EARLY IDENTIFICATION OF AN IMPENDING ROCKSLIDE LOCATION VIA A SPATIALLY-AIDED GAUSSIAN MIXTURE MODEL

By Shuo Zhou[1,*], Howard Bondell[1,†], Antoinette Tordesillas[1,‡], Benjamin I. P. Rubinstein[2,§] and James Bailey[2,¶]

[1]*School of Mathematics and Statistics, The University of Melbourne,* *[*]shuo.zhou@unimelb.edu.au;*
[†]*howard.bondell@unimelb.edu.au;* [‡]*atordesi@unimelb.edu.au*

[2]*School of Computing and Information Systems, The University of Melbourne,* [§]*brubinstein@unimelb.edu.au;*
[¶]*baileyj@unimelb.edu.au*

Movement of soil and rocks in an unstable slope under gravitational forces is an example of a complex system that is highly dynamic in space and time. A typical failure in such systems is a landslide. Fundamental studies of granular media failure combined with a complex network analysis of radar monitoring data show that distinct partitions emerge in the kinematic field in the early stages of the prefailure regime, and these patterns yield clues to the ultimate location of failure. In this study we address this partitioning of constituent units in complex systems by clustering the kinematic data, specifically, with a Gaussian mixture model. In addition, we assume that neighboring units should move together. As a result, spatial information is taken into account in our model so that spatial proximity is retained. Our case study of a rockslide from high resolution radar monitoring data shows that, by incorporating spatial information, our approach is more effective in revealing the dynamics of the system and detecting the location of a potential landslide, compared to the use of only the kinematics.

1. Introduction. Slope failures, such as landslides, are a class of geological hazards that involve the downslope movement of soil and rocks under gravitational forces. The most common and effective technique for determining slope stability (i.e., the resistance of a slope to failure) is by monitoring the small precursory movements of its surface prior to collapse. Modern early warning systems on both natural and man-made slopes that seek to detect these precursory movements have become key components of risk management programs. However, although these monitors produce a large amount of data, effective use of this data is still lagging behind the technology. For example, slope stability radar (SSR) can remotely scan a surface and detect movements with submillimeter precision at high spatial and temporal resolutions. Yet, typical use of this data is based on analysis of displacement and displacement rates at just one, or a few, suspected locations.

Slope failure is not instantaneous. There exists precursory failure regimes preceding failure where patterns emerge that give clues to the precise location of failure. Earlier study (Tordesillas, Zhou and Batterham (2018)) suggests that such data can be organized as a time-evolving network in the kinematic space and distinct patterns can be identified. Specifically, at eventual failure the slip region moves together in rigid body motion as it falls, while the nonslip region maintains its stable flow, with a narrow shear band that exhibits more chaotic behavior. With this in mind, using real-time measurements of the surface kinematics monitored on a spatial grid, the detection of the location of a potential landslide points to discovering these types of clustering patterns in the kinematics through space and time. As opposed to examining the displacement data via a network structure and partitioning in that space, we

propose to cluster based on these surface kinematic observations to form the basis for determining their behavior. Using clustering to reveal such behavior can then, potentially, lead to early prediction of the location of the eventual failure.

We examine landslide data from an open-pit mine, in which displacement was measured as a univariate quantity. For this dataset we anticipate the nonslip region to have small movements, and the slip region to have larger motion. We thus use a two-component Gaussian mixture to represent this phenomenon at each time step.

In addition to the kinematic data, spatial information is a readily easy to access data source and can assist in guiding the identification of clusters. The challenge here is to appropriately balance between kinematic similarity and spatial proximity in the clustering process. To incorporate spatial information into the clustering algorithm, we additionally model the spatial distribution of the pixels within each of the clusters. Given that these spatial distributions are complex, that is, the boundaries of landslides can take very complex shapes, we use an embedded Gaussian mixture model on the two-dimensional spatial coordinates within each of the first-level clusters. The reason for the mixture of Gaussians within each first-level cluster is to allow for the spatial regions to have more complex geometries, as a mixture of Gaussians can yield any desired shape as the number of mixture components within each first-level cluster increases. Such a mixture of mixture configuration is beneficial to produce locally-smooth yet not over-simplified clustering patterns in the spatial domain. In this work we are interested in the clustering pattern at each time stamp and how this varies over time. We do anticipate that the clusters should vary relatively slowly with time, and a direct approach to account for this without assuming a model for this behavior is to use a sliding window over the past time points. An alternative way would be to fit a more complex evolution model over time. However, in addition to having to choose a model assumption on the evolution, it would also require using all of the data at once in order to fit the temporal model which can be time consuming and infeasible. We show how combining kinematic and spatial data to build analytic tools for landslide monitoring, or more broadly speaking, how to discover patterns in one feature space while encouraging coherence in a separate domain.

We demonstrate our approach by providing a comprehensive characterization of patterns in ground motions from a high-resolution (spatially and temporally dense) dataset for a large, real-world cascading failure—a rockfall. Our approach resolves some of the outstanding challenges in the area of landslide hazard forecasting and early prediction raised by Wasowski and Bovenga (2014). As opposed to the prevailing approach of analyzing the time course of data from one or a few selected locations, our analysis makes full use of data from all monitored points (here, there are 1803 spatial locations). In addition, compared to existing studies where the complex spatial patterns are often overlooked, in our analysis we characterize the complex spatial patterns by performing a clustering of the ground surface deformation feature space which encourages spatial coherence via a flexible spatial distribution. Our analysis is able to capture the patterns of localized deformations preceding failure, even though these are small and vary nonlinearly with space and time. Crucially, we were able to extract useful insights that reliably predict the location of the yet-to-form failure and its complex dynamics in the lead up to the collapse. In particular, we were able to capture the interaction between multiple unstable zones (here, two zones) and distinguish which zone is likely to collapse based on the spatiotemporal persistence of its pattern. The presence of multiple "competing" unstable zones is a defining aspect of granular media failure in the precursory regime (Tordesillas et al. (2013)). In the field this can lead to forecasting errors with costly and, at times, fatal consequences: a collapse may be missed (false negative) with attention focused elsewhere while a false positive incurs the costs of a risk investigation. Thus, tools that can deliver accurate knowledge of where failure might occur and distinguish it from other unstable zones would help geologists and geotechnical engineers provide more reliable and effective risk assessment and prevention measures. Moreover, our proposed method can also assist decision

makers in: establishing emergency preparedness maps that indicate hazard zones, evacuation routes and safe shelters; designing appropriate remedial measures (e.g., removal of material, water drainage, adding retaining structures like barriers and walls in the hazard area); and potentially optimizing slope steepness for safety and economic performance in man-made slopes, such as those in open-pit mines where our studied landslide took place.

2. Related work. Statistical and machine learning approaches have a long history in the landslide literature. One area drawing increasing attention from statisticians and machine learning practitioners is landslide susceptibility modeling that aims to highlight potentially unstable slopes based on linking characteristics of a particular slope with past landslide events on similar slopes (Pourghasemi and Rahmati (2018)). Typically, a landslide inventory is maintained and terrain attributes, such as slope angle, elevation and profile curvature, are collected from one or few monitoring sites on the slope. Statistical and machine learning methods, such as generalized linear models (GLM) (Brenning (2005)), support vector machines (SVM) (Yao, Tham and Dai (2008), Bui et al. (2012)), random forests (RF) (Hong, Pourghasemi and Pourtaghi (2016)) and neural networks (NN) (Pham et al. (2017)), are applied to fit the historical data; the landslide susceptibility map can be generated as a classification problem. Compared to traditional physical models (Thanh and De Smedt (2014)), data driven approaches are less expensive and more applicable to large scale areas Pourghasemi and Rahmati (2018). More details and a comparison of approaches can be found in Goetz et al. (2015), Pourghasemi and Rahmati (2018) and a recent comprehensive review on this topic by Reichenbach et al. (2018). Our problem differs from a susceptibility study in that, rather than having location characteristics and determining if that location may be similar to those that had previous landslides, we focus on the next stage of management, that is, performing continuous monitoring of a previously identified susceptible site and examining the smaller scale spatial resolution characteristics over time. As a result, the aforementioned methods do not apply to this stage in the monitoring process. Instead, we address the problem with a clustering approach that reveals the dynamic changes of the system, given the high frequency measurements generated by our large amount of monitoring stations.

Another relevant area is landslide forecasting where time series generated by the monitoring stations are fitted and predicted in order to attempt to produce early warnings. Statistical and machine learning methods have been applied in this task. Examples include support sector regression (SVR) (Miao et al. (2018)), classification and regression tree (CART) (Chen et al. (2017)), kernel extreme learning machine (KELM) (Zhou et al. (2016)) and NN (Chen and Zeng (2013), Neaupane and Achet (2004), Yao et al. (2015)). However, the major limitation of such numerical approaches is that sudden events like landslides usually occur within a short period of time and may have significant deviations from the models fitted based on the historical data. Consequently, they are more often used to detect when the observed displacement exceeds the forecasted value while failing to give long-term forecasting of potential landslide (Zhou et al. (2016)).

The closest study to our work is that of Tordesillas, Zhou and Batterham (2018). Inspired by knowledge from fundamental studies of granular failure showing that a distinct kinematic pattern is initiated early on and well before ultimate failure (Tordesillas et al. (2013)), they argued that the ultimate failure pattern is encoded in the surface kinematics of a slope well before collapse (weeks in advance) for a landslide and sought to establish a reliable method for predicting the "strain-localization region," the boundary of the slip region. Specifically, k-nearest neighbor networks are generated based on features derived from the displacement data at each time stamp. Closeness centrality, a measurement that quantifies how close one node is to all others, is calculated for each of the nodes. As reported in their study, nodes with relatively high closeness centrality coincide with the shear band that divides the slip and

stable regions. As a result, monitoring this network connectivity over time enabled the early detection of the location of the landslide. However, unlike our proposed method, there is no use of spatial information in that work, and the regions that are found are not necessarily spatially contiguous or enclosed. In addition, their approach only learns the boundary between the slip and nonslip regions. In their study the identified boundary zone is sufficient to determine the location of the landslide, but in other datasets, like Mount Stromboli (Di Traglia et al. (2015, 2018)), the landslide might be associated with more complex patterns where the slip regions is consisted of two zones with irregular shape and disconnected to each other. In such cases knowledge of only the boundary may not be sufficient for detecting the landslide site and may even deliver misleading results. To address this issue, we propose to directly detect where the landslide is going to happen via a clustering technique which can encode complex spatial patterns.

3. Statistical modeling.

3.1. *sGMM*: *Spatially-aided Gaussian mixture model.* Let $X \in \mathbb{R}^{M \times D}$ be the *feature* data matrix comprising M instances with D features, $S \in \mathbb{R}^{M \times 2}$ be the *spatial* locations of these instances such that s_{m1} and s_{m2} represent the x and y coordinates of mth data point, respectively. We model the feature data by a K-component Gaussian mixture model (GMM) (Fraley and Raftery (2002)) with parameters $\Theta_f = \{\alpha_1, \ldots, \alpha_K, \mu_1, \ldots, \mu_K, \Sigma_1, \ldots, \Sigma_K\}$ as

$$(3.1) \qquad p(x_m|\Theta_f) = \sum_{k=1}^{K} \alpha_k \varphi(x_m|\mu_k, \Sigma_k),$$

where $\varphi(x_m|\mu_k, \Sigma_k)$ is the Gaussian density with mean vector μ_k and covariance matrix Σ_k and α_k is the mixing coefficient of the kth component. The mixture weights are nonnegative and sum up to 1.

Assuming g is a latent variable such that $g_m = k$ stands for x_m arising from the kth component, we have

$$(3.2) \qquad \begin{aligned} p(g_m = k) &= \alpha_k, \\ (x_m|g_m = k) &\sim \mathcal{N}(\mu_k, \Sigma_k). \end{aligned}$$

Given that we anticipate complex spatial patterns in our data (e.g., see Figure 4(b)), direct modeling of spatial coordinates with a mixture model would enforce inappropriate elliptical spatial distributions per cluster. Instead, we would like to model the spatial distribution within each cluster with a more flexible model, allowing for much more complex geometries. Inspired by Reich and Bondell (2011), we model the spatial distribution within each cluster as a nested GMM with L components.

In this way spatial information aids in the cluster allocation in that the number of subcomponents, L will control the role of the spatial information in cluster allocation. The larger the L, the less weight is placed on the spatial information, while as $L \to 1$ yields a restrictive spatial structure that is forced to be elliptical within each cluster.

Formally we have

$$(3.3) \qquad (s_m|g_m = k) \sim \sum_{l=1}^{L} \beta_{kl} \mathcal{N}(v_{kl}, \Lambda_{kl}),$$

where $0 \leq \beta_{kl} \leq 1$ is the mixing coefficient, v_{kl} is the mean vector, Λ_{kl} is the covariance matrix of the lth sub-component of the kth cluster, respectively, and $\sum_{l=1}^{L} \beta_{kl} = 1$.

Given the clustering assignment g_m, the distribution of the features are assumed to be homogeneous across the spatial locations within the cluster, that is, $(\boldsymbol{x}_m \perp \boldsymbol{s}_m)|g_m$ and, hence, the density of the mth instance can be factored as

$$(3.4) \qquad p(\boldsymbol{x}_m, \boldsymbol{s}_m|g_m) = p(\boldsymbol{x}_m|g_m)p(\boldsymbol{s}_m|g_m),$$

resulting in the following observed data log-likelihood:

$$(3.5) \qquad \begin{aligned} \ell(\boldsymbol{\Theta}|X, S) &= \sum_{\boldsymbol{g}} p(g_m)p(\boldsymbol{x}_m, \boldsymbol{s}_m|g_m) \\ &= \sum_{m=1}^{M} \log\left(\sum_{k=1}^{K} \alpha_k \varphi(\boldsymbol{x}_m|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \sum_{l=1}^{L} \beta_{kl}\varphi(\boldsymbol{s}_m|\boldsymbol{v}_{kl}, \boldsymbol{\Lambda}_{kl})\right). \end{aligned}$$

3.2. *EM algorithm for sGMM.* Directly maximizing (3.5) is difficult. Instead, the widely applied technique of Expectation-Maximization (EM) (Dempster, Laird and Rubin (1977)) is preferred for parameter estimation. Here, we introduce the EM algorithm for maximizing the log-likelihood of the sGMM.

Since the spatial distribution within a cluster is represented by its own GMM, we introduce a latent variable $\boldsymbol{h}$ for indicating this subcomponent membership conditional on the cluster assignment, that is, $(h_m = l)|(g_m = k)$ means that the mth instance is located in the lth sub-region of the kth cluster. As a result, (3.3) can be represented as

$$(3.6) \qquad \begin{aligned} p(h_m|g_m) &= \beta_{kl}, \\ (\boldsymbol{s}_m|g_m, h_m) &\sim \mathcal{N}(\boldsymbol{v}_{kl}, \boldsymbol{\Lambda}_{kl}). \end{aligned}$$

Given the feature vector and the location of the mth instance, let $r_{mkl}^{(t)}$ be the posterior probability of this instance belonging to the kth cluster and the lth sub-component within it at time t, we have

$$(3.7) \qquad \begin{aligned} r_{mkl}^{(t)} = p(g_m, h_m|\boldsymbol{x}_m, \boldsymbol{s}_m, \boldsymbol{\Theta}^{(t)}) &= \frac{p(\boldsymbol{x}_m, \boldsymbol{s}_m, g_m, h_m|\boldsymbol{\Theta}^{(t)})}{\sum_{\boldsymbol{g},\boldsymbol{h}} p(\boldsymbol{x}_m, \boldsymbol{s}_m, g_m, h_m|\boldsymbol{\Theta}^{(t)})} \\ &= \frac{p(\boldsymbol{x}_m|g_m, h_m, \boldsymbol{\Theta}^{(t)})p(\boldsymbol{s}_m|g_m, h_m, \boldsymbol{\Theta}^{(t)})p(h_m|g_m, \boldsymbol{\Theta}^{(t)})p(g_m|\boldsymbol{\Theta}^{(t)})}{\sum_{\boldsymbol{g},\boldsymbol{h}} p(\boldsymbol{x}_m, \boldsymbol{s}_m, g_m, h_m|\boldsymbol{\Theta}^{(t)})} \\ &= \frac{\alpha_k^{(t)} \varphi(\boldsymbol{x}_m|\boldsymbol{\mu}_k^{(t)}, \boldsymbol{\Sigma}_k^{(t)})\beta_{kl}^{(t)} \varphi(\boldsymbol{s}_m|\boldsymbol{v}_{kl}^{(t)}, \boldsymbol{\Lambda}_{kl}^{(t)})}{\sum_{k=1}^{K} \sum_{l=1}^{L} \alpha_k^{(t)} \varphi(\boldsymbol{x}_m|\boldsymbol{\mu}_k^{(t)}, \boldsymbol{\Sigma}_k^{(t)})\beta_{kl}^{(t)} \varphi(\boldsymbol{s}_m|\boldsymbol{v}_{kl}^{(t)}, \boldsymbol{\Lambda}_{kl}^{(t)})}. \end{aligned}$$

In the E-step, $Q(\boldsymbol{\Theta}, \boldsymbol{\Theta}^{(t)})$ is calculated as

$$(3.8) \qquad \begin{aligned} Q(\boldsymbol{\Theta}, \boldsymbol{\Theta}^{(t)}) &= \mathbb{E}_{\boldsymbol{g},\boldsymbol{h}|X,S,\boldsymbol{\Theta}^{(t)}} \log p(X, S, \boldsymbol{g}, \boldsymbol{h}|\boldsymbol{\Theta}) \\ &= \sum_{\boldsymbol{g},\boldsymbol{h}} p(\boldsymbol{g}, \boldsymbol{h}|X, S, \boldsymbol{\Theta}^{(t)}) \sum_{m=1}^{M} \log(p(\boldsymbol{x}_m, \boldsymbol{s}_m, g_m, h_m|\boldsymbol{\Theta})) \\ &= \sum_{m=1}^{M}\sum_{k=1}^{K}\sum_{l=1}^{L} r_{mkl}^{(t)} \log \alpha_k + \sum_{m=1}^{M}\sum_{k=1}^{K}\sum_{l=1}^{L} r_{mkl}^{(t)} \log \varphi(\boldsymbol{x}_m|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \\ &\quad + \sum_{m=1}^{M}\sum_{k=1}^{K}\sum_{l=1}^{L} r_{mkl}^{(t)} \log \beta_{kl} + \sum_{m=1}^{M}\sum_{k=1}^{K}\sum_{l=1}^{L} r_{mkl}^{(t)} \log \varphi(\boldsymbol{s}_m|\boldsymbol{v}_{kl}, \boldsymbol{\Lambda}_{kl}). \end{aligned}$$

In the M-step, the following closed-form updates can then be readily obtained as

$$\alpha_k^{(t+1)} = \frac{1}{M} \sum_{m=1}^{M} \sum_{l=1}^{L} r_{mkl}^{(t)},$$

$$\boldsymbol{\mu}_k^{(t+1)} = \frac{\sum_{m=1}^{M} \sum_{l=1}^{L} r_{mkl}^{(t)} \boldsymbol{x}_m}{\sum_{m=1}^{M} \sum_{l=1}^{L} r_{mkl}^{(t)}},$$

$$\boldsymbol{\Sigma}_k^{(t+1)} = \frac{\sum_{m=1}^{M} \sum_{l=1}^{L} r_{mkl}^{(t)} (\boldsymbol{x}_m - \boldsymbol{\mu}_k^{(t+1)})(\boldsymbol{x}_m - \boldsymbol{\mu}_k^{(t+1)})^\top}{\sum_{m=1}^{M} \sum_{l=1}^{L} r_{mkl}^{(t)}},$$

(3.9)

$$\beta_{kl}^{(t+1)} = \frac{\sum_{m=1}^{M} r_{mkl}^{(t)}}{\sum_{m=1}^{M} \sum_{l=1}^{L} r_{mkl}^{(t)}},$$

$$\boldsymbol{v}_{kl}^{(t+1)} = \frac{\sum_{m=1}^{M} r_{mkl}^{(t)} \boldsymbol{s}_m}{\sum_{m=1}^{M} r_{mkl}^{(t)}},$$

$$\boldsymbol{\Lambda}_{kl}^{(t+1)} = \frac{\sum_{m=1}^{M} r_{mkl}^{(t)} (\boldsymbol{s}_m - \boldsymbol{v}_{kl}^{(t+1)})(\boldsymbol{s}_m - \boldsymbol{v}_{kl}^{(t+1)})^\top}{\sum_{m=1}^{M} r_{mkl}^{(t)}}.$$

4. Simulation study. In this section we perform experiments on a synthetic dataset in order to illustrate the merits of the proposed sGMM algorithm compared to two baseline methods, K-Means and GMM.

4.1. *Data.* The simulated data consists of 400 data points which are evenly divided into two clusters, $\mathcal{C}_1$ and $\mathcal{C}_2$. Each point has a D-dimensional *feature* vector (we set $D = 2$ for ease of visualization) and another two-dimensional vector as its *spatial* coordinates. Feature values of points in $\mathcal{C}_1$ and $\mathcal{C}_2$ are sampled from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $\mathcal{N}(\boldsymbol{v}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu} = [0, 0]$, $\boldsymbol{v} = [1, 1]$ and $\boldsymbol{\Sigma} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, respectively. The visualization of the generated feature data can be found in Figure 1. The clustering pattern in the feature space is highly overlapped, with more than 20% of data points not correctly assigned by baselines to their actual clusters if this data is only used as input. The spatial coordinates of data points are uniformly distributed in a 2×2 squared space, and three different spatial patterns are simulated in our experiment: SP1, SP2 and SP3, as shown in the second column of Table 1.

4.2. *Settings.* Two baseline methods, K-Means and GMM, are used, and each of them has two variants, indicated by the (X), and (X + S) suffices followed by the algorithm name.
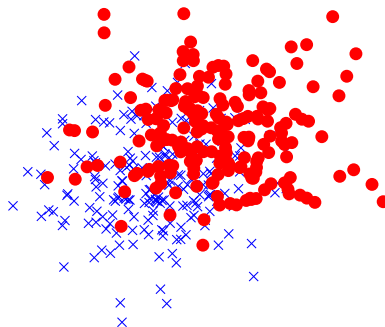

FIG. 1. *Ground truth clustering pattern in the feature space.*

TABLE 1
Visualization of clustering results of different methods in the spatial space

	Spatial pattern	K-Means (X)	K-Means (X + S)	GMM (X)	GMM (X + S)*	sGMM
SP1						
SP2						
SP3						

*GMM (X + S) is equivalent to sGMM with just one spatial component ($L = 1$) when diagonal covariance is used.

K-Means (X) and GMM (X) takes only the feature data as input, whereas the feature and spatial data are concatenated and fed to K-Means (X + S) and GMM (X + S). The proposed sGMM algorithm uses both the feature and spatial data. In addition, it is worth noting that when diagonal covariance is used, GMM (X + S) is equivalent to sGMM with just one spatial component.

The number of clusters is set to two for all algorithms. For our proposed method, an additional parameter, the *number of spatial components* for each cluster, L, is selected from one to four via Bayesian Information Criterion (BIC) (Schwarz (1978)) as

$$\text{BIC} = \ell(\mathbf{\Theta}|X, S) + |\mathbf{\Theta}|\log(M),$$

where $\ell(\mathbf{\Theta}|X, S)$ is the log likelihood, $|\mathbf{\Theta}|$ is the number of parameters and M is the total number of observations, respectively. In addition, to quantify the quality of prediction produced by each algorithm, we measure the similarity between the produced clustering and the ground truth by Normalized Mutual Information (NMI) (Vinh, Epps and Bailey (2010)) as

$$\text{NMI}(\tilde{Y}, Y) = \frac{I(\tilde{Y}; Y)}{\sqrt{H(\tilde{Y})H(Y)}},$$

where $\tilde{Y}$ is the prediction, Y is the ground truth, $I(\tilde{Y}; Y)$ is the mutual information between $\tilde{Y}$ and Y and $H(\cdot)$ is the entropy of the corresponding clustering partition.

4.3. *Results and discussion.* We visualize the clustering assignment produced by each algorithm in the spatial domain, as shown in Table 1, along side the of clustering quality results in Figure 2. The reported sGMM results (column 6 in Table 1) are chosen by the BIC results presented in Figure 3.

As shown in Figure 2, baselines that use only the feature data (K-Means (X) and GMM (X)) demonstrate poor performance with an NMI score under 0.3. Generally, incorporating spatial information helps the improvement in clustering quality. For example, for the trivial case like SP1, GMM (X + S) only mistakenly clusters a few observations and reaches a NMI score above 0.8 by simply concatenating feature and spatial data. However, this concatenating strategy can easily fail for the slightly more complex cases. Only minor performance boost can be found in SP3 for K-Means (X + S) and GMM (X + S) and the NMI score of GMM (X + S) even decreases for SP2, compared to the results of using solely the feature data.
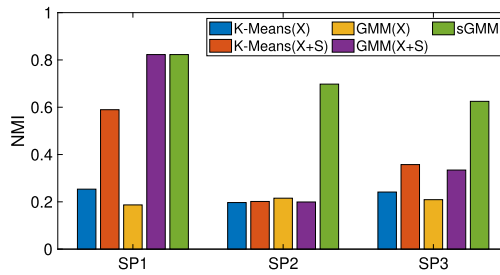
FIG. 2.   *Clustering quality compared to the ground truth* (*the higher the better*).

In contrast, by modeling the spatial coordinates of observations with an embedded mixture model, significant performance improvement can be found in our method. Compared to baselines, the clustering patterns produced by sGMM are less noisy and show much better similarity to the ground truth. In addition, it is clear that our model can easily accommodate different spatial patterns by simply choosing different number of spatial components, as demonstrated in Figure 3, where 1, 2, 3 spatial components are used for SP1, SP2 and SP3, respectively. Overall, it is clear to see the proposed method is a better way of using spatial information to aid the clustering problem, especially when data are hard to distinguish in the feature space while complex patterns are presented in the spatial domain.

5. Application to rock slope data.

5.1. *Rock slope dataset.*   Slope failures are a class of geological hazards that involve the downslope movement of soil and rocks under gravitational forces. The most common and effective technique for determining slope stability (i.e., the resistance of a slope to failure) is by monitoring the small precursory movements of its surface prior to collapse. Various radar interferometry systems have been developed for this purpose. One example is the slope stability radar (SSR). This monitoring technology can remotely scan a surface and detect movements with submillimeter precision at high spatial and temporal resolutions. Readers are referred to Casagli et al. (2010), Dick et al. (2014), Pienaar and Anton (2013) for details on: (a) the monitoring technology, (b) measurement methods and other data generated using SSR as well as (c) current best practice for prediction of slope failure from these datasets.

In this case study we analyze a rock slope dataset on which SSR was deployed to monitor the rock slope shown in Figure 4(a). Spatiotemporal raster data on movement were gathered for 1803 locations on the slope surface: a time series of the cumulative surface displacement along the line-of-sight from the stationary (ground-based) monitoring station to each observed location on the surface. This displacement is measured relative to a reference baseline observation and is updated at every six minute interval over a period of three weeks—10:07,
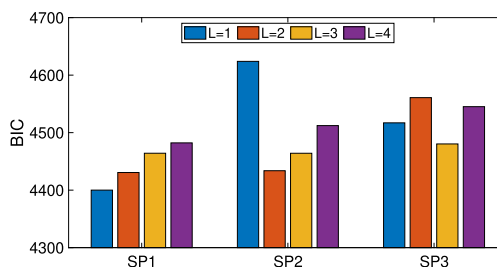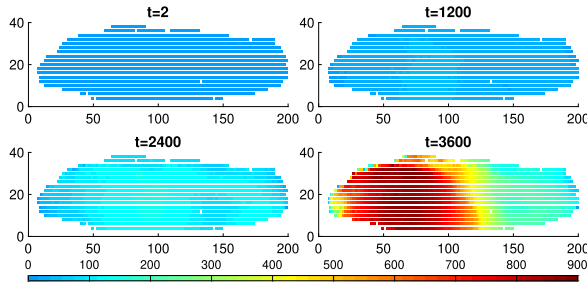

FIG. 3.   *BIC of sGMM with different number of spatial components* (*the lower the better*).
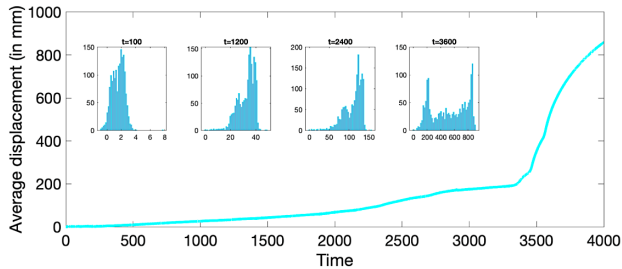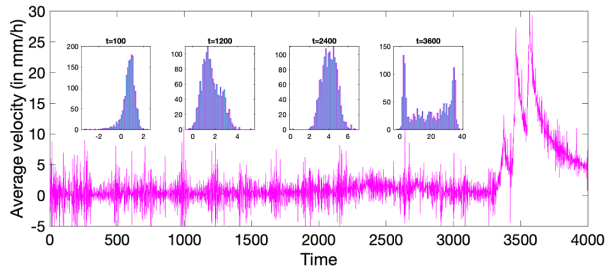
(a) Monitored site

(b) Cumulative displacement heatmaps of the rock slope at different time points.

(c) Average cumulative displacement over time (subplots are displacement distributions at different time stamps).

(d) Average velocity over time (subplots are velocity distributions at different time stamps).

FIG. 4.   *Rock slope dataset.*

May 31 to 23:55, June 21. A collapse occurred to the west in the early afternoon of June 15. Specifically, accelerating motions were recorded along the western side of the slope reaching peak velocity of 74.8 mm/h at location $(x, y) = (77m, 4m)$ at 1:10 p.m., June 15, which represents time point 3567 in this data. This coincided in time with the global average velocity reaching its peak of 33.6 mm/h (Figure 4(d)). Smaller surface movements were recorded in the south-eastern corner of the slope from June 1, but these stabilized on June 14. Two distinct subzones of the slope emerged in the failure. Large cumulative downward displacements characterize the rockslide region to the west, whilst relatively small movements developed in the stable region to the east.

5.2. *Settings*.   We apply sGMM on the first 3600 time stamps (10:07 May 31 to 16:25 June 15) of the data in a sliding window fashion. That is, given a particular window size W, at each time stamp a small snapshot of length W (taken backward from current time) is extracted from the displacement time series and clustered by sGMM, with the help of positional data of each radar monitor. Thus, 3600 clusterings from sGMM are produced. In order to demonstrate the value of incorporating spatial information in the clustering process, we use K-Means and GMM which use only the displacement of the given time range, as baselines.

NMI is used to quantify the quality of the produced clustering at each time stamp. Since no label information is presented in this data, we assume that the K-Means clustering of the displacement data at the final time stamp is the ground truth. This is a reasonable assumption, since at the last time stamp the rockslide has occurred and the groups (functional units) have become clearly defined.

In terms of parameter settings, the number of clusters is set to 2 for all algorithms, where one cluster indicates the active region and the other refers to the relative stable subzone. The number of iterations is set to 100 for the EM steps. Other parameters are selected via a grid search with *window size*, W, chosen from 1 (six minutes), 5 (30 minutes) and 10 (one hour), and the *number of spatial components* each cluster can have, L, is selected among 1, 2, 4, 8 and 16. The best parameters are chosen by the averaged BIC across all time steps.

In addition, for GMM and sGMM, we quantify the model uncertainty as the posterior probability for each point to be allocated to the cluster with relatively higher mean displacement in order to provide additional information regarding the certainty of the regions, that is, answering the question of whether a particular point will belong to the portion that collapses or to the portion that does not collapse. The other model parameters, such as the mixture component means and variances, are not themselves of interest but rather a way to obtain the clustering patterns. With this target in mind, we note that these posterior probabilities are functions of the parameters in the model and, hence, will have uncertainty associated with them as well. Specifically, we perform a weighted likelihood bootstrap as suggested for GMM (Newton and Raftery (1994), O'Hagan et al. (2019)). We use 1000 bootstrap samples at each time point. Based on these bootstrap samples, we then quantify the uncertainty in the posterior probability for each location to belong to the active region and demonstrate via plots (Figure 8) how the uncertainty shows up. We construct a 95% confidence interval for each probability, and our plots now show three sets of locations: those points whose confidence interval does not contain 0.5, which allows us to classify them to either the active or stable clusters (depending on whether they are above or below 0.5), and those points whose interval does contain 0.5. These are the points that we are uncertain as to which cluster they would belong, and we highlight these points separately.

The experiments are conducted on Spartan (Lafayette and Wiebelt (2017)), a research platform with multiple computing nodes, each having 12 CPU cores and 251 GB RAM. Due to the fact that Spartan is a highly utilized system, we limit the computing resources for each algorithm-dataset pair to two CPU cores, 16 GB memory and 12-hour maximum running time. The whole experiment is implemented in MATLAB,[1] and the code of GMM are adapted from Gebru et al. (2016).

5.3. *Results and discussion*.

5.3.1. *Overall performance of sGMM*.   According to Figure 5(a), the best parameter to fit this rock slope data is by using a time window with single time step and four spatial

---

[1]The MATLAB implementation of sGMM is available upon on request.

(a) Averaged BIC for different parameters (the lower the better).

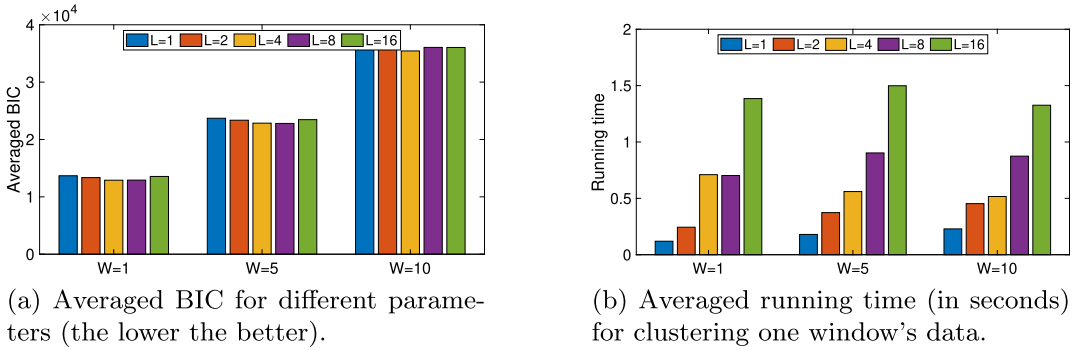(b) Averaged running time (in seconds) for clustering one window's data.

FIG. 5. *Performance of sGMM under different parameters.*

components for each cluster. In terms of running time, as shown in Figure 5(b), sGMM often takes under a second to cluster one time window's data, and generally, this goes up linearly with the increase of the number of spatial components. Considering the six minutes time interval in our data and usually even lower data generating frequency in landslide monitoring applications, such time consumption is insignificant and acceptable.

5.3.2. *Comparison to K-means.* Comparison between sGMM and K-Means can be found in Figures 6 and 7. We see that the major advantage of using spatial information in addition to kinematic features in clustering is that it allows sGMM to reveal more detail about the underlying dynamic changes of the system. Specifically, as shown in Figure 7, a key difference between the clusterings found by sGMM and K-Means is that the small region on the southeast side shows a "come-and-go" pattern in sGMM (i.e., it constantly joins and leaves the red cluster between time 700 to 2400 in Figure 7(a)). Unlike the "coactivate" pattern discovered by K-Means, sGMM detects that the ultimate failure is the outcome of a continuous competition between the west and southeast regions, with the western region ultimately winning the competition (being responsible for the failure). A similar type of competition phenomenon has been observed in previous small-scale laboratory studies, which observed that multiple strain localization regions initially compete with each other, with the winning region developing into the final shear band (see Vitone et al. (2013), Le Bouil et al. (2014), Karimi and Barrat (2018)). By incorporating spatial information into clustering, our method is able to detect the competition between the intermittent failure and winning failure regions in this large-scale real-world example. With this in mind, it is desirable to ignore the distraction of the southeast site, treating it as a false alarm occurring at an earlier stage, while
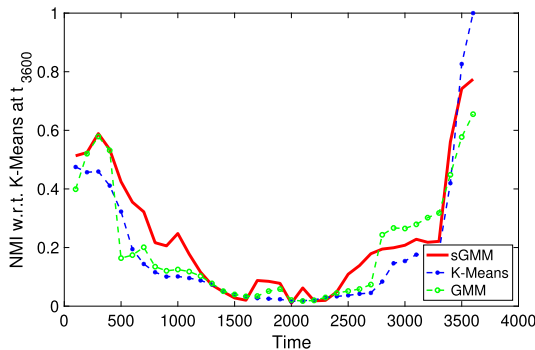

FIG. 6. *Changes of clustering similarities w.r.t. time. Larger NMI values indicate better similarity to the ground truth.*
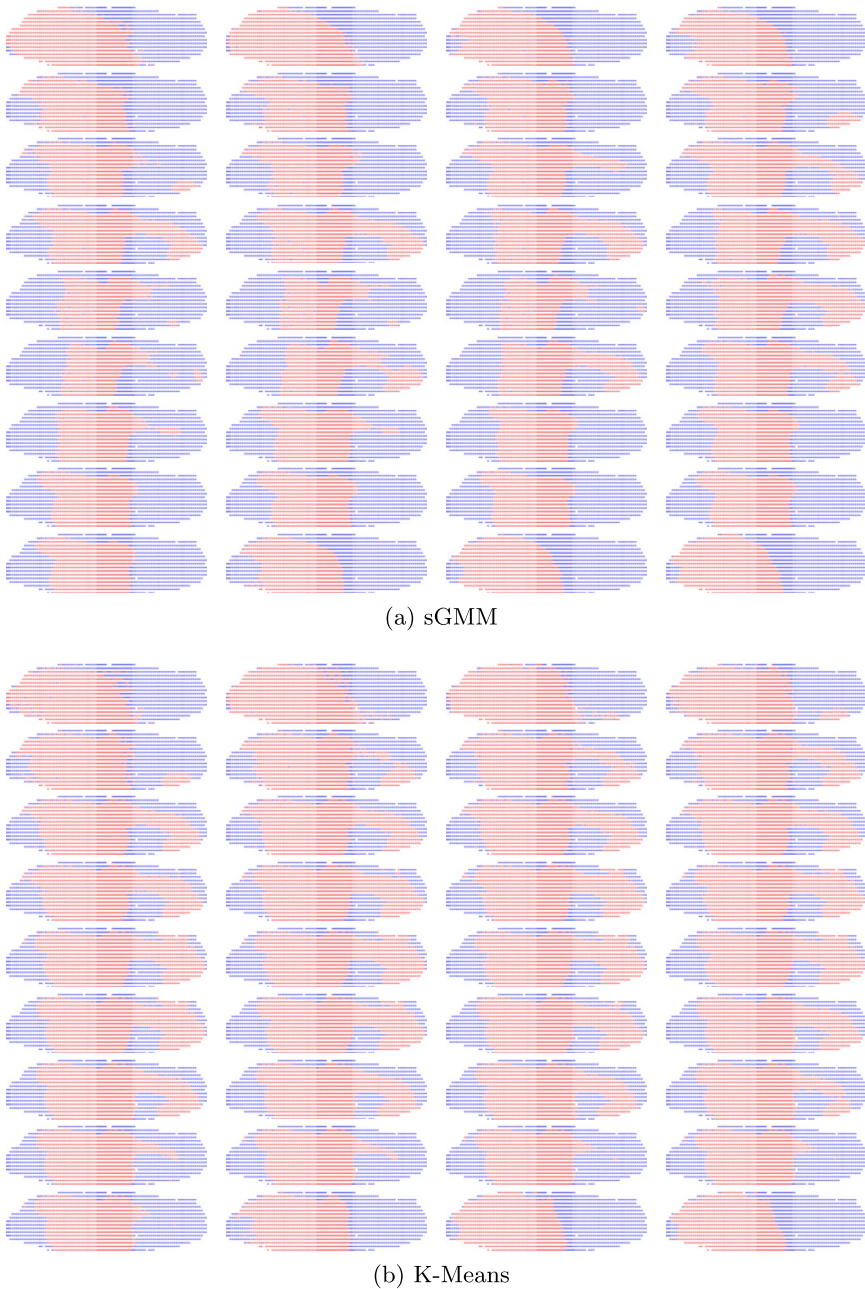
(a) sGMM

(b) K-Means

FIG. 7. *Clustering patterns at a sequence of time points. Results are from time* 100 *to* 3600, *with step size as* 100, *reading from left to right, top to bottom. Red cluster is the one with higher average displacement.*

focusing on the region that does actually collapse, the western wall, which is persistently active.

In fact, the incorporation of the spatial locality in our sGMM approach is able to automatically remove the southeast region from the active cluster 400 time stamps (∼40 hours, ∼five days prior to the eventually landslide) in advance, compared to K-Means, which consistently learns the arch that is bridging the west and southeast wall from time 700 to 2800 (∼three days before the eventually landslide). As a result, sGMM is more effective in terms of identifying the failure region over time (as shown in Figure 6, sGMM generates more accurate results that are similar to the final pattern formulated at the landsliding time). An additional

two days of warning on the landslide location is a significant gain and has the potential for a major impact for public safety.

Lastly, as an additional benefit of optimizing spatial locality, sGMM produces more spatially coherent clustering partitions than K-Means. This effect is especially noticeable in the early phase of the monitoring procedure. For example, clustering partitions produced by K-Means is more noisy than those of sGMM in the first 300 time stamps, where more isolated sensors can be found next to the boundary between two clusters. This appears to happen because in the early stages, the difference in the displacement data among radar sensors is subtle and highly concentrated within a small value range, as presented in Figure 4(c). Thus, if purely using the displacement data with K-Means, two sensors with only slightly different displacement records may be assigned to different clusters. In contrast, the clustering partitions produced by sGMM are generally smoother compared to K-Means, since spatial information has been incorporated in the clustering process, which is more interpretable and meaningful for understanding the evaluation of the rock slope.

5.3.3. *Comparison to GMM.* We report the posterior probability of a monitor being assigned to the active cluster by GMM and sGMM in Figure 8. The uncertain monitors found via the bootstrap test are also visualized by the black points. Additionally, the change of NMI scores of GMM can be found in Figure 6.

As can be seen in Figures 6 and 8(b), the clustering assignment of GMM is quite similar to K-Means in the first 1500 time stamps, where a large number of noisy points can be found in the early phase and a gradual development of the active region towards the east wall is presented. However, unlike K-Means that predicts the whole east wall as a part of slip region until time 2900, GMM shows a similar "come-and-go" pattern as our method inside the east wall and eliminates the distraction from this site 100 time stamps earlier than K-Means.

Compared to our proposed sGMM algorithm, the "come-and-go" pattern illustrated by GMM is less obvious and less frequent. One can easily see from Figure 8(a) that there are periodical expansion of uncertain monitors to the east wall from time 1600 to 2400 in sGMM. This means that, by taking spatial information into consideration, it is less confident to say that the east wall is of high likelihood to be a part of the active region, thus, the easier and earlier to narrow down the predicted active region to the the actual rockslide site. In contrast, the uncertain monitors found by GMM locate mainly in the boundary of the whole west and east sites that have relatively higher displacement, and it shows less clear and frequent signals in order to eliminate the east wall as earlier as sGMM can.

6. Conclusion. In this paper we study the failure detection problem in complex systems and propose a clustering algorithm that incorporates spatial information into a Gaussian mixture model. A case study on a rockslide monitoring dataset demonstrates that our model not only improves the spatial coherence of the clustering but also effectively identifies the correct failure region earlier, as much as two days in advance, compared to the baseline. More importantly, much richer patterns are discovered by the proposed model, such as the presence of an intermittent competing region of failure. Knowledge of such regions opens the door to deeper understanding of the dynamic of failure in complex systems.

More generally, the clustering method proposed in this paper is a unified approach for discovering patterns in one feature space while encouraging coherence in a separate feature space. This could potentially be applied to many real-world problems other than the rockfall forecasting we examine in our case study. Examples include: (1) multiview problems in cyber security, where one feature space reflects node behavior (patterns of API access or user application profiles) and the second space represents internode communication, such as connections endowed with statistics on port and protocols used; (2) traffic monitoring systems
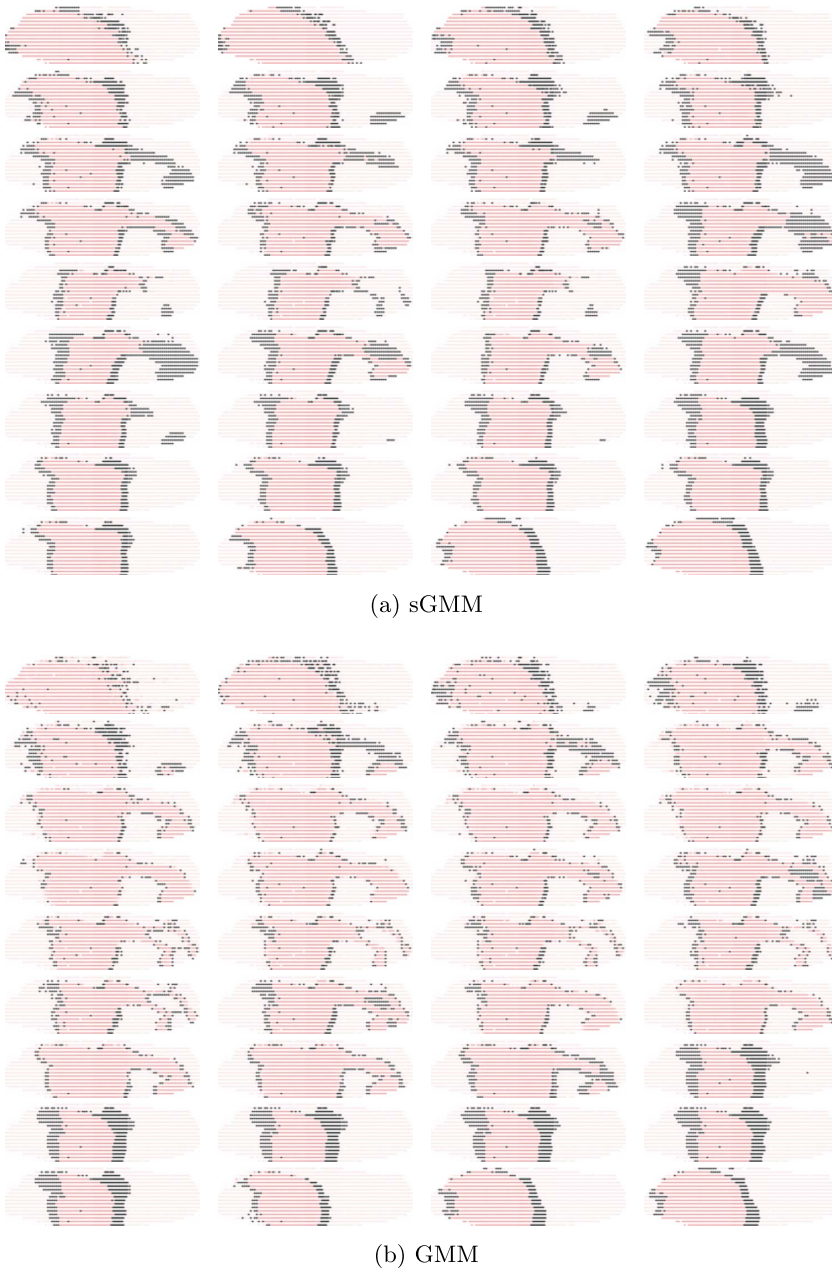
(a) sGMM

(b) GMM

FIG. 8. *Posterior probability of being in the active region. The darker the red, the closer the probability to* 1. *Black points are monitors that have* 0.5 *within its* 95% *confidence interval of the posterior.*

where the spatial proximity in the road network should be taken into account when clustering sensors based on the traffic volume; (3) recommender systems that partition users according to both their rating activities and their demographic information; and (4) social networks where communities or user groups are detected based not only on the linkage information among users but also their features such as age, occupation, and personal interest.

In terms of future work, one possible extension to our existing model is to incorporate domain knowledge as prior distributions on the parameters in a Bayesian modeling framework, similar to Reich and Bondell (2011). However, it is not entirely clear how best to incorporate this prior information on these model parameters which would represent things such as

the locations of the spatial clusters. In some sense this would then attempt to use the prior information to place a prior on the location of the rockslide, which we seek to find in the data-driven approach. Given that we have 1803 sensor locations within our framework, the data tends to identify the location well (as our results show), so a prior distribution on the location may not provide much additional information in this case. We also note that, in our example, we only have measurements of the kinematic data. Without additional measurements regarding characteristics of the slope, the experts are not typically able to give reliable information about the location of the eventual rockslide. In other cases there may perhaps be information about other characteristics of the slope, such as moisture readings, terrain density, etc. With that additional information, it would be potentially useful to include these readings as prior information on the suspected locations. For example, the spatial centers could be modeled via a hierarchical regression using these other characteristics as covariates.

REFERENCES

BRENNING, A. (2005). Spatial prediction models for landslide hazards: Review, comparison and evaluation. *Nat. Hazards Earth Syst. Sci.* **5** 853–862.

BUI, D. T., PRADHAN, B., LOFMAN, O., REVHAUG, I. and DICK, O. B. (2012). Landslide susceptibility mapping at Hoa Binh province (Vietnam) using an adaptive neuro-fuzzy inference system and GIS. *Comput. Geosci.* **45** 199–211.

CASAGLI, N., CATANI, F., DEL VENTISETTE, C. and LUZI, G. (2010). Monitoring, prediction, and early warning using ground-based radar interferometry. *Landslides* **7** 291–301.

CHEN, H. and ZENG, Z. (2013). Deformation prediction of landslide based on improved back-propagation neural network. *Cogn. Comput.* **5** 56–62.

CHEN, W., XIE, X., WANG, J., PRADHAN, B., HONG, H., BUI, D. T., DUAN, Z. and MA, J. (2017). A comparative study of logistic model tree, random forest, and classification and regression tree models for spatial prediction of landslide susceptibility. *Catena* **151** 147–160.

DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B* **39** 1–38. MR0501537

DI TRAGLIA, F. D., BATTAGLIA, M., NOLESINI, T., LAGOMARSINO, D. and CASAGLI, N. (2015). Shifts in the eruptive styles at Stromboli in 2010–2014 revealed by ground-based InSAR data. *Sci. Rep.* **5** Art. ID 13569. https://doi.org/10.1038/srep13569

DI TRAGLIA, F., CALVARI, S., D'AURIA, L., NOLESINI, T., BONACCORSO, A., FORNACIAI, A., ESPOSITO, A., CRISTALDI, A., FAVALLI, M. and CASAGLI, N. (2018). The 2014 effusive eruption at Stromboli: New insights from in situ and remote-sensing measurements. *Remote Sens.* **10** Art. ID 2035.

DICK, G. J., EBERHARDT, E., CABREJO-LIÉVANO, A. G., STEAD, D. and ROSE, N. D. (2014). Development of an early-warning time-of-failure analysis methodology for open-pit mine slopes utilizing ground-based slope stability radar monitoring data. *Can. Geotech. J.* **52** 515–529.

FRALEY, C. and RAFTERY, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *J. Amer. Statist. Assoc.* **97** 611–631. MR1951635 https://doi.org/10.1198/016214502760047131

GEBRU, I. D., ALAMEDA-PINEDA, X., FORBES, F. and HORAUD, R. (2016). EM algorithms for weighted-data clustering with application to audio-visual scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **38** 2402–2415. https://doi.org/10.1109/TPAMI.2016.2522425

GOETZ, J. N., BRENNING, A., PETSCHKO, H. and LEOPOLD, P. (2015). Evaluating machine learning and statistical prediction techniques for landslide susceptibility modeling. *Comput. Geosci.* **81** 1–11.

HONG, H., POURGHASEMI, H. R. and POURTAGHI, Z. S. (2016). Landslide susceptibility assessment in Lianhua County (China): A comparison between a random forest data mining technique and bivariate and multivariate statistical models. *Geomorphology* **259** 105–118.

KARIMI, K. and BARRAT, J.-L. (2018). Correlation and shear bands in a plastically deformed granular medium. *Sci. Rep.* **8** Art. ID 4021.

LAFAYETTE, L. and WIEBELT, B. (2017). Spartan and nemo: Two HPC-cloud hybrid implementations. In 2017 *IEEE* 13*th International Conference on E-Science* (*e-Science*) 458–459. IEEE, Piscataway, NJ.

LE BOUIL, A., AMON, A., SANGLEBOEUF, J.-C., ORAIN, H., BÉSUELLE, P., VIGGIANI, G., CHASLE, P. and CRASSOUS, J. (2014). A biaxial apparatus for the study of heterogeneous and intermittent strains in granular materials. *Granul. Matter* **16** 1–8.

MIAO, F., WU, Y., XIE, Y. and LI, Y. (2018). Prediction of landslide displacement with step-like behavior based on multialgorithm optimization and a support vector regression model. *Landslides* **15** 475–488.

NEAUPANE, K. M. and ACHET, S. H. (2004). Use of backpropagation neural network for landslide monitoring: A case study in the higher Himalaya. *Eng. Geol.* **74** 213–226.

NEWTON, M. A. and RAFTERY, A. E. (1994). Approximate Bayesian inference with the weighted likelihood bootstrap. *J. Roy. Statist. Soc. Ser. B* **56** 3–48. MR1257793

O'HAGAN, A., MURPHY, T. B., SCRUCCA, L. and GORMLEY, I. C. (2019). Investigation of parameter uncertainty in clustering using a Gaussian mixture model via jackknife, bootstrap and weighted likelihood bootstrap. *Comput. Statist.* **34** 1779–1813. MR4030442 https://doi.org/10.1007/s00180-019-00897-9

PHAM, B. T., BUI, D. T., POURGHASEMI, H. R., INDRA, P. and DHOLAKIA, M. B. (2017). Landslide susceptibility assesssment in the Uttarakhand area (India) using GIS: A comparison study of prediction capability of naïve Bayes, multilayer perceptron neural networks, and functional trees methods. *Theor. Appl. Climatol.* **128** 255–273.

PIENAAR, A. and ANTON, J. (2013). Guidelines for deriving alarm settings based on pre-determined criteria using the movement and surverying radar. *Aust. Cent. Geomech. Newsl.* **41** 7–12.

POURGHASEMI, H. R. and RAHMATI, O. (2018). Prediction of the landslide susceptibility: Which algorithm, which precision? *Catena* **162** 177–192.

REICH, B. J. and BONDELL, H. D. (2011). A spatial Dirichlet process mixture model for clustering population genetics data. *Biometrics* **67** 381–390. MR2829007 https://doi.org/10.1111/j.1541-0420.2010.01484.x

REICHENBACH, P., ROSSI, M., MALAMUD, B. D., MIHIR, M. and GUZZETTI, F. (2018). A review of statistically-based landslide susceptibility models. *Earth-Sci. Rev.* **180** 60–91.

SCHWARZ, G. (1978). Estimating the dimension of a model. *Ann. Statist.* **6** 461–464. MR0468014

THANH, L. N. and DE SMEDT, F. (2014). Slope stability analysis using a physically based model: A case study from A Luoi district in Thua Thien-Hue Province, Vietnam. *Landslides* **11** 897–907.

TORDESILLAS, A., ZHOU, Z. and BATTERHAM, R. (2018). A data-driven complex systems approach to early prediction of landslides. *Mech. Res. Commun.* **92** 137–141.

TORDESILLAS, A., WALKER, D. M., ANDÒ, E. and VIGGIANI, G. (2013). Revisiting localized deformation in sand with complex systems. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **469** Art. ID 20120606.

VINH, N. X., EPPS, J. and BAILEY, J. (2010). Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.* **11** 2837–2854. MR2738784

VITONE, C., COTECCHIA, F., VIGGIANI, G. and HALL, S. A. (2013). Strain fields and mechanical response of a highly to medium fissured bentonite clay. *Int. J. Numer. Anal. Methods Geomech.* **37** 1510–1534.

WASOWSKI, J. and BOVENGA, F. (2014). Investigating landslides and unstable slopes with satellite multi temporal interferometry: Current issues and future perspectives. *Eng. Geol.* **174** 103–138.

YAO, X., THAM, L. G. and DAI, F. C. (2008). Landslide susceptibility mapping based on support vector machine: A case study on natural slopes of Hong Kong, China. *Geomorphology* **101** 572–582.

YAO, W., ZENG, Z., LIAN, C. and TANG, H. (2015). Training enhanced reservoir computing predictor for landslide displacement. *Eng. Geol.* **188** 101–109.

ZHOU, C., YIN, K., CAO, Y. and AHMED, B. (2016). Application of time series analysis and PSO–SVM model in predicting the Bazimen landslide in the Three Gorges Reservoir, China. *Eng. Geol.* **204** 108–120.