

BAYESIAN HIERARCHICAL MODELING AND ANALYSIS FOR ACTIGRAPH DATA FROM WEARABLE DEVICES

BY PIERFRANCESCO ALAIMO DI LORO^{1,a}, MARCO MINGIONE^{2,b}, JONAH LIPSITT^{3,c},
CHRISTINA M. BATTEATE^{4,e}, MICHAEL JERRETT^{3,d} AND SUDIPTO BANERJEE^{5,f}

¹Department GEPLI, LUMSA, ^ap.alaimodiloro@lumsa.it

²Department of Political Sciences, Roma Tre University, ^bmarco.mingione@uniroma3.it

³Department of Environmental Health Sciences, University of California, Los Angeles, ^cjonahlipsitt@gmail.com,
^dmjerrett@ucla.edu

⁴Center of Occupational and Environmental Health, University of California, Los Angeles, ^ecbatteate@ucla.edu

⁵Department of Biostatistics, University of California, Los Angeles, ^fsudipto@ucla.edu

The majority of Americans fail to achieve recommended levels of physical activity, which leads to numerous preventable health problems, such as diabetes, hypertension, and heart diseases. This has generated substantial interest in monitoring human activity to gear interventions toward environmental features that may relate to higher physical activity. Wearable devices, such as wrist-worn sensors that monitor gross motor activity (actigraph units) continuously record the activity levels of a subject, producing massive amounts of high-resolution measurements. Analyzing actigraph data needs to account for spatial and temporal information on trajectories or paths traversed by subjects wearing such devices. Inferential objectives include estimating a subject's physical activity levels along a given trajectory, identifying trajectories that are more likely to produce higher levels of physical activity for a given subject, and predicting expected levels of physical activity in any proposed new trajectory for a given set of health attributes. Here, we devise a Bayesian hierarchical modeling framework for spatial-temporal actigraphy data to deliver fully model-based inference on trajectories while accounting for subject-level health attributes and spatial-temporal dependencies. We undertake a comprehensive analysis of an original dataset from the Physical Activity through Sustainable Transport Approaches in Los Angeles (PASTA-LA) study to ascertain spatial zones and trajectories exhibiting significantly higher levels of physical activity while accounting for various sources of heterogeneity.

1. Introduction. Promoting a healthy lifestyle continues to stoke substantial research activities in public health. The “Physical Activity Guidelines for Americans” (second edition) suggests that most individuals, depending on age and body composition, receive 150–300 minutes of moderate to vigorous physical activity (MVPA) weekly (Piercy et al. (2018)). In general, the scientific community agrees that regular physical activity (PA) can have immediate and long-term health benefits (Reiner et al. (2013), Bull et al. (2020)). Despite these well-known benefits, most Americans fail to meet recommended requirements (Piercy et al. (2018)). Specifically, only one in five high-school adolescents and one in four adults meet recommended levels of physical activity. Given the well-established relationships between lack of PA and several obesity-related chronic conditions, such as heart disease, type 2 diabetes, and cancer as well as many physical and mental health benefits, an urgent need exists to improve monitoring of PA and to establish public health programs that promote more PA.¹

Technologies for monitoring spatial energetics James et al. (2016), Drewnowski et al. (2020) and promoting physical activity continue to emerge. Among others, *actigraphy*

Received November 2022; revised February 2023.

Key words and phrases. Bayesian hierarchical models, directed acyclic graph, Gaussian processes, physical activity, sparsity, spatial-temporal statistics.

¹More details at <https://www.cdc.gov/chronicdisease/resources/publications/factsheets/physical-activity.htm>.

broadly refers to the monitoring of human rest and activity cycles using wearable devices. Actigraphy data are gathered directly from wearable sensors or indirectly through smartphone mobile applications and record repeated measurements at very high resolutions. In particular, accelerometers are motion sensors that measure acceleration along different axes and are able to collect large amounts of data (Plasqui and Westerterp (2007), Sikka et al. (2019)). They are increasingly conspicuous because of their affordability, accuracy, and availability in smart-phones, smart-watches, and other wearable devices. Many devices also include Global Positioning System (GPS) sensors that reference measurements with location tracking along trajectories, or paths, traversed by the subject. Collected data can be quickly downloaded and promptly analyzed to obtain insights into their pattern and structure.

We pursue a comprehensive analysis of an original actigraphy data set from the Physical Activity through Sustainable Transport Approaches in Los Angeles (PASTA-LA) study. Actigraphy and GPS data analysis customarily involve *idle* records that occur if a charged device does not detect acceleration over a specified time interval (e.g., 10 seconds). While idle records may correspond to periods of a subject's inactivity, they can also arise from other factors, including technical malfunctions or the subject not wearing the device. On the other hand, eliminating idle records does not exclude all inactive periods because the accelerometer still records minor movements from where it is worn while a subject may be mostly inactive. Attempting to account for idle records as representative of inactive periods are likely to confound assessments of a subject's activity levels with current technological capabilities. Hence, we do not consider idle records and focus on the following specific data analytic aims: (i) estimating a subject's physical activity levels along trajectories, (ii) identifying trajectories that are more likely to produce higher levels of physical activity for a given subject, and (iii) predicting expected levels of physical activity in any proposed new trajectory for a given set of health attributes. Researchers find actigraph tracking especially attractive, as it allows for a better understanding of what behavioral and environmental factors influence population and individual health and hence aid in public health recommendation and policy.

Actigraphs generate data evolving over space and time, which suggests rich classes of space-time models for analysis (Gelfand et al. (2010), Cressie and Wikle (2011)). In particular, actigraph analysis presents some notable challenges (Kestens et al. (2017)): the data sets are large, or even massive, as they are recorded at very high frequencies; they exhibit dependence along trajectories, which should be accounted for both explanation and prediction (Bai et al. (2018), Ray et al. (2018)). We argue against a customary spatial-temporal process over $\mathbb{R}^2 \times \mathbb{R}^+$ and propose disentangling spatial effects from temporal dependence along trajectories. The balance of the paper is organized as follows. Section 2 introduces the PASTA-LA dataset with insights into accelerometry data. The model for the temporal correlation is introduced in Section 3, while spatial effects are discussed in Section 3.4. An extensive simulation study validating our model is proposed in Section 3.5. Data analysis, model assessment, and comparisons are presented in Section 4. Finally, we conclude with a discussion in Section 5.

2. Data. Our data set is compiled from the original Physical Activity through Sustainable Transport Approaches in Los Angeles (PASTA-LA) study conducted on a cohort of 460 individuals monitored between May 2017 and June 2018. Data were collected through different sources: online questionnaires, a smartphone app named *MOVES*, a GPS device (GlobalSat DG-500), and a wearable actigraph unit (Actigraph GT3X+). Data collected through the *MOVES* app, whose reliability must still be verified and discussed, are not considered in this paper. While 460 is the sample size of the complete study, the GPS and actigraph devices were deployed only on a nested sample of 134 individuals due to cost considerations. We analyze data collected through these two devices that were supposed to be worn by the participants in the nested sample for two one-week periods (one in 2017 and

one in 2018). Study protocols for safeguarding participant information received necessary institutional review board (IRB) approval. The data were stored on a secure computer, and a redacted version was created for purposes of data sharing.

2.1. Questionnaires. The online questionnaires included two baseline and four follow-up surveys, i.e. one baseline and two follow-ups for each collection period of the actigraph and GPS data. Each survey consisted of responses pertaining to the participant's demographics and transportation habits. Here, we consider the *first baseline questionnaire*, which is the only one available for all the participants in the nested sample. Personal information and other socioeconomic factors have been encoded as follows for subsequent analysis:

- *Sex*: Female or Male;
- *Ethnicity*: Asian, Black/African/Caribbean, Latin-American, White, or Other (mixed multiple ethnic groups or prefer not to answer);
- *Age (years) class*: (0, 18], (18, 25], (25, 34], (34, 45], (45, 70];
- *BMI (kg/m^2) class*²: underweight if $\text{BMI} \in (15, 18.5]$, normal if $\text{BMI} \in (18.5, 24.5]$, overweight if $\text{BMI} \in (24.5, 30]$, and obese if $\text{BMI} > 30$;
- *Yearly Income Level* (in thousands \$): (0, 10], (10, 25], (25, 50], (50, 75], (75, 100], (100, 150], (150, $+\infty$], and Don't know/Prefer not to answer;
- *Educational attainment*: High-school diploma, College graduate, Associate degree, Graduate.

We filtered unreasonable values of the BMI, that is, $\text{BMI} < 10$, which was observed just for one individual, leaving 133 out of 134 individuals in the nested sample. A user ID was assigned to each survey response data, and a redacted master key was generated using all ID types for joining with other study data.

2.2. Actigraph. The Actigraph unit is an accelerometer roughly the same size and weight of the average wrist-watch. It can be worn on the wrist, hip, and thigh and measure the directional acceleration at a specified time frequency (up to 100Hz). The Actigraph GT3X+ model used for the PASTA-LA study can detect accelerations measured in gravitational units (G) with a sensitivity of ± 3 milligravity (mG) in the three orthogonal planes (anteroposterior, mediolateral, and vertical). Data are stored in an internal memory and can be downloaded to other hardware for analysis through a proprietary software. The participants were asked to start wearing the accelerometer on their dominant wrist as soon as it was handed to them, as the devices could have been properly calibrated at that time. The study protocol demanded that participants wear the Actigraph unit at all times other than during bathing and sleeping (awake time was assumed approximately from 7 a.m. to 11 p.m.). The sampling frequency has been set to 30Hz, and the *idle sleep* mode has been activated in order to save battery and memory. With this mode on, the device would go idle every time it records no acceleration ($< \pm 40mG$) for 10 consecutive seconds. The Actigraph GT3X+ grants access to the a `.gt3x` file with the raw acceleration measurements. It can be loaded in R, using the `read.gt3x` package, and contains the raw accelerations at each timestamp. Such accelerations comprise the basic ingredients to get a proxy for body movement from an accelerometer (Mathie et al. (2003), Migueles et al. (2019), Bammann et al. (2021)). There are substantial investigations into its statistical relationships with PA measures, such as *energy expenditure* measures (EE) (Crouter, Clowers and Bassett Jr (2006), Freedson et al. (2012), Taraldsen et al. (2012)) and the *Metabolic Equivalent of Task* (MET) (Lyden et al. (2014),

²According to standard guidelines of the Center of Disease Control and Prevention <https://www.cdc.gov/obesity/basics/adult-defining.html>.

Staudenmayer et al. (2015), Migueles et al. (2017), van Loo et al. (2018)). Among various metrics, we take the instantaneous body vector Magnitude of Acceleration (MAG) as the primary endpoint of our analysis (van Hees et al. (2011), White et al. (2016), Doherty et al. (2017)). Further discussion about the conversion of MAG into energy expenditure measures is reported in Section 2 of the Supplementary Material (Alaimo Di Loro et al. (2023)).

We were able to retrieve the Actigraph raw data only on $K = 97$ out of the 133 original individuals. Let x , y and z be the dynamic acceleration of the body of the k th individual. The pointwise MAG is defined as

$$(1) \quad \text{MAG}_{kt} = \sqrt{x_{kt}^2 + y_{kt}^2 + z_{kt}^2}, \quad k = 1, \dots, K.$$

However, the raw accelerations recorded by the accelerometer must be appropriately processed to assess body movement (Doherty et al. (2017)). Indeed, the raw acceleration recorded by each axle is the sum of both the *static* and the *dynamic* acceleration, but only the second is the effect of actual body movement. First, we remove *idle* records, that is, all the occasions in which the accelerometer recorded zero acceleration for longer than 15 seconds and the device went *idle*. It is very unlikely that these idle records with zero acceleration coincide with a subject's inactive periods because the accelerometer still records positive, albeit small, magnitudes of accelerations over inactive periods due to movements in the wrists, hips, and thighs. Idle records, on the other hand, are likely to arise from technical malfunctions or from a subject violating protocol and not wearing the device in the experimental time window.

Second, the raw accelerations, recorded by single axes, must be disentangled from unwanted static or nonstatic components: the effect of the earth's gravitational force and other external accelerations (e.g., car, bus, elevators) at low frequencies, machine noise, and vibrations at high frequencies. To address this issue, we adopted a Band-Pass Butterworth digital filter of order 4 with frequency window (0.25, 10) to clean the signals from these long and short waves (Mathie et al. (2003)). Indeed, most human activities result in signals with a frequency between 0.25Hz and 10Hz (Khusainov et al. (2013)). An example of how the raw signal is modified through this process is provided in Figure 1 of the Supplementary Material (Alaimo Di Loro et al. (2023)). We subsequently evaluate the pointwise MAG_{kt} using the filtered accelerations $(\tilde{x}_{kt}, \tilde{y}_{kt}, \tilde{z}_{kt})$. However, the instantaneous MAG, evaluated at the original 30Hz frequency, is extremely erratic, and the single value may not represent well the PA intensity of the participant at that time. For this reason it is usually averaged over 5 to 10 second epochs to acquire a more suitable measure of PA intensity (Migueles et al. (2017), Doherty et al. (2017)). Here, we perform a kernel smoothing of the 30Hz measurements, in order to be representative of the single time-point, and get a one second time resolution. The resulting vector magnitude is

$$(2) \quad \widetilde{\text{MAG}}_{kt} = \sum_j k_b(t - t_j) \cdot \text{MAG}_{kt_j},$$

where $k_b(\cdot)$ is a Gaussian kernel with bandwidth $b = 5$ seconds. This ensures that the impact of the neighboring points becomes negligible for $|t - t_j| > 10$ seconds; see Figure 2 of the Supplementary Material (Alaimo Di Loro et al. (2023)) as an example. Finally, we removed all the observations recorded outside of the prespecified daily time window, that is, from 7 a.m. to 11 p.m.

2.3. GPS. The GPS device *GlobalSat DG-500* recorded the subject's location (latitude and longitude) roughly every 5 seconds, together with date, time, and speed (*km/h*, measured as distance over time through linear interpolation). This work restricts the attention to 93 out of 97 subjects living and working in the Westwood neighborhood of Los Angeles in order

to avoid a geographical imbalance that could bias and invalidate the model estimates. This area hosts the university campus of UCLA, and it includes the largest part of all the available observations. Westwood is a walk-friendly neighborhood with a lot of green areas, parks, and major roads with shops and amenities. People were free to move inside and between buildings (e.g., people at the gym, office, etc), and we are interested in quantifying their movement in all these settings.

However, GPS measurements can be affected by possible inaccuracies, especially around buildings, that may cause unreasonable jumps in a very small time span. We note that most of these issues are already mitigated by an automatic filtering process of the GPS device that would drop records for which the signal is not strong enough. Nevertheless, to further enhance the cleaning process of GPS measurements, we removed all data points for which the computed average speed between two subsequent points was larger than 100 km/h . We picked such a high threshold, as we do not want to drop observations related to individuals standing or sitting in a bus or car.

2.4. External covariates. PA levels are not only affected by individual characteristics but also can be fostered by specific features of the surrounding area they are navigating. Therefore, we included three external covariates to account for some of the built-in environment features of the Westwood area. In particular, we used the following:

(i) The weighted overlay distance to parks (in km) with a spatial resolution of 23×23 , which can be downloaded from <https://egis2.lacounty.gov/arcgis>. It represents a weighted distance of each point from officially recognized parks, and it can be seen as a proxy of the *green area density* (see Figure 1(a); darker shades indicate proximity to parks).

(ii) The Normalized Difference Vegetation Index (NDVI), which is available with a spatial resolution of 30×30 , can be downloaded from <https://earthexplorer.usgs.gov/> and provides a measure of the *greenness* of the patch itself (see Figure 1(b); darker shades depict more greenness).

(iii) The slope (azimuth), with a spatial resolution of 23×23 , is computed from the digital elevation model (DEM) downloaded from http://www.webgis.com/terr_pages/CA/dem1/losangeles.html. It represents the average angular inclination of the ground patch with respect to the horizon line (see Figure 1(c); darker shades depict higher slopes).

While previous studies (e.g., Maddison et al. (2009)) have reported on these variables affecting PA levels, they usually consider the environmental impact on the average PA level through a buffer around the home location of the participant and not on its instantaneous PA level. With our current work, we want to discover and establish direct associations between PA levels of a subject and these spatially-indexed covariates along trajectories.

2.5. Joining. GPS and accelerometer data were all assigned a participant ID aligned with the questionnaires' master key to facilitate joining across all ID types (including email) while redacting and encrypting user data. The first baseline questionnaire, Actigraph, and GPS were available for the aforementioned group of $K = 93$ individuals. Henceforth, we refer to this specific group of units. The joining of different data sources follows these steps:

1. Actigraph data are joined to the first baseline questionnaire using the individual master-key. The resulting data set includes the physical activity endpoint, recorded by the Actigraph at the different timestamps, and all the available individual information but no spatial information.

2. GPS data are joined to the dataset obtained in (1) using the individual master-key and the timestamp. Following the processing in Section 2.2, the Actigraph data are available at the same time resolution of GPS data as long as the subject wore them simultaneously.

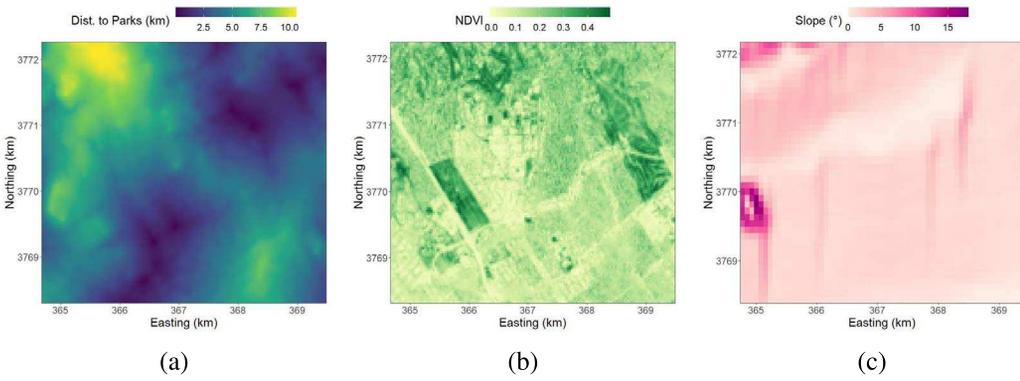


FIG. 1. *Spatial-varying covariates: (a) distance to parks, (b) NDVI, and (c) slope.*

Therefore, we decided to use the GPS as the leading table in the joining process. This avoids use of artificial data (e.g., interpolating GPS locations).

3. Spatial covariates are joined to the data set obtained in (2) through the *minimum distance* criteria; that is, each location is assigned the value of the closest point on the grid for each spatial-varying covariate.

The temporal coverage is not balanced across individuals because: (i) subjects move around Westwood in different segments of the overall time window, and (ii) some participants violated the study protocol. Indeed, not all the participants were available for both of the one-week surveillance periods in 2017 and 2018. In fact, only 58 out of the 93 participants have data for the first week only and missed the follow-up survey. In the end we go from the least represented individuals, having $\approx 5 \times 10^2$ observations (≈ 2 hours of data), to the most represented ones with more than $\approx 5 \times 10^5$ observations (≈ 14 days of data). Considering all the *five* second time segments between the first and last observed point of each individual in each day as the potential observation window, the proportion of missing measurements ranges between $\approx 31\%$ and $\approx 97\%$. The overall proportion of missing measurements in the entire database (based upon aggregated *five*-second time segments for all individuals) is $\approx 83\%$. Figure 2 shows the number of observations for each individual $k = 1, \dots, K$ in each hour $h = 7, \dots, 22$, where the y-axis has been ordered in ascending order according to proportion of missing measurements for each subject. Overall, we can state that all hours are well represented, but only a few individuals have data for the whole daily time window. Summing up, all subsequent analysis will refer to the final data set consisting of $N \simeq 7 \times 10^5$ measurements across $K = 93$ individuals, scattered over Westwood (see Figure 3).

3. The model. The outcomes corresponding to the $K = 93$ subjects are referenced with respect to the time at which they are recorded and the position in the trajectory. While it is tempting to work with a spatiotemporal process, dependence introduced by such processes may not be appropriate. An individual can visit the same location numerous times in his/her trajectory. These revisits need not occur at regular intervals and can be at distant time points. This suggests that proximity of two spatial locations in a trajectory need not result in strongly dependent MAGs recorded there. It appears more reasonable to model dependence among MAG measurements through a temporal process. In fact, such temporal processes can be motivated by the position vectors defining the trajectories, as we describe below.

Let $Z_k(\cdot) : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a spatial process corresponding to individual k . The domain of $Z_k(\cdot)$ is restricted to the trajectories $\gamma_k(t) = (\gamma_k^x(t), \gamma_k^y(t))$, where $k = 1, \dots, K$ and $t \in \mathbb{R}^+$, which defines the movements of the k th individual along time. As shown in Figure 4, the process actually belongs to a one-dimensional space for which we define a proper distance

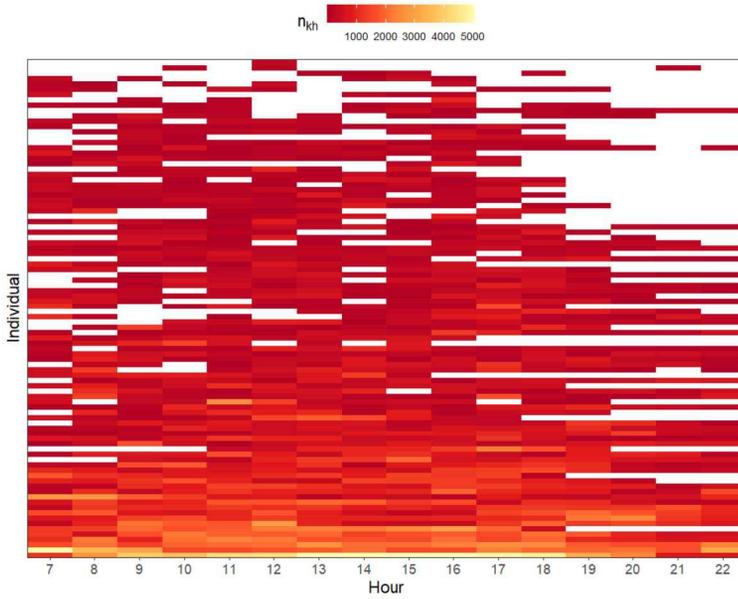


FIG. 2. Missing data pattern by hour of the day and individual.

measure $d(t_{ki}, t_{kj}) = \|\gamma_k(t_{kj}) - \gamma_k(t_{ki})\|$, where t_{ki} is the i th recorded time point from individual k . We approximate such distances as the elapsed time between the two points $d(t_{ki}, t_{kj}) = |t_{kj} - t_{ki}|$, which would result in a good approximation of the spatial distance (especially if the subject is moving at constant speed). More generally, the elapsed separation across time will reflect dependence better than the spatial distance. The faster an individual is moving from one point to the other, the shorter the time elapsed, and higher the correlation between the two measurements. Hence, we model our measurements as $Y_k(\cdot) \equiv Z_k \circ \gamma_k(\cdot) : \mathbb{R}^+ \rightarrow \mathbb{R}$, which, by construction, is a valid stochastic process.

This will form the edifice of the model in Section 3.1, where we are modeling the dependence by solely considering stochastic evolution through time. How should spatial information be introduced in the model? Two individuals at the same spatial coordinate experience the same spatial effect but different temporal effects because their physical activities are a function of their trajectory’s temporal evolution. An added complication is that trajectories

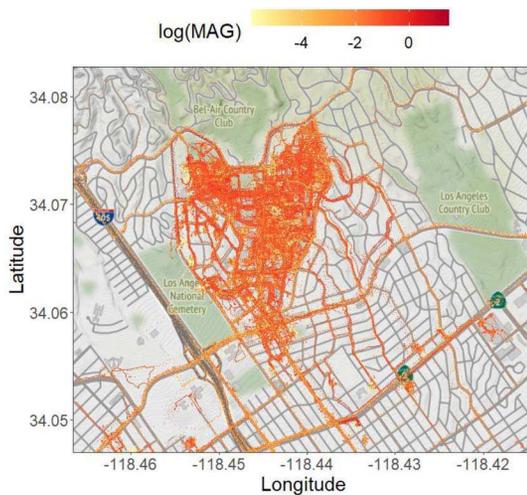


FIG. 3. Observed locations over the Westwood area.

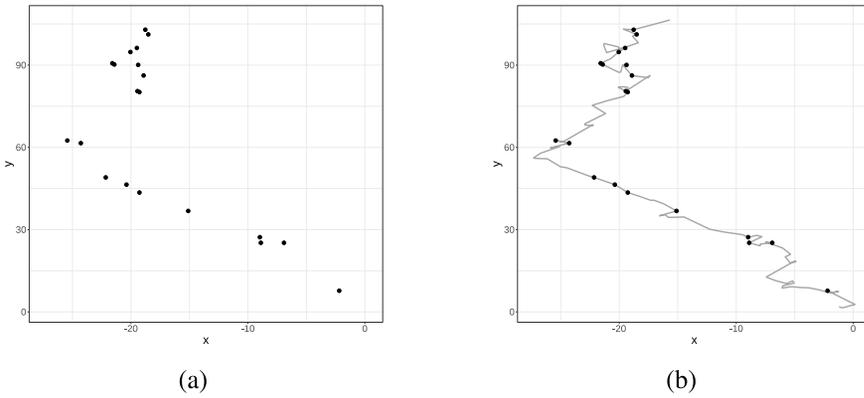


FIG. 4. Example of observed points (a) and trajectory (b): black dots are realizations, grey line is the domain of the process.

intersect and overlap and, in practice, can have multiple observations at the same location. Even more flexible spatiotemporal covariance kernels (e.g., nonseparable or nonstationarity) will struggle to recognize the above features. Hence, we introduce the spatial effect in the mean using spline regression (see Section 3.4).

3.1. *Temporal model.* Let $\mathcal{T} = \bigcup_{k=1}^K \mathcal{T}_k$, where $\mathcal{T}_k = \{t_{ki}\}_{i=1}^{T_k}$ and $t_{ki} \in \mathbb{R}^+$ be the set of the $n = \sum_{k=1}^K T_k$ observed time points. We model $Y(\mathcal{T})$ as the finite realization of a K -variate process $Y(\cdot)$ over \mathbb{R}^+ ,

$$(3) \quad Y(t) = X(t, \boldsymbol{\gamma}(t))^\top \boldsymbol{\beta} + \boldsymbol{w}(t) + \boldsymbol{\varepsilon}(t), \quad t \in \mathbb{R}^+,$$

where $Y(t) = (Y_1(t), Y_2(t), \dots, Y_K(t))^\top$ is a $K \times 1$ vector of measurements at time t on the K individuals, $X(t, \boldsymbol{\gamma}(t))$ is a $p \times K$ matrix, each row being the values of a covariate for the K individuals, $\boldsymbol{w}(t) = (w_1(t), w_2(t), \dots, w_K(t))^\top$ is a $K \times 1$ vector comprising a temporal process for each individual, and $\boldsymbol{\varepsilon}(t) \sim \mathcal{N}_K(0, \tau^2 \mathbf{I}_K)$, $\tau^2 \in \mathbb{R}^+$ is a white noise process for measurement error. Each element of $\boldsymbol{w}(t)$ is specified as $w_k(t) \stackrel{\text{ind}}{\sim} \mathcal{GP}(0, c_\theta(\cdot, \cdot))$, where $c_\theta(\cdot, \cdot)$ is a covariance function with parameters $\boldsymbol{\theta} \in \Theta$.

Let y_{ki} and \boldsymbol{x}_{ki} be the outcome and covariates for individual k at time point t_{ki} , respectively, so $\{(y_{ki}, \boldsymbol{x}_{ki}) : k = 1, \dots, K, i = 1, \dots, T_k\}$ is the observed data. Let \boldsymbol{y}_k and \boldsymbol{w}_k be $T_k \times 1$ vectors comprising all measurements and random effects on patient k , respectively. Forming the $n \times 1$ vectors $\boldsymbol{y} = [y_1^\top, y_2^\top, \dots, y_K^\top]^\top$ and $\boldsymbol{w} = [w_1^\top, w_2^\top, \dots, w_K^\top]^\top$ and the $n \times p$ matrix $\boldsymbol{X} = [X_1^\top, X_2^\top, \dots, X_K^\top]^\top$, where X_k is the $T_k \times p$ matrix of predictors corresponding to \boldsymbol{y}_k , we extend (3) to a hierarchical model with posterior distribution

$$(4) \quad p(\boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{\theta}, \tau^2 | \boldsymbol{y}) \propto p(\boldsymbol{\theta}, \tau^2) \times N(\boldsymbol{\beta} | \boldsymbol{\mu}_\beta, \boldsymbol{V}_\beta) \times N(\boldsymbol{w} | \mathbf{0}, \boldsymbol{C}_\theta) \times N(\boldsymbol{y} | \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{w}, \tau^2 \mathbf{I}_n).$$

The covariance matrix $\boldsymbol{C}_\theta = \text{diag}(\boldsymbol{C}_{\theta_{1,1}}, \boldsymbol{C}_{\theta_{2,2}}, \dots, \boldsymbol{C}_{\theta_{K,K}})$ is $n \times n$ block-diagonal with $\boldsymbol{C}_{\theta_{k,k}} = [c_\theta(t_{ki}, t_{kj})]$ as the $T_k \times T_k$ temporal covariance matrix corresponding to individual k . Each individual is allowed its own covariance parameters, $\boldsymbol{\theta}_k$, and $\boldsymbol{\theta} = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_K\}$ in (4) is the collection of all the covariance kernel parameters. Applying (4) involves the determinant and inverse of \boldsymbol{C}_θ , which require $\mathcal{O}(n^2)$ storage space and $\mathcal{O}(n^3)$ floating point operations (flops). The block-diagonal structure of \boldsymbol{C}_θ considerably alleviates this burden since $\det(\boldsymbol{C}) = \prod_{k=1}^K \det(\boldsymbol{C}_{\theta_{k,k}})$ and $\boldsymbol{C}^{-1} = \text{diag}(\boldsymbol{C}_{\theta_{1,1}}^{-1}, \boldsymbol{C}_{\theta_{2,2}}^{-1}, \dots, \boldsymbol{C}_{\theta_{K,K}}^{-1})$. This reduces the flop count from $\mathcal{O}(n^3) = \mathcal{O}((\sum_{k=1}^K T_k)^3)$ to $\mathcal{O}(K \sum_{k=1}^K (T_k)^3)$ with a significant saving of calculations, especially when the T_k 's are reasonably small ($< 10^4$). Furthermore, each $\boldsymbol{C}_{\theta_{k,k}}$ can be computed in parallel rendering further scalability to the algorithm.

However, analyzing the Actigraph data in Section 2 will involve $T_k > 10^5$ measurements from some individuals. Full inference will be impractical without any exploitable structure for each $\mathbf{C}_{\theta,k,k}$. Analyzing massive spatiotemporal data has witnessed burgeoning interest, and a comprehensive review is beyond the scope of this work (see, e.g., Banerjee (2017), Heaton et al. (2019), and references therein). We will pursue an approximation, due to Vecchia (Vecchia (1988)), that has generated substantial recent interest (Datta et al. (2016a, 2016b), Katzfuss et al. (2020), Katzfuss and Guinness (2021), Peruzzi, Banerjee and Finley (2022)) in scalable Bayesian modeling.

3.2. *Independent DAG models over individuals.* We adapt Vecchia’s likelihood approximation (Vecchia (1988)) to the random effects \mathbf{w}_k for each $k = 1, 2, \dots, K$. Beginning with the observed time points $\{t_{k1} < t_{k2} < \dots < t_{kT_k}\}$ for individual k and the directed acyclic graphical (DAG) representation $p(\mathbf{w}_k) = p(w_{k1}) \prod_{i=2}^{T_k} p(w_{ki}|w_{k1}, \dots, w_{k(i-1)})$, we define

$$(5) \quad p(\mathbf{w}_k) \approx \tilde{p}(\mathbf{w}_k) = p(w_{k1}) \prod_{i=2}^{T_k} p(w_{ki}|\mathbf{w}_{k,N(i)}),$$

where $\tilde{p}(\cdot)$ is the joint density derived from $p(\mathbf{w}_k)$ by restricting the parents (conditional sets) of each w_{ki} in the DAG to a set $\mathbf{w}_{kN(i)} = \{w_{kj} : j \in N(i)\}$, where $N(i)$ is a set of prefixed size m comprising the m nearest neighbors of t_{ki} from the past. Thus, $N(i) = \{t_{k(i-m)}, \dots, < t_{k(i-1)}\}$ for $i > m$ and $N(i) = \{t_{k1}, \dots, t_{k(i-1)}\}$ for $i \leq m$. Such approximations yield valid probability likelihoods (Lauritzen (1996), Stein, Chi and Welty (2004), Murphy (2012)) and can be extended to stochastic processes (Datta et al. (2016a)) for inference on arbitrary time points.

The connection between sparsity and conditional independence follows by writing (5) as a linear model $\mathbf{w}_k = \mathbf{A}_k \mathbf{w}_k + \boldsymbol{\eta}_k$, where \mathbf{A}_k is a $T_k \times T_k$ strictly lower triangular matrix, $\boldsymbol{\eta}_k \sim \mathcal{N}_{T_k}(\mathbf{0}, \mathbf{D}_k)$ and \mathbf{D}_k is the $T_k \times T_k$ diagonal matrix such that $[\mathbf{D}_k]_{ii} = d_{ii} = \text{Var}(w_{ki}|\{w_{kj}, j < i\})$ for $i = 1, \dots, T_k$. The DAG imposes the lower-triangular structure on \mathbf{A}_k , and its (i, j) th entry is allowed to be nonzero only for $j \in N(i)$. Therefore, each row of \mathbf{A}_k has at most m nonzero entries so that $\tilde{\mathbf{C}}_k^{-1} = (\mathbf{I}_{T_k} - \mathbf{A}_k)^\top \mathbf{D}_k^{-1} (\mathbf{I}_{T_k} - \mathbf{A}_k)$ is sparse, where $\tilde{\mathbf{C}}_k^{-1}$ is the precision matrix corresponding to $\tilde{p}(\mathbf{w}_k)$. Replacing \mathbf{C} with $\tilde{\mathbf{C}}$ in (4) yields a computationally efficient hierarchical model with $N(\prod_{k=1}^K N(\mathbf{w}_k|\mathbf{0}, \tilde{\mathbf{C}}_k))$ as the prior on \mathbf{w} .

The key observation is that the nonzero elements of the i th row of \mathbf{A}_k is the solution \mathbf{a}_k of the $m \times m$ linear system $\mathbf{C}_{\theta,k}[N(i), N(i)]\mathbf{a}_k = \mathbf{C}_{\theta,k}[N(i), i]$, where $[\cdot, \cdot]$ indicates submatrices defined by the given row and column index sets. Obtaining the nonzero elements of \mathbf{A}_k and \mathbf{D}_k costs $\mathcal{O}(T_k m^3)$ (scales linearly with T_k), instead of $\mathcal{O}(T_k^3)$, as would have been without sparsity. This cheaply delivers the quadratic form $\mathbf{w}_k^\top \tilde{\mathbf{C}}_k^{-1} \mathbf{w}_k$ in terms of \mathbf{A}_k and \mathbf{D}_k and the determinant $\det(\tilde{\mathbf{C}}_k) = \prod_{i=1}^{T_k} d_{ii}$ at almost no additional cost. The lower triangular matrix \mathbf{A}_k is not just sparse but also banded and with a lower bandwidth equal to m . Consequently, $\tilde{\mathbf{C}}_k^{-1}$ is also banded with lower and upper bandwidth equal to m . This leads to further accrual of computational benefits. The overall cost is $\mathcal{O}(\sum_{k=1}^K T_k m^3) = \mathcal{O}(nm^3)$ (linear in n) for computing the posterior for any given values of the parameters.

3.3. *Implementation using collapsed models.* The Bayesian hierarchical model in (4), either with \mathbf{C}_θ or with $\tilde{\mathbf{C}}_\theta$ in the prior for \mathbf{w} , allows full posterior inference for $\{\boldsymbol{\beta}, \mathbf{w}, \boldsymbol{\theta}, \tau^2\}$ using Markov chain Monte Carlo (MCMC). Gibbs sampling with random walk Metropolis steps provide full conditional distributions in closed form for $\{\boldsymbol{\beta}, \mathbf{w}\}$ and also for τ^2 with an $\mathcal{IG}(a_\tau, b_\tau)$ prior. However, this convenience is nullified, in practice, by strong autocorrelation

and poor mixing of the chains (Liu, Wong and Kong (1994)). Samplers based on spatial DAG-based models have been devised, explored, and compared in Finley et al. (2019). Instead of (4), we sample from

$$(6) \quad p(\boldsymbol{\beta}, \boldsymbol{\theta}, \tau^2 | \mathbf{y}) \propto p(\boldsymbol{\theta}, \tau^2) \times N(\boldsymbol{\beta} | \boldsymbol{\mu}_\beta, \mathbf{V}_\beta) \times N(\mathbf{y} | \mathbf{X}\boldsymbol{\beta}, \tilde{\mathbf{C}}_\theta + \tau^2 \mathbf{I}_n),$$

which is derived from (4) by integrating out \mathbf{w} , thereby “collapsing” the parameter space to a much smaller domain without \mathbf{w} . This considerably improves mixing and convergence.

We will need to compute the inverse and determinant of $\tilde{\boldsymbol{\Lambda}} = \tilde{\mathbf{C}}_\theta + \tau^2 \mathbf{I}_n$, which is $n \times n$. While $\tilde{\boldsymbol{\Lambda}}^{-1}$ does not share the same convenient factorization of $\tilde{\mathbf{C}}^{-1}$ and is also not guaranteed to be sparse, the Sherman–Woodbury–Morrison formulas reveal

$$(7) \quad \tilde{\boldsymbol{\Lambda}}^{-1} = \tau^{-2} \mathbf{I} - \tau^{-4} \boldsymbol{\Omega}^{-1} \quad \text{with } \boldsymbol{\Omega} = \tilde{\mathbf{C}}^{-1} + \tau^{-2} \mathbf{I},$$

where $\boldsymbol{\Omega}$ enjoys the same sparsity as \mathbf{C}^{-1} . Moreover, $\det(\tilde{\boldsymbol{\Lambda}}) = \tau^{2n} \det(\tilde{\mathbf{C}}) \det(\boldsymbol{\Omega})$. The core of the algorithm is, therefore, to compute $\tilde{\boldsymbol{\Lambda}}^{-1}$ through $\boldsymbol{\Omega}$. In our application the random effect is assumed to be the realization of K independent temporal processes. As discussed in Section 3.2, this implies a block-diagonal structure for $\tilde{\mathbf{C}}$ that can be shown to be shared also by $\boldsymbol{\Omega}$ (see equation (7)). Each block $\boldsymbol{\Omega}_k$ of $\boldsymbol{\Omega}$ can be computed independently for each individual, and the same holds for its inverse and its determinant. This means that the body of the algorithm will consist of a loop over all the individuals, which allows for straightforward parallelization; see Algorithm 1 in the Supplementary Material (Alaimo Di Loro et al. (2023)). Unlike in spatial DAGs (Datta et al. (2016a), Finley et al. (2019)), we do not need fill-reducing permutation methods since neighbors sets for temporal processes consist of contiguous observations and $\{\boldsymbol{\Omega}_k\}_{k=1}^K$ are banded matrices with no gaps.

We devised a Gibbs sampler with Metropolis random walk updates for (6), where $\boldsymbol{\beta}$ is updated from its full conditional distribution, while $\{\boldsymbol{\theta}, \tau^2\}$ are updated using an adaptive Metropolis step based on Haario, Saksman and Tamminen (2001). Here, after the first few iterations, a new proposal covariance matrix is regularly computed on the run according to the empirical covariance of the current chain. Subsequently, a mixture of the original and adaptive proposal is used as the new proposal. Convergence toward the desired acceptance rate is assured for an appropriate choice of the variance terms and of the adaptation rule (Roberts and Rosenthal (2009)). The algorithm has been coded using the R 4.0.5 statistical environment. All expensive computations are managed by the Eigen library (version 3.3.7), which provides efficient routines for numerical linear algebra with an emphasis on sparse matrices. Our implementation of (6) outperforms the algorithms that update \mathbf{w} in terms of computational speed, as it is implemented in the `spNNGP` package (Finley, Datta and Banerjee (2017)). We present these comparisons in the Supplementary Material (Alaimo Di Loro et al. (2023)), including a link to the GitHub repository hosting codes to implement the models.

3.4. Including spatial effects. Accounting for spatial information in our Actigraph dataset presents some new considerations. As mentioned in Section 1, spatial information is available to us in terms of the physical location along the trajectory as well as through covariates that are functions of space. Considering the discussion in Section 3, the analytical goals of this dataset suggest accounting for spatial heterogeneity. Here, as argued earlier, modeling $\mathbf{w}(\cdot)$ in (3) as a spatiotemporal process, including scalable versions, has challenges, given that: (i) the trajectory’s domain does not have a positive area; and (ii) associations among the measurements are more amenable to the temporal scale. Therefore, we introduce spatial effects into the mean employing a smooth function of space, $f_S(\cdot) : \mathbb{R}^2 \rightarrow \mathbb{R}$, approximated by a spline basis representation (see, e.g., Goodman and

Hardin (2006), Ramsay and Silverman (2007)). For instance, if J_x and J_y are the dimensions of independently defined B-spline basis expansions on the x and y coordinates, respectively, then $f_S((x, y)) \approx \tilde{f}_S((x, y)) = \sum_{j_x=1}^{J_x} \sum_{j_y=1}^{J_y} \beta_{S,(j_x, j_y)} B_{x, j_x}(x) B_{y, j_y}(y)$, where $B_{x, j_x} = [B_x]_{j_x}$ and $B_{y, j_y} = [B_y]_{j_y}$ are the j_x th and j_y th element of the B-spline basis along the two axis. For any location $(x, y) \in \mathbb{R}^2$, the elements of the previous sum can be more compactly expressed through the tensor product basis $B_S(x, y) = (B_x \otimes B_y)(x, y)$. The size of this basis is $J_S = J_x \cdot J_y$ and depends on the size of the two original spline basis, which, in turn, depends on the chosen number of knots $\text{knots}_x, \text{knots}_y$ and degree $\text{deg}_x, \text{deg}_y$ (namely, $J_c = \text{knots}_c + \text{deg}_c$ for $c = x, y$). We now modify (3) to include the spline,

$$(8) \quad Y(t) = X(t, \boldsymbol{\gamma}(t))\boldsymbol{\beta} + B_S(\boldsymbol{\gamma}(t))\boldsymbol{\beta}_S + \mathbf{w}(t) + \boldsymbol{\varepsilon}(t), \quad t \in \mathbb{R}^+,$$

where $\boldsymbol{\gamma}(t) = \{\boldsymbol{\gamma}_1(t), \boldsymbol{\gamma}_2(t), \dots, \boldsymbol{\gamma}_K(t)\}$, $\boldsymbol{\gamma}_k(t) = (\gamma_{k,x}(t), \gamma_{k,y}(t)) : \mathbb{R}^+ \rightarrow \mathbb{R}^2$ is the trajectory function mapping time t for individual k to its position and $B_S(\boldsymbol{\gamma}(t))$ is the $K \times J_S$ matrix with row k corresponding to the J_S basis elements for the coordinates at time point t for individual k . A proper choice of J_S (i.e., knots and degree) is required to fit a spline surface flexible enough to describe the spatial variations at the scale of interest without incurring over-fitting. Let $B = B_S(\boldsymbol{\gamma}(\mathcal{T}))$ be the $n \times J_S$ matrix containing the B-spline basis elements evaluated at the observed location of each individual $\boldsymbol{\gamma}(\mathcal{T}) = \{\gamma_1(t_{11}), \gamma_1(t_{12}), \dots, \gamma_K(t_{K T_k})\}$. Following equation (6), we sample from

$$(9) \quad p(\boldsymbol{\beta}, \boldsymbol{\beta}_S, \boldsymbol{\theta}, \tau^2 | \mathbf{y}) \propto p(\boldsymbol{\theta}, \tau^2) \times p_S(\boldsymbol{\beta}_S) \times N(\boldsymbol{\beta} | \boldsymbol{\mu}_\beta, \mathbf{V}_\beta) \times N(\mathbf{y} | X\boldsymbol{\beta} + B\boldsymbol{\beta}_S, \tilde{C}_\theta + \tau^2 \mathbf{I}_n),$$

where the prior $p_S(\cdot)$ needs to be specified. The Actigraph data includes millions of observations in a limited study area of which some assume different values in the same location (or in its immediate vicinity), so over-fitting will not be an issue. However, some areas present sparsely observed points (trajectories are not uniformly distributed, as shown in Figure 3). This may cause coefficients corresponding to those regions to be weakly identified. To control for the balance of all these components, we may assign ad hoc priors to the spatial spline regression coefficients (Eilers and Marx (1996)) for penalizing deviation from a certain degree of smoothness and favoring identifiability. This behavior suggests the Bayesian P-Spline (Hastie and Tibshirani (2000), Lang and Brezger (2004)). While keeping the Gaussian priors, we effectuate shrinkage by choosing a suitable precision matrix P and introducing a shrinkage parameter λ at a deeper level of the hierarchy. To be precise, $\boldsymbol{\beta}_S | \lambda \propto \exp\{-\frac{\lambda}{2} \cdot \boldsymbol{\beta}_S P \boldsymbol{\beta}_S^T\}$ and $\lambda \sim \mathcal{G}(\alpha_\lambda, \beta_\lambda)$. We consider two possible forms for P , which imply different penalization for the coefficients:

- Ridge-like prior, which is to say $P = P_{RL} = \mathbf{I}_{J_S}$;
- First-order random walk prior, which is to say,

$$P = P_{RW} : [P_{RW}]_{ij} = \begin{cases} n_i & i = j, \\ -1 & i \sim j, \\ 0 & \text{otherwise,} \end{cases}$$

where n_i is the number of neighbors of knot i and $i \sim j$ denotes a neighboring relationship between the knots.

Both precision matrices provide a multivariate Gaussian prior distribution on the coefficients. However, the latter is improper since $\text{rank}(P_{RW}) < J_S$. Nevertheless, if we collect the B-Spline basis elements with the other covariates, as $X^* = [X, B]$, and stack the corresponding coefficients into the joint vector $\boldsymbol{\psi} = [\boldsymbol{\beta}, \boldsymbol{\beta}_S]$, then the posterior distribution of the latter is

a proper multivariate Gaussian with full conditional distribution $\boldsymbol{\psi}|\cdot \propto \mathcal{N}_J(\boldsymbol{\psi}|\mathbf{G}^{-1}\mathbf{g}, \mathbf{G}^{-1})$, where $\mathbf{G} = \mathbf{X}^{*\top}\tilde{\boldsymbol{\Lambda}}^{-1}\mathbf{X}^* + \mathbf{V}_{\boldsymbol{\psi}}^{-1}$ and $\mathbf{g} = \mathbf{X}^{*\top}\tilde{\boldsymbol{\Lambda}}^{-1}\mathbf{y} + \mathbf{V}_{\boldsymbol{\psi}}^{-1}\boldsymbol{\mu}_{\boldsymbol{\psi}}$ with $\mathbf{V}_{\boldsymbol{\psi}}^{-1} = \text{diag}(\mathbf{V}_{\boldsymbol{\beta}}^{-1}, \lambda \cdot \mathbf{P})$ and $\mathbf{g} = [\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\mu}_{\boldsymbol{\beta}_S}]^\top = \mathbf{0}^\top$. Moreover, the Gamma prior on λ implies a Gamma full-conditional distribution $\lambda|\cdot \propto \mathcal{G}(\lambda|\alpha_\lambda + 1/2, \beta_\lambda + \boldsymbol{\beta}_S^\top \mathbf{P} \boldsymbol{\beta}_S)$.

Estimating the model in (8) is achieved through a straightforward extension of Algorithm 1. We jointly update $\boldsymbol{\psi}$ and λ from their full conditional distributions. In particular, the Gibbs' sampling step of Algorithm 1 can be modified to get full inference also on the spline coefficients $\boldsymbol{\beta}_S$ and the shrinkage parameter λ ; see Algorithm 2 in the Supplementary Material Alaimo Di Loro et al. (2023). In practical terms this requires J_S additional linear coefficients to be estimated, whose size $p^* = p + J_S$ may undermine the efficiency of the algorithm. For example, calculations in Step 1b are quadratic w.r.t. $p^* \rightarrow \mathcal{O}(np^{*2})$. Steps 1a and 1b (i.e., the most expensive in p^*) are executed in the first iteration and, subsequently, only in those iterations where new values of $\boldsymbol{\theta}$ are accepted. When $\boldsymbol{\theta}$ is rejected, we retain in memory the previously computed value (which would stay unchanged). Thus, if we attain an optimal acceptance rate of $\approx 20\% - 30\%$ in the Metropolis–Hastings step on $\boldsymbol{\theta}$, the computation is avoided in the majority of cases with a sensible improvement in computation time and speed.

3.5. Simulations. We conducted simulation experiments to evaluate the model described in Section 3.4 and compared the performance of our algorithm in terms of fitting, prediction error, and computational speed with other routines available from the `spNNGP` package (Finley, Datta and Banerjee (2017)). Additional comparative experiments are provided in the Supplementary Material (Alaimo Di Loro et al. (2023)). We executed our MCMC algorithms on a computing environment equipped with 12 modern computational nodes with 32 cores each, roughly equivalent to three TeraFlop/sec and 256 Gb of RAM. Each of the presented applications have been executed on a single node exploiting the computational power of all cores. The results presented are based upon posterior samples that were retained after diagnosing convergence using visual tools (e.g., traceplots, autocorrelation), effective sample sizes, Monte Carlo standard errors (MCSE), and other diagnostics offered by the `coda`, `mcse` and `bayesplot` packages in the R computing environment; the Supplementary Material (Alaimo Di Loro et al. (2023)) includes specific details.

We first generated $T_k = 2 \times 10^5$ time points for $K = 5$ individuals, where each time point t_{ki} followed exponential waiting times between observations, that is, $t_{ki} = \sum_{h=1}^{i-1} \delta_h$, and $\delta_h \stackrel{\text{iid}}{\sim} \text{Exp}(5)$. Given the time points, we constructed spatial trajectories $\boldsymbol{\gamma}_k(\cdot)$, $k = 1, \dots, K$, by simulating $\mathbf{s}_k = [\gamma_k(t_{k1}), \dots, \gamma_k(t_{kT_k})]^\top$, where subsequent components were independent Gaussian random walks over the square $\mathcal{S} = (1, 10) \times (1, 10)$ with the variance of each step along the horizontal and vertical axis proportional to the elapsed time between two subsequent observations. If the trajectory left the square, it was projected onto the border, and the next step would resume from there. The simulated trajectories are shown in Figure 5(a).

Given the time points and positions (Figure 5(b)), we generated the latent temporal Gaussian processes $w_k(\cdot) \stackrel{\text{iid}}{\sim} \mathcal{GP}(0, c_\theta(\cdot, \cdot))$ with an exponential covariance $c_\theta(t, t') = \sigma^2 \exp\{-\phi \cdot |t - t'|\}$, where $\sigma^2 > 0$ represents the variance of the process, $\phi > 0$ is the decay in temporal correlation (range), and $\tau^2 > 0$ the residual variance (nugget). The spatial effects are then introduced through $f_S(\cdot) : \mathcal{S} \rightarrow \mathbb{R}$ by considering a tensor product spline basis of degree 2 and with *nine* knots over the square domain (including boundary knots), where the spline coefficients $\boldsymbol{\beta}_S$ have been fixed to randomly generated values from $\mathcal{N}_{81}(\mathbf{0}, \lambda \mathbf{I}_{81})$ with $\lambda = 0.5$. The model also included individual-specific intercepts $\{\beta_{0k}\}_{k=1}^5$ and the effect of three covariates with random values drawn independently at each location from a $\mathcal{N}(0, 1)$ distribution, leading to covariate vectors $\{\mathbf{x}_{ki}\}_{i=1}^{T_k}$, $k = 1, \dots, K$. The effect of the covariates is assumed common across individuals and set to be determined by slopes $\boldsymbol{\beta} = [\beta_1, \beta_2, \beta_3]^\top$.

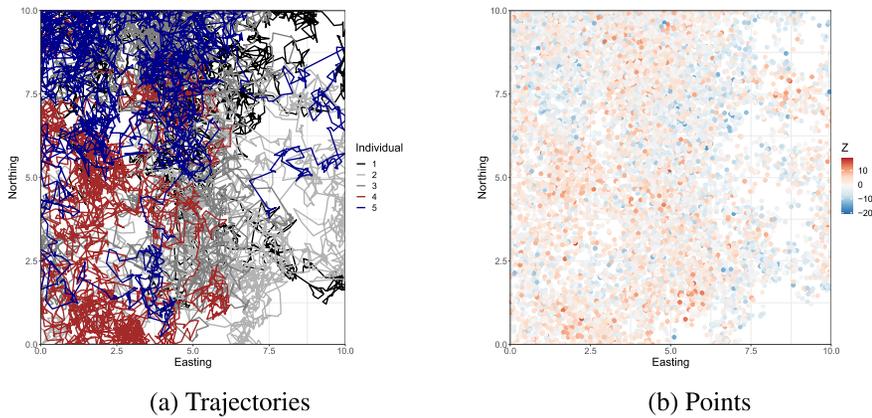


FIG. 5. Observed trajectories (a) and observed points (b) for the simulated dataset.

We generated values of the outcome for individual k at time t_{ki} and location $s_{ki} = \gamma_k(t_{ki})$, according to the generative process, defined by (8), with parameters fixed as above. This yielded a simulated dataset $D_{\text{sim}} = \{(\text{Ind}_j, t_j, s_j, y_j, \mathbf{x}_j^\top)\}_{j=1}^n$ with $n = 10^5$ observations, where Ind_j denotes the individual corresponding to row j . Then we fit the model in (9) on 70% of the total observations in D_{sim} . The remaining 30% were held out to assess out-of-sample predictive performances in terms of *Relative and Root Mean Squared Prediction Error* (RMSPE), *Coverage*, and *Predictive Interval Width* (PIW). Intercept and slope regression parameters were assigned $\mathcal{N}(0, 10^6)$ priors; the variance components, σ^2 and τ^2 , were both assigned inverse Gamma $\mathcal{IG}(2, 2)$ priors, and the decay parameter ϕ received a Gamma prior $\mathcal{G}(1, 1)$. For the spline coefficients, we considered both the penalized versions in Section 3.4. The first is referred to as an S-Spline (shrinking splines) and the second as P-Spline (penalized splines).

Table 1 presents the posterior estimates. We also included the deviance information criterion (DIC) for both models. Performances in the two settings are almost identical, but the DIC favors the S-Spline model. This is not surprising, as the data were generated using an analogous shrinkage prior for the β_S 's. Further details, including the estimates of the spline coefficients, are provided in the Supplementary Material (Alaimo Di Loro et al. (2023)). Figure 6 presents the posterior estimate of the spatial surface. We compare the true latent surface with just the S-Splines, as it performs slightly better with respect to the DIC, but notice that P-Splines provides practically identical estimates.

4. Application. We apply the proposed model in (8) to estimate the MAG (measured in G) for participants in the study, accounting for subject-specific features and spatial effects on the mean while modeling the latent temporal dependence, as described in Section 3.1. We split the data into training (70%) and testing (30%) subsets, where the records have been allocated to each subset according to a random sample stratified by individual. The testing set is used to assess the out-of-sample predictive performances in terms of *relative mean squared prediction error* (RMSPE), *root mean squared prediction error* (rMSPE), *coverage* (Cov), and *predictive interval width* (PIW). Posterior inferences are based on 5,000 samples retained after diagnosed convergence from 10,000 MCMC iterations.

4.1. Model specification. Spatial effects are introduced by considering the tensor product of two analogous univariate B-spline basis on each spatial axis. After a preliminary validation through the DIC (see Figure 6 of the Appendix), we choose two bases of degree 3 with 12 equally spaced knots over a square encompassing Westwood. This sums up to

TABLE 1

Parameter estimates, predictive validation, and fitting times (hours) on the simulated dataset for all of the considered models

Param. (True)	S-Spline		P-Spline	
	Point	Interval	Point	Interval
$\beta_{01}(-3.76)$	-3.799	(-3.846, -3.752)	-3.797	(-3.844, -3.75)
$\beta_{02}(0.65)$	0.572	(0.523, 0.62)	0.575	(0.526, 0.623)
$\beta_{03}(-0.60)$	-0.649	(-0.697, -0.6)	-0.646	(-0.693, -0.598)
$\beta_{04}(2.36)$	2.326	(2.277, 2.374)	2.328	(2.28, 2.376)
$\beta_{05}(-0.33)$	-0.359	(-0.408, -0.31)	-0.356	(-0.404, -0.308)
$\beta_1(2.59)$	2.599	(2.59, 2.608)	2.599	(2.59, 2.608)
$\beta_2(2.70)$	2.691	(2.683, 2.7)	2.691	(2.683, 2.7)
$\beta_3(-0.58)$	-0.586	(-0.595, -0.577)	-0.586	(-0.595, -0.577)
$\sigma^2(1)$	1.001	(0.973, 1.032)	0.993	(0.965, 1.023)
$\phi(1)$	0.994	(0.948, 1.04)	1.01	(0.964, 1.063)
$\tau^2(1)$	1.001	(0.984, 1.018)	1.001	(0.984, 1.018)

Metric	Out-of-sample	In-sample	Out-of-sample	In-sample
Coverage	0.95	0.99	0.95	0.99
RMSPE (r)	0.07 (1.18)	0.03 (0.84)	0.07 (1.19)	0.03 (0.84)
PIW	4.66	4.44	4.66	4.44
DIC		115'543		115'556
Fitting time (h)		2.18		2.2

$J_S = (9 + 3) \times (9 + 3) = 144$ terms for our complete spline basis, including the boundary knots. All numerical variables in $X(\cdot, \cdot)$ have been standardized for improving the efficiency of the MCMC sampling (Gilks and Roberts (1996)). The presence of temporal dependence in individual trajectories was investigated through an individual-specific exploratory analysis on the residuals from a standard linear regression, and an Ornstein–Uhlenbeck process (GP with an exponential covariance function) was specified to capture temporal dependence as a parsimonious and effective model for the behavior of the underlying residual process.

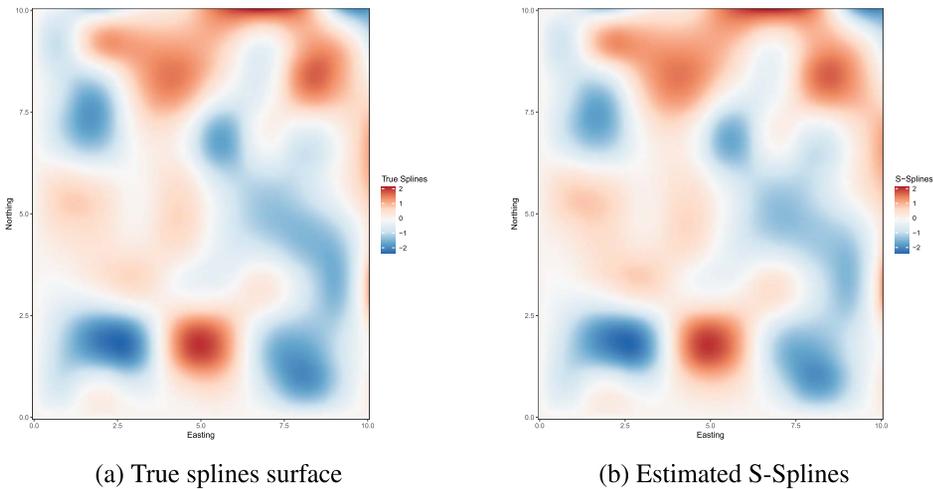


FIG. 6. True and estimated spline surfaces using S-Splines.

Finally, the outcome is log-transformed in order to comply with the Gaussianity assumption of the model. We denote the parameter associated with variable “varname” as β_{varname} and the levels of each categorical covariate as $\text{varname}_{(j)}$ for $j = 1, \dots, J_{\text{varname}}$. Hence,

$$\begin{aligned}
 & \mathbb{E}[\log(\text{MAG}_k(t))] \\
 &= \beta_0 + \sum_{j=2}^{J_{\text{BMI}}} \beta_{\text{BMI},j} \cdot \mathbb{I}(\text{BMI}_k = \text{BMI}_{(j)}) + \sum_{j=2}^{J_{\text{Sex}}} \beta_{\text{Sex},j} \cdot \mathbb{I}(\text{Sex}_k = \text{Sex}_{(j)}) \\
 &+ \sum_{j=2}^{J_{\text{Age}}} \beta_{\text{Age},j} \cdot \mathbb{I}(\text{Age}_k = \text{Age}_{(j)}) + \sum_{j=2}^{J_{\text{Eth}}} \beta_{\text{Eth},j} \cdot \mathbb{I}(\text{Eth}_k = \text{Eth}_{(j)}) \\
 (10) \quad &+ \beta_{\text{distHome}} \cdot \text{distHome}_k(\gamma_k(t)) + \beta_{\text{NDVI}} \cdot \text{NDVI}(\gamma_k(t)) \\
 &+ \beta_{\text{distParks}} \cdot \text{distParks}(\gamma_k(t)) + \beta_{\text{Slope}} \cdot \text{Slope}(\gamma_k(t)) \\
 &+ \sum_{j=1}^{J_S} \beta_{S,j} B_{S,j}(\gamma_k(t)),
 \end{aligned}$$

where $\mathbb{I}(\cdot)$ denotes the indicator function, $w_k(\cdot)$ is the DAG-based approximation (Section 3.2) for $\mathcal{GP}(0, c_\theta(\cdot, \cdot))$, and $\epsilon_k(t) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \tau^2)$. The baseline subject represents an underweight Asian female less than 18 years of age. Other socioeconomic factors (e.g., education and income level) have been excluded from the analysis, as they are strongly associated with ethnicity and age.

4.2. *Prior distributions.* The prior choices for each set of parameters and/or coefficients followed ad hoc strategies. We incorporated priors, such as $\beta \sim \mathcal{N}_J(\mathbf{0}, 10^6 \cdot \mathbf{I}_J)$, $\sigma^2 \sim \mathcal{IG}(2, 2)$, and $\tau^2 \sim \mathcal{IG}(2, 2)$ with J being the total number of β coefficients. The high spatial density of observations in several areas of the map enables robust estimation of the spatial effects. However, overfitting may emerge from the the high dimension of the spline basis. Furthermore, there are areas in Westwood that present sparsely observed data points, and the model could struggle to identify the spline coefficients referred to those areas and jeopardize convergence of the MCMC algorithm. Hence, we consider the S-Spline (Ridge-like prior), described in Section 3.4, to mitigate these potential issues, where the shrinkage parameter λ has been assigned a $\mathcal{G}(1, 1)$ prior.

4.3. *Results.* Fitting the model required ≈ 7 hours on a computer equipped with two processors AMD EPYC 7452, each one having 32 cores for a maximum of 256 Gb of RAM. The acceptance rate obtained is $\approx 28\%$, supporting the consistency of our adaptive strategy. Table 2 presents parameter estimates and performance metrics for the model in (8) with the explanatory variables, specified in (10), alongside estimates from a Bayesian linear regression model that includes the spatial spline terms but neglects the temporal dependence structure.

The estimates from the two models are largely consistent with each other, although accounting for temporal dependence tends to somewhat mitigate the effects of some predictors. We anticipate the temporal process to absorb the impact of certain predictors, especially when their relationship with the dependent variable is complex and nonlinear; this appears to be the case with “Slope”, which loses its significant positive impact on MAG once the temporal process is incorporated. All other variables seem to retain the nature of their impact on MAG. These coefficients are interpreted with respect to the baseline measure of an underweight 18-year-old Asian female.

The intercept represents the natural logarithm of the MAG for the baseline subject and reckons with both active as well as inactive time points for the high-resolution digitally filtered data. Since we model continuously over time, the resulting value of the MAG, which is about $\exp(-2.9) \approx 0.06$, corresponds to the low region of moderate physical activity level; see the Supplementary Material, [Alaimo Di Loro et al. \(2023\)](#). Other gleanings from Table 2 indicate that MAGs vary by ethnicity in the study cohort, as Whites tend to record lower MAGs while Latin Americans and African Americans tend to register larger MAGs compared to Asians. Impact of Age groups on MAGs also tend to vary with the groups of 25–34 and 45–70-year-old subjects tending to register lower MAGs than the baseline (young), while the middle-aged group tends to be higher than the baseline. This is not entirely surprising because subjects in the 25–34-year-old group tend to exercise less than the younger (baseline) and middle-aged groups with 25–34-year-old subjects having less time as they embark on their careers (less free time), while those in the 45–70 range also tend to follow a less vigorous lifestyle regimen due to their age. The effect of Body Mass Index (BMI) is also seen to vary based upon the categories of weight. While all three categories indicate a significantly higher MAG, compared to the baseline, the impact of the overweight, but not obese, category seems higher than the other two. We do not fully know the extent to which larger body weights affect accelerometer readings, but this variation must also account for the fact that higher BMI may also correspond to muscular (not unfit) individuals engaging in more vigorous lifestyle regimens. The spatially-indexed predictors indicated the expected positive impact of NDVI (more greenness encourages more outdoor activities and exercise) while it is also expected, especially in Westwood, that subjects tend to exercise along paths closer to their home thereby explaining the negative coefficient for the weighted distance to home.

The estimate of the temporal decay parameter ϕ implies that the temporal correlation drops to 0.05 in about $3/\hat{\phi} \approx 4.3$ minutes, where $\hat{\phi} \approx 0.7$ is the posterior median of ϕ . Unsurprisingly, including the spatial effect and the temporal process improves predictive performances (RMSPE or PIW in Table 2) over a model including only spatial effects (excluding the temporal process). The model incorporating the temporal process delivers satisfactory coverage and outperforms its competitor in all of the other indices for the training and testing data.

Figure 7(a) shows the estimated spatial surface, while Figure 7(b) presents the width of the posterior predictive intervals. The map clearly evinces zones (darker shades of red highlighted with white contours) that tend to depict high levels of physical activity. For example, the largest dark red blob in the north center-left almost perfectly tracks the UCLA campus boundary, reflecting a campus environment with active mobility (walking, running, biking). Other zones of high activity identify with locations where more participants in the study live, including those residing in student dorms (northwest corner) and residential areas immediately around and in the predefined Westwood/UCLA study area (such as the south central zone) or Century City shopping center (to the east). Lighter shades (orange) correspond to areas that are less developed (open space), such as the areas in the north east, or they are areas with a high degree of transportation infrastructure and traffic (e.g., toward the western boundary). These correspond to highways (such as the Interstate 405 highway or other vehicular transportation corridors) that often have lower levels of activity because they inhibit outdoor physical activities due to noise, pollution, safety, etc. Our analysis reveals three additional high-activity areas that are not gleaned from nonspatial models: the Los Angeles National Veteran Park; the Century City shopping center, and the Stone Canyon Park. The color gradient closely follows the spatial characteristics of the Westwood neighborhood and reveal how spatial patterns can impact physical activity behavior after accounting for variation attributable to known explanatory variables.

Figure 8 shows two examples of observed (left) and reconstructed (right) MAGs along trajectories carved out by two subjects. We find a good degree of agreement between the two

TABLE 2

Parameter estimates and model performance metrics for model (8) with and without the temporal process

Param.	Model (8) without temporal process		Model (8) with temporal process	
	Point	Interval	Point	Interval
Intercept	-2.750	(-2.754, -2.746)	-2.92	(-2.94, -2.91)
Eth. White	-0.128	(-0.146, -0.111)	-0.190	(-0.258, -0.125)
Eth. Other	0.122	(0.110, 0.134)	0.128	(0.077, 0.178)
Eth. Latin-American	0.259	(0.247, 0.271)	0.314	(0.264, 0.362)
Eth. Black/African/Caribbean	0.263	(0.248, 0.278)	0.400	(0.340, 0.461)
Sex Male	-0.348	(-0.358, -0.338)	-0.298	(-0.338, -0.258)
Normal weight	0.121	(0.110, 0.132)	0.297	(0.252, 0.343)
Over weight	0.351	(0.330, 0.372)	0.482	(0.398, 0.566)
Obese	0.220	(0.181, 0.258)	0.401	(0.241, 0.560)
Age (25-34]	-0.387	(-0.398, -0.377)	-0.320	(-0.362, -0.279)
Age (34-45]	0.080	(0.064, 0.097)	0.125	(0.057, 0.191)
Age (45-70]	-0.105	(-0.132, -0.079)	-0.091	(-0.192, 0.006)
Dist. from home	-0.135	(-0.142, -0.128)	-0.074	(-0.102, -0.046)
Slope	0.052	(0.047, 0.056)	-0.003	(-0.12, 0.005)
Dist. to parks	-0.221	(-0.227, -0.214)	-0.066	(-0.089, -0.043)
NDVI	0.226	(0.221, 0.231)	0.010	(0.004, 0.015)
σ^2			2.266	(2.237, 2.297)
ϕ			0.718	(0.704, 0.731)
τ^2	2.10	(2.08, 2.13)	0.050	(0.048, 0.053)

Metric	Out-of-sample	In-sample	Out-of-sample	In-sample
DIC		17'588'058		973'329
Coverage	0.95	0.95	0.93	0.99
RMSPE (r)	1.44 (0.68)	1.44 (0.68)	0.55 (0.09)	0.1 (0.003)
PIW	5.69	5.69	2	1.18

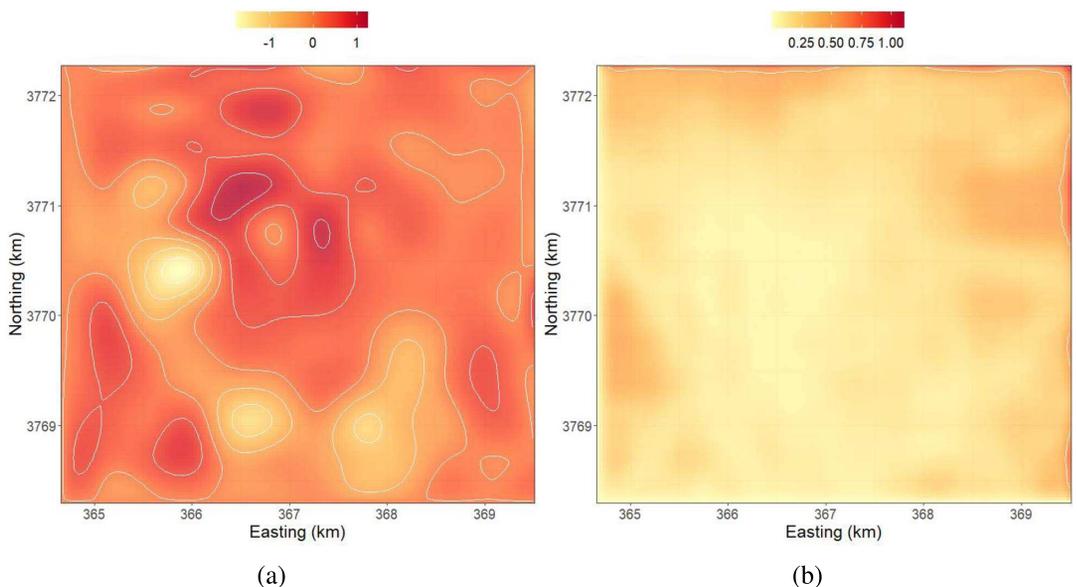


FIG. 7. (a) Spatially smoothed estimates from a shrinkage spline over Westwood, Los Angeles; (b) Standard deviation for the shrinkage spline.

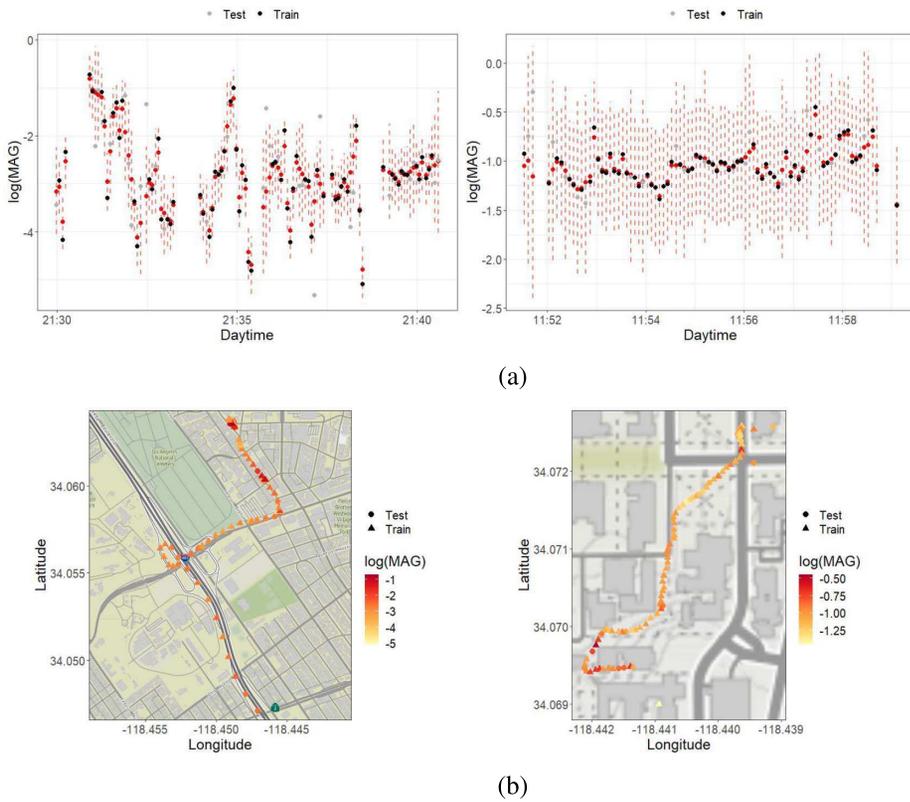


FIG. 8. Estimated $\log(\text{MAG})$ for two randomly selected individuals: (a) estimated $\log(\text{MAG})$ (red points) and 95% prediction intervals (red dashed line) for each point within the observed time-windows; (b) including the location.

plots, and the ability of our model to recover the $\log(\text{MAG})$ in locations where it has not been observed. The reliability of the predictions can be demonstrated through different metrics and, unsurprisingly, accommodating spatial effects and the temporal process improves predictive performances, as measured by MSPE or PIW. We deliver these personalized trajectory plots for every subject in the study and also predict personalized MAGs for each subject along any new trajectory. This enables personalized recommendations based upon an individual's health attributes, including suggestions for more effective paths to follow for optimal physical activities, while also informing community level interventions in the built environment.

5. Discussion. We have devised a Bayesian modeling framework to conduct fully model-based inference for high-resolution accelerometer data over trajectories compiled from the PASTA-LA study. Our key data analytic developments included: (i) modeling dependence over trajectories, (ii) accounting for subject-specific spatial-temporal variation for daily mobility, (iii) predicting or interpolating PA levels across trajectories and (iv) identifying zones of high physical activity in Westwood, Los Angeles. Our spatiotemporal analysis offers richer inference and reveals relationships between physical activity levels and a variety of factors, both at the subject level (e.g., personal attributes) and as a function of space and time. The temporal process was able to effectively extract the features of the data at finer resolutions, while the spatial splines accounted for residual spatial heterogeneity. Accommodating both temporal dependence and spatial heterogeneity demonstrably improved predictive ability and enabled us to effectively delineate zones of high physical activity. Furthermore, the ability of the model to pool information across individuals at all time points allows us to infer about those who present sparsely observed space-time points (due to technical issues or protocol

violation). In particular, we can interpolate and infer about PA levels with full uncertainty quantification and ensure the desired coverage by our prediction intervals. The methods we develop can be adapted to model animal tracking and be compared to existing spatial models (see, e.g., [Hedley and Buckland \(2004\)](#)).

Recent public health reviews call for interdisciplinary technological advances to more effectively measure spatiotemporal energetics of activity spaces in obesity and chronic disease research ([James et al. \(2016\)](#), [Kestens et al. \(2017\)](#), [Drewnowski et al. \(2020\)](#)). Individual-level data, at aggregate, can be used to identify anchor points for physical activity and reveal causal pathways between built environment exposures and health. Our work is a novel contribution demonstrating methodologies to answer these pressing research questions.

Our analysis also resolves practical difficulties in using actigraph data. It is not cost-effective to deploy research-grade GlobalSat GPS and Actigraph units, as they are very expensive and continued usage requires heavy staff involvement. Our methods can be applied to analogous, but less complete, data derived from smart phones and smart watches, then such devices could be deployed in much larger studies with much larger sample sizes at a fraction of the cost. Given the spatiotemporal nature of outdoor PA research, our ability to predict in areas of data missingness drastically improve inference related to the impacts of the built and natural environments on physical activity and active mobility.

We recognize that there are several avenues for further research. Substantive investigations pertaining to the PASTA-LA study will focus on the impact of intervention schemes designed to promote physical activities and ask questions related to controlling for weather while estimating the impact of the intervention. Our DAG-based approach for scalable temporal processes can be further enriched with recent developments ([Katzfuss and Guinness \(2021\)](#), [Peruzzi, Banerjee and Finley \(2022\)](#)), although any of the methods reviewed and evaluated by [Heaton et al. \(2019\)](#) can be incorporated into our framework. We also recognize a wealth of future research surrounding wearable devices and actigraphy data. Examples include methodological advancements in clustering of trajectories, according to different levels of physical activity, and creating personalized health recommendation systems for patients with regard to trajectories (e.g., walking, running, or biking routes) that will be most appropriate for them. Related to the clustering of trajectories, one can also pursue model-based learning about individual effects from the extent of (appropriately quantifies) spatial overlap in trajectories and discerning them from spatial effects. Finally, there is possible merit in modeling both the nonidle and idle times with a more comprehensive hierarchical specification ([Bai et al. \(2018\)](#)). The joint modeling could be achieved using Mixture Models, Hidden Markov Models, or the modeling of multivariate Gaussian censored outcomes ([De Oliveira \(2005\)](#), [Molstad, Hsu and Sun \(2021\)](#)). Combining such approaches with efficient estimation strategies is the major challenge, which will be tackled in future developments of this work.

Acknowledgments. The authors thank the Editor, Associate Editor, and two anonymous reviewers for several helpful comments and suggestions.

Funding. Sudipto Banerjee was supported, in part, by National Science Foundation (NSF) under grants DMS-2113778, DMS-1916349, and IIS-1562303. Sudipto Banerjee and Michael B. Jerrett have been supported by the National Institute of Environmental Health Sciences (NIEHS) under grants R01ES030210 and 5R01ES027027. The authors also acknowledge support from the NIOSH Education Research Center, the Center for Occupational and Environmental Health, and the UCLA Department of Transportation for this work. Finally, the authors acknowledge support for the survey administration from the EU Physical Activity through Sustainable Transport Approaches (PASTA) team members <https://www.pastaproject.eu/>. Pierfrancesco Alaimo Di Loro was supported, in part, by the Fulbright Scholarship Program, category *Graduate Study - Visiting Student Researcher*

SUPPLEMENTARY MATERIAL

Algorithms, simulations and data analysis (DOI: [10.1214/23-AOAS1742SUPPA](https://doi.org/10.1214/23-AOAS1742SUPPA); .pdf). Supplementary Material is available in [Alaimo Di Loro et al. \(2023\)](#). It includes further information about data processing, relationships between vector magnitude of acceleration and metabolic equivalent of task, technical details on some of the algorithms using the temporal NNGP, additional simulation experiments and some further analysis of the actigraph data.

Computer programs (DOI: [10.1214/23-AOAS1742SUPPB](https://doi.org/10.1214/23-AOAS1742SUPPB); .zip). Computer programs developed for implementing the models in the paper for the R statistical computing environment are available as Supplementary Material in the form of a compressed folder `EfficientTNNGPforActigraph-main.zip`. This can also be downloaded from a GitHub repository <https://github.com/minmar94/EfficientTNNGPforActigraph>.

REFERENCES

- ALAIMO DI LORO, P., MINGIONE, M., LIPSITT, J., BATTEATE, C. M., JERRETT, M. and BANERJEE, S. (2023). Supplement to “Bayesian hierarchical modeling and analysis for actigraph data from wearable devices.” <https://doi.org/10.1214/23-AOAS1742SUPPA>, <https://doi.org/10.1214/23-AOAS1742SUPPB>
- BAI, J., SUN, Y., SCHRACK, J. A., CRAINICEANU, C. M. and WANG, M.-C. (2018). A two-stage model for wearable device data. *Biometrics* **74** 744–752. MR3825361 <https://doi.org/10.1111/biom.12781>
- BAMMANN, K., THOMSON, N. K., ALBRECHT, B. M., BUCHAN, D. S. and EASTON, C. (2021). Generation and validation of ActiGraph GT3X+ accelerometer cut-points for assessing physical activity intensity in older adults. The OUTDOOR ACTIVE validation study. *PLoS ONE* **16** e0252615. <https://doi.org/10.1371/journal.pone.0252615>
- BANERJEE, S. (2017). High-dimensional Bayesian geostatistics. *Bayesian Anal.* **12** 583–614. MR3654826 <https://doi.org/10.1214/17-BA1056R>
- BULL, F. C., AL-ANSARI, S. S., BIDDLE, S., BORODULIN, K., BUMAN, M. P., CARDON, G., CARTY, C., CHAPUT, J.-P., CHASTIN, S. et al. (2020). World Health Organization 2020 guidelines on physical activity and sedentary behaviour. *Br. J. Sports Med.* **54** 1451–1462.
- CRESSIE, N. and WIKLE, C. K. (2011). *Statistics for Spatio-Temporal Data*. Wiley Series in Probability and Statistics. Wiley, Hoboken, NJ. MR2848400
- CROUTER, S. E., CLOWERS, K. G. and BASSETT JR, D. R. (2006). A novel method for using accelerometer data to predict energy expenditure. *J. Appl. Physiol.* **100** 1324–1331.
- DATTA, A., BANERJEE, S., FINLEY, A. O. and GELFAND, A. E. (2016a). Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *J. Amer. Statist. Assoc.* **111** 800–812. MR3538706 <https://doi.org/10.1080/01621459.2015.1044091>
- DATTA, A., BANERJEE, S., FINLEY, A. O., HAMM, N. A. S. and SCHAAP, M. (2016b). Nonseparable dynamic nearest neighbor Gaussian process models for large spatio-temporal data with an application to particulate matter analysis. *Ann. Appl. Stat.* **10** 1286–1316. MR3553225 <https://doi.org/10.1214/16-AOAS931>
- DE OLIVEIRA, V. (2005). Bayesian inference and prediction of Gaussian random fields based on censored data. *J. Comput. Graph. Statist.* **14** 95–115. MR2137892 <https://doi.org/10.1198/106186005X27518>
- DOHERTY, A., JACKSON, D., HAMMERLA, N., PLÖTZ, T., OLIVIER, P., GRANAT, M. H., WHITE, T., VAN HEES, V. T., TRENELL, M. I. et al. (2017). Large scale population assessment of physical activity using wrist worn accelerometers: The UK biobank study. *PLoS ONE* **12**.
- DREWNOWSKI, A., BUSZKIEWICZ, J., AGGARWAL, A., ROSE, C., GUPTA, S. and BRADSHAW, A. (2020). Obesity and the built environment: A reappraisal. *Obesity* **28** 22–30.
- EILERS, P. H. C. and MARX, B. D. (1996). Flexible smoothing with *B*-splines and penalties. *Statist. Sci.* **11** 89–121. MR1435485 <https://doi.org/10.1214/ss/1038425655>
- FINLEY, A., DATTA, A. and BANERJEE, S. (2017). Spnngp: Spatial regression models for large datasets using nearest neighbor Gaussian processes. R Package Version 0.1.1.
- FINLEY, A. O., DATTA, A., COOK, B. D., MORTON, D. C., ANDERSEN, H. E. and BANERJEE, S. (2019). Efficient algorithms for Bayesian nearest neighbor Gaussian processes. *J. Comput. Graph. Statist.* **28** 401–414. MR3974889 <https://doi.org/10.1080/10618600.2018.1537924>
- FREEDSON, P., BOWLES, H. R., TROIANO, R. and HASKELL, W. (2012). Assessment of physical activity using wearable monitors: Recommendations for monitor calibration and use in the field. *Med. Sci. Sports Exerc.* **44** S1–S4. <https://doi.org/10.1249/MSS.0b013e3182399b7e>
- GELFAND, A. E., DIGGLE, P., GUTTORP, P. and FUENTES, M. (2010). *Handbook of Spatial Statistics*. CRC press.

- GILKS, W. R. and ROBERTS, G. O. (1996). Strategies for improving MCMC. *Markov Chain Monte Carlo in Practice* **6** 89–114.
- GOODMAN, T. and HARDIN, D. (2006). Refinable multivariate spline functions. In *Topics in Multivariate Approximation and Interpolation. Stud. Comput. Math.* **12** 55–83. Elsevier, Amsterdam. MR2410695 [https://doi.org/10.1016/S1570-579X\(06\)80005-4](https://doi.org/10.1016/S1570-579X(06)80005-4)
- HAARIO, H., SAKSMAN, E. and TAMMINEN, J. (2001). An adaptive Metropolis algorithm. *Bernoulli* **7** 223–242. MR1828504 <https://doi.org/10.2307/3318737>
- HASTIE, T. and TIBSHIRANI, R. (2000). Bayesian backfitting. *Statist. Sci.* **15** 196–223. MR1820768 <https://doi.org/10.1214/ss/1009212815>
- HEATON, M. J., DATTA, A., FINLEY, A. O. et al. (2019). A case study competition among methods for analyzing large spatial data. *J. Agric. Biol. Environ. Stat.* **24** 398–425. MR3996451 <https://doi.org/10.1007/s13253-018-00348-w>
- HEDLEY, S. L. and BUCKLAND, S. T. (2004). Spatial models for line transect sampling. *J. Agric. Biol. Environ. Stat.* **9** 181–199.
- JAMES, P., JANKOWSKA, M., MARX, C., HART, J. E., BERRIGAN, D., KERR, J., HURVITZ, P. M., HIPPI, J. A. and LADEN, F. (2016). “Spatial energetics”: Integrating data from GPS, accelerometry, and GIS to address obesity and inactivity. *Am. J. Prev. Med.* **51** 792–800. <https://doi.org/10.1016/j.amepre.2016.06.006>
- KATZFUSS, M. and GUINNESS, J. (2021). A general framework for Vecchia approximations of Gaussian processes. *Statist. Sci.* **36** 124–141. MR4194207 <https://doi.org/10.1214/19-ST575>
- KATZFUSS, M., GUINNESS, J., GONG, W. and ZILBER, D. (2020). Vecchia approximations of Gaussian-process predictions. *J. Agric. Biol. Environ. Stat.* **25** 383–414. MR4139037 <https://doi.org/10.1007/s13253-020-00401-7>
- KESTENS, Y., WASFI, R., NAUD, A. and CHAIX, B. (2017). “Contextualizing context”: Reconciling environmental exposures, social networks, and location preferences in health research. *Curr. Environ. Health Rep.* **4** 51–60. <https://doi.org/10.1007/s40572-017-0121-8>
- KHUSAINOV, R., AZZI, D., ACHUMBA, I. E. and BERSCH, S. D. (2013). Real-time human ambulation, activity, and physiological monitoring: Taxonomy of issues, techniques, applications, challenges and limitations. *Sensors* **13** 12852–12902. <https://doi.org/10.3390/s131012852>
- LANG, S. and BREZGER, A. (2004). Bayesian P-splines. *J. Comput. Graph. Statist.* **13** 183–212. MR2044877 <https://doi.org/10.1198/1061860043010>
- LAURITZEN, S. L. (1996). *Graphical Models. Oxford Statistical Science Series* **17**. The Clarendon Press, New York. MR1419991
- LIU, J. S., WONG, W. H. and KONG, A. (1994). Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes. *Biometrika* **81** 27–40. MR1279653 <https://doi.org/10.1093/biomet/81.1.27>
- LYDEN, K., KEADLE, S. K., STAUDENMAYER, J. and FREEDSON, P. S. (2014). A method to estimate free-living active and sedentary behavior from an accelerometer. *Med. Sci. Sports Exerc.* **46** 386–397. <https://doi.org/10.1249/MSS.0b013e3182a42a2d>
- MADDISON, R., HOORN, S. V., JIANG, Y., MHURCHU, C. N., EXETER, D., DOREY, E., BULLEN, C., UTTER, J., SCHAAF, D. et al. (2009). The environment and physical activity: The influence of psychosocial, perceived and built environmental factors. *Int. J. Behav. Nutr. Phys. Act.* **6** 19.
- MATHIE, M. J., COSTER, A. C. F., LOVELL, N. H. and CELLER, B. G. (2003). Detection of daily physical activities using a triaxial accelerometer. *Med. Biol. Eng. Comput.* **41** 296–301. <https://doi.org/10.1007/BF02348434>
- MIGUELES, J. H., CADENAS-SANCHEZ, C., EKELEND, U., NYSTRÖM, C. D., MORA-GONZALEZ, J., LÖF, M., LABAYEN, I., RUIZ, J. R. and ORTEGA, F. B. (2017). Accelerometer data collection and processing criteria to assess physical activity and other outcomes: A systematic review and practical considerations. *Sports Med.* **47** 1821–1845. <https://doi.org/10.1007/s40279-017-0716-0>
- MIGUELES, J. H., CADENAS-SANCHEZ, C., ROWLANDS, A. V., HENRIKSSON, P., SHIROMA, E. J., ACOSTA, F. M., RODRIGUEZ-AYLLON, M., ESTEBAN-CORNEJO, I., PLAZA-FLORIDO, A. et al. (2019). Comparability of accelerometer signal aggregation metrics across placements and dominant wrist cut points for the assessment of physical activity in adults. *Sci. Rep.* **9** 1–12.
- MOLSTAD, A. J., HSU, L. and SUN, W. (2021). Gaussian process regression for survival time prediction with genome-wide gene expression. *Biostatistics* **22** 164–180. MR4207149 <https://doi.org/10.1093/biostatistics/kxz023>
- MURPHY, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press, Cambridge.
- PERUZZI, M., BANERJEE, S. and FINLEY, A. O. (2022). Highly scalable Bayesian geostatistical modeling via meshed Gaussian processes on partitioned domains. *J. Amer. Statist. Assoc.* **117** 969–982. MR4436326 <https://doi.org/10.1080/01621459.2020.1833889>

- PIERCY, K. L., TROIANO, R. P., BALLARD, R. M., CARLSON, S. A., FULTON, J. E., GALUSKA, D. A., GEORGE, S. M. and OLSON, R. D. (2018). The physical activity guidelines for Americans. *JAMA* **320** 2020–2028.
- PLASQUI, G. and WESTERTERP, K. R. (2007). Physical activity assessment with accelerometers: An evaluation against doubly labeled water. *Obesity* **15** 2371–2379. <https://doi.org/10.1038/oby.2007.281>
- RAMSAY, J. O. and SILVERMAN, B. W. (2007). *Applied Functional Data Analysis: Methods and Case Studies*. Springer, Berlin.
- RAY, E. L., SASAKI, J. E., FREEDSON, P. S. and STAUDENMAYER, J. (2018). Physical activity classification with dynamic discriminative methods. *Biometrics* **74** 1502–1511. MR3908166 <https://doi.org/10.1111/biom.12892>
- REINER, M., NIERMANN, C., JEKAUC, D. and WOLL, A. (2013). Long-term health benefits of physical activity—a systematic review of longitudinal studies. *BMC Public Health* **13** 1–9.
- ROBERTS, G. O. and ROSENTHAL, J. S. (2009). Examples of adaptive MCMC. *J. Comput. Graph. Statist.* **18** 349–367. MR2749836 <https://doi.org/10.1198/jcgs.2009.06134>
- SIKKA, R. S., BAER, M., RAJA, A., STUART, M. and TOMPKINS, M. (2019). Analytics in sports medicine: Implications and responsibilities that accompany the era of big data. *JBJS* **101** 276–283.
- STAUDENMAYER, J., HE, S., HICKEY, A., SASAKI, J. and FREEDSON, P. (2015). Methods to estimate aspects of physical activity and sedentary behavior from high-frequency wrist accelerometer measurements. *J. Appl. Physiol.* **119** 396–403.
- STEIN, M. L., CHI, Z. and WELTY, L. J. (2004). Approximating likelihoods for large spatial data sets. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **66** 275–296. MR2062376 <https://doi.org/10.1046/j.1369-7412.2003.05512.x>
- TARALDSEN, K., CHASTIN, S. F., RIPHAGEN, I. I., VEREIJKEN, B. and HELBOSTAD, J. L. (2012). Physical activity monitoring by use of accelerometer-based body-worn sensors in older adults: A systematic literature review of current knowledge and applications. *Maturitas* **71** 13–19.
- VAN HEES, V. T., RENSTRÖM, F., WRIGHT, A., GRADMARK, A., CATT, M., CHEN, K. Y., LÖF, M., BLUCK, L., POMEROY, J. et al. (2011). Estimation of daily energy expenditure in pregnant and non-pregnant women using a wrist-worn tri-axial accelerometer. *PLoS ONE* **6** e22922.
- VAN LOO, C. M., OKELY, A. D., BATTERHAM, M. J., HINKLEY, T., EKELUND, U., BRAGE, S., REILLY, J. J., TROST, S. G., JONES, R. A. et al. (2018). Wrist acceleration cut-points for moderate-to-vigorous physical activity in youth. *Med. Sci. Sports Exerc.* **50** 609.
- VECCHIA, A. V. (1988). Estimation and model identification for continuous spatial processes. *J. Roy. Statist. Soc. Ser. B* **50** 297–312. MR0964183
- WHITE, T., WESTGATE, K., WAREHAM, N. J. and BRAGE, S. (2016). Estimation of physical activity energy expenditure during free-living from wrist accelerometry in UK adults. *PLoS ONE* **11** e0167472. <https://doi.org/10.1371/journal.pone.0167472>