Denoising ambient seismic field correlation functions with convolutional autoencoders

Loïc Viens[1] and Chris Van Houtte[2]

[1] *Disaster Prevention Research Institute, Kyoto University, Japan*

[2] *GNS Science, Lower Hutt, New Zealand*

1 July 2019

SUMMARY

This paper is a non-peer reviewed preprint submitted to EarthArXiv and has been submitted to Geophysical Journal International for peer review.

Seismic interferometry is an established method for monitoring the temporal evolution of the Earth's physical properties. We introduce a new technique to improve the precision and temporal resolution of seismic monitoring studies based on deep learning. Our method takes advantage of a convolutional denoising autoencoder, called ConvDeNoise, to denoise ambient seismic field correlation functions. The technique can be applied to traditional two-station cross-correlation functions but this study focuses on single-station cross-correlation (SC) functions. SC functions are computed by cross correlating the different components of a single seismic station and can be used to monitor the temporal evolution of the Earth's near surface. We train and apply our algorithm to SC functions computed with a time resolution of 20 minutes at a seismic station in the Tokyo metropolitan area, Japan. We show that the relative seismic velocity change ($dv/v(t)$) computed from SC functions denoised with ConvDeNoise has less variability than that calculated from raw SC functions. Compared to other denoising methods such as the SVD-based Wiener Filter method developed by Moreau et al. (2017), the dv/v results obtained after using our algorithm have similar precision. The advantage of our technique

is that once the algorithm is trained, it can be apply to denoise near-real-time SC functions. The near-real-time aspect of our denoising algorithm may be useful for operational hazard forecasting models, for example when applying seismic interferometry at an active volcano.

Key words: Seismic interferometry – Seismic noise – Computational seismology – Hydrogeophysics – Machine learning

1   INTRODUCTION

In the last decade, seismic interferometry has become a very popular technique for extracting useful geophysical information from continuous seismic signals. By cross-correlating two continuous data streams of the ambient seismic field, the response of the medium between the streams' recording locations can be calculated (Shapiro & Campillo 2004). Under the assumption that the ambient seismic wavefield is equipartitioned, mathematical models demonstrated that the correlation function computed between two seismic stations should yield the elastodynamic Green's function (Weaver & Lobkis 2001; Fichtner & Tsai 2019). However, such a condition is rarely fulfilled on Earth, as the ambient seismic field is generated by uneven distributions of sources, for example from ocean waves at low frequencies (< 1 Hz) and human activities at high frequencies. This uneven source distribution can corrupt the correlation functions with spurious, non-physical, signals (Shapiro et al. 2006; Zeng & Ni 2010) and lead to errors in the travel-time measurements (Tsai 2009; Yao & Van Der Hilst 2009).

Despite the anisotropic nature of the ambient seismic wavefield, correlation functions have been heavily utilised to image the geophysical properties of the Earth from the shallow subsurface down to the core (Shapiro et al. 2005; Lin et al. 2008; Boué et al. 2013; Mordret et al. 2014). Correlation functions have also been widely used for seismic velocity monitoring, where changes to arrival times of coda wave-packets are tracked through time and used to infer relative seismic velocity (dv/v) changes in the Earth. This technique has been used to monitor geophysical changes in a wide range of environments such as volcanoes (Brenguier et al. 2008b), landslides (Mainsant et al.

2012), aquifer water levels (Voisin et al. 2017; Clements & Denolle 2018), glaciers (Mordret et al. 2016), and active faults (Brenguier et al. 2008a; Gassenmeier et al. 2016).

The quality of the correlation functions used for both tomographic and monitoring applications, depends significantly on the processing of the ambient seismic field signals. One way to improve the quality of correlation functions is to pre-process ambient seismic field records with techniques such as one-bit normalisation and/or pre-whitening (Bensen et al. 2007). Individual correlation functions are then calculated from short duration time-windows ranging from a few minutes to a few hours. Despite the pre-processing, correlation functions computed over these short time windows are often corrupted by the non-uniform distribution of ambient seismic field sources and tend to have a low signal to noise ratio (SNR). For imaging purposes, common mitigation practice consists of stacking correlation functions over a longer period of time (e.g., from a few days to a few years) to increase the SNR, and decrease the influence of the non-uniform distribution of the noise sources. In monitoring applications, correlation functions can only be stacked over a short period of time to maintain temporal resolution, meaning that correlation functions tend to have lower SNR. The lower SNR makes monitoring applications of seismic interferometry much more sensitive to incoherent signals in the ambient seismic wavefield than tomographic applications.

To avoid stacking correlation functions over long periods of time, several techniques have been developed to 'denoise' correlation functions and improve their SNR while maintaining temporal resolution. A first set of methods involves applying a wavelet-type transform to a stacked correlation function, and then using the transformed data to design data-adaptive filters. These filters can then be applied to individual, unstacked correlation functions to remove incoherent energy (Baig et al. 2009; Hadziioannou et al. 2011; Stehly et al. 2015). However, these filter-based techniques require the individual correlation functions to remain relatively similar to the stacked correlation function, otherwise the filters have the potential to remove real signal. Therefore, their application is essentially limited to detecting relatively small seismic velocity changes in the medium. An alternative method was suggested by Durand et al. (2011), who applied a singular value decomposition (SVD) to their correlation functions to isolate undesirable signals, in their case seasonal changes

in the noise source distribution, from the signal of interest. The eigenvector that corresponds to the seasonal variations could then be removed from their analysis. Moreau et al. (2017) recently extended the approach of Durand et al. (2011) by combining a singular value decomposition with a Wiener filter (SVDWF) to denoise correlation functions. In this approach, the SVD of a matrix of correlation functions is first calculated. Then, a Wiener filter is individually applied to a selected number of singular vectors, before the matrix of correlation functions is reconstructed using the filtered singular vectors. The reconstructed matrix of correlation functions is finally filtered again with the same Wiener filter. The Wiener filter works by minimising signal incoherence across a two-dimensional filter kernel. In the case of ambient noise monitoring, the two dimensions are the correlation lag time and the observation period. The SVDWF technique is a powerful method for denoising correlation functions, and has the advantage over the S-transform and curvelet methods in that it does not require a reference waveform.

While very useful for retrospective analyses, the SVDWF technique has the disadvantage of being rather impractical to apply in near-real-time monitoring. The technique is relatively inefficient computationally, because when new correlation functions are calculated in near-real-time, it is necessary to recalculate the SVD of an ever-increasing matrix of correlation functions. Additionally, because the Wiener filter kernel requires a value in the observation time dimension, the correlation function at a given timestamp (t_i) depends on the surrounding correlation functions at t_{i-k} and t_{i+k}, where k is the one-sided dimension of the Wiener filter kernel. This creates practical issues for the correlation at t_n, where n is the time 'now', because the denoised correlation function will be different when a new correlation function arrives and the correlation becomes t_{n-1}, and again at t_{n-2}, etc. The changes may affect previously-calculated dv/v values backward in time from t_n, meaning that these values would need to be recalculated.

To increase the temporal resolution of continuous seismic monitoring without the drawbacks of the previously mentioned techniques, we propose to use denoising autoencoders (DAEs, Vincent et al. 2008, 2010; Goodfellow et al. 2016) to denoise correlation functions in an unsupervised manner. An autoencoder compresses some input data into a low-dimensional representation and

reconstructs the input data from that representation (Hinton & Salakhutdinov 2006; Bengio 2009). A denoising autoencoder is a specific type of autoencoder that is trained on pairs of partially-corrupted and clean data. During the training phase, the autoencoder learns the salient features of an underlying signal within some corrupted data, and how the noise can be removed from the data to reveal a denoised signal. DAEs have been used to denoise a wide range of data, such as images (Xie et al. 2012; Shen et al. 2017) and speech (Lu et al. 2013; Grais & Plumbley 2017). Not only can DAEs be effective at denoising correlation functions for more precise monitoring, they are also a practical, efficient solution for denoising correlation functions in near-real-time data products. Once trained, DAE can denoise correlation functions for dv/v calculations with minimal required computation, and unlike the SVDWF method, no recalculations of dv/v for previous time steps are necessary.

In this study, we aim to demonstrate that a DAE, specifically a convolutional denoising autoencoder (CDAE), can effectively denoise single-station cross-correlation (SC) functions to more precisely monitor changes in the physical properties of the shallow subsurface. Unlike two-station cross-correlation functions, SC functions are computed by cross correlating signals from the different components of a single station. SC functions have been previously used to reveal clear geophysical changes in subsurface media as a result of volcanic activity (De Plaen et al. 2016) and strong ground motions (Hobiger et al. 2014, 2016; Viens et al. 2018a). We first introduce the computation of SC functions and the stretching technique that is used to compute relative seismic velocity changes. We then present the architecture of our CDAE and apply it to denoise SC functions computed from data recorded at a seismic station in the Tokyo metropolitan area in Japan.

2 METHODS

2.1 Single-station cross-correlation functions and relative seismic velocity changes

In this section, we briefly introduce the computation of single-station cross-correlation functions and the stretching method used to calculate relative seismic velocity changes, and refer the reader to Viens et al. (2018a) for further details.

Single-station cross-correlation functions are computed by cross-correlating continuous signals recorded by the three-component sensors at a seismic station. In this study, we use the continuous data recorded by the NS7M station of the Metropolitan Seismic Observation network (MeSO-net, Kasahara et al. 2009; Sakai & Hirata 2009). The NS7M station is located 5.25 km away from the Funabashi weather station which is operated by the Japan Meteorological Agency (JMA, Figure 1). We select 20 months of data recorded between April 1, 2017 and December 31, 2018 by the three-component accelerometers with a sampling rate of 200 Hz and correct the records for their instrumental response. As MeSO-net sensors are located in 20-m deep boreholes, the horizontal components are generally not aligned along the north-south and east-west directions. We rotate the horizontal components using the orientations determined by Kano et al. (2015). We then divide the dataset into five-minute-long time series and apply the one-bit normalisation method to reduce the unwanted effect of transient signals (Bensen et al. 2007). The SC functions between the vertical (Z) and north-south (N) components and between the vertical and east-west (E) components are finally computed in frequency domain as

$$\mathrm{SC}_{Z,H}(t) = \mathfrak{F}^{-1}\left(\frac{\hat{v}_H(\omega)\hat{v}_Z^*(\omega)}{\{|\hat{v}_Z(\omega)|\}^2}\right),\tag{1}$$

where $\hat{v}_Z^*(\omega)$ is the Fourier transform of a 5-minute vertical acceleration record and $\hat{v}_H(\omega)$ is the Fourier transform of the 5-minute acceleration time window recorded by either the N or E component (H). The asterisk symbol (*) represents the complex conjugate, $|\cdot|$ denotes the absolute value, and $\{\cdot\}$ is a smoothing of the denominator term to stabilise the deconvolution. The inverse Fourier transform ($\mathfrak{F}^{-1}$) of the deconvolution output is computed to retrieve time-domain SC functions. Then, the causal part of the SC functions is selected, tapered and bandpass filtered between 1 and 20 Hz using a two-pass, four-pole Butterworth filter. Finally, the SC functions are stacked over 20 minutes (e.g., stack of four 5-minute waveforms) to increase the SNR.

Relative seismic wave velocity variations through time ($dv/v(t)$) are calculated with respect to a reference waveform. The reference waveforms for the Z–N and Z–E components are calculated by stacking the respective SC functions from April 1 to December 31, 2017. The relative seismic wave velocity variations are finally computed between each current (20-minutes) SC function and

its corresponding reference waveform using the stretching method (Lobkis & Weaver 2003; Sens-Schönfelder & Wegler 2006). Under the assumption of homogeneous seismic velocity changes in the medium, the relative velocity is related to the relative time shift between the waveforms as

$$\epsilon = -\frac{dt}{t} = \frac{dv}{v},$$ (2)

where dt/t is the relative time shift and dv/v is the relative velocity change.

The computation of the relative velocity change is performed in two steps. First, we stretch and compress the first second of each current SC function between $-\epsilon_{max}$ and ϵ_{max}, with ϵ_{max} being equal to 5%, every 0.5%. For each pair of stretched-reference waveforms, we compute a correlation coefficient and select the stretched waveform that maximises it. Second, we refine the measurements around the ϵ with the previously-maximised correlation coefficient, by interpolating the stretched waveforms 500 times between the neighbouring ϵ values. The final relative velocity change is obtained by selecting the refined ϵ value with the maximum correlation coefficient.

We finally average the dv/v measurements computed from the $Z–N$ and $Z–E$ components following the weighting average scheme proposed by Hobiger et al. (2012). This averaging can be summarised as

$$\frac{dv}{v}(t) = \frac{\sum\limits_{j=1}^{2} c_j^2(t) \cdot \frac{dv_j}{v}(t)}{\sum\limits_{j=1}^{2} c_j^2(t)},$$ (3)

where $c_j(t)$ is the correlation coefficient after stretching for the jth component ($Z–N$ or $Z–E$). The variable t represents the time for which the velocity variations are computed (e.g., every 20 minutes). The correlation coefficients after stretching are also averaged following Hobiger et al. (2012) as

$$c(t) = \frac{\sum\limits_{j=1}^{2} c_j^3(t)}{\sum\limits_{j=1}^{2} c_j^2(t)}.$$ (4)

2.2   Autoencoders, denoising autoencoders, and convolutional denoising autoencoders

Autoencoders are types of neural networks that are composed of an encoding part and a decoding part (Vincent et al. 2008). The simplest autoencoder is composed of an input layer, one hidden layer, and an output layer. For a given input vector $\mathbf{x}$ of length d such as $\mathbf{x} \in [0,1]^d$, the hidden layer encodes the input vector $\mathbf{x}$ to a latent representation $\mathbf{y} \in [0,1]^{d'}$. This is done by using a $d' \times d$ matrix of weights $\mathbf{W}$ and an activation function f as

$$\mathbf{y} = f(\mathbf{W}\mathbf{x} + \mathbf{b}), \tag{5}$$

where $\mathbf{b}$ is a bias vector. The decoder then maps the encoded latent representation y back into a reconstruction $\hat{\mathbf{x}} \in [0,1]^d$ of the same shape as $\mathbf{x}$ as

$$\hat{\mathbf{x}} = f\left(\mathbf{W}'\mathbf{y} + \mathbf{b}'\right). \tag{6}$$

The parameters of the network, $\theta = \{\mathbf{W}, \mathbf{b}\}$ and $\theta' = \{\mathbf{W}', \mathbf{b}'\}$, are optimised by minimising the reconstruction error

$$\theta^{\star}, \theta'^{\star} = \underset{\theta, \theta'}{\arg\min}\left(\mathcal{L}\left(\mathbf{x}, \hat{\mathbf{x}}\right)\right), \tag{7}$$

where $\mathcal{L}$ is a loss function such as the squared error $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) = \|\mathbf{x} - \hat{\mathbf{x}}\|^2$.

The basic goal of this simple autoencoder is to obtain an output $\hat{\mathbf{x}}$ that is almost identical as the input $\mathbf{x}$. However, the mapping of $\mathbf{x}$ to the hidden representation $\mathbf{y}$ can be viewed as a lossy compression of $\mathbf{x}$, which implies that the reconstructed signal $\hat{\mathbf{x}}$ is degraded compared to the original input. For a singular hidden layer with linear activation functions f, the mapping of $\mathbf{x}$ to a lower representation is equivalent to the dimensionality reduction obtained with the Principal Component Analysis (PCA, Hinton & Salakhutdinov 2006). However, for an autoencoder composed of multiple hidden layers and non-linear activation functions, such as the rectified linear unit function (ReLU, Nair & Hinton 2010) or the sigmoid function, the autoencoder behaves differently from the PCA and can capture non-linear features of the input distribution.

Denoising autoencoders work in a similar way as autoencoders, but are trained with pairs of a corrupted ($\mathbf{x_n}$) and clean ($\mathbf{x_c}$) signals. The corrupted signal $\mathbf{x_n}$ is fed to the network which tries to

optimise its weights and biases by minimising the reconstruction error between its output $\hat{\mathbf{x}}$ and the clean signal $\mathbf{x_c}$ (e.g., $\mathcal{L}(\mathbf{x_c}, \hat{\mathbf{x}})$).

Classic autoencoders are generally composed of fully-connected layers, which means that each neuron of a given hidden layer is connected with every neuron of the previous layer. Such a large number of connections leads to a total number of parameters that increases significantly with the number of layers. To reduce the number of parameters and thus reduce memory usage and training time of the network, convolutional layers can be used instead of fully connected layers. For such a network, each neuron is only connected to nearby neurons of the previous layer with a set of small filters. This allows the network to be deeper with fewer parameters compared to fully-connected layers.

2.3 ConvDeNoise: a convolutional denoising autoencoder for seismic monitoring

The goal of our CDAE, called ConvDeNoise, is to denoise SC functions to increase the precision and the temporal resolution of relative velocity change measurements. To train a CDAE to denoise SC functions, examples of paired 'noisy' and 'clean', noise-free, SC waveforms are required. The noisy SC functions are simply the 20-minute $Z–E$ and $Z–N$ SC functions, but the corresponding clean SC functions are unknown. To approximate a clean SC function at each time step, we first stack the noisy SC functions over two hours around the considered 20-minute time window (i.e., for a SC function calculated from the data recorded between 5:00 and 5:20 of a given day, the corresponding clean waveform is the stack of the SC functions between 4:10 and 6:10). The stack of the waveforms minimises the incoherent energy that corrupts the observation of the SC function, and by having a large training set of paired noisy and clean SC functions, the CDAE is able to learn the characteristics of the noise and how to suppress them.

2.3.1 Training set

For the NS7M station, the training set is composed of pairs of noisy and clean SC functions computed from April 2, 2017 to March 30, 2018 (e.g., 363 days). We show the pairs of noisy and

clean SC functions for the $Z-N$ and $Z-E$ components at this station in Figure 2. Several wave arrivals can be observed for both noisy and clean SC functions, but appear more clearly in the clean SC functions. Note that the data are normalised between 0 and 1 because it helps to speed up the training process (LeCun et al. 2012).

To train the network, a training set and a validation set are required. The training set is used to train the model parameters/weights and the validation set is used to evaluate the performance of the model during the training phase and to decide when to stop it. From April 2, 2017 to March 30, 2018, we have 26,136 pairs of noisy/clean SC functions for each component (e.g., $Z-E$ and $Z-N$, total 52,272). As our algorithm takes the time-dependent $Z-E$ and $Z-N$ SC functions as input, we randomly divide the dataset into a training set and a validation set but keep the $Z-E$ and $Z-N$ SC functions at each given time together. The training set is composed of 90% of the 26,136 pairs of noisy/clean SC functions for each component (e.g., $23,515 \times 2$ SC functions) and the validation set is composed of the remaining 10% (e.g., $2,621 \times 2$ SC functions). The SC functions calculated from the data recorded after April 1, 2018 are not used to train the algorithm but to assess its performance (e.g., generalisation) on an unseen test data set.

2.3.2   Architecture of ConvDeNoise

ConvDeNoise is composed of (1) an input layer, which takes both noisy $Z-E$ and $Z-N$ SC functions at a given station, (2) six hidden layers which compose the encoder and the decoder, and (3) an output layer which computes denoised $Z-E$ and $Z-N$ SC functions. A schematic representation of the architecture of the network is shown in Figure 3.

The input layer of the network is composed of 2 channels with 200 samples (or features), which corresponds to the 1-s duration $Z-E$ and $Z-N$ noisy SC functions sampled at 200 Hz. The input data are first passed to the encoder, which is composed of three hidden layers. Each hidden layer is the result of the data going through a one-dimensional (1D) convolution operation, a ReLU activation function, and a max-pooling operation. The 1D convolution takes the data and convolves them with a set of filters. For the first layer, the size of the kernel, which is the window sliding through

the data, is set to 130 samples, and we use 40 filters. All convolutions have a stride length of 1 and use zero padding as the boundary condition, which implies that the output of the convolution has the same length as its original input. A bias term per filter is added to the sum of the convolved data. The first convolution has 10440 trainable parameters ($40 \times 130 \times 2$ weights and 40 biases). The output of the 1D convolution is then passed to a ReLU activation function, which introduces non-linearity to the data representation, and to a max-pooling operation, which reduces the number of features by a factor of two by keeping the maximum value within a window of two features. The first hidden layer is composed of 40 channels with 100 features (Figure 3). The second and third 1D convolution of the encoder both use 40 filters with kernel sizes of 65 and 40 features, respectively. The bottleneck of the network, which is the latent representation of our input data, is composed of 40 channels each with 25 features.

The encoded representation is then passed to the decoder which is composed of three hidden layers. The layers are obtained using a 1D convolution, a ReLU activation function, and an up-sampling operation. The up-sampling of the data increases the number of features in the representation maps to generates new ones with higher dimensions. The output layer is finally obtained by a 1D convolution and a sigmoid activation function. This layer outputs the denoised $Z{-}E$ and $Z{-}N$ SC functions which have the same dimension as the initial input.

This architecture may not be appropriate for every set of correlation functions, as the architecture might depend on the quality and quantity of training data as well as the dimensions of the input data. Nevertheless, the number of layers, filter sizes, and kernel sizes can easily be changed.

2.3.3   *Training the network*

We optimise the 389,042 parameters of the network (e.g., 388,800 weights and 242 biases) by minimising a mean squared error loss function between the clean SC functions and the denoised SC functions from the output of ConvDeNoise. We use the Adam optimiser (Kingma & Ba 2014) with all parameters set to the default values recommended by the authors to update the network parameters. To reduce the memory usage during training, we use a batched stochastic gradient

descent algorithm. We first randomly shuffle the data of the training set and form batches of 256 windows. For each training step, a batch is fed to the network. The loss function is used to calculate the loss of the batch and the network parameters are updated using back-propagation (LeCun et al. 2015). We train the network over 200 epochs, with an early stopping if the loss of the validation set does not improve for 20 epochs. The network has been developed in Python using the Keras library (Chollet et al. 2015).

To train the network on 35 CPUs (Intel Xeon Gold 6154 at 3.00 GHz), the average time to complete one epoch is around 20 s, while training the network on an Nvidia Quadro P6000 GPU has a computation time of around 2.5 s per epoch. For our training data set, 80 to 150 epochs are generally required to train the network. Note that if the network is trained multiple times, the results change slightly as the initial weights of each convolutional operation are randomly set with a Xavier normal initializer (Glorot & Bengio 2010). Moreover, the results will also change if another training set data is used. Finally, the weights trained for a specific station cannot be used to denoise SC functions from another station as their characteristics (e.g., frequency content or noise level) might be different.

3   RESULTS

3.1   Denoising SC functions at the NS7M station

Figure 4 compares raw SC functions from April 1, 2017 to December 31, 2018 with those denoised using ConvDeNoise, and those denoised with the SVDWF technique. For the SVDWF approach, we keep the first 25 singular values and set the K and L parameters, which define the size of the Wiener filter, to 3 and 5, respectively. The $K = 3$ parameter corresponds to a 1-h window in the vertical direction in Figure 4c and the $L = 5$ parameter is a 0.025 s window in the horizontal direction. We only keep the first 25 singular values which contain most of the information (Supplementary material Figure S1). For the SC functions denoised using ConvDeNoise, the data between April 1, 2018 and December 31, 2018 have not been previously seen by the algorithm during the training period.

For all panels in Figure 4, consistent wave arrivals can be observed through the considered time period in the Z–E and Z–N SC functions. Small variations of wave-packet arrival times can also be observed, but these variations are clearer in the denoised SC functions than in the raw waveforms. For example, the wave packet arriving at 0.25 s on the Z–E component fluctuates seasonally, but this is much more apparent on the denoised SC functions.

3.2   Relative velocity change (dv/v) measurements

For each 20-minute SC function, we compute the relative velocity change (dv/v) with respect to a reference waveform that is defined as the stack of the raw SC functions computed over 2017. dv/v is calculated independently for the raw SC functions, those denoised with ConvDeNoise and those denoised with SVDWF, using technique- and component-specific reference waveforms. For both the Z–E and Z–N components, we show the computed relative velocity changes in Figure 5. For all panels in Figure 5, clear temporal variations of dv/v can be observed over the 20-month observation period. However, the dv/v measurements computed from the raw SC functions have more variability than those calculated from the denoised SC functions. The dv/v measurements computed from the SC function denoised with ConvDeNoise show even less variability than with the SVDWF method.

To evaluate the accuracy of the dv/v measurements calculated from the denoised SC functions, we stack the dv/v measurements obtained from the raw, denoised with ConvDeNoise and with the SVDWF method over 1 day (e.g., stack of 72 dv/v measurements) and average the Z–E and Z–N components using the weighting average scheme from Equation 3. The averaging over a longer period of time decreases the temporal resolution of the dv/v measurements but increases their stability. The averaging over the two components also increases the stability of the dv/v measurements. Therefore, we expect the stacked and averaged dv/v measurements from the three methods to be similar.

We show the stacked and averaged $dv/v(t)$ measurements over the 20-month observation period in Figure 6. The $dv/v(t)$ obtained from SC functions denoised with ConvDeNoise and SVDWF

reproduce well the $dv/v(t)$ measurements obtained from raw waveforms, as most variations are accurately reproduced. This indicates that the 20-minute dv/v measurements calculated from the denoised waveforms are relatively accurate. More interestingly, the dv/v calculated using the SC functions denoised with ConvDeNoise after April 1, 2018, which are not shown during the training phase, show similar variations as those calculated from the raw SC functions. This indicates that our network generalises well.

4 DISCUSSION

4.1 Effect of earthquake ground motions

To investigate the potential effect of earthquake strong ground motions on the $dv/v(t)$ variations, we show the peak ground acceleration (PGA) through time at the NS7M station in Figure 6d. The PGA is the geometric mean of the peak ground accelerations recorded by the two horizontal components over 1-day periods. Over the 20-month observation period, we do not observe any clear drop of dv/v related to the PGA peaks. This is most likely because the largest PGA is smaller than 0.1 m/s^2, which is too low to cause any damage to the shallow subsurface. For comparison, the PGA during the 2011 M_w 9.0 Tohoku-Oki earthquake at the NS7M station exceeded 1 m/s^2 and caused a dv/v drop of $\sim 7\%$ (Viens et al. 2018a).

4.2 Effect of temperature variations

We show the 1-day average temperature between April 1, 2017 and December 31, 2018 at the Funabashi weather station in Figure 6c. While Viens et al. (2018b) showed that there is a positive correlation between dv/v variations and temperature fluctuations at the NS7M station between 2010 and 2017, such a correlation does not appear clearly over our 20-month observation period. Nevertheless, the highest values of dv/v are observed during the periods where the temperatures are the highest in Figure 6.

4.3 Effect of precipitation events

The main feature in Figure 6 is that rapid dv/v drops seem to occur during precipitation events. Sensitivity of dv/v to precipitation events has already been observed in the past (Sens-Schönfelder & Wegler 2006; Viens et al. 2018a), however, the 1-day resolution of Figure 6 does not allow us to accurately measure the duration over which the changes occur.

In Figure 7, we focus on 16 days between September 25 and October 11, 2017 and show the $dv/v(t)$ measurements with a 20-minute time resolution calculated from the raw and denoised SC functions. We also present the precipitation data with a 20-minute time resolution at the Funabashi weather station. The first rainfall event started on September 27 at 23:40 and ended on September 28 at 8:20. The cumulative precipitation during this period is 69 mm with some heavy rain between 4:00 and 7:00 (e.g., average of 14.5 mm per hour). For the second event shown in Figure 7d, the rain started on October 6 at 16:40 and ended around 5:20 the next day. The cumulative precipitation for this event is 54.5 mm and some heavy rain was recorded between 2:00 and 4:00 on October 7 with a precipitation rate of 8 mm per hour.

The $dv/v(t)$ calculated from the raw SC functions shows a relatively continuous decrease during the two precipitation events, which starts with the rain and seems to last for some time after the rain stops. The $dv/v(t)$ calculated from the SC functions denoised with ConvDeNoise and the SVDWF method, however, show slightly different behaviour. For the first event, the dv/v decreases slowly during the first hours where the precipitation is less than 4 mm/20-min. The increase of the precipitation up to 12 mm/20-min triggers a sudden drop of dv/v which can only be observed for the dv/v calculated from denoised waveforms. The change in slope of the dv/v measurements coincident with the heaviest rainfall, is shown in more detail in Figure 8. For the second precipitation event, the precipitation rate is relatively constant, and $dv/v(t)$ decreases smoothly with time until the end of the precipitation event. These changes in dv/v with rainfall events are likely to be structural changes of the subsurface rather than due to variations to the source of the ambient seismic wavefield, because dv/v does not return to the pre-rain level at the end of each rainfall event.

5   CONCLUSIONS

We introduced a new technique based on deep learning to denoise correlation functions and improve the precision and temporal resolution of seismic monitoring measurements. After training, we showed that our CDAE, called ConvDeNoise, can be used to denoise SC functions at a MeSO-net station located in the Tokyo metropolitan area. The relative velocity changes (dv/v) calculated from the denoised SC functions are more precise than those obtained from the raw SC functions, and provide comparable results to the SVDWF denoising method.

An ultimate goal of ambient noise monitoring is to develop operational algorithms and data products that can be used to monitor environmental and hazard events in near-real-time. The main advantage of ConvDeNoise is that the algorithm can be applied efficiently in near-real-time. In this study, we have detected relative velocity variations related to precipitation. There are many other applications where having near-real-time relative velocities from denoised correlation functions will also have step-change improvements. For example, applying the technique to a seismograph network at an active volcano could improve operational forecasting of volcanic unrest and eruption.

REFERENCES

Baig, A. M., Campillo, M., & Brenguier, F., 2009. Denoising seismic noise cross correlations, *J. Geophys. Res. Solid Earth*, **114**(B8).

Bengio, Y., 2009. Learning deep architectures for AI, *Foundations and Trends in Machine Learning*, **2**, 1–127, Also published as a book. Now Publishers, 2009.

Bensen, G. D., Ritzwoller, M. H., Barmin, M. P., Levshin, A. L., Lin, F., Moschetti, M. P., Shapiro, N. M., & Yang, Y., 2007. Processing seismic ambient noise data to obtain reliable broad-band surface wave dispersion measurements, *Geophys. J. Int.*, **169**, 1239–1260.

Boué, P., Poli, P., Campillo, M., Pedersen, H., Briand, X., & Roux, P., 2013. Teleseismic correlations of ambient seismic noise for deep global imaging of the Earth, *Geophys. J. Int.*, **194**, 844–848.

Brenguier, F., Campillo, M., Hadziioannou, C., Shapiro, N. M., Nadeau, R. M., & Larose, E., 2008a. Post-seismic relaxation along the San Andreas fault at Parkfield from continuous seismological observations, *Science*, **321**, 1478–1481.

Brenguier, F., Shapiro, N. M., Campillo, M., Ferrazzini, V., Duputel, Z., Coutant, O., & Nercessian, A., 2008b. Towards forecasting volcanic eruptions using seismic noise, *Nature Geoscience*, **1**, 126–130.

Chollet, F. et al., 2015. Keras, `https://keras.io`.

Clements, T. & Denolle, M. A., 2018. Tracking groundwater levels using the ambient seismic field, *Geophysical Research Letters*, **45**(13), 6459–6465.

De Plaen, R. S. M., Lecocq, T., Caudron, C., Ferrazzini, V., & Francis, O., 2016. Single-station monitoring of volcanoes using seismic ambient noise, *Geophys. Res. Lett.*, **43**, 8511–8518.

Durand, S., Montagner, J. P., Roux, P., Brenguier, F., Nadeau, R. M., & Ricard, Y., 2011. Passive monitoring of anisotropy change associated with the parkfield 2004 earthquake, *Geophys. Res. Lett.*, **38**(13), L13303.

Fichtner, A. & Tsai, V. C., 2019. *Theoretical foundations of noise interferometry*, Cambridge University Press.

Gassenmeier, M., Sens-Schönfelder, C., Eulenfeld, T., Bartsch, M., Victor, P., Tilmann, F., & Korn, M., 2016. Field observations of seismic velocity changes caused by shaking-induced damage and healing due to mesoscopic nonlinearity, *Geophys. J. Int.*, **204**, 1490.

Glorot, X. & Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, vol. 9 of **Proceedings of Machine Learning Research**, pp. 249–256, PMLR, Chia Laguna Resort, Sardinia, Italy.

Goodfellow, I., Bengio, Y., & Courville, A., 2016. *Deep Learning*, MIT Press, `http://www.deeplearningbook.org`.

Grais, E. M. & Plumbley, M. D., 2017. Single channel audio source separation using convolutional de-

noising autoencoders, *CoRR*, **abs/1703.08019**.

Hadziioannou, C., Larose, E., Baig, A., Roux, P., & Campillo, M., 2011. Improving temporal resolution in ambient noise monitoring of seismic wave speed, *J. Geophys. Res. Solid Earth*, **116**(B7).

Hinton, G. E. & Salakhutdinov, R. R., 2006. Reducing the dimensionality of data with neural networks, *Science*, **313**(5786), 504–507.

Hobiger, M., Wegler, U., Shiomi, K., & Nakahara, H., 2012. Coseismic and postseismic elastic wave velocity variations caused by the 2008 iwate-miyagi nairiku earthquake, japan, *J. Geophys. Res. Solid Earth*, **117**, B09313.

Hobiger, M., Wegler, U., Shiomi, K., & Nakahara, H., 2014. Single-station cross-correlation analysis of ambient seismic noise: application to stations in the surroundings of the 2008 iwate-miyagi nairiku earthquake, *Geophys. J. Int.*, **198**, 90.

Hobiger, M., Wegler, U., Shiomi, K., & Nakahara, H., 2016. Coseismic and post-seismic velocity changes detected by passive image interferometry: comparison of one great and five strong earthquakes in japan, *Geophys. J. Int.*, **205**, 1053.

Hunter, J. D., 2007. Matplotlib: A 2d graphics environment, *Computing In Science & Engineering*, **9**(3), 90–95.

Kano, M., Nagao, H., Sakai, S., Nakagawa, S., Mizusako, S., Hori, M., Hirata, N., Shiomi, K., & Honda, R., 2015. Azimuth verification of the meso-net seismographs, *Zisin (Journal of the Seismological Society of Japan. 2nd ser.)*, **68**, 31–44.

Kasahara, K., Sakai, S., Morita, Y., Hirata, N., Tsuruoka, H., Nakagawa, S., Nanjo, K. Z., & Obara, K., 2009. Development of the Metropolitan Seismic Observation network (MeSO-net) for Detection of Mega-thrust beneath Tokyo Metropolitan Area, *Bull. Earthq. Res. Inst.*, **84**, 71–88.

Kingma, D. P. & Ba, J., 2014. Adam: A method for stochastic optimization, *CoRR*, **abs/1412.6980**.

LeCun, Y., Bengio, Y., & Hinton, G., 2015. Deep learning, *Nature*, **521**, 436 EP –.

LeCun, Y. A., Bottou, L., Orr, G. B., & Müller, K.-R., 2012. *Efficient BackProp*, pp. 9–48, Springer Berlin Heidelberg, Berlin, Heidelberg.

Lin, F.-C., Moschetti, M. P., & Ritzwoller, M. H., 2008. Surface wave tomography of the western United States from ambient seismic noise: Rayleigh and Love wave phase velocity maps, *Geophys. J. Int.*, **173**, 281–298.

Lobkis, O. I. & Weaver, R. L., 2003. Coda-wave interferometry in finite solids: Recovery of p-to-s conversion rates in an elastodynamic billiard, *Phys. Rev. Lett.*, **90**, 254302.

Lu, X., Tsao, Y., Matsuda, S., & Hori, C., 2013. Speech enhancement based on deep denoising autoencoder.

Mainsant, G., Larose, E., Brönnimann, C., Jongmans, D., Michoud, C., & Jaboyedoff, M., 2012. Ambient seismic noise monitoring of a clay landslide: Toward failure prediction, *J. Geophys. Res. Earth Surface*,

117, F01030.

Mordret, A., Landès, M., Shapiro, N. M., Singh, S. C., & Roux, P., 2014. Ambient noise surface wave tomography to determine the shallow shear velocity structure at Valhall: depth inversion with a Neighbourhood Algorithm, *Geophys. J. Int.*, **198**, 1514–1525.

Mordret, A., Mikesell, T. D., Harig, C., Lipovsky, B. P., & Prieto, G. A., 2016. Monitoring southwest greenland's ice sheet melt with ambient seismic noise, *Science Advances*, **2**(5).

Moreau, L., Stehly, L., Boué, P., Lu, Y., Larose, E., & Campillo, M., 2017. Improving ambient noise correlation functions with an svd-based wiener filter, *Geophys. J. Int.*, **211**, 418–426.

Nair, V. & Hinton, G. E., 2010. Rectified linear units improve restricted boltzmann machines, in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pp. 807–814, Omnipress, USA.

Sakai, S. & Hirata, N., 2009. Distribution of the Metropolitan Seismic Observation network, *Bull. Earthq. Res. Inst.*, **84**, 57–69.

Sens-Schönfelder, C. & Wegler, U., 2006. Passive image interferometry and seasonal variations of seismic velocities at merapi volcano, indonesia, *Geophys. Res. Lett.*, **33**, L21302.

Shapiro, N. M. & Campillo, M., 2004. Emergence of broadband Rayleigh waves from correlations of the ambient seismic noise, *Geophys. Res. Lett.*, **31**, L07614.

Shapiro, N. M., Campillo, M., Stehly, L., & Ritzwoller, M. H., 2005. High-resolution surface-wave tomography from ambient seismic noise, *Science*, **307**, 1615–1618.

Shapiro, N. M., Ritzwoller, M. H., & Bensen, G. D., 2006. Source location of the 26 sec microseism from cross-correlations of ambient seismic noise, *Geophys. Res. Lett.*, **33**, L18310.

Shen, D., Wu, G., & Suk, H.-I., 2017. Deep learning in medical image analysis, *Annual Review of Biomedical Engineering*, **19**, 221–248, PMID: 28301734.

Stehly, L., Froment, B., Campillo, M., Liu, Q. Y., & Chen, J. H., 2015. Monitoring seismic wave velocity changes associated with the mw 7.9 wenchuan earthquake: increasing the temporal resolution using curvelet filters, *Geophys. J. Int.*, **201**, 1939–1949.

Tsai, V. C., 2009. On establishing the accuracy of noise tomography travel-time measurements in a realistic medium, *Geophys. J. Int.*, **178**, 1555–1564.

Viens, L., Denolle, M. A., Hirata, N., & Nakagawa, S., 2018a. Complex near-surface rheology inferred from the response of greater tokyo to strong ground motions, *J. Geophys. Res. Solid Earth*, **123**, 5710–5729.

Viens, L., Jiang, C., Denolle, M. A., & Hirata, N., 2018b. Long-term seismic monitoring of an urban sedimentary basin, *Abstract S31B-06*, **presented at 2018 Fall Meeting, AGU, Washington DC, 10-14 dec**.

Vincent, P., Larochelle, H., Bengio, Y., & Manzagol, P.-A., 2008. Extracting and composing robust features

with denoising autoencoders, in *Proceedings of the 25th International Conference on Machine Learning*, ICML '08, pp. 1096–1103, ACM, New York, NY, USA.

Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., & Manzagol, P.-A., 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion, *J. Mach. Learn. Res.*, **11**, 3371–3408.

Voisin, C., Guzmán, M., Réfloch, A., Taruselli, M., & Garambois, S., 2017. Groundwater monitoring with passive seismic interferometry, *Journal of Water Resource and Protection*, **9**, 1414–1427.

Weaver, R. L. & Lobkis, O. I., 2001. Ultrasonics without a Source: Thermal Fluctuation Correlations at MHz Frequencies, *Phys. Rev. Lett.*, **87**, 134301.

Xie, J., Xu, L., & Chen, E., 2012. Image denoising and inpainting with deep neural networks, in *Advances in Neural Information Processing Systems 25*, pp. 341–349, eds Pereira, F., Burges, C. J. C., Bottou, L., & Weinberger, K. Q., Curran Associates, Inc.

Yao, H. & Van Der Hilst, R. D., 2009. Analysis of ambient noise energy distribution and phase velocity bias in ambient noise tomography, with application to SE Tibet, *Geophys. J. Int.*, **179**, 1113–1132.

Zeng, X. & Ni, S., 2010. A persistent localized microseismic source near the kyushu island, japan, *Geophys. Res. Lett.*, **37**, L24307.
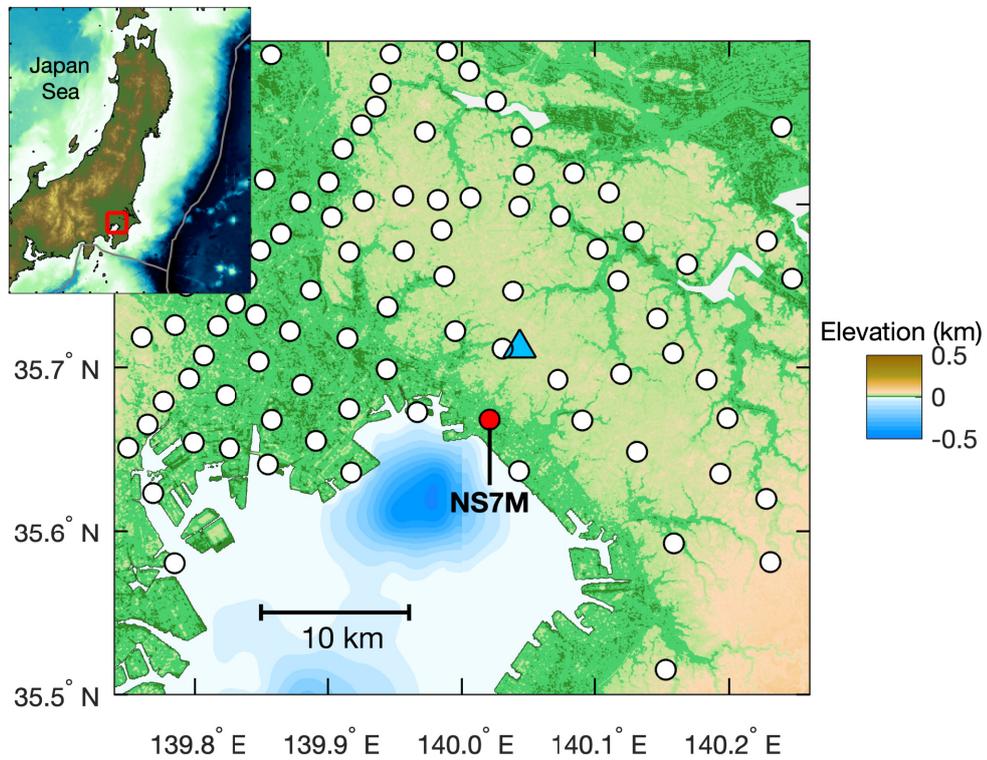
Figure 1. Topographic map of a part of the greater Tokyo area, Japan. The white circles represent the MeSO-net stations and the red circle is the NS7M station which is used in this study. The blue triangle represents the Funabashi weather station operated by JMA. The inset map shows a part of Honshu Island and the red rectangle is the area of interest. Plate boundaries are indicated by the grey lines.
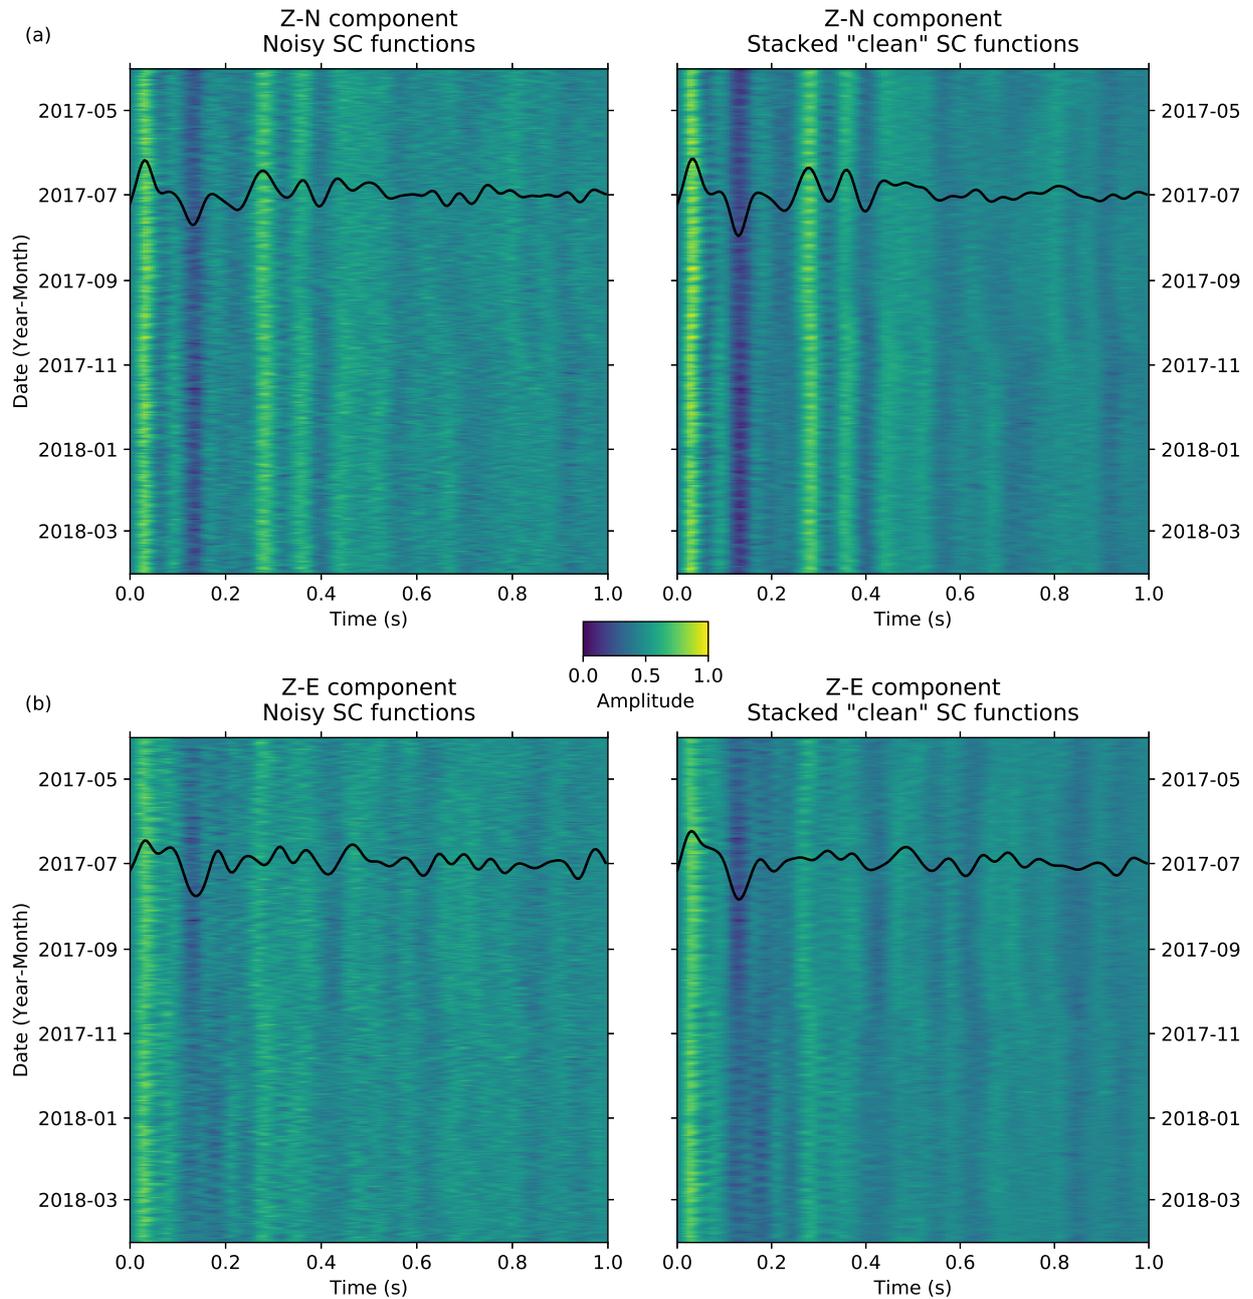
Figure 2. (a) 20-min noisy SC functions for the $Z-N$ component between April 2, 2017 and March 30, 2018 and their corresponding stacked 'clean' SC functions (e.g., SC functions stacked over 2 hours around each 20-min time window). (b) Same as (a) for the $Z-E$ SC functions. For each panel, the first trace on July 1, 2017 is also shown. All the SC functions are bandpass filtered between 1 and 20 Hz and their amplitudes are normalised between 0 and 1. These waveforms are used to train the algorithm.
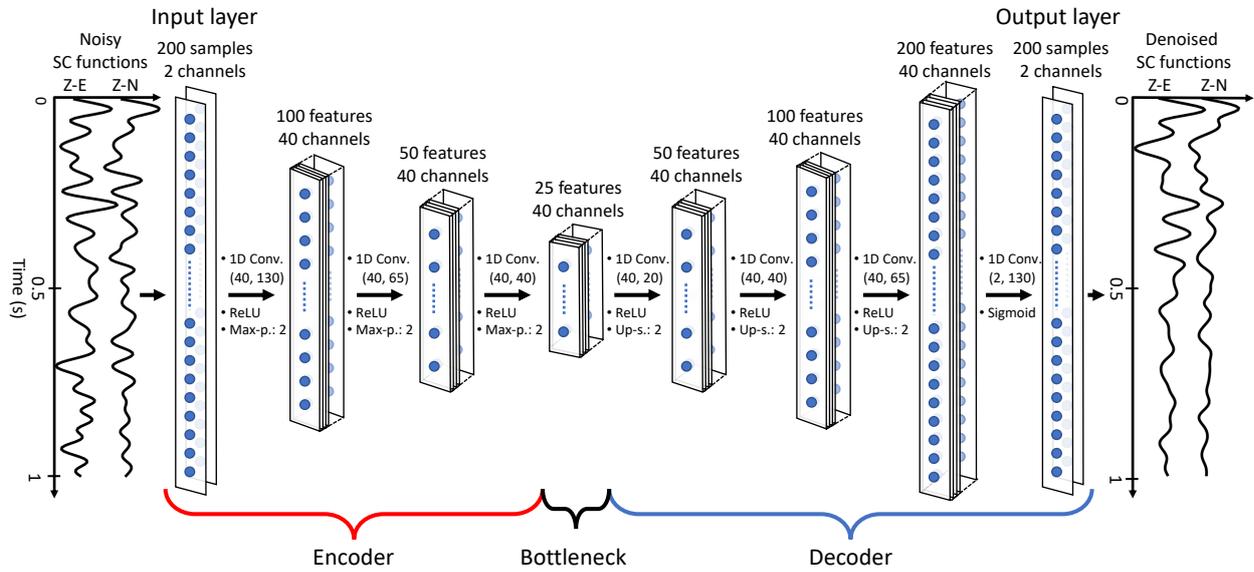
Figure 3. Architecture of ConvDeNoise. The $Z-N$ and $Z-E$ noisy SC functions are fed to the network, which outputs the corresponding denoised SC functions via six hidden layers. The layers of the network are shown by the rectangles together with their dimensions. The number of samples/features is shown on top of each layer and represents the number of points in the vertical direction. The number of channels represents the depth direction of each layer. The operations applied to the data are indicated between the layers. For each 1D convolutional operation (1D Conv.), two numbers are given between parenthesis. The first number is the number of filters and corresponds to the number of sliding windows applied to the data. The second number is the size of the kernel and represents the size of the sliding window that convolves across the data. After each 1D convolution, a non-linear activation function is applied to the data (e.g., ReLU or sigmoid). The max-pooling (Max-p.) and up-sampling (Up-s.) operations are also indicated and the numbers following each operation represent the size of the max-pooling window or the up-sampling factor.
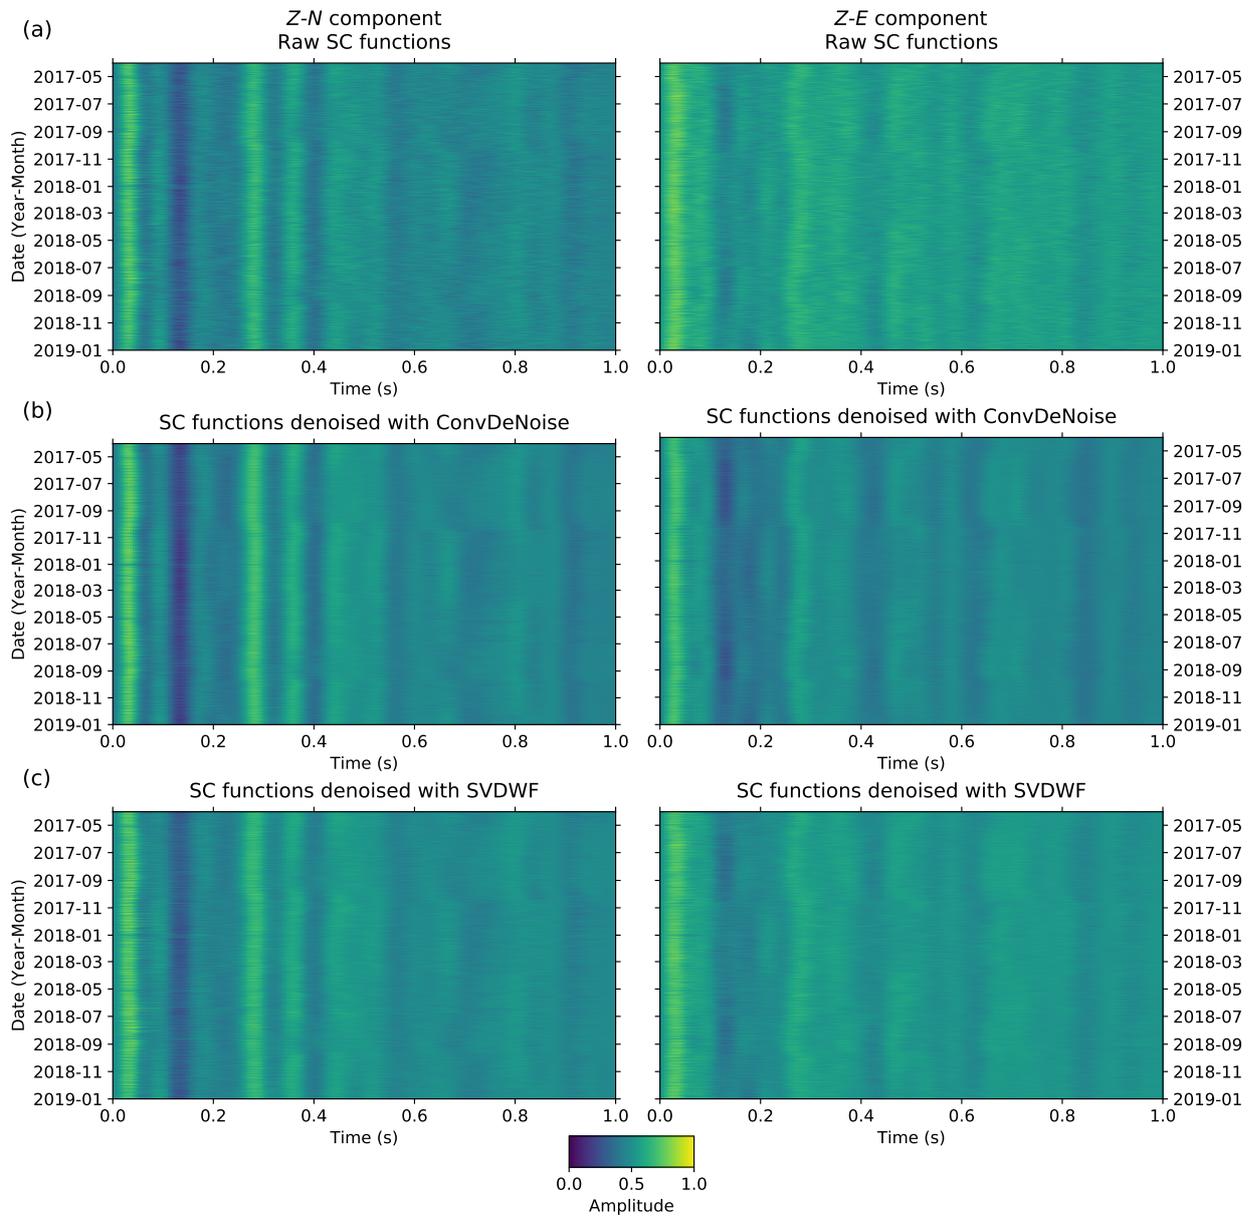
Figure 4. (a) 20-min noisy SC functions at the NS7M station for the $Z–N$ and $Z–E$ components calculated over 20 months between April 1, 2017 and December 31, 2018. (b) 20-min $Z–N$ and $Z–E$ SC functions denoised with ConvDeNoise and (c) denoised with the SVDWF (first 25 singular values, $K = 3$ and $L = 5$) method. All the waveforms are bandpass filtered between 1 and 20 Hz and the amplitudes are normalised between 0 and 1.
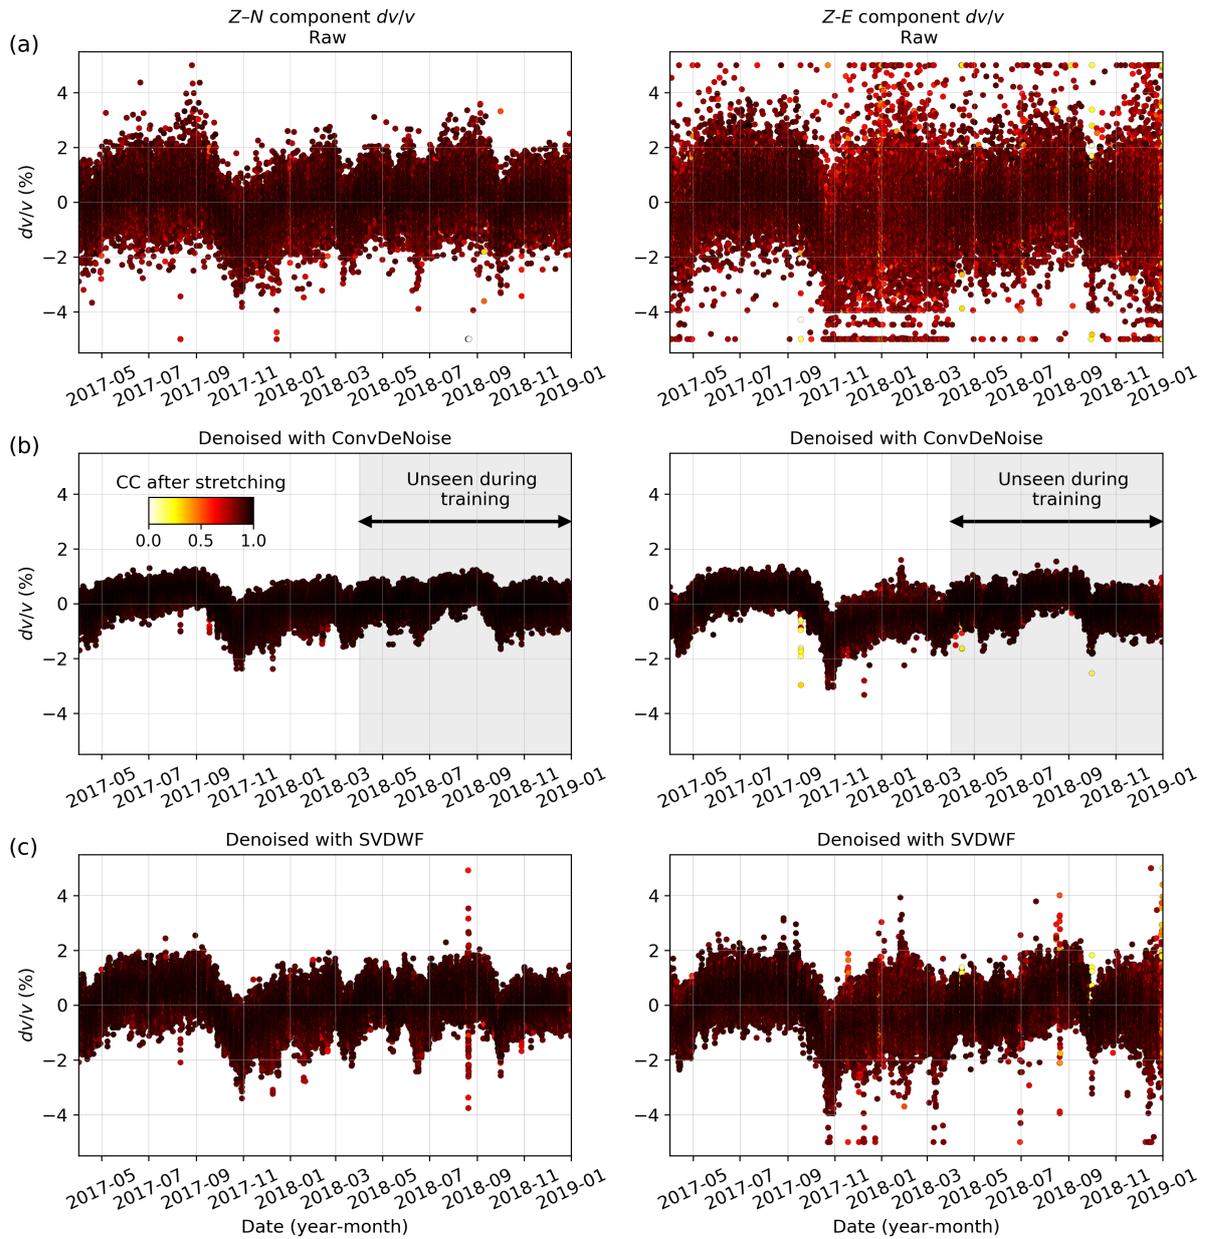
Figure 5. Relative velocity change $dv/v(t)$ computed between April 1, 2017 and December 31, 2018 from the (a) raw, (b) denoised with ConvDeNoise, and (c) denoised with SVDWF ($K = 3$ and $L = 5$; see text) 20-min SC functions for both the Z–N and Z–E components at the NS7M station. The colour of the circles represents the correlation coefficient after stretching. The grey area for both panels in (b) represents dv/v measurements calculated from SC functions that are not used to train ConvDeNoise.
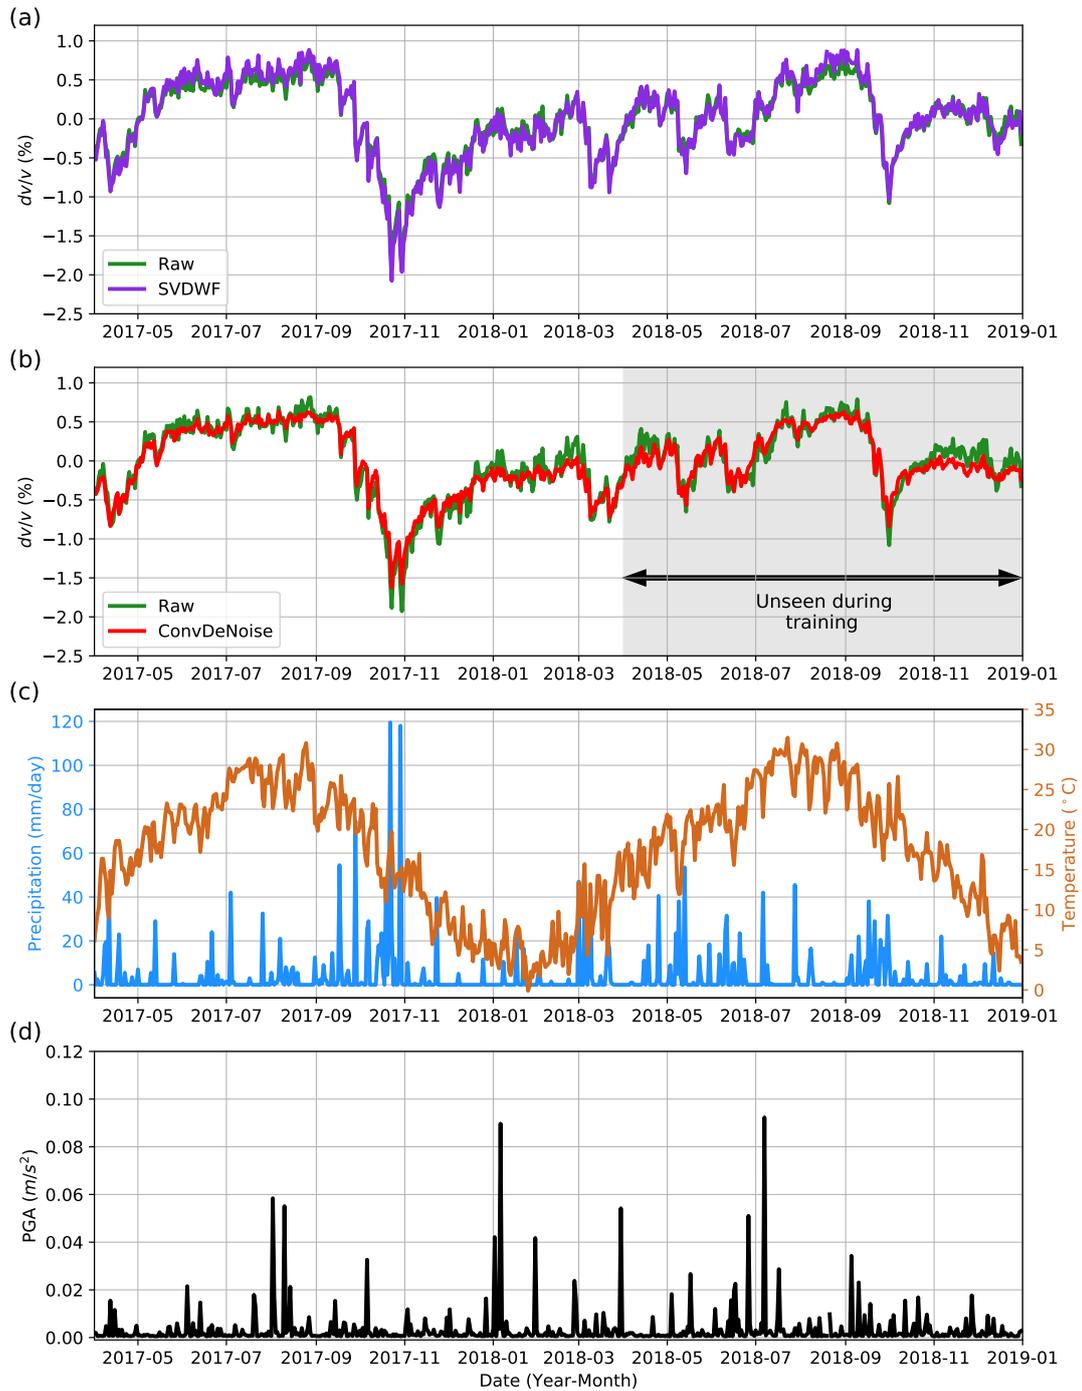
Figure 6. Relative seismic velocity change $dv/v(t)$ stacked over 1 day from April 1, 2017 to December 31, 2018 computed from (a) the raw and denoised with SVDWF ($K = 3$ and $L = 5$) SC functions, and from (b) the raw and denoised with ConvDeNoise SC functions at the NS7M station. The dv/v measurements for the $Z–N$ and $Z–E$ components are averaged using the weighting scheme of Equation 3. dv/v measurements calculated from SC functions unseen during the training phase of ConvDeNoise are highlighted by the grey background in (b). (c) Cumulative daily precipitation (blue) and average temperature (brown) at the Funabashi weather station located 5.25 km away from the NS7M station. (d) Peak ground acceleration (PGA) calculated over 20-minute time windows using the geometric mean of the peak of the two horizontal components.
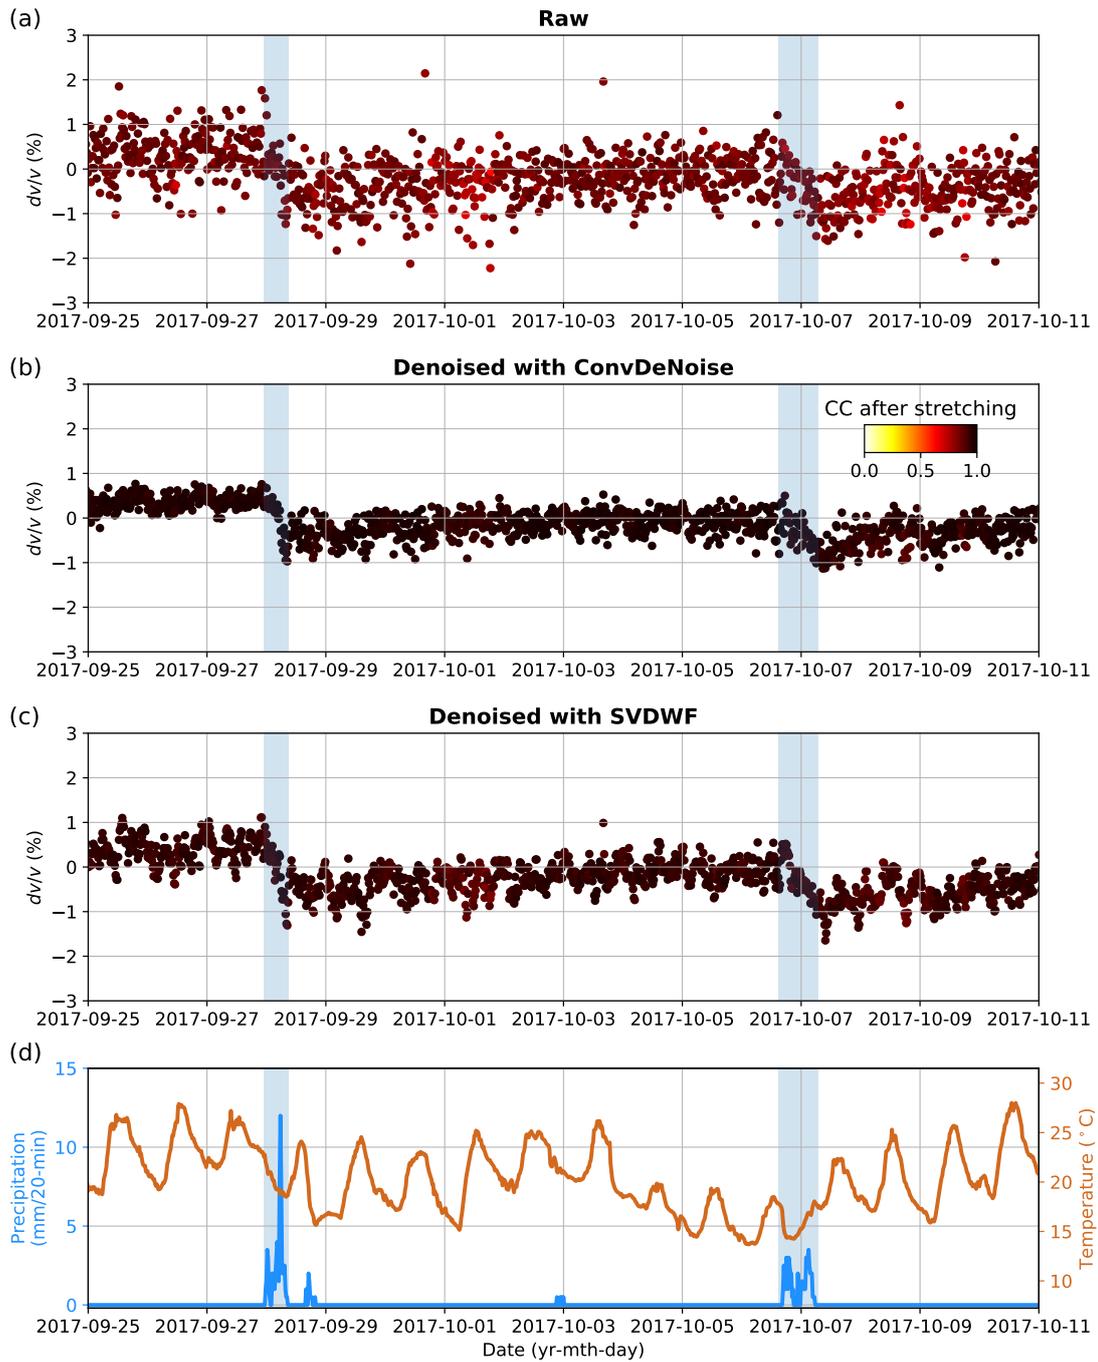
Figure 7. 20-min resolution $dv/v(t)$ calculated from the (a) raw, (b) denoised with ConvDeNoise, and (c) denoised with the SVDWF method ($K = 3$ and $L = 5$; see text) SC functions at the NS7M station. (d) 20-min cumulative precipitation (blue) and 20-min average temperature (brown) at the Funabashi weather station, located 5.25 km away from the NS7M station, from September 25 to October 11, 2017.
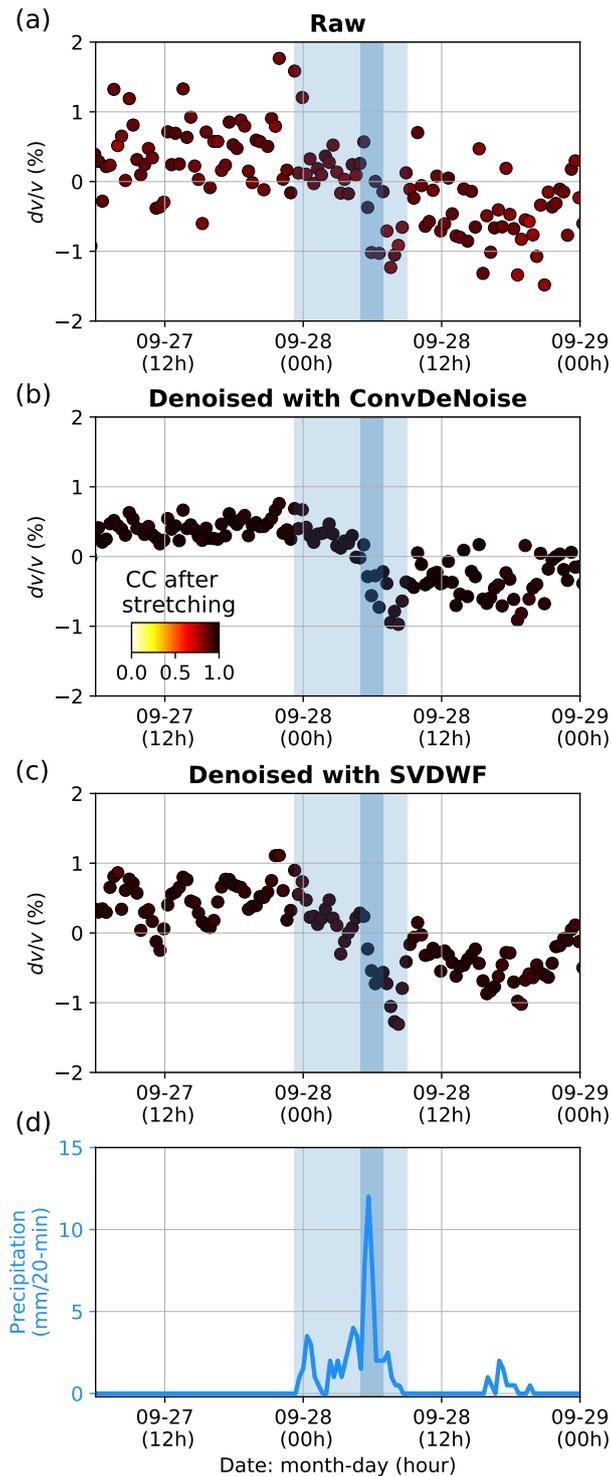
Figure 8. 20-min resolution $dv/v(t)$ calculated from the (a) raw, (b) denoised with ConvDeNoise, and (c) denoised with the SVDWF method ($K = 3$ and $L = 5$; see text) SC functions at the NS7M station. (d) 20-min cumulative precipitation at the Funabashi weather station, located 5.25 km away from the NS7M station, from September 27, 2017 at 6:00 to September 29, 2017 at 00:00. For each subplot, the duration of the precipitation event is highlighted by the light blue area and the heavy rain period is shown by the dark blue area.