. Although not usually described as such, these estimators can be understood in Bayesian terms as finding the maximum posterior likelihood value of $Z(p)$ from a number of known neighbouring values $Z(p_1), \ldots, Z(p_n)$ by combining priors $P_i(Z(p)|Z(p_i), d(p, p_i))$ based on the variogram, and the distance to each neighbouring point $d(p, p_i)$.

The proposed technique can in fact be viewed as related to Kriging, albeit applied in the topological space of the road network, not in the Euclidean space of the DEM. Kriging has even been applied on networks before [30]. The difference is that in the current case, estimates of each unmeasured value are informed by a physical model of prior expectations on road gradient, and crucially, output estimates are not independent: each is affected by estimation of its neighbours (as determined by network topology) and priors on the gradient between them (as determined by network shape).

The other option for DEM interpolation is DEM super-resolution through deep learning [31], [32], [33], [34]. A recent review of deep learning in DEM processing [5] summarizes advantages of deep learning as improved classification accuracy, autonomy and efficiency, and ability to handle heterogeneous data; and the disadvantages as data and resource requirements, dependence on representative training data and interpretability. The proposed approach differs in a number of ways. It aims only to determine elevation of the vector features, rather than reconstruct a complete DEM (this is essential in order to avoid the model error illustrated in Figure 1). It makes use of the road centreline network as auxiliary information: while some super-resolution approaches [33] make use of gridded auxiliary information, none make use of the information contained in a topological network, so this is a unique contribution of the current study. Finally, the parameters of the likelihood model used in the current study are derived analytically and hence interpretable [35] without need for ongoing research in explainable AI [36].

Although the proposed technique is demonstrated as an enhancement of naïve interpolation, it could also be applied to the output of either Kriging, or deep learning approaches. Such combinations are outside the scope of this study. It should be noted that the Kriging results presented in this study are not applying the proposed Bayesian correction to the output of Kriging, but compare the proposed Bayesian correction of naïve drape, to the results achieved using both naïve drape, and ordinary Kriging, to interpolate the DEM.

III. METHODS

The method used is based on maximum a posteriori likelihood estimation. It is Bayesian in the sense that priors are used to quantify our pre-existing belief in the expected distribution of feature gradients and error terms, and then these beliefs are

updated by the terrain and network data we see, to produce an estimate of the true heights. Section III-A formalizes the error model informally introduced in the introduction. Section III-B focuses on the structure of the likelihood definition used, while Section III-C explores choice of form for the prior. Section III-D discusses implementation details, and Section III-E the choice of parameters and methods for validating the model.

A. ERROR MODEL

Equation (2) defines the normal GIS drape operation, computing height z_i for a vertex i with coordinates x_i, y_i on a digital elevation model **DEM**. In the current study INTERPOLATE is a bilinear interpolation function as used in naïve drape [1], however without loss of generality we can consider other forms of interpolation such as bicubic, Kriging, or DEM super-resolution.

$$z_i^{\text{draped}} = \text{INTERPOLATE}(\mathbf{DEM}, x_i, y_i) \tag{2}$$

The true heights can be related to z_i^{draped} as follows:

$$z_i^{\text{true}} = z_i^{\text{draped}} + \epsilon_i \tag{3}$$

where ϵ_i is composed of the following sources of error: (1) errors arising from the measurement of the DEM; (2) errors arising from measurement of horizontal positions of the draped features; (3) sampling error, attributable to the loss of information that occurs because the DEM does not record the elevation of all points but only a sampled grid; (4) model misspecification error inherent in the assumption that elevation of a feature can always be correctly recovered by draping a DEM (that assumption being violated in the case of vertical or overhanging cliffs where the true terrain has multiple elevations for the same x, y).

The current model attempts a better estimate of z_i^{true} by applying a vertical offset Δz_i to z_i^{draped}:

$$z_i^{\text{estimated}} = z_i^{\text{draped}} + \Delta z_i \tag{4}$$

Hutchinson [37] presents a simple model for the standard deviation of sampling error when interpolating a DEM tile of uniform gradient, related to the maximum change in elevation across the tile Δz^{tile} by assuming a point sampled at random from the tile:

$$\sigma_z = \Delta z^{\text{tile}} / \sqrt{12} \tag{5}$$

More recent research points at the possibility of developing more sophisticated models of sampling error σ_z by dispensing with the assumption of a uniform DEM tile [38]. However, noting that in the current model Δz_i is intended to correct for all four sources of error in any case, for simplicity Δz_i is taken to be drawn from a normal distribution with heteroscedastic variance from the Hutchinson sampling error model. A homoscedastic measurement error $\sigma_z^{\text{measurement}}$ is also added:

$$\Delta z_i \sim \mathcal{N}(0, (\frac{1}{\sqrt{12}} \Delta z_i^{\text{tile}} \sigma_{z\prime})^2 + \sigma_z^{\text{measurement}^2}) \tag{6}$$

where $\sigma_{z\prime}$ is an empirically determined scaling factor that allows the model to be adapted to the task at hand, in this case draping roads over terrain. In the current study it is assumed that sampling errors will dominate, hence $\sigma_z^{\text{measurement}}$ is set to zero.

B. LIKELIHOOD MODEL

The unknown parameters Δz_i in (4) are computed by maximum a posteriori likelihood estimation based on the input data $\mathbf{x}$, $\mathbf{y}$, **ADJ**, **DEM**. As the estimated maximum likelihood elevation of any vertex is dependent on the estimated elevations of neighbouring vertices on any polylines it is part of, with this dependency propagating through the network, all Δz_i must be estimated simultaneously as $\Delta \mathbf{z}\prime$:

$$\Delta \mathbf{z}\prime = \text{argmax}_{\Delta \mathbf{z}} L(\Delta \mathbf{z} | \mathbf{x}, \mathbf{y}, \mathbf{ADJ}, \mathbf{DEM}) \tag{7}$$

In addition to the DEM, the posterior likelihood $L(\Delta \mathbf{z} | \mathbf{x}, \mathbf{y}, \mathbf{ADJ}, \mathbf{DEM})$ is based on data consisting of a set of 2-d polylines for which we wish to estimate elevation. In practice these are represented as

- arrays of vertex coordinates $\mathbf{x}, \mathbf{y}$. The $\mathbf{x}, \mathbf{y}$ pairs are unique: even though an individual vertex may appear on multiple different polylines, it appears only once in these arrays which contain all point coordinates without reference to their containing polylines.
- a vertex adjacency matrix **ADJ** indicating which vertices are neighbours on at least one polyline. Note that where polylines share vertices they are considered to be connected in the network. This is reflected by the shared vertices having additional neighbours in the adjacency matrix.

The posterior likelihood function L of a particular value of all vertex offset parameters $\Delta \mathbf{z}$ given the input data, is the product of the likelihood of the resulting gradients (based on our prior beliefs about gradients), and the prior likelihood of the sampling and random errors that $\Delta \mathbf{z}$ would correct:

$$\begin{aligned} &L(\Delta \mathbf{z} | \mathbf{x}, \mathbf{y}, \mathbf{ADJ}, \mathbf{DEM}) \\ &\quad = L_{\text{gradients}}(\mathbf{x}, \mathbf{y}, \mathbf{z}^{\mathbf{draped}} + \Delta \mathbf{z}, \mathbf{ADJ}) \\ &\quad \times L_{\text{offset}}\,(\Delta \mathbf{z} | \mathbf{x}, \mathbf{y}, \mathbf{DEM}) \end{aligned} \tag{8}$$

where $L_{\text{gradients}}(\mathbf{x}, \mathbf{y}, \mathbf{z}^{\mathbf{draped}} + \Delta \mathbf{z}, \mathbf{ADJ})$ is the likelihood of the gradients $\mathbf{g}$ resulting from elevations $\mathbf{z}$ (specifically in this case $\mathbf{z} = \mathbf{z}^{\mathbf{draped}} + \Delta \mathbf{z}$) between all adjacent points on the road network data, given our prior beliefs about the gradients a road network should have, encoded in the probability density function of the grade distribution $\text{PDF}_G(g | \sigma_G, \alpha)$ (the distribution G is discussed in the following section III-C):

$$\begin{aligned} &L_{\text{gradients}}(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{ADJ}) \\ &\quad = \prod_{\texttt{adjacent}} \text{PDF}_G(\text{abs}(g(n_1, n_2)) | \sigma_G, \alpha)^{d_{n_1, n_2}} \end{aligned} \tag{9}$$

where $\prod_{\texttt{adjacent}}$ denotes the product over all adjacent vertices $n_1, n_2 . \mathbf{ADJ}(n_1, n_2) = 1$, and g is the grade arising

from $\mathbf{z} = \mathbf{z}^{\mathbf{draped}} + \Delta\mathbf{z}$,

$$g(n_1, n_2) = (z_{n_2} - z_{n_1})/d_{n_1,n_2} \tag{10}$$

with d_{n_1,n_2} being the horizontal distance between vertices n_1 and n_2:

$$d_{n_1,n_2} = \sqrt{(x_{n_1} - x_{n_2})^2 + (y_{n_1} - y_{n_2})^2} \tag{11}$$

Note that gradient probability densities are weighted per unit length of network as we require likelihood to be unaffected by the modifiable unit problem. That is, we should be indifferent to how many vertices, and with what spacing, are used to digitize any given road feature. Length weighting as applied in equation (9) achieves this by ensuring that for any two line segments a, b of equal gradient g and lengths l_a, l_b, the combined likelihood of both line segments is the same as that of a single segment ab of gradient g and length $a + b$:

$$\begin{aligned}\forall a, b :\\ L_{\text{gradients}}(x_a, y_a, z_a, \text{ADJ}_a) L_{\text{gradients}}(x_b, y_b, z_b, \text{ADJ}_b)\\ = L_{\text{gradients}}(x_{ab}, y_{ab}, z_{ab}, \text{ADJ}_{ab})\end{aligned} \tag{12}$$

Likewise, $L_{\text{offset}}\ (\Delta\mathbf{z}|\mathbf{x}, \mathbf{y}, \mathbf{DEM})$ is the prior likelihood of all offsets $\Delta\mathbf{z}$ together given the combined sampling, model and interpolation error models:

$$L_{\text{offset}}\ (\Delta\mathbf{z}|\mathbf{x}, \mathbf{y}, \mathbf{DEM}) = \prod_n L_{\text{offset}}(z_n|x_n, y_n, \mathbf{DEM}) \tag{13}$$

where $L_{\text{offset}}(z_n|x_n, y_n, \mathbf{DEM})$ is the likelihood of the offset for a single vertex

$$L_{\text{offset}}(z_n|x_n, y_n, \mathbf{DEM}) = \text{PDF}_{\Delta z}(\Delta z_n|\sigma_{z'}, \Delta z_n^{\text{tile}})^{d_n} \tag{14}$$

based on the probability density $\text{PDF}_{\Delta z}(\Delta z_n|\sigma_{z'}, \Delta z_n^{\text{tile}})$ of the distribution of the error term Δz defined above in (6), and weighted by d_n, the total length of road segments closest on the network to vertex n (equivalent to half of the total length of road segments attached to vertex n, as each segment is shared between two vertices):

$$d_n = \frac{1}{2} \sum_{i.\mathbf{ADJ}(n,i)=1} d_{n,i} \tag{15}$$

For vertices on bridges and tunnels (henceforth brunels) the height is not estimated by an offset from the DEM, but rather is estimated directly, allowing them to 'decouple' from the terrain. In the case of a brunel consisting of a single link, heights could be computed by linear interpolation between its endpoints (the endpoints being estimated in turn according to the likelihood defined in (8)). However this does not address the general case in which junctions between roads may occur on/in a brunel, nor does it capture the likelihood of the brunel's gradient. The unknown brunel vertex heights are therefore estimated directly, by adding special cases to the above definitions: in (8), $\mathbf{z}^{\mathbf{draped}} = 0$, and in (14), $L_{\text{offset}}(z_n|x_n, y_n, \mathbf{DEM}) = 1$ for decoupled vertices only. This allows decoupled vertices to be estimated simultaneously with normal vertices, but for these vertices we interpret Δz to represent the absolute elevation and maximise only the likelihood of gradients. The offset from the DEM is not directly considered but the DEM will influence these points indirectly via the estimates of Δz where brunels join ordinary links at their endpoints.

Finally the model allows for fixed features for which the height is set directly by input data. In these cases nothing is estimated for the fixed vertex itself, but gradients and angles of segments joining to fixed vertices are included in likelihood computations for their neighbours.

C. FORM OF PRIORS

Real world distribution of road gradients is close to exponential, however, the exponential prior on gradients has undesirable properties if the aim is to smooth elevation data. Here we consider these properties. An exact exponential grade prior defined per unit length of road network, is indifferent to the shape of a slope profile, provided the slope is conservative - that is, for a road that eventually rises by a given height h then all slope profiles that eventually rise to h are equally likely, so long as they never lose height (Fig. 5). Conversely, for a prior with downward curving log likelihood (in contrast to the 'straight line' log likelihood of the exponential distribution), uniform slopes are preferred (Fig. 6).

As discontinuity in interpolated slopes is potentially undesirable (particularly so in the case of cycling models), the prior distribution for grades G is thus designed to give control over both steepness of slopes, and uniformity of slopes, by parameterizing the space between an exponential and half-normal distribution (with scale parameter σ_G defining the mean steepness of roads):

$$G \sim \begin{cases} \text{Exp}(\sigma_G), & \text{if } \alpha = 0 \\ \mathcal{N}_{1/2}(\sigma_G), & \text{if } \alpha = 1 \end{cases} \tag{16}$$

where α specifies the mixture of distributions, with $\alpha = 0$ giving exponential, and $\alpha = 1$ giving half normal. Interpolation for any $0 \leq \alpha \leq 1$ is achieved efficiently in computing log probability density for slope grade x as

$$\log(\text{PDF}_G(x)) = k - \frac{1-\alpha}{\sigma_G} x - \frac{\alpha}{\pi\sigma_G^2} x^2 \tag{17}$$

(with normalizing constant k set to zero in practice). Note that this prior is defined per unit length of road network to eliminate dependency on the number of vertices used to encode any given feature.

D. IMPLEMENTATION

The proposed approach uses local likelihood maximization, conducted by minimization of the negative log likelihood function using the L-BFGS-B algorithm [39]. The implementation used is part of Scipy [40]. L-BFGS-B runs much faster if provided with the gradient of the function to be minimized as well as the function itself, as it need not then use p samples of the function to compute the gradient at each step, where

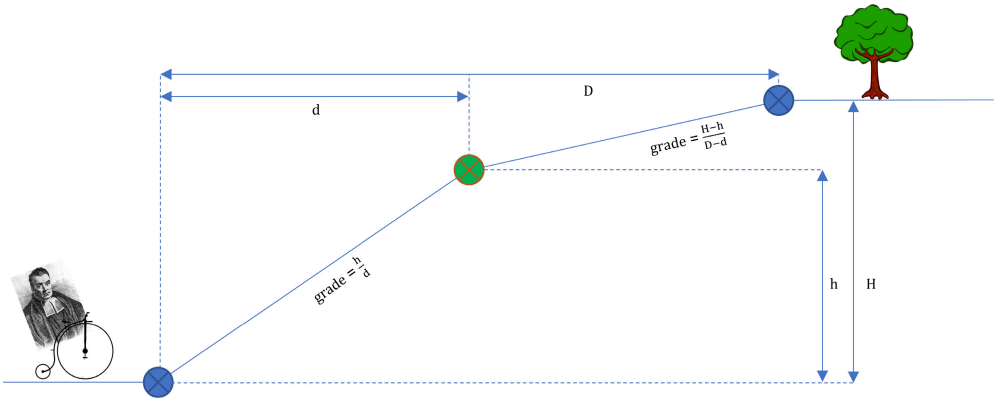

FIGURE 5. If applying an exponential prior, the length weighted grade likelihood of any conservative slope is $(\lambda e^{-\lambda h/d})^d (\lambda e^{-\lambda \frac{H-h}{D-d}})^{D-d} = \lambda^D e^{-\lambda H}$, i.e. depends only on the total height gained and total distance (λ is inverse scale of the exponential distribution, all other terms defined in figure).

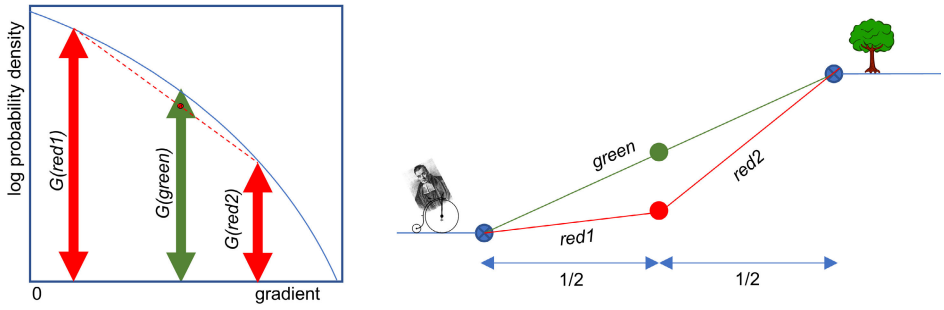

FIGURE 6. Without loss of generality, consider two slopes covering unit horizontal distance, each divided into two segments as shown. The total likelihood for the non-uniform (red) slope is the geometric mean of the priors of its component parts, hence the arithmetic mean of log likelihoods. If the log prior has negative curvature (unlike the exponential log prior which is linear), a uniform slope (green) is always deemed more likely.

p is the number of parameters. To this end, differentiable programming is used to define a negative log likelihood gradient function. This tracks the operations performed during computation of the function itself, then backtraces them to compute the gradient using the chain rule. The implementation used is from the PyTorch framework [41] which, although initially developed for neural networks, is capable of automatic differentiation tasks in general.

The steps of the implementation can therefore be summarized as follows:

1) Insert extra vertices on polylines where they intersect DEM cell boundaries
2) Break down polylines to individual vertices (each classified as estimated or decoupled) and build network adjacency matrix of vertices
3) Define parameter vector $\Delta\mathbf{z}$ for likelihood model, which is interpreted as
   a) offsets from draped elevation for ordinary vertices, all initialized to 0
   b) direct estimates of z for decoupled vertices, initialised to an inverse network distance weighted mean of their respective naïvely draped brunel endpoint elevations (see end of section III-D)
4) Perform naïve interpolation of DEM to obtain $\mathbf{z_{draped}}$ for all ordinary vertices
5) Define function to be minimized f(params) as negative posterior log likelihood as defined in Section III-B
6) Define gradient function f'(params) using automatic differentiation of f (enable gradient tracking on params, call f(params) and backtrace)
7) Call L-BFGS-B minimizer to find optimal $\Delta\mathbf{z}$ which minimizes f given f, f'
8) Assign elevations of $\mathbf{z_{draped}} + \Delta\mathbf{z}$ to all ordinary vertices
9) Assign elevations of $\Delta\mathbf{z}$ to all decoupled vertices
10) Rebuild polylines from vertices and adjacency matrix

The question arises of how frequently, on the feature polylines, we should sample an elevation from the interpolation model. Within a single terrain tile, in most cases, smaller sampling distances on the polyline will tend to produce vertex elevations which are close to the linear interpolation of elevations derived from larger sampling distances (moreso under the influence of priors which encourage smooth

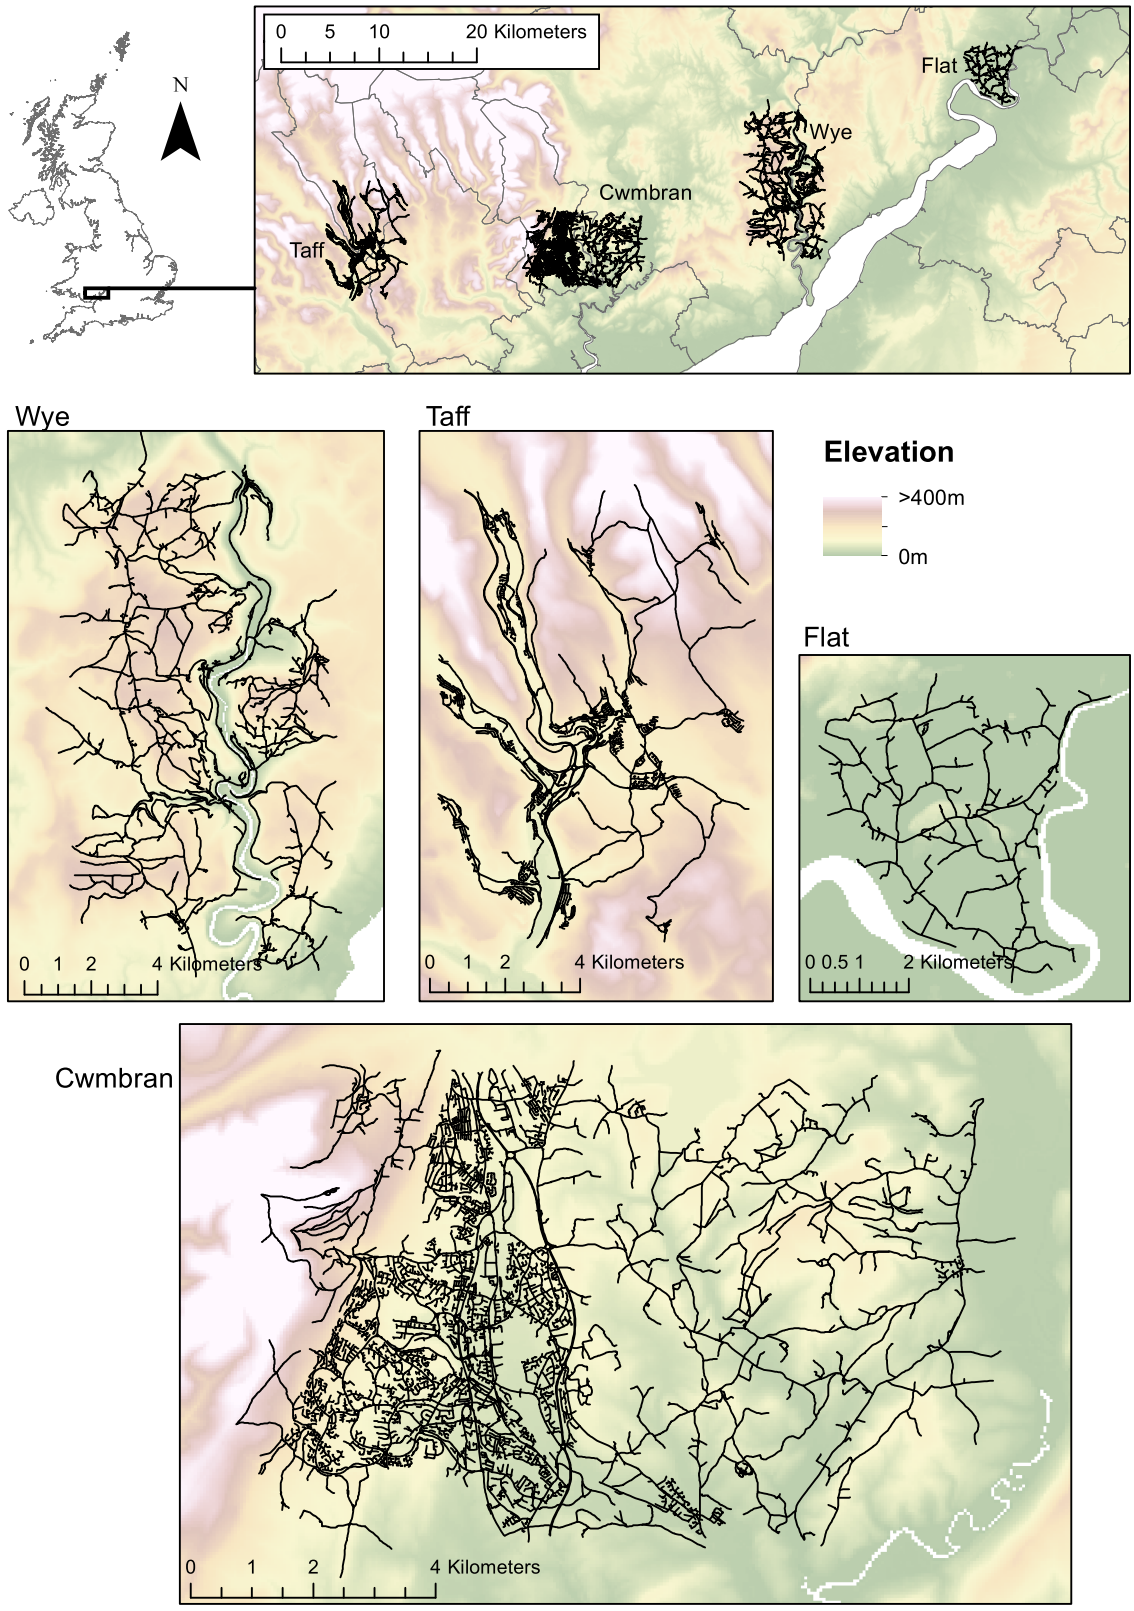

FIGURE 7. Maps showing the four road networks of test study areas. Boundaries shown in top figure are UK local authorities (Office for National Statistics licensed under the Open Government Licence v.3.0). Contains OS data © Crown copyright and database right 2021-2024.

changes in elevation). As the polyline data model itself implies straight line segments between vertices, the sampling distance therefore makes little difference, provided the following are sampled: (i) existing vertices in the polyline as it may change direction; (ii) points where a polyline crosses a DEM cell boundary, where it is possible that the DEM gradient will exhibit a large change. The only exception to this could be where DEM tiles deviate substantially from being planar, in which case finer sampling distances could better reflect the change in gradient within the tile. However for the cases we are trying to correct, we cannot expect the interior of a DEM tile to be well represented interpolation anyway. The approach taken for sampling roads, therefore, is only to add extra vertices to input features where they cross cell boundaries in the DEM (unless a vertex already exists within 1/100th the DEM resolution of the cell boundary). DEM gradient computation is smoothed for vertices close to cell boundaries, to ensure an average gradient of both cells is used. Draped elevations are not smoothed.

The optimizer's starting guess used for ordinary vertices ($\Delta z = 0$) corresponds to a naïve drape operation, so we can guarantee that results will be at least as good as naïve drape. To derive a starting guess for brunels, note that for any $\sigma_G > 0$ and $0 < \alpha \leq 1$, the maximum likelihood elevation profile of a brunel with two fixed endpoints corresponds to a linear interpolation (straight line) between the endpoint elevations. To accommodate decoupled junctions this is generalized to > 2 endpoints using an inverse network distance weighted mean of endpoint elevations (interpolated from terrain as above with $\Delta z = 0$). In the general case this does not equate to maximum likelihood but serves as a suitable initial guess.

E. CHOICE OF PARAMETERS AND VALIDATION

The prior scales and shapes σ_G, $\sigma_Z\prime$, α of the Bayesian model can be calibrated by comparing its output against known good height data in e.g. a Bayesian Monte Carlo framework. However, this approach is of little use for end users of the model who don't have access to sufficiently good height data. Instead, this section reports tests of a range of the unitless metaparameters $\sigma_Z\prime$, α on four study areas with the aim of finding values likely to be suited to general use. It is assumed that mean slope used to determine the scale parameter σ_G can be measured based on a simple drape of road network data over terrain: although this draped data will by definition include the elevation errors we aim to correct, these errors are outliers which have little effect on the mean, so we can estimate σ_G using uncorrected data.

To assess the accuracy of the model's output it is compared to elevation changes recorded on the Ordnance Survey (OS) Mastermap Highways Network. The OS derive such estimates by combining a drape of the road network (ranging from 0.4 to 4.1m planimetric accuracy) over their own detailed height content, for which the resolution is unspecified, though assumed at least equal to their best publicly released model which has 5m DEM resolution. The OS also supplement this with additional height data where deemed necessary e.g. for features which traverse hillsides. Features also receive extra annotation where height data is "expected to be low quality" [42].

TABLE 1. Performance characteristics for key optimizer runs.

Dataset	α	$\sigma_z\prime$	$\frac{\Delta loglik}{param}$	Iterations	Time (s)	Links	Params
Wye 50m	0	0.25	1.70	1,816	13.7	1,400	24,114
	0	0.4	2.40	2,483	19.0		
	0.5	0.25	1.99	3,845	31.0		
	0.5	0.4	2.70	1,832	14.9		
	1	0.25	2.65	1,420	11.8		
	1	0.4	3.43	1,402	11.7		
Cwmbran	0	0.25	0.80	1,394	21.0	5,707	49,240
	0	0.4	1.23	1,400	21.2		
	0.5	0.25	0.66	1,282	21.0		
	0.5	0.4	1.03	1,330	21.4		
	1	0.25	0.71	1,734	29.1		
	1	0.4	1.05	1,661	27.9		
Taff	0	0.25	1.90	2,413	17.0	1,712	19,908
	0	0.4	2.75	2,568	18.1		
	0.5	0.25	1.74	2,845	19.2		
	0.5	0.4	2.51	1,959	13.3		
	1	0.25	1.97	3,055	21.2		
	1	0.4	2.73	2,616	18.2		
Flat	0	0.25	0.28	1,916	4.3	299	5,311
	0	0.4	0.48	5,200	11.7		
	0.5	0.25	0.25	2,659	6.3		
	0.5	0.4	0.40	2,393	5.6		
	1	0.25	0.29	1,930	4.7		
	1	0.4	0.41	1,723	4.2		
Wye 5m	0	0.25	0.09	1,982	62.5	1,400	97,659
	0	0.4	0.14	2,220	69.2		
	0.5	0.25	0.12	3,476	123.0		
	0.5	0.4	0.18	2,338	83.6		
	1	0.25	0.19	3,825	136.0		
	1	0.4	0.27	3,428	125.0		

All runs converged with no further significant change to log likelihood. CPU is single core of an Intel(R) Xeon(R) Gold 6148 @ 2.40GHz. 'Links' is the number of links in each network dataset, 'Params' is number of parameters in the model (equivalent to number of estimated vertices in the network).

Test areas are shown in Figure 7. The primary area is a small region of the Lower Wye Valley between Monmouth and Chepstow, which due to its topography contains numerous roads traversing steep hillsides likely to generate problematic draped elevations. OS MasterMap metadata shows that supplemental height data has been used to compute elevations and none of these are designated low quality, with all tested roads having 1.1m planimetric accuracy or better. It is assumed that MasterMap elevation changes for the test data are accurate for the purpose of testing the Bayesian model, which uses as input the much lower resolution OS Terrain 50 DEM [7]: the model is working with terrain data of at least 10x lower resolution than the validation data. The test network is extracted from OS MasterMap Highways [42], and includes two bridges exceeding 50m length which are marked for decoupling.

Three further study areas are included for validation; in each case using data from the same sources. Cwmbran is chosen as an urban hillside setting, in contrast to the primary test area which is rural. The Taff valley is another steep sided valley albeit with a dual carriageway running along its length (in addition to local and residential streets); the dual

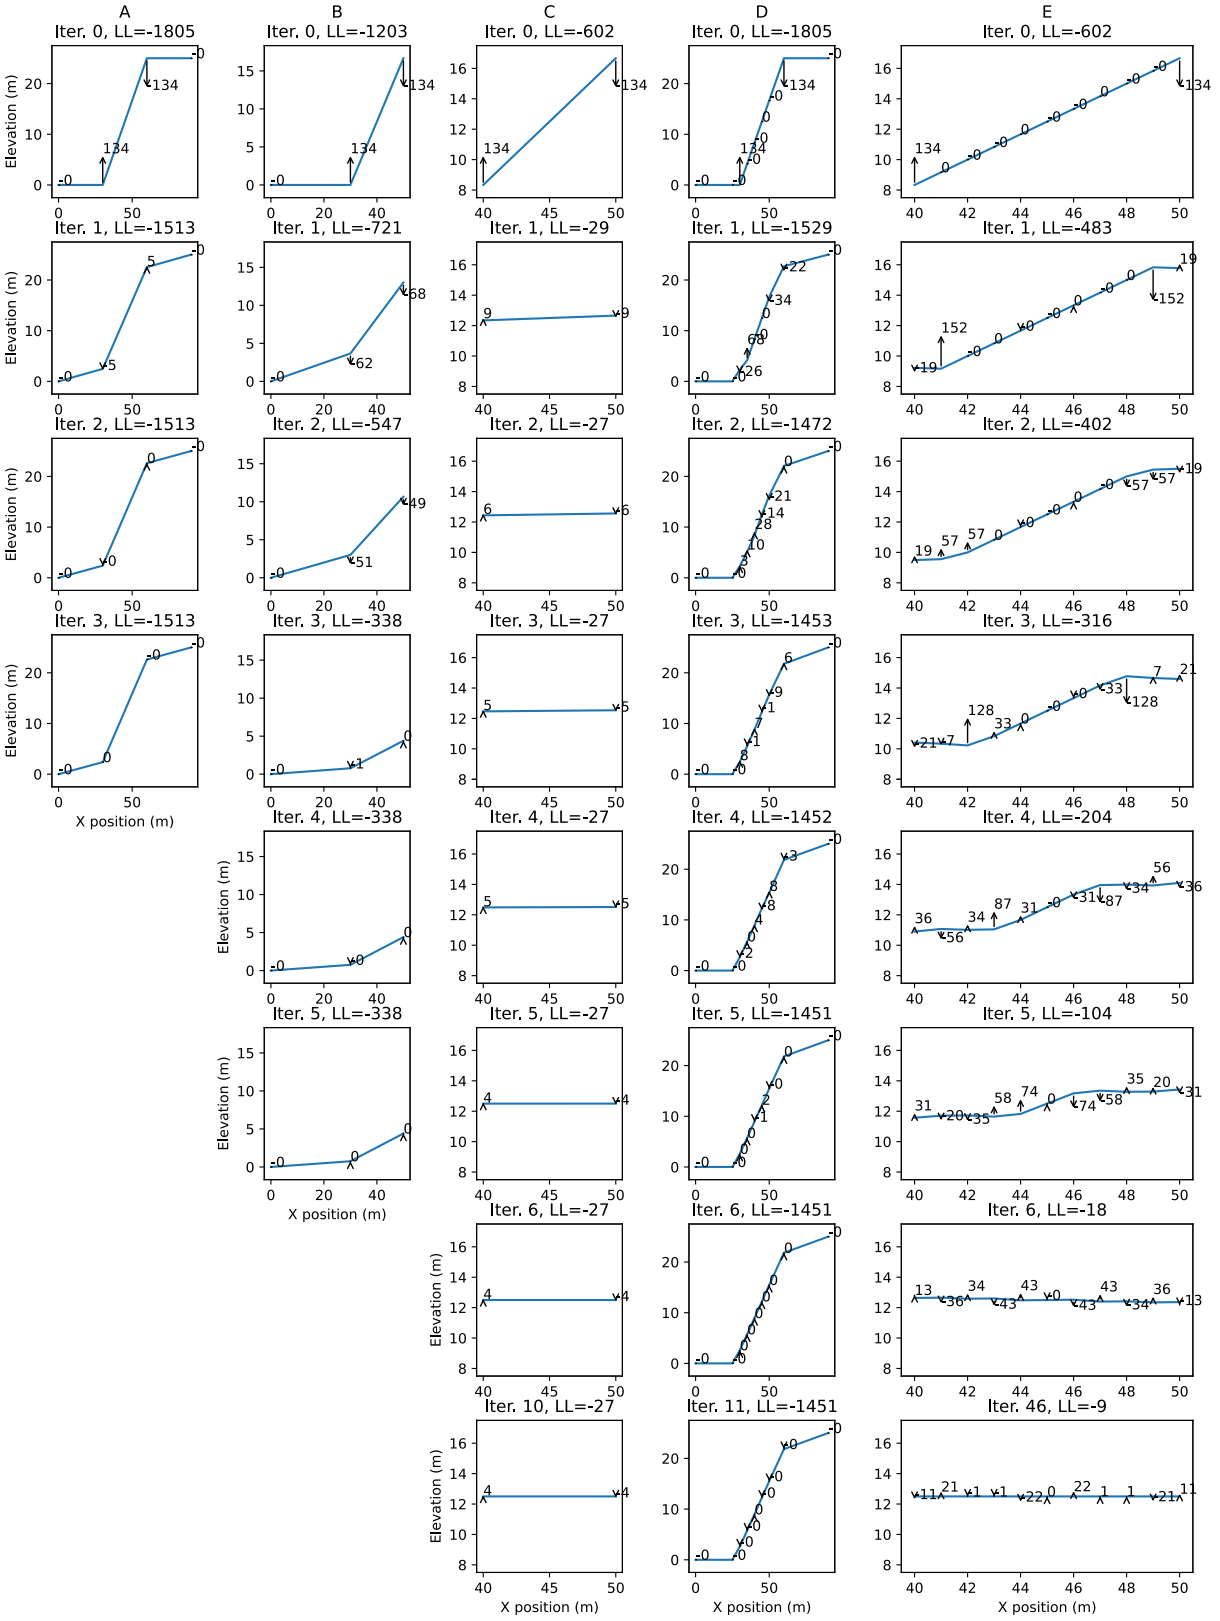

FIGURE 8. Convergence of elevations on simulated test cases, showing a single road in each column ($\alpha = 0.5, \sigma_{Z'} = 0.25, \sigma_G = 2.66$). Iteration 0 (top row) is equivalent to a naïve drape of the terrain in each case. Iter = iteration; LL = log likelihood. Arrows show log likelihood partial derivative with respect to each elevation point.

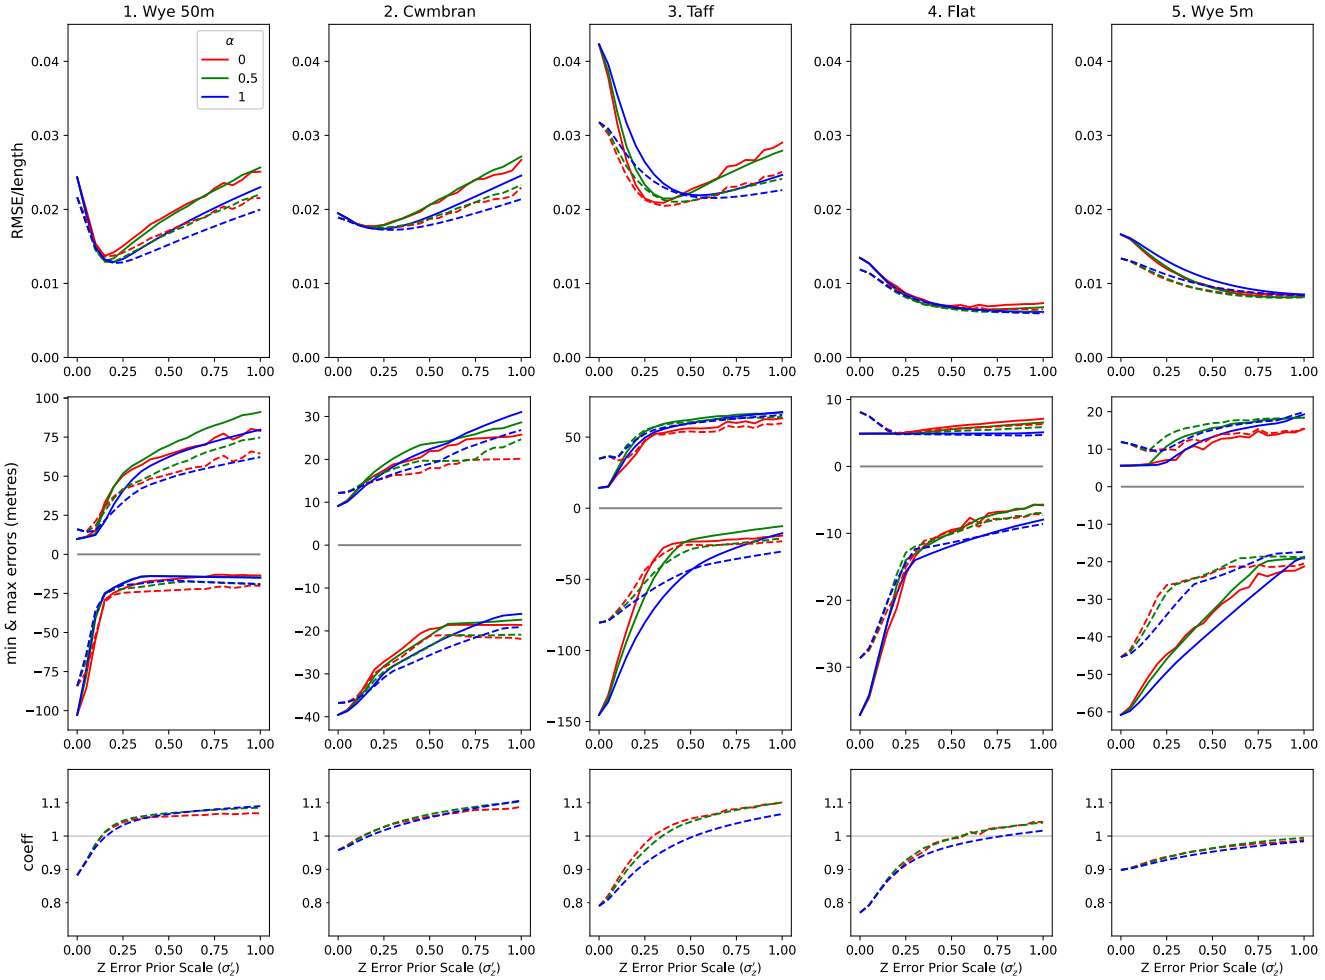

FIGURE 9. Performance of the Bayesian model on each geographical region, as a function of z error prior scale $0 \leq \sigma_{z'} \leq 1$ and for three values of α (0, 0.5, 1). Solid lines show direct application of the model; dotted lines show outputs with bias corrected by a further regression model. Leftmost points on each plot ($\sigma_{z'} = 0$) are equivalent to a naïve drape. Note that minimum and maximum error plots have different scale for each geographical region.

carriageway crosses large viaducts. A much flatter region near to Westbury-on-Severn, Gloucestershire, is included for comparison. Finally to show performance of the model where higher resolution terrain data is already available, the analysis of the primary area (Lower Wye Valley) is repeated using a 5m (OS Terrain 5) DEM in place of the 50m DEM.

In evaluating outputs from different sets of parameters, the primary measures for comparison are RMS error per unit length (to allow comparison between geographic regions with different total road length), alongside maximum and minimum errors (differences between the true elevation change and raw outputs from Bayesian model). Bearing in mind that it is possible to over-smooth the data, but that this may be desirable in some circumstances (where reducing minimum errors is more important and bias introduced by over-smoothing can be corrected), a second set of outputs is shown in which a regression model has been used to remove systematic bias from over-smoothing:

$$\Delta E_{\mathrm{TRUE}} = k + \beta \Delta E_{\mathrm{BD}} + \epsilon \tag{18}$$

where ΔE_{TRUE} is the true cumulative absolute elevation change (i.e., elevation gains and losses do not cancel one another), and ΔE_{BD} is the elevation change derived from the Bayesian model. The error term is heteroscedastic so the model is fitted using the HC3 estimator [43]. Similar outputs are reported from this regression model (RMS per unit length, minimum and maximum errors) and also the fitted coefficient value β to quantify the bias.

In addition, draped profiles for two manually selected problematic links are displayed (both of which are outliers in terms of spurious elevation change), to show the effect of different prior combinations on the shape of output networks. The links are:

1) A riverside track in the primary data set for which OS elevation data is available. This is an extremely difficult link to drape correctly as it traverses steep terrain and for many of the tested models is the outlier with the largest overpredicted elevation change in the study area. The link is approximately 1km in length, which allows us to visualise the qualitative

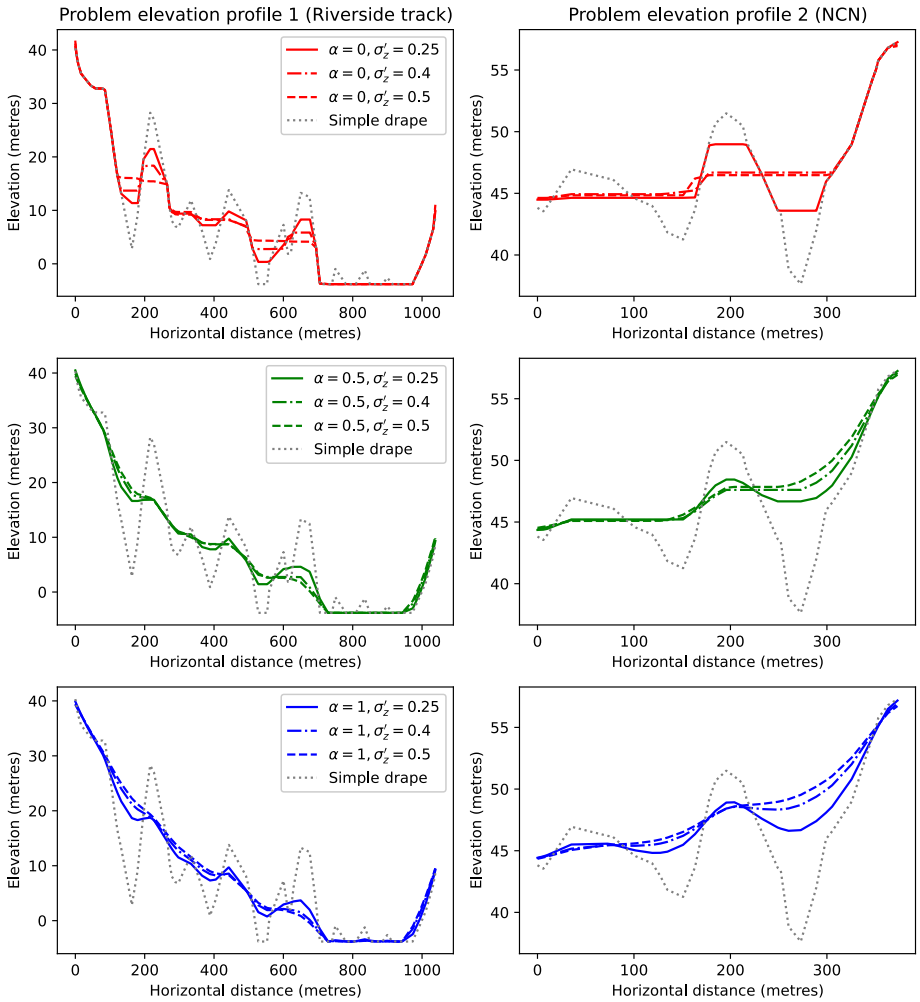

FIGURE 10. Elevation profiles output by the Bayesian model for manually selected problematic links.

behaviour of different models, as well as ensuring the central parts of the link are relatively unaffected by estimated elevations of neighbouring links. Ground truth inspection of this link showed it to consist of a level central section with conservative slopes at either end.

2) The National Cycle Network (NCN) link shown in Fig. 2(c) which caused issues in modelling by Lovelace and Cooper [24]. The OS does not provide elevation data for this link, however, ground truth inspection reveals it to have conservative gradient i.e. there are no "bumps" in the true elevation profile.

IV. RESULTS AND DISCUSSION

This section first presents some simulated test cases, then discusses the fit and sensitivity of a variety of real world models as a function of their prior metaparameters $\sigma_Z\prime, \alpha$.

Convergence of the simulated test cases is shown in Fig. 8. All cases converge with no further significant change to log likelihood. Column A shows the case of a road climbing a slope with flat terrain tiles above and below. The endpoints of the road are thus fixed in elevation as their $\sigma_Z = 0$, so limited smoothing of the slope occurs. Column B shows half of the same road; as the rightmost endpoint is now on a sloping terrain tile, its elevation is not fixed and is adjusted downwards to reduce elevation change to the fixed leftmost endpoint. Column C shows the case where both endpoints are on sloping tiles; both are adjusted to remove any gradient. Column D shows a similar situation to column A, albeit with more vertices used to encode the leftmost half of the road. An initial log likelihood equal to that in column A shows the invariance of log likelihood on different encodings of the same shape. However the extra vertices on flat ground fix the elevation at those points, so the final estimated profile is slightly different. Column E shows the same situation as column C, albeit with more vertices along the entire length of the road. It can be seen that elevation changes must propagate inwards from the line's endpoints (the only places for which partial derivative of the log likelihood is initially nonzero) until the elevation profile becomes flat.

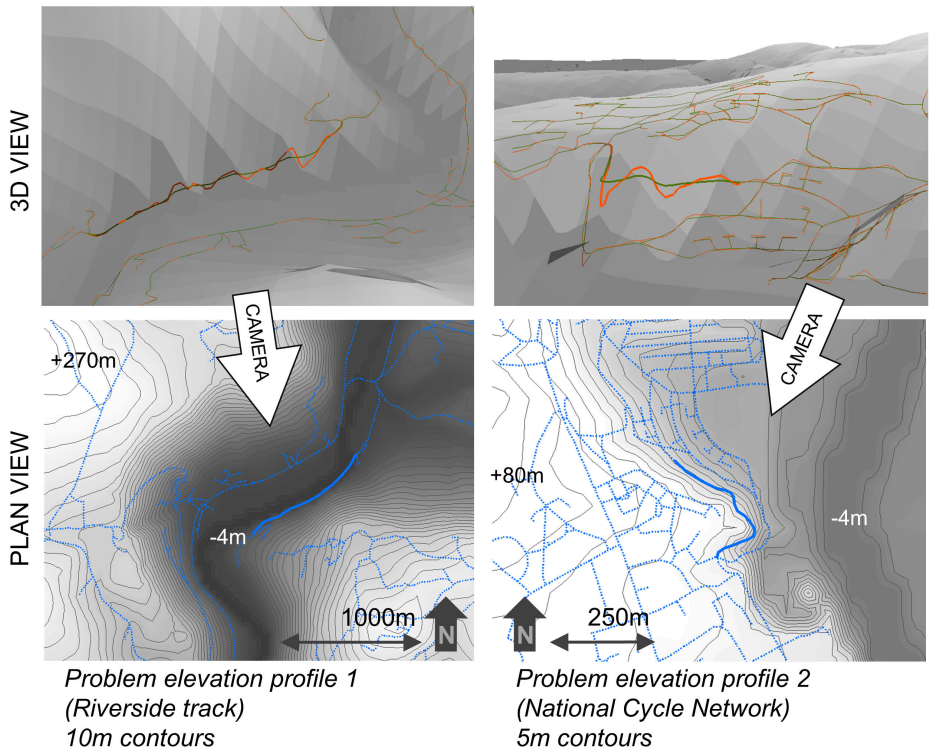

FIGURE 11. Further views of the two manually selected problematic outlier links: 2-d surrounding road/path network (blue), and 3-d views of the naïve drape (orange) and Bayesian correction (green, using the recommended prior $\sigma_{z'} = 0.25$, $\alpha = 0.5$). Vertical exaggeration of 5 is used for the 3-d views. Greyscale gradient represents height in plan view, height with added hillshade in 3-d view.

Results for all real terrain models together are plotted in Fig. 9. Elevation profiles output by each model for the two manually selected outlier links are shown in Fig. 10, and further views of these links are shown in Fig. 11. Table 1 shows dataset characteristics and optimizer performance.

Looking at the difference in performance of the model across all five study areas, we see that choosing an appropriate value of $\sigma_{z'}$ reduces absolute error and minimum errors (overpredicted elevation change), while increasing maximum errors to a lesser extent (therefore giving an improvement in errors overall). Substantial improvement in minimum errors is seen even in the flat study area, and high resolution terrain model.

The mean slope measured from naïve drape of road network in the primary study area was 2.66° and this value is used in all model runs.

Considering the choice of α, Fig. 10 (top plots in red) shows as expected that setting $\alpha = 0$ produces link profiles consisting of level sections with rapid changes in elevation between them. $\alpha = 1$ produces the smoothest link profiles, but at the cost of substantially worse mean error in two of the five study areas (Fig. 9). $\alpha = 0.5$ is therefore recommended for general usage as it strikes a balance between these two extremes.

Across all five datasets, without correcting Bayesian model outputs for bias, setting $\sigma_{z'} = 0.25$ (with $\alpha = 0.5$) reduces absolute errors in elevation, compared to a naïve drape, by an average of 32%. Minimum errors reduce in magnitude by an average of 51%. Maximum errors increase, but to a lesser extent (on average only 41% of the corresponding absolute reduction in minimum errors for the same dataset). $\sigma_{z'} = 0.25$ is therefore recommended for use in cases where bias must be avoided.

The alternative, if the user wishes to achieve greater smoothing at the cost of increased bias, is to set a higher $\sigma_{z'}$. Using $\sigma_{z'} = 0.4$, then correcting Bayesian model outputs for bias (18), reduces absolute error by an average (across all 5 datasets) of 40%. Minimum errors reduce in magnitude by an average of 63%. Again, maximum errors increase, but to a lesser extent (on average 33% of the corresponding absolute reduction in minimum errors for the same dataset).

Testing the model on hilly regions shows broadly comparable results in both urban and rural settings. Initially it may seem surprising that a substantial improvement in accuracy is seen even in the flat region. Inspection of residuals reveals that this is due to a single highway cutting alongside a minor hill (Fig. 12). Features of similar height, such as railway embankments, can occur even in regions typically thought of as flat. The errors introduced by such embankments are not limited by the height of the embankment, as a naïvely draped polyline can appear to climb and descend the embankment

FIGURE 12. The embankment corrected by the Bayesian model in the flat region (© Google, 2023).

TABLE 2. Comparing Bayesian correction of naïve drape, to naïve drape and Kriging.

Dataset	Model	RMSE/length	Min error (m)	Max error (m)
Wye 50m	Krige	0.018	-57	10
	Naïve	0.024	-103	10
	Bayesian	0.014	-18	52
Cwmbran	Krige	0.016	-25	9
	Naïve	0.019	-40	9
	Bayesian	0.018	-28	19
Taff	Krige	0.033	-94	14
	Naïve	0.042	-145	14
	Bayesian	0.023	-62	54
Flat	Krige	0.011	-30	5
	Naïve	0.013	-37	5
	Bayesian	0.008	-14	5

Kriging is conducted using parameters derived from Gaussian semivariogram fitting. Bayesian model has recommended universal metaparameters $\sigma_z\prime = 0.25, \alpha = 0.5$.

multiple times. It is therefore recommended to use the Bayesian model even on relatively flat areas.

The high resolution terrain model shows (as expected) reduced sampling error for the naïve drape, but the Bayesian model still improves on this substantially. Notably, in both the flat and high-resolution models it is possible to choose a higher $\sigma_z\prime$ (up to $\sigma_z\prime = 1$) without any increase in absolute error due to over-smoothing.

For comparison, Table 2 shows results from Kriging (with Gaussian semivariogram also used by [38]), and naïve drape, alongside Bayesian correction of naïve drape. In three out of four areas, Kriging shows improvement over naïve drape, but is still substantially outperformed by the Bayesian correction of naïve drape. This is especially true when considering reduction of minimum errors - the design goal of the Bayesian model. This is understandable as Kriging cannot make use of priors on the road network adjacency data to avoid spurious elevation change.

There is one exception in the case of Cwmbran, in which Kriging performs better overall. Inspection of the differences between model outputs in this area, reveals that Kriging is working better on a few long road segments with small (hence, not a priori unlikely) road gradients. In these cases an isolated DEM measurement causes spurious draped elevation change in naïve interpolation and hence the Bayesian output, while the Kriging output is smoother due to the influence of neighbouring terrain measurements. Most of the network length falls within an urban area, which contains fewer sudden changes in gradient requiring Bayesian correction. In this case therefore, the advantage of introducing gradient priors is outweighed by the more spatially dispersed sampling error model inherent in Kriging. Although the Bayesian model outperforms Kriging in the other areas, this anomaly suggests that use of naïve drape as an input for the Bayesian model may impact robustness.

V. CONCLUSION

This paper has presented a Bayesian approach to correcting heights on vector features interpolated over a terrain model. Using the naïve drape as a basis for correction, the Bayesian model has been shown to substantially improve the accuracy of estimated heights on road networks, in a variety of geographical areas of varying hilliness, and on DEMs of both 5m and 50m resolution. The Bayesian Drape code is open source and available both at https://github.com/fiftysevendegreesofrad/BayesianDrape, and in a reproducible capsule attached to this paper.

If using the model to improve height data to road networks in the absence of ground truth data for calibration, one of the following approaches to choosing priors is advised, based on results of the current study:

1) Determine an estimate of mean slope. Either
   a) use a simple drape of 2-d road network over terrain (complete with the errors the model aims to correct) to estimate this, or
   b) use the figure measured in this study (2.66°) which is likely to be an overestimate for regions

less hilly than the study area, however, the results show improvements from using this value on a flatter region in any case.

2) Determine the acceptable level of bias. Note there is no guarantee that the bias figures shown in the results of this study will be usable for correction of bias in other study areas, as they may depend on the mean hilliness of the terrain in the study area. In the absence of ground truth data the bias also cannot be corrected by linear regression as shown here. So if no data is available for bias correction, then lower bias is likely preferable. However, if a containing model is able to correct bias based on other data, then a consistent bias can be compensated: the containing model will adjust its coefficient for elevation change in the generalized cost equation. Therefore for applications such as a transport model calibrated to cyclist counts, a higher level of elevation bias may be allowed.
    a) If reduced bias is preferred, use $\sigma_z\prime = 0.25$, $\alpha = 0.5$.
    b) If stronger outlier correction at the expense of slightly higher bias is desired, use $\sigma_z\prime = 0.4$, $\alpha = 0.5$.

Although the approach above is satisfactory for the stated aim, as with any model, there is room for improvement.

A more empirical approach to ensuring continuity of slopes would be to introduce a prior for curvature, albeit with substantially greater computation and data requirements.

Although the proposed model is shown to improve on results from naïve drape, its robustness may be limited by its use of heights derived from naïve drape as a basis for Bayesian correction. Future implementations may therefore benefit from using Kriging output, as the input for the Bayesian model. This conclusion is supported by the fact that Kriging performed better on one of the four test areas. Further enhancements may be possible by combining the Kriging and Hutchinson error models. Alternatively, given the potential for improved DEM interpolation accuracy available from deep learning approaches, a DEM super-resolution model could provide the input for the Bayesian model.

Optimizer performance is likely to be adequate for most applications, however if performance is an issue then the option remains to use GPU acceleration. Initial exploration using PyTorch shows that GPUs provide approximately an 8x speed increase for datasets of half a million vertices, i.e. approximately 20x the size of the test data used in the current study. It is likely that this advantage will increase with larger datasets, and that performance could be further improved by reducing CPU-GPU communication.

It is likely that the Bayesian approach will be useful for interpolating other types of data over terrain models, for example, GPS traces recorded by fitness tracking devices. Broader applications such as improving existing river gradient estimations in hydrology [28] may also be possible using revised priors. There will be other data sets where knowledge on desired feature shapes can also be encoded as priors, for example, 3-d models of buildings in which heights of all entrances and exits must match terrain outside the building.

ACKNOWLEDGMENT

Study Based on Ordnance Survey data © Crown copyright and database right 2021. This research was undertaken using the supercomputing facilities at Cardiff University operated by Advanced Research Computing at Cardiff (ARCCA) on behalf of the Cardiff Supercomputing Facility and the HPC Wales and Supercomputing Wales (SCW) projects. The author would like to thank Thomas Green and Jose Criollo for supercomputing advice and support, Jonathan Gillard for advice on describing mathematics, and James Sample for introducing him to Bayesian modelling.

REFERENCES

[1] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, "Numerical recipes in C," in *The Art of Scientific Computing*, 2nd ed. Cambridge, U.K.: Univ. Press, Nov. 1992.
[2] *Geographic Resources Analysis Support System (GRASS GIS) Software*, GRASS Develop. Team, Open Source Geospatial Found., Chicago, IL, USA, 2020.
[3] *Interpolate Shape (3D Analyst)—ArcMap*, ESRI, Redlands, CA, USA, 2021.
[4] D. G. Krige, "A statistical approach to some basic mine valuation problems on the Witwatersrand," *J. Southern Afr. Inst. Mining Metall.*, vol. 52, no. 6, pp. 119–139, 1951.
[5] J. J. Ruiz-Lendínez, F. J. Ariza-López, J. F. Reinoso-Gordo, M. A. Ureña-Cámara, and F. J. Quesada-Real, "Deep learning methods applied to digital elevation models: State of the art," *Geocarto Int.*, vol. 38, no. 1, Dec. 2023, Art. no. 2252389.
[6] *Shuttle Radar Topography Mission (SRTM) Data*, Jet Propuls. Lab., CA, USA, 2014.
[7] *OS Terrain 50*, Ordnance Surv., Southampton, U.K., 2021.
[8] T. Tadono, H. Nagai, H. Ishida, F. Oda, S. Naito, K. Minakawa, and H. Iwamoto, "Initial validation of the 30 m-mesh global digital surface model generated by ALOS PRISM," in *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Prague, Czech Republic. Hanover, Germany: Institute of Photogrammetry and GeoInformation, Jul. 2016.
[9] *OS Terrain 5*, Ordnance Surv., Southampton, U.K., 2021.
[10] A. Aoun, J. Bjornstad, B. DuBose, M. Mitman, and M. Pelon, *Bicycle and Pedestrian Forecasting Tools: State of the Practice*. Washington, DC, USA: Federal Highway Administration, 2015.
[11] C. H. V. Cooper, I. Harvey, S. Orford, and A. J. F. Chiaradia, "Using multiple hybrid spatial design network analysis to predict longitudinal effect of a major city centre redevelopment on pedestrian flows," *Transportation*, vol. 48, no. 2, pp. 643–672, Dec. 2019.
[12] R. Ewing, G. Tian, J. Goates, M. Zhang, M. J. Greenwald, A. Joyce, J. Kircher, and W. Greene, "Varying influences of the built environment on household travel in 15 diverse regions of the United States," *Urban Stud.*, vol. 52, no. 13, pp. 2330–2348, Dec. 2014.
[13] J. B. Griswold, A. Medury, R. J. Schneider, D. Amos, A. Li, and O. Grembek, "A pedestrian exposure model for the California state highway system," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2673, no. 4, pp. 941–950, Apr. 2019.
[14] National Academies of Sciences, Engineering, and Medicine, *Estimating Bicycling and Walking for Planning and Project Development: A Guidebook*. Washington, DC, USA: The National Academies Press, 2014, doi: 10.17226/22330.
[15] F. Martinez-Gil, M.-F. I. Lozano, and F. Fernandez, "Modeling, evaluation, and scale on artificial pedestrians: A literature review," *ACM Comput. Surv.*, vol. 50, no. 5, p. 72, 2017.
[16] S. Munira and I. N. Sener, "Use of direct-demand modeling in estimating nonmotorized activity: A meta-analysis," Safety Through Disruption (Safe-D) Nat. Univ. Transp. Center (UTC) Program, Texas A&M Transp. Inst., Bryan, TX, USA, Tech. Rep. 01-003, 2017.

[17] S. Turner, *Synthesis of Methods for Estimating Pedestrian and Bicyclist Exposure to Risk at Area Wide Levels and on Specific Transportation Facilities*. Washington, DC, USA: Federal Highway Administration, Office of Safety, 2017.
[18] P. Arnesen, O. K. Malmin, and E. Dahl, "A forward Markov model for predicting bicycle speed," *Transportation*, vol. 47, no. 5, pp. 2415–2437, Oct. 2020.
[19] E. Y. C. Chan and C. H. V. Cooper, "Using road class as a replacement for predicted motorized traffic flow in spatial network models of cycling," *Sci. Rep.*, vol. 9, no. 1, pp. 1–12, Dec. 2019.
[20] C. H. V. Cooper, "Predictive spatial network analysis for high-resolution transport modeling, applied to cyclist flows, mode choice, and targeting investment," *Int. J. Sustain. Transp.*, vol. 12, no. 10, pp. 714–724, Nov. 2018.
[21] R. Lovelace, A. Goodman, R. Aldred, N. Berkoff, A. Abbas, and J. Woodcock, "The propensity to cycle tool: An open source online system for sustainable transport planning," *J. Transp. Land Use*, vol. 10, no. 1, pp. 505–528, Jan. 2017.
[22] J. D. D. Ortuzar and L. G. Willumsen, *Modelling Transport*, 4th ed. Chichester, U.K.: Wiley-Blackwell, Mar. 2011.
[23] Sustrans. (2021). *Route 42*. [Online]. Available: https://www.sustrans.org.uk/find-a-route-on-the-national-cycle-network/route-42
[24] R. Lovelace and C. H. V. Cooper, "CycleMon: Automated generation of basic walking and cycling analysis from open data," Tech. Rep., May 2021. [Online]. Available: https://github.com/Robinlovelace/cyclemon
[25] *London Residents Cycling to Work Doubles in 10 Years*, Office Nat. Statist., Newport, U.K., 2014.
[26] *Design Guidance, Active Travel (Wales) Act 2013*, Welsh Government, Cardiff, U.K., 2014.
[27] R. Lovelace and R. Felix, "Slopes R package," Tech. Rep., 2021. [Online]. Available: https://github.com/ropensci/slopes
[28] S. Cohen, T. Wan, M. T. Islam, and J. P. M. Syvitski, "Global river slope: A new geospatial dataset and global-scale analysis," *J. Hydrol.*, vol. 563, pp. 1057–1067, Aug. 2018.
[29] S. Sarkka, *Bayesian Filtering and Smoothing: 3*. Cambridge, U.K.: Cambridge Univ. Press, Sep. 2013.
[30] N. Shiode and S. Shiode, "Street-level spatial interpolation using network-based IDW and ordinary kriging," *Trans. GIS*, vol. 15, no. 4, pp. 457–477, Aug. 2011.
[31] B. Z. Demiray, M. Sit, and I. Demir, "D-SRGAN: DEM super-resolution with generative adversarial networks," *Social Netw. Comput. Sci.*, vol. 2, no. 1, p. 48, Jan. 2021.
[32] D. Jiao, D. Wang, H. Lv, and Y. Peng, "Super-resolution reconstruction of a digital elevation model based on a deep residual network," *Open Geosci.*, vol. 12, no. 1, pp. 1369–1382, Nov. 2020.
[33] C. Kirkwood, T. Economou, N. Pugeault, and H. Odbert, "Bayesian deep learning for spatial interpolation in the presence of auxiliary information," *Math. Geosci.*, vol. 54, no. 3, pp. 507–531, Apr. 2022.
[34] X. Han, X. Ma, H. Li, and Z. Chen, "A global-information-constrained deep learning network for digital elevation model super-resolution," *Remote Sens.*, vol. 15, no. 2, p. 305, Jan. 2023.
[35] D. Castelvecchi, "Can we open the black box of AI?" *Nature*, vol. 538, no. 7623, pp. 20–23, Oct. 2016.
[36] G. Vilone and L. Longo, "Notions of explainability and evaluation approaches for explainable artificial intelligence," *Inf. Fusion*, vol. 76, pp. 89–106, Dec. 2021.
[37] M. F. Hutchinson, "A locally adaptive approach to the interpolation of digital elevation models," in *Proc. 3rd Int. Conf./Workshop Integrating GIS Environ. Model.* Santa Barbara, CA, USA: Univ. California, 1996, pp. 21–26.
[38] C. Wang, Q. Yang, D. Jupp, and G. Pang, "Modeling change of topographic spatial structures with DEM resolution using semi-variogram analysis and filter bank," *ISPRS Int. J. Geo-Inf.*, vol. 5, no. 7, p. 107, Jun. 2016.
[39] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, "Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization," *ACM Trans. Math. Softw.*, vol. 23, no. 4, pp. 550–560, Dec. 1997.
[40] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, and S. J. Van Der Walt, "SciPy 1.0: Fundamental algorithms for scientific computing in Python," *Nature Methods*, vol. 17, pp. 261–272, Feb. 2020.
[41] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, and A. Desmaison, "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alche-Buc, E. Fox, and R. Garnett, Eds. New York, NY, USA: Curran Associates, 2019, pp. 8024–8035.
[42] *OS Mastermap Highways Network—Roads—Technical Specification V2.3*, Ordnance Surv., Southampton, U.K., Mar. 2021.
[43] J. G. MacKinnon and H. White, "Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties," *J. Econ.*, vol. 29, no. 3, pp. 305–325, Sep. 1985.

CRISPIN H. V. COOPER received the B.A. degree in computer science from the University of Cambridge, England, U.K., in 2002, and the M.A. (Cantab.) and Ph.D. degrees in computer science and geography from Cardiff University, Wales, in 2005 and 2010, respectively.

From 2003 to 2005, he was a Research Assistant with the Department of Electronics, University of York, U.K. From 2010 to 2021, he was a Research Associate and then a Research Fellow with the Sustainable Places Research Institute, Cardiff University, where he is currently a Lecturer with the School of Computer Science and Informatics. He is a Lead Developer of the sDNA software for spatial network analysis used by both academics and planners worldwide and leads a research program on sustainable transport, including the application of spatial network analysis to modeling walking and cycling. His whole system modal shift model formed the basis of the report "Moving On: Greener Travel for the U.K." (Green Alliance, 2023).