

OPEN

Forecasting high-dimensional dynamics exploiting suboptimal embeddings

Shunya Okuno ^{1,2*}, Kazuyuki Aihara ^{1,3,5} & Yoshito Hirata ^{3,4,5}

Delay embedding—a method for reconstructing dynamical systems by delay coordinates—is widely used to forecast nonlinear time series as a model-free approach. When multivariate time series are observed, several existing frameworks can be applied to yield a single forecast combining multiple forecasts derived from various embeddings. However, the performance of these frameworks is not always satisfactory because they randomly select embeddings or use brute force and do not consider the diversity of the embeddings to combine. Herein, we develop a forecasting framework that overcomes these existing problems. The framework exploits various “suboptimal embeddings” obtained by minimizing the in-sample error via combinatorial optimization. The framework achieves the best results among existing frameworks for sample toy datasets and a real-world flood dataset. We show that the framework is applicable to a wide range of data lengths and dimensions. Therefore, the framework can be applied to various fields such as neuroscience, ecology, finance, fluid dynamics, weather, and disaster prevention.

Forecasting a future system state is an important task in various fields such as neuroscience, ecology, finance, fluid dynamics, weather, and disaster prevention. If accurate system equations are unknown but a time series is observed, we can use model-free forecasting approaches. Although there are various model-free methods, if the time series is assumed to be nonlinear and deterministic, one of the most common approaches is delay embedding. According to embedding theorems^{1,2}, we can reconstruct the underlying dynamics by delay embedding (see Methods for details). These theorems ensure that the map from the original attractor to the reconstructed attractor has a one-to-one correspondence. These embedding theorems have also been extended to multivariate data and nonuniform embeddings³. In addition, multivariate embeddings have been empirically used^{4–7} before the extended theorem was proposed.

When multivariate time series are observed, various embeddings can be obtained on the basis of the extended embedding theorem³. Ye and Sugihara⁸ exploited this property; they developed multiview embedding (MVE), which combines multiple forecasts based on the top-performing multiple embeddings scored by the in-sample error. MVE can yield a more accurate forecast than the single best embedding, especially for short time series⁸. Although MVE works well with low-dimensional data, it cannot be simply applied to high-dimensional data. This is because MVE requires forecasts derived from all possible embeddings, whose total number combinatorially increases with the number of variables. Although we can approximate MVE by randomly chosen embeddings if the total number of possible embeddings is beyond the computable one⁹, it is difficult to find high-performance embeddings for high-dimensional data. Recently, Ma *et al.*¹⁰ presented an outstanding framework, randomly distributed embedding (RDE), to tackle these high-dimensional data. Their key idea is to combine forecasts yielded by randomly generated “nondelay embeddings” of target variables. These “nondelay embeddings” can also reconstruct the original state space for some cases and drastically reduce the possible number of embeddings. Ma *et al.*¹⁰ also showed that small embedding dimensions worked fine, even for high-dimensional dynamics, and successfully forecasted short-term high-dimensional data using the RDE framework. Although RDE showed outstanding results, there is room for improvement for some specific tasks. For example, “nondelay embedding”

¹Institute of Industrial Science, The University of Tokyo, 4-6-1 Komaba, Meguro-ku, Tokyo, 153-8505, Japan. ²Disaster Reduction & Environmental Engineering Department, Kozo Keikaku Engineering Inc., 4-5-3 Chuo, Nakanoku, Tokyo, 164-0011, Japan. ³International Research Center for Neurointelligence (WPI-IRCIN), The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033, Japan. ⁴Mathematics and Informatics Center, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033, Japan. ⁵These authors contributed equally: Kazuyuki Aihara and Yoshito Hirata. *email: okuno@sat.t.u-tokyo.ac.jp

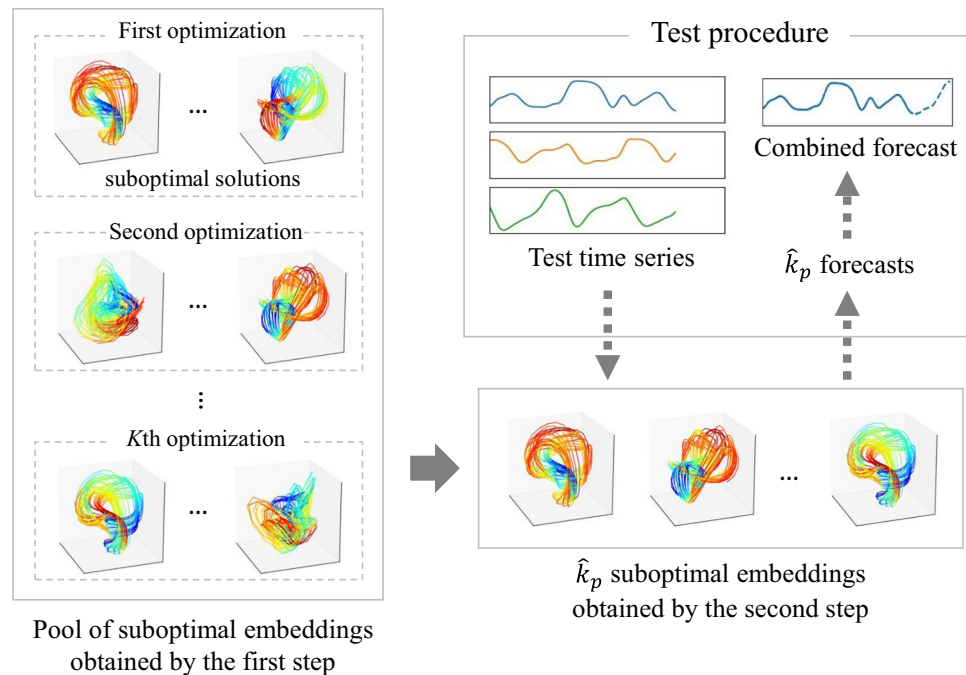


Figure 1. Schematic of the proposed forecasting procedure. We prepare a pool of suboptimal embeddings in the first step. We solve K combinatorial optimization problems to obtain various embeddings in this step. Next, we pick \hat{k}_p embeddings to minimize the error of the combined forecast in the second step. We combine the forecasts obtained by the \hat{k}_p embeddings to test the time series.

may overlook important pieces of information if delayed observations carry such pieces of information, e.g., the upstream river height for flood forecasting. A random distribution may also be a problem for the case where only partial variables are valid system variables; in this case, most of the random embeddings are invalid and not suitable to forecast. In addition, several important hyperparameters—the embedding dimension and the number of embeddings to combine—must be manually or empirically selected. Although ensembles tend to yield better results with significant diversity among its members^{11,12}, RDE only aggregates forecasts for a fixed embedding dimension with a fixed number of embeddings to combine, regardless of their forecast performance.

In this study, we propose another forecasting framework that overcomes the disadvantages of MVE and RDE. Our key idea is to prepare diverse “suboptimal embeddings” via combinatorial optimization. We combine the optimal number of these embeddings to maximize the performance of the combined forecast. In contrast to the existing frameworks, the suitable embedding dimensions and the number of embeddings to combine are automatically determined through this procedure. Our proposed framework achieves better forecast performance than the existing frameworks for high-dimensional toy models and a real-world flood dataset.

Results

In this section, we introduce a forecasting framework that overcomes some drawbacks of existing frameworks in the first subsection. Then, we numerically validate the performance of the proposed framework by using toy models as well as a real dataset in the second subsection.

Proposed forecasting framework. Our proposed framework yields a single forecast to combine multiple ones, which are obtained by multiple embeddings. Although the principal concept is the same as those of MVE and RDE, the procedure to yield the multiple embeddings differs from the existing ones. These embeddings are obtained through two steps. First, we obtain suboptimal embeddings, which are suboptimal solutions to minimize the forecast error. The aim of this procedure is to yield diverse embeddings to improve the performance of the combined forecast because such diversity is important for the performance of the ensemble^{11,12}. Second, we make the in-sample forecasts based on all suboptimal embeddings following MVE and RDE but pick the optimal number of embeddings to minimize the error of the combined forecasts. A schematic of the procedure is illustrated in Fig. 1. We describe each step in detail in the following sub-subsections.

First step: Preparing suboptimal embeddings. In our framework, we first obtain diverse suboptimal embeddings to minimize the in-sample forecast error. With the observed time series $y(t) \in \mathbb{R}^n$, a set of possible embeddings \mathcal{E} , and a corresponding delay vector $v_e(t) \in \mathbb{R}^E$ with $e \in \mathcal{E}$, we can obtain the p -steps-ahead forecast at time t by a map ψ_e as follows:

$$\hat{y}_f^e(t+p|t) = \psi_e(v_e(t)), \quad (1)$$

where y_f is the target variable to be forecasted. We can obtain such maps by suitable fitting algorithms, e.g., the conventional method of analogues¹³, as in MVE, and Gaussian process regression¹⁴, as in RDE. Throughout this paper, we apply the same algorithm, a variation of the method of analogues (see Supplementary Information), to all forecasting frameworks including the proposed one. This is because we need to compare the performance of the frameworks on a fair basis that is not affected by the performance of each fitting algorithm.

To obtain high-performance embeddings in an efficient manner, we solve combinatorial optimization problems instead of random selection or using brute force. We define the time indices used for training as $t < 0$, and suppose the forecasting of unobserved samples for $t \geq 0$. We split the set of training time indices $\mathcal{T}_{train} = \{t | t < 0\}$ into K datasets. Then, we solve the following optimization problem for each $K' \in \{1, 2, \dots, K\}$ to minimize the error:

$$\text{minimize}_{e \in \mathcal{E}} \sum_{t \in \mathcal{T}_{train}^{K'}} \sum_p \|\hat{y}_f^e(t+p|t) - y_f(t+p)\|, \quad (2)$$

where $\mathcal{T}_{train}^{K'}$ is the K' th set of training time indices and $\|\cdot\|$ is an appropriate norm. We use the L^2 norm throughout this paper. Note that \hat{y}_f^e is computed in an in-sample manner, similar to MVE; we use all of \mathcal{T}_{train} , except for the current time, as its library (leave-one-out cross-validation). When we consider embeddings up to l lags to embed the latest observation of the target variable, the number of possible embeddings is 2^{nl-1} . If nl is sufficiently large, it is almost impossible to compute all possible embeddings because of combinatorial explosion. Here, we take a straightforward approach to minimize the forecast error: an evolution strategy that interprets an embedding as a one-dimensional binary series^{15,16}. See Methods for details. Throughout this paper, we applied the $(\mu + \lambda)$ evolution strategy¹⁷.

In the process of each optimization, we store not only the best embedding but also the ‘‘Hall of Fame,’’ which preserves the best individuals in each generation. From the Hall of Fame, we select the best M solutions that satisfy the condition that the Hamming distance $d(\cdot, \cdot)$ is larger than or equal to a certain threshold value θ as follows:

$$d(e_i, e_j) \geq \theta, \quad \forall i, j; i \neq j, \quad (3)$$

where e_i is the i th embedding of the sorted Hall of Fame. Namely, we sort the Hall of Fame by the fitness (the objective function of Eq. 2) and pick the best embeddings in order from the sorted Hall of Fame to satisfy Eq. (3). Note that the Hamming distance is defined according to the binary expressions of the embeddings. This condition prunes similar embeddings and keeps diverse embeddings for the next step. We can explicitly control the diversity of the embeddings by changing the threshold value θ , while the existing approaches do not include such a mechanism. The diversity of embeddings can affect the performance of the combined forecast. To apply this procedure, we obtain $K \cdot M$ suboptimal embeddings in total. See also the first part of Algorithm 1 in Methods.

The procedure of the K times optimizations yields diverse but high-performance embeddings. If a dataset contains a large number of variables, it is extremely difficult to obtain the exact solution by evolution strategies or any other metaheuristic algorithms. These algorithms find suboptimal solutions for most cases, and these solutions can be easily changed by slight changes in the optimization conditions. The procedure rather exploits this property; the K times optimization can find K sets of diverse suboptimal solutions. Instead of minimizing the whole in-sample error with K different conditions, we minimize each K split dataset and utilize the history of solutions to significantly reduce the computational time. We numerically show the effect of K times optimization in Supplementary Information.

To create diverse embeddings and improve the performance of the combined forecast, we can optionally consider L finite impulse response (FIR) filters for each suboptimal embedding; namely, $K \cdot L \cdot M$ embeddings are obtained in total. A study has shown that the combination of various forecasts via linear filters improves the forecast result in some cases¹⁸. These filtered embeddings are justified by the filtered delay embedding prevalence theorem², which ensures the embedding with finite filtered delay coordinates. Although we can compute suboptimal embeddings for each filter, we simply consider L filters based on the obtained $K \cdot M$ embeddings. Empirically, the optimization of the embeddings for each filter does not realize a significant improvement in performance to justify the additional computational cost. See Methods for the details of forecasting filtered time series.

Second step: combining forecasts to minimize the in-sample error. Next, we pick the optimal number of embeddings from the $P = K \cdot L \cdot M$ embeddings to minimize the error of the combined forecast. We compute in-sample forecasts by the embeddings for all of \mathcal{T}_{train} , as done in the MVE or RDE aggregation schemes. We determine the optimal number of embeddings to integrate forecasts, which is neither the square root of the possible number of embeddings (as MVE does) nor the embedding dimension (as the RDE aggregation scheme does). Here, we define a tuple of the forecast indices sorted by the in-sample error for step p as \mathcal{J}_p , and the i th best embedding is written as $\mathcal{J}_p(i)$. Note that \mathcal{J}_p is separately computed for each forecast horizon because accurate embeddings may differ according to the forecast horizons, as discussed later. We obtain \hat{k}_p to minimize the combined forecast as follows:

$$Y_k(t+p|t) = 1/k \sum_{i=1}^k \hat{y}_f^{\mathcal{J}_p(i)}(t+p|t), \quad (4)$$




Figure 2. Forecast performance with the Lorenz'96I model: comparisons of the performance (a) up to 10 steps ahead with the fixed data length (4000) without noise, (b) with different values of the data length, and (c) with different scales of observational noise. We computed the RMSE of the five-steps-ahead forecasts with randomly distributed embedding (RDE), multiview embedding (MVE), state-dependent weighting (SDW), single-best embedding based on the $(\mu + \lambda)$ -ES algorithm (SBE), single-variable embedding (SVE), all-variable embedding (AVE), and the proposed framework. These tests were carried out with 20 datasets generated with different random initial conditions and noise. The median, upper quartile, and lower quartile are shown.

$$\hat{k}_p = \arg \min_{k=1,2,\dots,P} \sum_{t \in \mathcal{F}_{train}} [Y_k(t + p|t) - y_f(t + p)]^2, \tag{5}$$

$$\hat{y}_f(t + p|t) = Y_{\hat{k}_p}(t + p|t). \tag{6}$$

We can obtain \hat{k}_p in a brute-force manner with a very small computational cost because the k th combined forecast $Y_k(t + p|t)$ can be recursively computed as follows:

$$Y_{k+1}(t + p|t) = [kY_k(t + p|t) + \hat{y}_f^{\mathcal{J}_p(k+1)}(t + p|t)] / (k + 1), \tag{7}$$

where $k = 1, 2, \dots, P - 1$, $Y_1(t + p|t) = \hat{y}_f^{\mathcal{J}_p(1)}(t + p|t)$. We compute Y_k and the corresponding error for all k and pick the optimal value to minimize the squared error. Note that \hat{k}_p can be different for each p ; empirically, \hat{k}_p tends to be larger for a longer step.

For unobserved samples $t \geq 0$, the forecast $\hat{y}_f(t + p|t)$ is computed as follows:

$$\hat{y}_f(t + p|t) = \sum_{k=1}^{\hat{k}_p} \hat{y}_f^{\mathcal{J}_p(k)}(t + p|t) / \hat{k}_p. \tag{8}$$

Note that only \hat{k}_p forecasts are necessary to compute unobserved samples (test data) rather than $K \cdot L \cdot M$ forecasts. We summarize the whole algorithm in Algorithm 1 in Methods.

Numerical experiments. Toy models. We demonstrate our framework with toy models. We first tested with the 10-dimensional Lorenz'96I model¹⁹ using forcing term 8.0. The model is described as 10-dimensional differential equations, and we used five of their system variables and five random walks as a dataset—namely, Lorenz'96I's x_0, x_1, x_2, x_3, x_4 and random walks x_5, x_6, x_7, x_8, x_9 . We forecasted x_0 up to 10 steps. The single time step corresponds to the 0.5 non-dimensional unit of the original system. Note that the 0.05 non-dimensional unit corresponds to 6 hours with the forcing 8.0 in terms of the error growth rate under the conditions of the original study¹⁹. See Supplementary Information for the detailed setup. We compared the proposed framework with existing frameworks—namely, randomly distributed embedding with an aggregation scheme (RDE), multiview embedding (MVE), a variation of MVE (state-dependent weighting⁹ (SDW)), and the single best embedding via the $(\mu + \lambda)$ -ES algorithm. In addition, we also compared the proposed framework with a conventional single-variable embedding (SVE) and an all-variable embedding (AVE). See Supplementary Information for the detailed conditions.

We first tested the performance with the data length of 4000 for the database and 500 for the evaluation. Figure 2(a) shows that the proposed forecast achieved the best performance among the existing frameworks up to 10 steps ahead with the length of 4000.

We conducted a sensitivity analysis with the data length. We set the data length, i.e., the size of the training dataset, to {100, 200, 500, 1000, 2000, 4000} and tested the performance with 500 samples for all cases. Figure 2(b) shows that the proposed five-steps-ahead forecast yielded the minimum error among the existing frameworks for all the data lengths.

We also checked the robustness against observational noise. We added Gaussian observational noise scaled by the standard deviation of the target variable and tested the performance with a scale range of {0.0, 0.1, 0.2, 0.3}. We set the training data length as 500 and evaluated with 500 samples. Figure 2(c) shows that a lower noise level results in better performance for the proposed framework compared with the other frameworks. Although




Figure 3. Forecast profile of the Lorenz'96I model with the data length of 4000. Panel (a) shows the proportion of embedding of the proposed forecasts. The color indicates the proportion of embedding for each variable averaged over the number of combined forecasts for each step. Note that for each prediction step, the sum of the proportions of all variables is one. Variables x_0, \dots, x_4 are the variables of the 10-dimensional Lorenz'96I equations, and the other variables are random walks. Panel (b) shows the averaged filter coefficient ρ and embedding dimension E over the combined forecasts for each step. Note that the averaged dimensions can be a decimal fraction since they are averaged over the number of combined forecasts for each step. Panel (c) shows the relation between the number of combined forecasts and the errors. The solid line shows the in-sample error, and the dashed line shows the test error. The vertical solid line shows the selected number of combined forecasts k_p to minimize the in-sample error.

the proposed framework achieved the best performance up to 10% noise, the performance was degraded by high-level noise and is the second best at most for 30% noise.

We profiled which variables were embedded for the combined forecasts for the case where the data length is 4000 without noise. We summed \hat{k}_p embeddings for each variable, where 1 means embedded and 0 means not embedded. Then, we computed the proportion of embedding of each variable, normalizing the sum to one for each p . Figure 3(a) shows that random walks x_5, x_6, x_7, x_8 and x_9 were seldom embedded with the proposed framework. This result suggests that our proposed framework can select the appropriate variables. The results also show that the proportion of embedding of the target variable x_0 is large for short-term prediction and small for longer-term prediction. This result is consistent with that of an existing study²⁰.

We also profiled the filter coefficient ρ and the embedding dimension E of the proposed combined forecasts. Details regarding the definition of ρ are presented in the Methods section. We averaged ρ and E over the \hat{k}_p combined members for each step. Figure 3(b) shows that smaller dimensions are selected for shorter steps, and larger dimensions are selected for longer predictions. The figure also shows that smaller and larger values of ρ are selected for shorter and longer steps, respectively. Because we employ high-pass filtering ($\rho \leq 0$) with this model, a smaller ρ focuses on a smaller time period and vice versa.

We investigated the number of combined forecasts and the forecast errors to check whether the proposed integration scheme works properly. Although the proposed framework determines \hat{k}_p on the basis of the in-sample error, \hat{k}_p also minimizes the test error in this example (see Fig. 3(c)).

We also demonstrate our framework with the Kuramoto–Sivashinsky equations^{21,22}. We generated 20-dimensional time series: the values of the first 10 grids of the Kuramoto–Sivashinsky equations x_0, x_1, \dots, x_9 and 10 random walks $x_{10}, x_{11}, \dots, x_{19}$. Note that a single time step corresponds to 1.0 of the original system. See Supplementary Information for the detailed setup. We forecasted x_0 up to 10 steps using the same parameter values, data length, and noise scales as used in the Lorenz'96I example. See Supplementary Information for the detailed conditions.

Despite the increases in the variables, Fig. 4(a,b) show that the proposed forecast yielded the minimum error among the existing frameworks, as in the Lorenz'96I example. Regarding the noise sensitivity, a lower noise level results in better performance for the proposed framework compared to the other frameworks. The proposed framework achieved the best performance for a wide data range and relatively small observational noise.

Finally, we tested the sensitivity to the number of variables. We changed the dimensions of Lorenz'96I to {10, 20, 40, 80, 160, 320}, and we substituted a half of each dataset with random walks. We set the training data length to 500 and evaluated with 500 samples. For all ranges of variables, the proposed framework achieved the best performance (Fig. 5(a)). We also tested the case where all variables are available and the datasets do not contain random walks; this is the case in which delay embedding is not needed because all system variables are observed. In this case, the performance difference is negligible for all the top four frameworks, and RDE achieved the best score for the case where the number of variables is 320. Although this condition is too ideal because the observation function is an identity map, we can expect RDE to perform well for such cases that include many informative variables.

Flood dataset. We tested the proposed framework with real-world data, namely the flood forecasting competition dataset at “Artificial Neural Network Experiment (ANNEX 2005/2006)”²³. This dataset contains nine variables: the river stages of the target site (Q), and three upstream sites ($US1, US2$, and $US3$) and five rain gauges ($RG1, RG2, RG3, RG4$, and $RG5$) for three periods. We forecasted 6, 12, 18, and 24h ahead of the target river stage using all nine variables. See Supplementary Information for the detailed conditions. We also compared our framework with three other conventional machine-learning methods used in an existing study⁹: a recurrent neural network with long short-term memory (LSTM)²⁴, support vector regression (SVR)²⁵, and random forest regression²⁶. See the detailed conditions in ref. ⁹.




Figure 4. Forecast performance with the Kuramoto–Sivashinsky equations: comparisons of the performance (a) up to 10 steps ahead with the fixed data length (4000) without noise, (b) with different values of the data length, and (c) with different scales of observational noise. We computed the RMSE of the five-steps-ahead forecasts with randomly distributed embedding (RDE), multiview embedding (MVE), state-dependent weighting (SDW), single-best embedding based on the $(\mu + \lambda)$ -ES algorithm (SBE), single-variable embedding (SVE), all-variable embedding (AVE), and the proposed framework. These tests were carried out with 20 datasets generated with different random initial conditions and noise. The median, upper quartile, and lower quartile are shown.

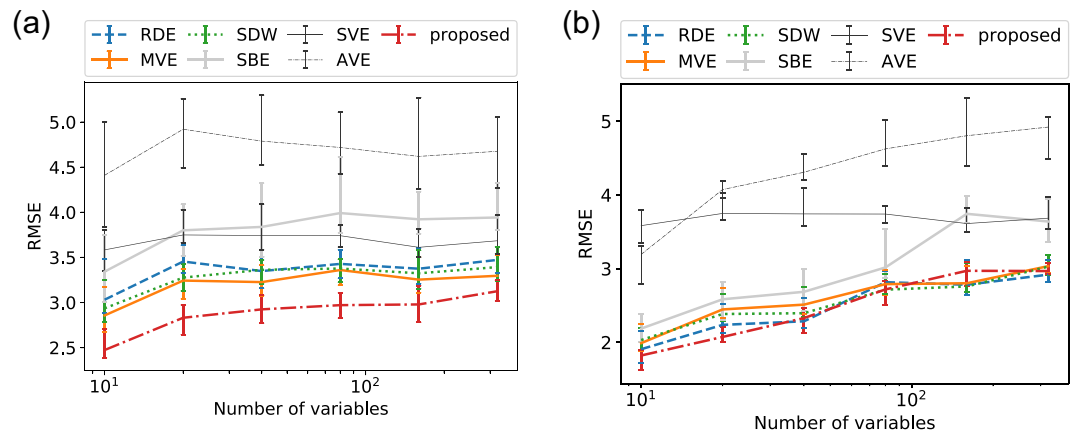


Figure 5. Forecast performance with the Lorenz'96I model for various numbers of variables: cases where (a) a half of the variables are substituted with random walks and (b) all variables are available. We computed the RMSE of the five-steps-ahead forecasts with randomly distributed embedding (RDE), multiview embedding (MVE), state-dependent weighting (SDW), single-best embedding based on the $(\mu + \lambda)$ -ES algorithm (SBE), single-variable embedding (SVE), all-variable embedding (AVE), and the proposed framework. These tests were carried out with 20 datasets generated with different random initial conditions and noise. The median, upper quartile, and lower quartile are shown.

Prediction steps [h]	Proposed	RDE	MVE	LSTM	SVR	Random Forest
6	0.068	0.069	0.060	0.123	0.073	0.079
12	0.167	0.174	0.171	0.213	0.177	0.185
18	0.247	0.278	0.273	0.291	0.268	0.275
24	0.349	0.389	0.373	0.372	0.365	0.363

Table 1. RMSEs of the flood dataset computed using the proposed framework, randomly distributed embedding (RDE), and multiview embedding (MVE) and also compared with those obtained by long short-term memory (LSTM), support vector regression (SVR), and the random forest. The proposed framework achieved the best accuracy for the 12-, 18-, and 24-h-ahead forecasts. Although the result obtained by the proposed framework for the 6-h-ahead forecast is the second best, the difference is slight.

The results in Table 1 show that the proposed framework yielded the best root-mean-square error (RMSE) through 12–24 h ahead, and the differences compared to the other methods were expanded over the forecast horizon. Although the 6 h (one step)-ahead forecast was not the best, the difference compared to the best score was negligible. The dataset contains only 1460 points, and long-term memory is not required for this task; hence, LSTM did not perform well. Although the test data include the largest river stage on 1995-02-01, Fig. 6(a) shows that the proposed framework properly forecasted the inexperienced river stage.

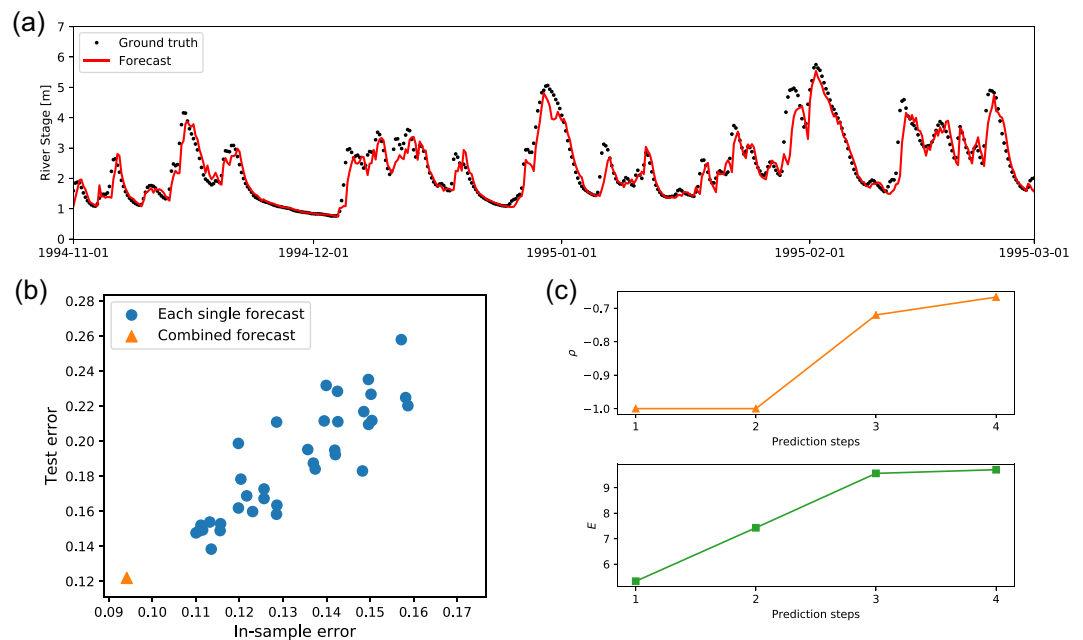


Figure 6. Forecast results for the flood dataset. Panel (a) shows a comparison of the ground truth and the proposed 24-h-ahead forecast. The proposed forecast did not underestimate the maximum river stage, which is the maximum value of the whole dataset. Panel (b) shows a comparison of the in-sample and test errors of the ensemble members. The results demonstrate the difficulty of selecting the best forecast because the best in-sample forecast does not always perform the best. In contrast, the combined forecast performed the best by far for both training and test data. Panel (c) shows the filter coefficient ρ and embedding dimension E averaged over the combined forecasts for each step.

We compared the combined forecast with its individual members obtained with the $(\mu + \lambda)$ -ES algorithm. Regarding the individual members, the order of the in-sample score was different from the order of the test (see Fig. 6(b)). In contrast, the combined forecast achieved the minimum error by far for the both the in-sample and test errors. This forecast stability is one of the largest advantages of combining forecasts. Simultaneously, we profiled the best embeddings for combination, as shown in Fig. 7. Although each embedding performs competitively, they exhibit vastly different embeddings. Interestingly, the embedding dimensions are approximately 10 for all embeddings. This suggests that the high-performing embedding dimensions are approximately identical although there are various high-performing embeddings for the forecasts.

We also profiled ρ and E averaged over the k_p combined members for each step. The results in Fig. 6(c) suggest the same trends as the toy models, even for the real-world data. Specifically, smaller dimensions and smaller values of ρ were selected for shorter steps.

Discussion

Our numerical experiments suggest that our proposed framework works well for a wide range of the data length, including short time series such as 100. In contrast, the framework did not appear to work fine with very noisy data, and MVE achieved the best performance, as shown in the noise sensitivity analyses (Figs. 2(c) and 4(c)). This suggests that noise prevents the framework from selecting optimal embeddings on the basis of the squared errors, and it results in obtaining invalid embeddings. MVE can still be a good option for very noisy data. Note that the performance of RDE was not satisfactory in some of our numerical experiments, but this is because the RDE framework was originally proposed for extremely high-dimensional data, e.g., hundreds of variables. As shown in Fig. 5(b), we can expect RDE to perform well for cases that include many informative variables, and delay embedding is not needed.

The diversity of suboptimal embeddings is deeply linked to the performance of the proposed framework. As shown in the numerical experiments, the proposed framework worked well for high-dimensional dynamics and datasets containing high-dimensional variables. On the other hand, it is not suitable to apply it to relatively low-dimensional data. We tested the framework with low-dimensional datasets, and the results suggest that the proposed framework does not always yield the best performance, especially for low-dimensional datasets such as the Rössler equations²⁷ (see Supplementary Information). This is because it is easy to find suitable embeddings to forecast in a brute-force manner, especially for low-dimensional datasets, while it becomes combinatorially difficult as the number of variables increases. In addition, the aim of K times optimization is to yield diverse embeddings (see the effect of K times optimization in Supplementary Information). However, if a dataset contains a small number of variables, almost the exact solutions are obtained for every optimization. We numerically ensured this behavior in the “Diversity of embeddings for split dataset” section in the Supplementary Information. The result shows that high-dimensional datasets with smaller data

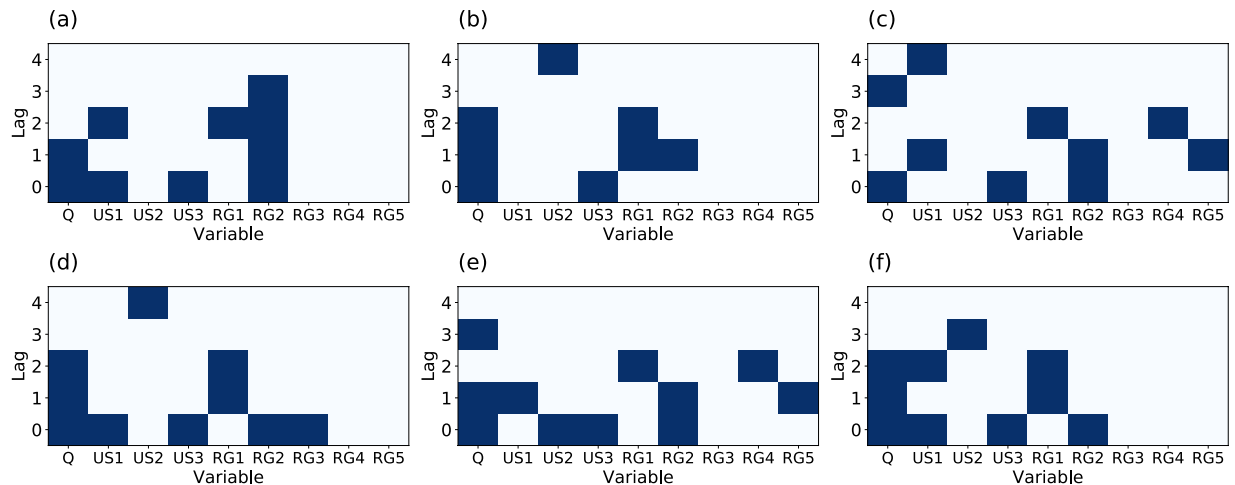


Figure 7. The top six embeddings to be combined for a 24-h-ahead forecast. Dark blue cells represent embeddings, and white cells indicate no embeddings. The horizontal axis indicates variables, and the vertical axis denotes the lags.

length produced diverse embeddings, but low-dimensional datasets with larger data length obtained similar or common embeddings across segments. This property is consistent with the performance of high- and low-dimensional datasets.

The proposed method combines forecasts via the simple average, optimizing the number of forecasts. The simple average works fine for most cases because suboptimal embeddings can equally forecast well thanks to solving the combinatorial optimization problem, and the test performance order is not always the same as the corresponding training data, especially for real-world data (see Fig. 6(b)). If suboptimal embeddings show quite different forecast performances, the weighted average based on the error distribution¹⁰ may work better. We can also apply the expert advice framework²⁸ if better forecasts are assumed to be changed over time.

We need to select the values of several parameters to apply the proposed framework, i.e., the number of split datasets K , the number of the suboptimal embeddings for each batch M , the criteria of the Hamming distance θ of Eq. (3), the number of filters L , and their components. We can control the variations in forecasts with these parameters. In order to set K , we consider the number of samples in each divided dataset; empirically, the framework works well with more than 50 samples for each divided dataset. Note that these samples are used only for the evaluation of an objective function, and the corresponding forecasts are yielded by using whole training data in an in-sample manner. Samples that are too small prevent the framework from finding good embeddings, but samples that are too large reduce the diversity of the embeddings. Although M can be any large integer (up to the total number of individuals of the evolution strategy), $M \leq 5$ is sufficient for most cases. Note that the value of M that is too large is overkill and unnecessarily increases the computational cost. The criteria for the Hamming distance θ directly controls the similarity of suboptimal embeddings. We set $\theta = 3$ in this study, and it worked fine for all cases. Note that conventional MVE and RDE guarantee that the Hamming distance among all embeddings is larger than two (i.e., $\theta = 2$) for a fixed embedding dimension. In practice, appropriate filters depend on the prediction steps. If we predict shorter time periods, high-pass filters will work fine; if we predict longer time periods, low frequencies are needed to obtain fine forecasts, as the numerical experiments suggest. Although we treat high-pass filters in this study, low-pass filters may work for extremely noisy data. We need to understand the effects of these parameters and set the appropriate values within reason, $P = KLM$, which is usually much smaller than the total number of embeddings evaluated in the first step.

Our numerical experiments suggest that the appropriate embeddings vary with the prediction steps. The results show that a smaller E with a smaller ρ is preferable for a short-term forecast, and a larger E with a larger ρ is preferable for a longer-term forecast. These results are consistent with those of an existing study²⁰ and with the fact that the longer-step forecasts need more complex map from the current inputs to the forecasts. Therefore, more information is needed to express the complex map.

Although we applied the $(\mu + \lambda)$ -ES algorithm to obtain the suboptimal embeddings, it is worth considering other methods to solve the combinatorial optimization problem. We can apply other evolution algorithms, e.g., simulated annealing²⁹ and ant colony optimization³⁰. Any algorithm may work as long as it can store the suboptimal embeddings through its optimization procedure. Testing and evaluating other algorithms are open problems.

Throughout this study, we used a variation of a local nonlinear prediction method to forecast the target variable. One of the most significant advantages of these local prediction methods is that we can compute in-sample forecasts very easily with a small computational cost; all we have to do is to exclude the current query from the nearest neighbors. In addition, owing to recent progress in approximate neighboring search algorithms (e.g., refs. 31,32), these algorithms are fast enough to apply the evolution algorithms. Moreover, it is also suitable to obtain suboptimal embeddings through the forecast errors because local nonlinear prediction methods are sensitive to the false nearest neighbors derived from invalid features. However, we can apply other regression methods, e.g., Gaussian process regression¹⁴, as done in RDE¹⁰, for further improvements. Note that these methods incur a very

large computational cost to yield in-sample forecasts, as mentioned in ref.¹⁰. To reduce the computational cost, we can apply these methods to only the process in the second step and can approximate the in-sample error by the k-fold cross-validation error.

Methods

Delay embedding. Delay embedding is a method of reconstructing the original state space using the delay coordinates $v(t)$. Here, with the n -dimensional observed time series $y(t) = [y_1(t), y_2(t), \dots, y_n(t)]$, variable $\sigma_i \in \{1, 2, \dots, n\}$, and time lag $\tau_i \in \{0, 1, \dots, l - 1\}$, $v(t) \in \mathbb{R}^E$ is defined as follows:

$$v(t) = [y_{\sigma_1}(t - \tau_1), y_{\sigma_2}(t - \tau_2), \dots, y_{\sigma_E}(t - \tau_E)] \quad (9)$$

where E is the embedding dimension, and no duplication is allowed for any element. According to embedding theorems^{1,3}, an attractor reconstructed by $v(t)$ is an embedding with the appropriate E . Note that the number of possible reconstructions increase combinatorially.

Optimizing embedding by an evolution strategy. We optimize embedding via combinatorial optimization to minimize forecast errors. Based on existing studies^{15,16}, we treat embeddings as a binary series, i.e., 1 denoting embedded and 0 denoting not embedded, and solve Eq. (2) by using a combinatorial optimization problem. For example, with $n = 3$ and $l = 2$, the embedding $[y_1(t), y_2(t), y_3(t - 1)]$ can be expressed by the $nl = 6$ binary code $e = [1, 0, 1, 0, 0, 1]$: the first element 1 corresponds to $y_1(t)$ to embed, the second element 0 corresponds to $y_1(t - 1)$ to embed, the third element 1 corresponds to $y_2(t)$ to embed, and so on. The optimization of a binary series is a typical combinatorial optimization problem. In this paper, we apply the $(\mu + \lambda)$ evolution strategy¹⁷. The parameter μ denotes the number of parents, λ denotes the number of offspring, and $+$ denotes *plus selection*, which means that the next parents are selected from both the previous parents and offspring. Precisely, the algorithm yields λ new offspring from the top-performing μ individuals selected from the previous parents and offspring. The algorithm is an elitist method since the algorithm retains the best individuals unless they are replaced by superior individuals. The algorithm tends to converge fast, but it is likely to become trapped at local optima. Therefore, the algorithm is suitable for our proposed framework. However, it is worth considering the application of other existing methods^{33–35} to obtain suboptimal embeddings.

Forecasting through FIR filters. We apply FIR filters to the original time series, and create embeddings using the filtered time series to obtain various embeddings and forecasts. Note that this procedure can also be written using the delay coordinates composed of the original time series with the corresponding filter matrix. We apply the following linear transformation for each variable:

$$z_i(t) = \left[\sum_{k=0}^{N-1} h_i(k) y_i(t - k) - \mu_i \right] / \sigma_i \quad (10)$$

where h_i 's are filter coefficients with N elements, $h_i(0) = 1 \quad \forall i$, and μ_i and σ_i are the coefficients for standardization. When we apply the filter to the target variable y_f , we obtain filtered forecast $\hat{z}_f(t)$. After forecasting $z_f(t)$, we need to restore $\hat{z}_f(t)$ for the original time series. For $p = 1$, we restore $\hat{y}_f(t + p|t)$ using $\hat{z}_f(t + p|t)$ as follows:

$$\hat{y}_f(t + p|t) = \sigma_f \hat{z}_f(t + p|t) + \mu_f - \sum_{k=1}^{N-1} h_f(k) y_f(t + p - k). \quad (11)$$

For $p \geq 2$, we obtain the p -steps-ahead forecast $\hat{y}_f(t + p|t)$ by substituting forecasts for the unobserved $y_f(t \geq t + 1)$ as follows:

$$\hat{y}_f(t + p|t) = \sigma_f \hat{z}_f(t + p|t) + \mu_f - \sum_{k=1}^{p-1} \left(h_f(k) \hat{y}_f(t + p - k|t) \right) + \sum_{k=p}^{N-1} \left(h_f(k) y_f(t + p - k) \right). \quad (12)$$

We can restore $\hat{y}_f(t + p|t)$ for any p to apply Eq. (12) recursively.

We can apply FIR filters such as low-pass filters to suppress noise. In contrast, an existing study¹⁸ focuses on the filtered delay coordinates constructed by blending the derivatives of time series and the original one to search the nearest neighbors. Although these are a type of high-pass FIR filters inspired by the difference of time series, such filters improve the prediction accuracy in less noisy time series. In this paper, we set the filter coefficients as $h_i(0) = 1$, $h_i(1) = \rho \quad \forall i$ with the number of taps $N = 2$ for all numerical examples, based on an existing study¹⁸, although there are many possible filters to improve forecast. More suitable filters can be designed in a future work.

Overall algorithm of the proposed framework. We denote a multivariate time series as $y(t) \in \mathbb{R}^n$ and the target variable to be forecasted as y_f . The proposed forecasting framework is presented in Algorithm 1 in terms of the user-defined number of split training datasets K , the number of top-performing embeddings chosen for each split M , the linear filters, and the number of filters L .

Algorithm 1. Proposed forecasting framework.

```

1: Start the algorithm by setting the set of suboptimal embeddings  $\mathcal{S}$  to the empty
2: for each  $K'$  in  $\{1, 2, \dots, K\}$  do ▷ First step: Obtaining suboptimal embeddings
3:   Initialize a set of the Hall of Fame embeddings  $\mathcal{H}$ 
4:   Initialize a set of embeddings  $\mathcal{P}$  (population)
5:   for each generation do
6:     Compute  $\hat{y}_f^e(t+p|t) \forall t \in \mathcal{T}_{train}^{K'}$  by map  $\psi_e \forall e \in \mathcal{P}$ 
7:     Evaluate the fitness by  $\sum_{t \in \mathcal{T}_{train}^{K'}} \sum_p \|\hat{y}_f^e(t+p|t) - y_f(t+p)\|$ 
8:     Append the elements of  $\mathcal{P}$  to  $\mathcal{H}$ 
9:     Create a new  $\mathcal{P}$  for the next generation using an evolution strategy
10:   end for
11:   Sort  $\mathcal{H}$  by fitness
12:   for each  $e_h$  in  $\mathcal{H}$  do
13:     if  $d(e_h, e_s) \geq \theta \forall e_s \in \mathcal{S}$  or  $\mathcal{S}$  is empty then Append  $e_h$  to  $\mathcal{S}$  with  $L$  filters
14:     if  $M \cdot L$  embeddings are appended to  $\mathcal{S}$  then break
15:   end for
16: end for
17: for each  $e$  in  $\mathcal{S}$  do ▷ Second step: Combining forecasts to minimize the in-sample error
18:   Compute  $\hat{y}_f^e(t+p|t) \forall t \in \mathcal{T}_{train}$  by map  $\psi_e$ 
19:   Evaluate the in-sample error of  $\hat{y}_f^e(t+p|t)$ 
20: end for
21:  $\mathcal{S}_p(i) :=$  Sort  $\mathcal{S}$  by the in-sample error of the  $p$ th step
22: for each  $k$  in the indices of  $\mathcal{S}_p$  do
23:   Compute the combined forecast  $Y_k(t+p|t)$  by a recursive formula
24:   Evaluate the in-sample error  $Err(k)$  of  $Y_k(t+p|t)$ 
25:   Obtain the optimal number to combine  $\hat{k}_p = \arg \min_{k=1,2,\dots,p} Err$ 
26: end for
27: while  $t \geq 0$  do ▷ Forecast unobserved samples  $t \geq 0$  by  $\hat{k}_p$  embeddings
28:   for  $k = 1$  to  $\hat{k}_p$  do
29:     Compute  $\hat{y}_f^{\mathcal{S}_p(k)}(t+p|t)$ 
30:   end for
31:   Yield the combined forecast by  $\sum_{k=1}^{\hat{k}_p} \hat{y}_f^{\mathcal{S}_p(k)}(t+p|t) / \hat{k}_p$ 
32: end while

```

Data availability

To obtain the flood dataset, please make a contact with Professor Christian W. Dawson (Loughborough University) directly. The other data related to this paper are available from the corresponding author on reasonable request.

Received: 22 July 2019; Accepted: 20 December 2019;

Published online: 20 January 2020

References

1. Takens, F. Detecting strange attractors in turbulence. *Lect. Notes Math. Berlin Springer Verlag* **898**, 366, <https://doi.org/10.1007/BFb0091924> (1981).
2. Sauer, T., Yorke, J. A. & Casdagli, M. Embedology. *J. Stat. Phys.* **65**, 579–616, <https://doi.org/10.1007/BF01053745> (1991).
3. Deyle, E. R. & Sugihara, G. Generalized theorems for nonlinear state space reconstruction. *PLoS ONE* **6**, e18295, <https://doi.org/10.1371/journal.pone.0018295> (2011).
4. Garcia, S. P. & Almeida, J. S. Multivariate phase space reconstruction by nearest neighbor embedding with different time delays. *Phys. Rev. E* **72**, 27205, <https://doi.org/10.1103/PhysRevE.72.027205> (2005).
5. Hirata, Y., Suzuki, H. & Aihara, K. Reconstructing state spaces from multivariate data using variable delays. *Phys. Rev. E* **74**, 26202, <https://doi.org/10.1103/PhysRevE.74.026202> (2006).
6. Pecora, L. M., Moniz, L., Nichols, J. & Carroll, T. L. A unified approach to attractor reconstruction. *Chaos: An Interdiscip. J. Nonlinear Sci.* **17**, 13110, <https://doi.org/10.1063/1.2430294> (2007).
7. Vlachos, I. & Kugiumtzis, D. Nonuniform state-space reconstruction and coupling detection. *Phys. Rev. E* **82**, 16207, <https://doi.org/10.1103/PhysRevE.82.016207> (2010).
8. Ye, H. & Sugihara, G. Information leverage in interconnected ecosystems: Overcoming the curse of dimensionality. *Sci.* **353**, 922–925, <https://doi.org/10.1126/science.aag0863> (2016).
9. Okuno, S., Aihara, K. & Hirata, Y. Combining multiple forecasts for multivariate time series via state-dependent weighting. *Chaos: An Interdiscip. J. Nonlinear Sci.* **29**, 33128, <https://doi.org/10.1063/1.5057379> (2019).
10. Ma, H., Leng, S., Aihara, K., Lin, W. & Chen, L. Randomly distributed embedding making short-term high-dimensional data predictable. *Proc. Natl. Acad. Sci.* **115**, E9994–E10002, <https://doi.org/10.1073/pnas.1802987115> (2018).
11. Sollich, P. & Krogh, A. Learning with ensembles: how over-fitting can be useful. In *Advances in neural information processing systems*, 190–196 (1996).
12. Kuncheva, L. I. & Whitaker, C. J. Measures of Diversity in Classifier Ensembles and Their Relationship with the Ensemble Accuracy. *Mach. Learn.* **51**, 181–207, <https://doi.org/10.1023/A:1022859003006> (2003).
13. Lorenz, E. N. Atmospheric predictability as revealed by naturally occurring analogues. *J. Atmospheric Sci.* **26**, 636–646 (1969).
14. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
15. Vitrano, J. B., Povinelli, R. J., B Vitrano, J. & Povinelli, R. J. Selecting dimensions and delay values for a time-delay embedding using a genetic algorithm. In *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation*, GECCO'01, 1423–1430 (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2001).
16. Small, M. Optimal time delay embedding for nonlinear time series modeling arXiv:nlin/0312011 (2003).

17. Schwefel, H.-P. *Numerical Optimization of Computer Models* (John Wiley & Sons, Chichester, 1981).
18. Okuno, S., Takeuchi, T., Horai, S., Aihara, K. & Hirata, Y. Avoiding underestimates for time series prediction by state-dependent local integration. *Math. Eng. Tech. Reports METR 2017–22*, The University of Tokyo (2017).
19. Lorenz, E. N. Predictability: a problem partly solved. In *Seminar on Predictability*, 1–18 (ECMWF, Reading, England, 1996).
20. Chayama, M. & Hirata, Y. When univariate model-free time series prediction is better than multivariate. *Phys. Lett. A* **380**, 2359–2365, <https://doi.org/10.1016/j.physleta.2016.05.027> (2016).
21. Kuramoto, Y. & Tsuzuki, T. Persistent Propagation of Concentration Waves in Dissipative Media Far from Thermal Equilibrium. *Prog. Theor. Phys.* **55**, 356–369, <https://doi.org/10.1143/PTP.55.356> (1976).
22. Sivashinsky, G. I. Nonlinear analysis of hydrodynamic instability in laminar flames-I. Derivation of basic equations. *Acta Astronaut.* **4**, 1177–1206, [https://doi.org/10.1016/0094-5765\(77\)90096-0](https://doi.org/10.1016/0094-5765(77)90096-0) (1977).
23. Dawson, C. *et al.* A comparative study of artificial neural network techniques for river stage forecasting. In *Proceedings of the International Joint Conference on Neural Networks*, vol. 4, 2666–2670, <https://doi.org/10.1109/IJCNN.2005.1556324> (IEEE, Montreal, Canada, 2005).
24. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780, <https://doi.org/10.1162/neco.1997.9.8.1735> (1997).
25. Boser, B. E., Guyon, I. M. & Vapnik, V. N. A Training Algorithm for Optimal Margin Classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory, COLT' 92*, 144–152, <https://doi.org/10.1145/130385.130401> (ACM, New York, NY, USA, 1992).
26. Breiman, L. Random Forests. *Mach. Learn.* **45**, 5–32, <https://doi.org/10.1023/A:1010933404324> (2001).
27. Rössler, O. E. An equation for continuous chaos. *Phys. Lett. A* **57**, 397–398 (1976).
28. Cesa-Bianchi, N. & Lugosi, G. *Prediction, Learning, and Games* (Cambridge University Press, 2006).
29. Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by simulated annealing. *Sci.* **220**, 671–680 (1983).
30. Dorigo, M. & Stützle, T. *Ant Colony Optimization* (Bradford Company, Scituate, MA, USA, 2004).
31. Muja, M. & Lowe, D. G. Scalable nearest neighbor algorithms for high dimensional data. *IEEE Transactions on Pattern Analysis Mach. Intell.* **36**, 2227–2240, <https://doi.org/10.1109/TPAMI.2014.2321376> (2014).
32. Fu, C. & Cai, D. EFANNA: An extremely fast approximate nearest neighbor search algorithm based on kNN graph *arXiv:1609.07228* (2016).
33. Runge, J., Donner, R. V. & Kurths, J. Optimal model-free prediction from multivariate time series. *Phys. Rev. E* **91**, 1–14, <https://doi.org/10.1103/PhysRevE.91.052909> *arXiv:1506.05822* (2015).
34. Vlachos, I. & Kugiumtzis, D. State space reconstruction from multiple time series. In *Topics on Chaotic Systems: Selected Papers from Chaos 2008 International Conference*, 378–387, https://doi.org/10.1142/9789814271349_0043 (World Scientific, 2009).
35. Chen, Y. & Wong, M. L. An ant colony optimization approach for stacking ensemble. *Second. World Congr. on Nat. Biol. Inspired Comput. (NaBIC)* 146–151, <https://doi.org/10.1109/NABIC.2010.5716282> (2010).

Acknowledgements

We thank Professor Christian W. Dawson for permission to use the flood dataset. This research is partially supported by Kozo Keikaku Engineering Inc., JSPS KAKENHI 15H05707, WPI, MEXT, Japan, AMED JP19dm0307009, and JST CREST JPMJCR14D2, Japan. We thank Dr. Takahiro Omi for fruitful discussions.

Author contributions

S.O., K.A. and Y.H. conceived the original design of this study. S.O. and K.A. designed the study of the toy models, and S.O. and Y.H. designed the study of the flood data. S.O. conducted the numerical experiment. All authors contributed to interpreting the results and writing the manuscript. All authors checked the manuscript and agreed to submit the final version of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-019-57255-4>.

Correspondence and requests for materials should be addressed to S.O.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020