

# Natural climate reconstruction in the Norwegian stave churches through time series processing with variational autoencoders

Noemi Manara, Lorenzo Rosset, Francesco Zambelli and  
Andrea Zanola

*Department of Physics and Astronomy, University of Padua, Padua, Italy, and  
America Califano*

*Department of Physics and Astronomy, University of Padua, Padua, Italy and  
Department of Mechanical and Industrial Engineering,  
Norwegian University of Science and Technology, Trondheim, Norway*

## Abstract

**Purpose** – In the field of heritage science, especially applied to buildings and artefacts made by organic hygroscopic materials, analyzing the microclimate has always been of extreme importance. In particular, in many cases, the knowledge of the outdoor/indoor microclimate may support the decision process in conservation and preservation matters of historic buildings. This knowledge is often gained by implementing long and time-consuming monitoring campaigns that allow collecting atmospheric and climatic data.

**Design/methodology/approach** – Sometimes the collected time series may be corrupted, incomplete and/or subjected to the sensors' errors because of the remoteness of the historic building location, the natural aging of the sensor or the lack of a continuous check of the data downloading process. For this reason, in this work, an innovative approach about reconstructing the indoor microclimate into heritage buildings, just knowing the outdoor one, is proposed. This methodology is based on using machine learning tools known as variational auto encoders (VAEs), that are able to reconstruct time series and/or to fill data gaps.

**Findings** – The proposed approach is implemented using data collected in Ringebu Stave Church, a Norwegian medieval wooden heritage building. Reconstructing a realistic time series, for the vast majority of the year period, of the natural internal climate of the Church has been successfully implemented.

**Originality/value** – The novelty of this work is discussed in the framework of the existing literature. The work explores the potentials of machine learning tools compared to traditional ones, providing a method that is able to reliably fill missing data in time series.

**Keywords** Wood, Microclimate, Heritage buildings, Machine learning, Variational auto encoders

**Paper type** Research paper

---

© Noemi Manara, Lorenzo Rosset, Francesco Zambelli, Andrea Zanola and America Califano. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) licence. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this licence may be seen at <http://creativecommons.org/licenses/by/4.0/legalcode>

This work is part of the Symbol Research Project no. 274749 and of the Spara Och Bevara Project no. 50049-1 and has been funded by the Norwegian Research Council and the Swedish Energy Agency, respectively. This work has been possible thanks to the collaboration between the University of Padua (Italy) and the Norwegian University of Science and Technology – NTNU (Norway). The ownership of microclimatic data collected and used in this work is of the Principal Investigator of the Symbol Project. The author(s) give special thanks to Prof. Marco Baiesi and Prof. Chiara Bertolin for their precious guidance and support.



## 1. Introduction

Studying the microclimate of heritage buildings has always been fundamental for assessing the impact that temperature, relative humidity, mixing ratio and other atmospheric variables have on historic materials conservation. In particular, the variations of these atmospheric variables affect organic hygroscopic materials and may trigger deterioration mechanisms, potentially catastrophic (Camuffo, 2019). For this reason, it is mandatory to identify the microclimate conditions that are desirable and suitable for the studied heritage and historical buildings, in order to slow down their decay and deterioration. In this framework, being able to gather as much information as possible on the ambient physical conditions surrounding a given building or building asset may be of extreme help. Generally, the microclimatic situation is depicted by recovering and/or collecting time-series of temperature and relative humidity of the selected environment; however, this often implies to carry out extensive and time-consuming on-the-field monitoring campaigns (Frasca *et al.*, 2017; Olstad *et al.*, 2020; Varas-Muriel *et al.*, 2014). To avoid too long and expensive monitoring campaigns, exploiting tools of Building Simulation (Kramer *et al.*, 2013; de Wit, 2006) could be of help, for instance. These kinds of tools allow to compute and forecast the indoor environmental conditions by simulating the building behavior, subjected to the external climate. In addition, the modern approaches of ML have been recently gaining more and more interest due to their exceptional capability of handling huge datasets, recognizing anomalies and, in some cases, filtering them out (Wen and Keyes, 2019). Among the several ML techniques, the VAEs (Mehta *et al.*, 2019) are appealing as they are able to reconstruct realistic time-series and to reliably fill data gaps. This capability may be groundbreaking in the field of Heritage Science that is often based on studies that either need too many and too expensive experimental observations and, sometimes, rely on corrupted data.

In general, time-series forecasting and reconstruction is a task common to many fields, such as IoT (Han *et al.*, 2021), solar physics (Arslan and Sekertekin, 2019), building energy minimization (Liguori *et al.*, 2021) and cultural heritage (Bertolin *et al.*, 2015). In literature, three main different approaches are proposed: missing data are reconstructed through parametric models based on *a priori* knowledge about the system (American Society of Heating Refrigerating and Air-Conditioning Engineers, 2011; Ciulla *et al.*, 2010; Giaconia and Orioli, 2000), through data-driven approaches (such as deep learning) (Arslan and Sekertekin, 2019; Liguori *et al.*, 2021), or through a mixture of the two (Han *et al.*, 2021).

The first method (Ciulla *et al.*, 2010; Giaconia and Orioli, 2000) is consolidated in the field, but it does not generalize to problems with non-negligible stochasticity (as, for example, the cases in which the impact of human actions is evident and not trivial to include in a deterministic model). Therefore, fitting a polynomial model with such spoiled data would not give good results (Liguori *et al.*, 2021). Even so, the application of pure deep learning could not be advisable as, in case of lack of data, any effective learning may be undermined.

The third approach (Han *et al.*, 2021), through an initial pre-processing and augmentation of the data, makes it possible to build a sufficiently large dataset to feed a machine learning algorithm able to model the system, including its stochastic parts. In this way, good use is made of both clean and spoiled data, which still carry meaningful information and should not thus be simply discarded. Even though the pre-processing phase makes this approach less generalizable to other kind of tasks, it seems to be the only way to tackle a problem with high stochasticity and low data availability. However, in literature the focus is on the forecast of a single measurement, given its history (Arslan and Sekertekin, 2019; Han *et al.*, 2021).

The current work is based on reconstructing the whole natural climate time-series of a wooden heritage building located in Norway: the Ringebu Stave Church. Based on the literature framework, the existing methods seem to be not directly applicable to the current case study as they mainly focus on one-step-ahead predictions (Arslan and Sekertekin, 2019; Han *et al.*, 2021) or short intervals reconstruction (less than 20 h) (Liguori *et al.*, 2021). In the case of this work, instead, the measurements cover one-year period and data gaps that last some days, up to a week, are

present as well. To the authors' knowledge, the task of reconstructing such a wide time interval is not common. For this reason, a different approach from the one recently tested by (Liguori *et al.*, 2021) has been implemented: the choice is to exploit the powerful representation and regularization capabilities of Variational Auto Encoders (VAEs) (Mehta *et al.*, 2019). This has been carried out by analyzing microclimatic data collected in Ringebu Stave Church, a Norwegian medieval wooden building. The study case is presented in Section 2; then, the time series reconstruction has been carried out in Section 3, by using a simple transfer function (Section 3.1) and the VAEs (Section 3.2). The main results are discussed in Section 4 and the main conclusions are highlighted in Section 5.


## 2. Materials

The data used in this work were collected during a monitoring campaign conducted within the framework of the symbol – Sustainable Management of Heritage Building in a long-term perspective project in the Ringebu Church, a stave church located in the region of Innlandet, Norway. Stave churches are medieval Christian churches entirely made of wood, built from the 11th century onward. In particular, Ringebu Church dates back to the 13th century, coming to the present days with only a few modifications. Stave churches are a great heritage for Norwegian culture, and among the over 1,000 churches that were built in the middle age, only 28 still survive. The oldest ones collapsed because the wood was dug directly into the soil, a fact that led to a fast deterioration. In the following centuries new techniques were developed, and the wood was put above a stone base, allowing the remaining stave churches to last in time. However, the risk of losing these masterpieces in recent years is getting stronger and stronger. This is due to the fact that in order to use the churches for celebration, heating is used to create a comfortable internal climate for the churchgoers. This significantly threatens the wood composing the churches architecture which arrived till nowadays through many decades of history and through adaptation to the natural climate of Norway. The work of reconstruction of the natural climate of the church has been performed by using time-series of temperature and relative humidity collected in Ringebu Church using three data loggers installed during a monitoring campaign that started in 2019. Two of these sensors (denoted as DL1 and DL2, with a reading acquisition time of 5 min) are located inside the church, at different positions, whereas the third one (DL3 with a reading acquisition time of 15 min) is placed outside the building, providing external climate measurements. To quantify the similarity between the DL1 and DL2 indoor temperature series, the normalized root mean square error (NRMSE) was used

$$NRMSE = \frac{RMSE}{IQR} \quad (1)$$

where RMSE is the root mean squared error between the two series and IQR is the interquartile range, i.e., the range between the 0.25 and 0.75 quantiles. Since an NRMSE of 0.087 was obtained, proving that DL1 and DL2 time series are very similar to each other, it is reasonable to consider only DL1 as representative of the indoor measures and to compare it with the DL3 (outdoor). The data loggers were installed in March 2019 and, since then, they have been collecting microclimatic data. The current work is focused on a time window of one calendar year, from March 2019 to March 2020. Figure 1 shows the temperature and relative humidity profiles collected from DL1 and DL3 over the observed period. In detail, Figure 1a represents the indoor (red curve) and the outdoor (blue curve) temperature, while Figure 1b represents the indoor (red curve) and the outdoor (blue curve) relative humidity.

As easily noticeable from Figure 1, during the considered year there were sporadic peaks of the internal temperature that lay beyond the natural (external) variability, particularly during the cold season. These events are classified as *anomalous*, as they introduce perturbations in the microclimatic time-series, and they are caused by the artificial heating



**Figure 1.** Temperature (a) and Relative Humidity (b) collected from March 2019 to March 2020 in Ringebu stave Church, Norway

**Note(s):** The red curves represent the indoor environmental variables, while the blue curves represent the outdoor ones. The inset in Figure 1a highlights the long period during which the heating system was kept on

that is activated in view of celebrations. In particular, there was a week in which the heating was kept on without interruptions, causing the huge *plateau* that can be spotted in the inset of Figure 1a. Besides, additional data (*metadata*) carrying information about the use of artificial heating are available as well, allowing to precisely identify anomalous (i.e., artificial) events.

### 3. Methods

Clearly, using a machine learning-based strategy for reconstructing the natural climate of the Church by having at disposal a unique and flawed time series might be tricky and may lead to unreliable results, if a proper training phase is not carried out. A possible solution could be based on the assumption that the internal temperature of the building is easily obtained from the external one (Giaconia and Orioli, 2000) by means of a deterministic transfer function plus a sub-leading noise contribution. In other words, the building's indoor temperature is inevitably dependent on the outdoor one and it is influenced by the presence of the building envelope acting as a buffer between the outdoor and indoor environments. Indeed, high-resolution historical external temperature and relative humidity time series are provided by the weather forecast station E6 Frya, which is located 5 km away from the Ringebu Church. In particular, it was possible to gather data collected from the weather forecast station in the period September 2014–April 2020, although with some missing data.

At this point, the temperature time-series recorded by the station was compared to the one collected by DL3. An NRMSE of 0.21 was obtained, confirming that the two series are indeed very similar to one another. Hence, the idea is to collect the available data from the station and to apply a properly calibrated transfer function that returns a synthetic internal climate of the church, as it would have been without any human intervention. Eventually, the weather


station's time series, along with the reconstructed one, can be used as a *simulated* DL3 and DL1 series in order to train the VAE. From now on, this dataset will be referred to as "station dataset".

The same procedure was implemented for the relative humidity as well, because adding this information can be reasonably beneficial for the VAE, in order to recognize possible connections between temperature and relative humidity time series.

### 3.1 A transfer function for simulating the indoor environment

As previously mentioned, a transfer function applied to the external time series is able to reconstruct the internal one. This idea is based on the fact that there exist relations, in terms of amplitude and time delay, between the temperature outside a building and that inside (Bertolin *et al.*, 2015; Camuffo *et al.*, 2014). Moreover, this concept of thermal phase shift is supported also by data, as visible in Figure 2. In Figure 2 the measured (red curve) and the rebuilt (blue curve) temperature profiles of DL1 (excluding the trendline and, so, the average behavior in order to better see the fluctuations) are shown, together with the actual indoor (red dots) and outdoor (blue dots) measured data and the datalogger error bar (red dashed curve). It can be easily noticed that the maxima of the internal temperature suddenly follow that of the external temperature. This phenomenon is obviously a consequence of the fact that the outdoor environment drives the temperature inside the building. As a first attempt, a simple attenuation and time shifting mechanism is applied. Applying the corresponding transformations to the external temperature time series, a time series for the indoor temperature, approximately equivalent to the one recorded by DL1, can be obtained. Conceptually, this is not completely correct, because this simple assumption would mean that the transfer function of the walls has a constant response for all frequencies, which is far from being realistic. This reasoning leads to the idea that this simple mechanism could be extended specifically to all the frequencies that compose the spectrum of the signal.

A possible issue consists in guessing the correct transfer function of the building walls,  $G(s)$ . Potentially, any  $G(s)$  can be modeled as an infinite combination of single zero-pole



**Figure 2.**  
Example of DL1 reconstruction from DL3 (blue continuous line) in an example from 20 to 23 March 2020

**Note(s):** Internal measured temperature (red dots), external measured temperature (blue dots), internal measured temperature without trendline (red curve), reconstructed internal temperature from the external one without trendline (blue curve) and sensor error bar (red dashed curve)

functions with arbitrarily close zeros and poles; the easiest thing, however, is to consider just one zero-pole and check the results. The following transfer function has been considered:

$$G(s) = K_{stat} \frac{1 + s\tau_{zero}}{1 + s\tau_{pole}} \quad (2)$$

where  $K_{stat}$  is a static gain,  $\tau_{zero}$  and  $\tau_{pole}$  are the zeros and the poles of  $G(s)$  respectively, and  $s$  is the independent variable in the Laplace domain. With MATLAB Simulink® a grid search approach has been implemented in order to find the best value of  $\tau_{zero}$  and  $\tau_{pole}$ , where the grid search is based on an Level 2 (L2)-loss minimization (Goodfellow *et al.*, 2016) between the reconstructed “synthetic” time series and the measured one in the spring-summer period, namely when the heating was turned off. An eventual static gain  $K_{stat}$  has been added, whose best value, however, turns out to be near to one ( $1.0810 \pm 0.005$ ). The main drawback in this procedure is the temperature offset resulting between the average year temperature and the average warm season’s temperature. However, this offset is harmless since the dataset has to be standardized before being fed to the VAE. The Bode diagram of the transfer function shown in Eq. (2) is represented in Figure 3 showing the magnitude (Figure 3a), the phase (Figure 3b) and the temporal shift (Figure 3c).


Every sinusoidal component of the signal has an argument

$$\omega t + \varphi = \omega(t + t_{shift}) \quad (3)$$

where  $t_{shift}$  was introduced in order to recover temporal information from the phase diagram. A property of the zero-pole structure is that it gives a finite attenuation to high frequencies, and this is desirable because the high-frequency fluctuations need to be attenuated or neglected as they are too quick to be able to propagate through the walls of the Church. It is also true that the procedure of minimizing the L2 loss brings the risk of overfitting the data. However, the chosen structure of  $G(s)$  is sufficiently simple to reasonably generalize the result, acting as a regularizer. The consequence is that the values of  $\tau_{zero}$  ( $6.00 \pm 0.05$  h) and  $\tau_{pole}$  ( $27.25 \pm 0.05$  h) are chosen in such a way that they will automatically dampen the contribution of the high frequencies, as can be seen in Figure 3a.

### 3.2 Variational auto encoders (VAEs) for reconstructing the indoor environment

An innovative approach for reconstructing the indoor environment may be seen in exploiting the capabilities of VAEs. In order to train this model, the whole dataset has been sliced into samples representing a short time-window of one-week. The time span of one week has been chosen to allow the VAE to recognize short and intermediate-range correlations between data points, such as the daily periodicity. Nearby samples are intentionally correlated, as the series has been sliced by using a one-day-long stride. This has been based on the heuristic assumption that, during the training, this strategy would enforce the VAE to learn temporal correlations between subsequent slices. Another good reason for this choice is that, when it comes to reconstructing the complete year time series, adjacent samples would tend to display abrupt discontinuity between one another. Thus, overlapping the samples allow to carry out averaging operations. Finally, slicing the dataset in this way allows having a larger training set, hence improving the training phase. Each sample embodies eight different features: the internal and external temperature and relative humidity time series, and four additional features to encode the temporal information, representing the sine and the cosine of the daily and yearly periodicity, respectively. To fill the windows with missing data, the extremal points have been interpolated with a straight line. Finally, before being fed to the VAE, samples have been standardized in order to prevent gradients from exploding during the training.

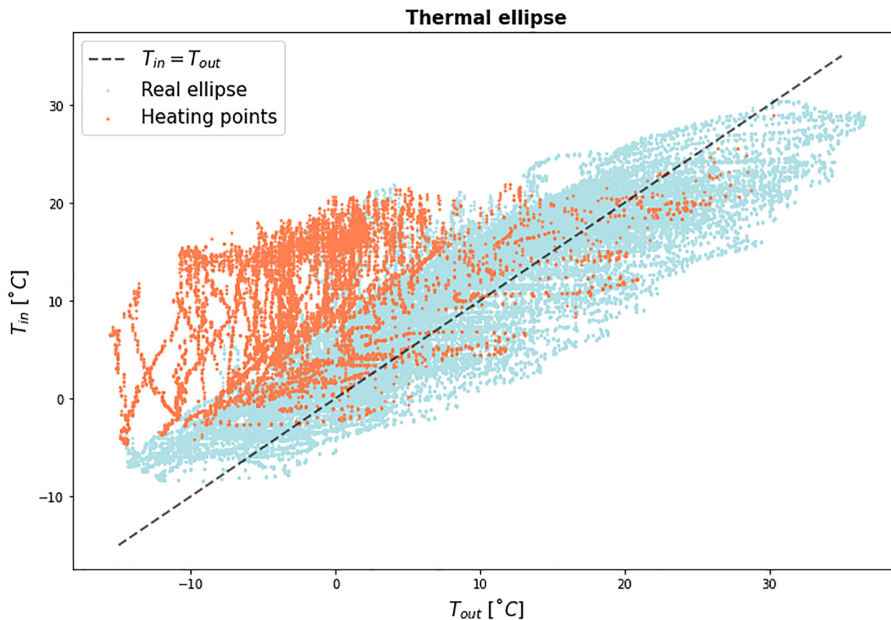


**Figure 3.**  
Bode diagrams for the tuned transfer function  $G(s)$ : magnitude (a), phase (b) and temporal shift (c)

The different tested architectures have been trained over the whole station dataset, without using any test set (more details in Section 4); in addition, the reconstruction of each of those has been assessed either qualitatively, comparing the result with the Ringebu series, and with the support of a properly defined score. The idea is that, on one hand, a reliable reconstruction of the original time series is desirable, and on the other hand, it is also important to quantify the ability of the VAE to cut the heating peaks out. In principle, a possible measure of how well the original time series has been reconstructed, regardless of the presence of heating peaks in the reconstruction, could have been the NRMSE between the two series. In practice, though, it turns out that, for this particular task, a very regularized (i.e. less detailed) reconstruction typically outperforms the NRMSE of a more detailed one. This kind of assessment, namely whether the reconstruction is too mild or too detailed has been carried on in a graphical way. The score that has been implemented, instead, accounts only for the fact that it is not desirable to reconstruct anomalous peaks. Plotting the internal temperature versus the external one (Figure 4), it can be seen that data points (light blue dots in Figure 4) are distributed following what is called the *thermal ellipse* (TE) (Camuffo, 2019), which essentially represents the thermal inertia due to the walls of the Church. The key observation here is that most of the points that are known from metadata to correspond to the heating (orange dots in Figure 4) fall outside this region. Hence, the idea is to fit the uncorrupted data with a multinomial Gaussian ( $N$ ), determining the mean and standard deviation ( $\mu_{TE}$ ,  $\Sigma_{TE}$ ) and the reconstruction score  $J_{TE}(X)$  as

$$J_{TE}(X) = -\langle \log(N(x|\mu_{TE}, \Sigma_{TE})) \rangle - J_{TE}^* \quad (4)$$

where the average is computed over the whole dataset and  $J_{TE}^*$  is the score of the Ringebu series without the heating events. Having subtracted the baseline corresponding to the



**Note(s):** Red points correspond to the events that are recorded when the heating system was turned on, while the cyan points to the natural climate

**Figure 4.**  
Thermal ellipse of the church's dataset




cleaned original data, the resulting score will end up being either positive or negative. In the first case, the thermal ellipse of the reconstruction is broader than the ideal one, and this can be a signal that the heating peaks have not been properly eliminated. In the second case, instead, the thermal ellipse of the reconstruction is narrower than the original one. This is also not advisable, because it means that the natural thermal inertia of the building has not been properly reconstructed. In conclusion, a good TE score is the one that has a small absolute value.

Having defined the TE score, it is possible to go on with the definition of the VAEs. A VAE (Goodfellow *et al.*, 2016; Mehta *et al.*, 2019) is an unsupervised machine learning model that is able to learn and approximate the probability distribution of training data, allowing then the sampling of newly generated data. The architecture is made of three main components: the encoder, the bottleneck and the decoder. The encoder is a neural network in charge of compressing the input data into a low-dimensional representation. In the present setting, this representation consists of the mean and variance vectors of a symmetric multinomial Gaussian, which represents the probability distribution of the latent variables,  $z$ . The dimensionality of the latent space is a hyper parameter to be tuned, and it corresponds to the amount of information that has to be retained from the input data. In the bottleneck, it is possible to sample latent variables and feed them into the decoder neural network, which produces in output a sample resembling the original data.

In principle, after the architecture has been trained, by randomly sampling latent variables in the latent space it is possible to generate completely new data. In this work, though, it was important to sample the latent variables precisely and sequentially in order to reconstruct the temperature time series over the whole year. For this reason, in the reconstruction phase, the original time series has been fed into the VAE's encoder, thus producing the correct sequence of latent variables to get a filtered version of the original signal in the output.

Among all the tried architectures, the final selected model is described as follows and shown in Figure 5.

The encoder of the VAE is made of two sections. The first section alternates 1D convolutional layers with a pooling layer and a dropout layer for regularization. This is meant



**Figure 5.**  
Architecture of  
the VAE

**Note(s):** Each layer type is represented with a unique color. Their activation functions are in the right side of the box: in orange the ReLu and in dark blue the sigmoid. At the end of the first conv block, Max pooling (red) is performed. After conv2 and ConvT blocks, dropout (purple) is inserted

to detect some local patterns of the series and to select higher-level features to be fed into the second section, represented by a Long-Short-Term-Memory (LSTM) layer (Goodfellow *et al.*, 2016). This layer should be able to recognize both short and long-range temporal correlations between data points. The output of the encoder is the couple of parameters  $(\mu, \Sigma)$  that allows sampling the latent variables through the reparameterization trick. In particular, it has been found that a 17-dimensional latent space was able to provide a good reconstruction. The decoder's design is simpler, and it consists of a dense layer followed by a 1D convolutional transposed layer and a dropout. In the following, this architecture will be referred to as "VAE 1".

Another relevant architecture that has been tested consists of two VAEs concatenated. The first one is made of two LSTM layers, one in the encoder neural network and the other in the decoder neural network. This architecture proved to be very efficient in understanding the yearly periodicity of the data, although the fine structure of the time series was completely eliminated. For this reason, this architecture was trained on the original Ringebu's time series in order to learn the temporal pattern of the data, and its output has been fed into the second VAE, trained on the station dataset. In the following, this architecture is denoted as "VAE 2", and the labels "VAE 2a" and "VAE 2 b" are used to indicate the two components of the architecture.

The supervised training has been performed through the minimization of the (negative) Evidence Lower Bound (ELBO) (Goodfellow *et al.*, 2016):

$$-ELBO = L2(x_{input}, x_{rec}) + \beta \cdot KL(q(z|x) || p(z)) \quad (5)$$

where  $L2$  is the L2-loss between the input ( $x_{input}$ ) time-series and its reconstruction ( $x_{rec}$ ),  $KL$  is the Kullback Leibler divergence between the posterior distribution of the latent variables,  $q(z|x)$ , and the prior standardized normal distribution  $p(z)$ , and  $\beta = 270$  is a regularization coefficient, fine-tuned for the two contributions to both be effective.

#### 4. Results and discussion

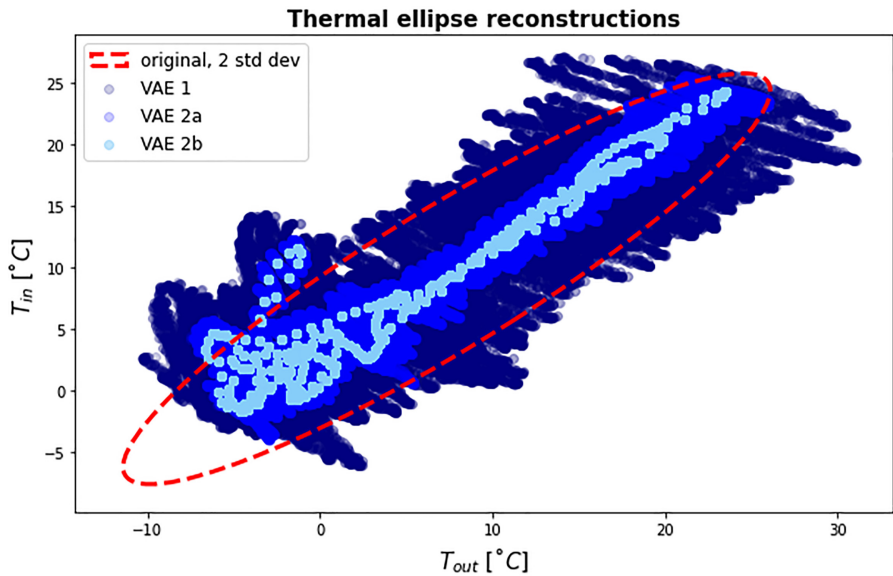
After 50 epochs (i.e., the times that the learning algorithm will work through the entire training dataset) of training, a stabilization of the training error was noticed, thus the training phase was stopped in order to move to the performance testing phase.

The architecture VAE 1 had a TE score of 0.28, whereas VAE 2 scored  $-0.32$ . As discussed above, the interpretation is that the thermal ellipse reproduced by VAE 1 is broader than the original one, whereas VAE 2 tends to underestimate the thermal inertia of the building by generating a narrow thermal ellipse. Figure 6 shows the thermal ellipses of VAE 1 and the two stages of VAE 2 compared with the fitted original one.

As for the single sample processing, shown in Figure 7, the daily periodicity is present in the reconstruction (red curve) of each feature (internal humidity, Figure 7a, internal temperature, Figure 7b, external humidity, Figure 7c, and external temperature, Figure 7d) and the signals (blue dashed curve) of both humidity and temperature are filtered, resulting in more regular curves that still have the local unique variations typical of a natural climate. In Figure 7 it can be noticed how the reconstructed signals tend to be more regularized than the original ones, avoiding abrupt deviations from the average pattern. This is because, ideally, the reconstruction has to be interpreted as the average of many realizations of a stochastic process.

Moving to the whole year reconstruction, it can be subdivided into two sections: cold seasons (autumn and winter) and warm seasons (spring and summer).

The months of cold seasons are affected by the use of artificial heating, which corrupts the original signal. By feeding those samples to the trained VAE, as shown in Figure 8 (Goodfellow *et al.*, 2016), it recognizes that the heating peaks are anomalous and thus filters



**Figure 6.**  
Thermal ellipses  
generated by the  
different VAE  
architectures

**Note(s):** VAE 1 (purple dots), VAE 2a (blue dots) and VAE 2b (light blue dots). The TE extrapolated by the data is shown with the red dashed curve

them, retrieving a time series that shows a strong and clean daily periodicity, resembling the one of a natural climate.

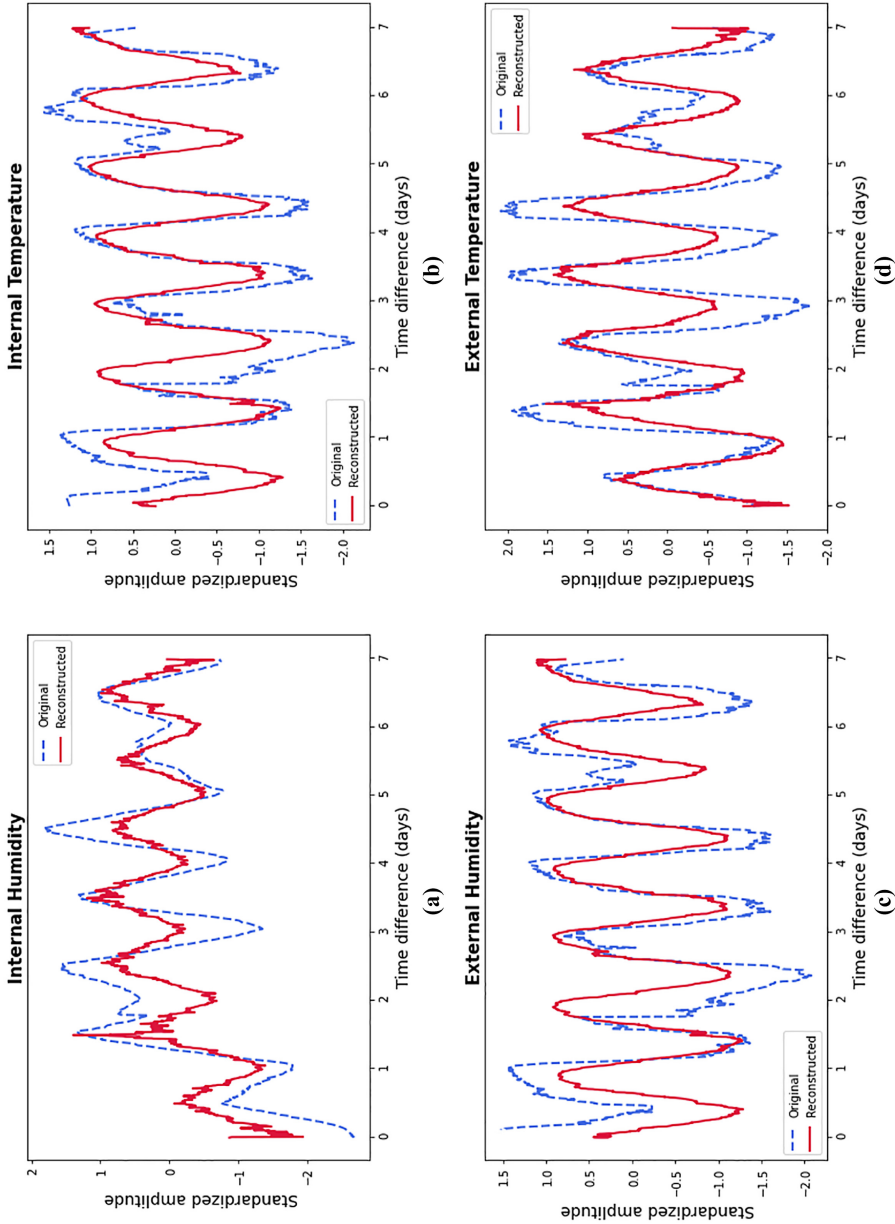
Instead, warm seasons are less affected by artificial heating, but still some effects due to the presence of visitors within the church are present. As shown in [Figure 8](#), when feeding the VAE with these kinds of time series, the natural behavior is respected.

The auto encoder is also able to reconstruct the behavior of the temperature when some data are missing (or removed because too corrupted): up to a week of observations can be artificially restored by feeding the VAE with the values corresponding to a linear interpolation of the extremes of the missing interval, as can be seen in [Figure 9](#).

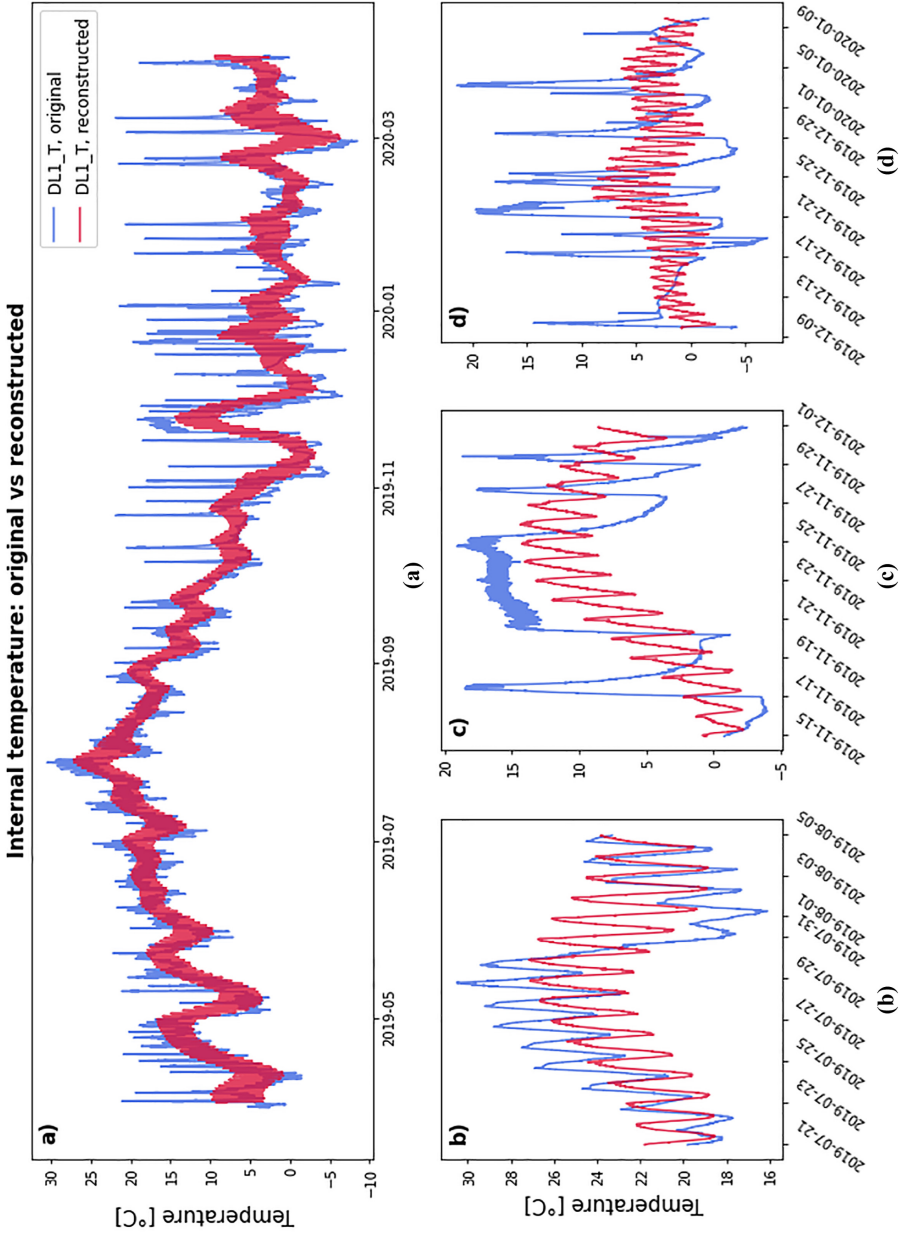
Longer missing intervals should be tackled with a different approach being the climate variations substantial and a simple linear interpolation not sufficient for a proper reconstruction.

Finally, a simple assessment of the latent space may be carried out by plotting the scatterplots of the training samples encodings in [Figure 10](#). The most important aspect is the presence of seasonal clusters, as it is an index of the correct encoding of the yearly periodicity.

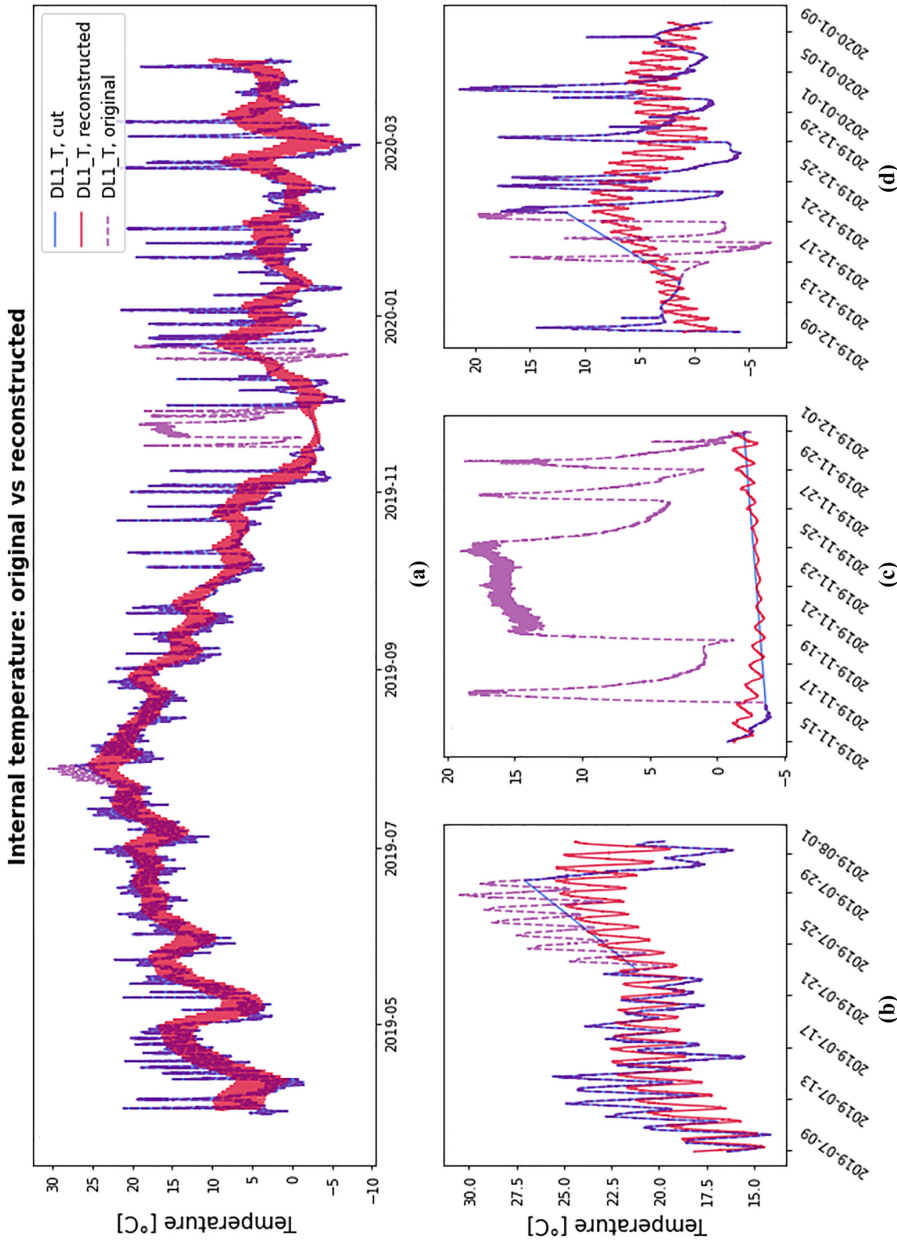
An important remark has to be done concerning the training procedure. Typically, the machine learning way to proceed is to update the weights of the model based on the training samples' reconstruction and to monitor overfitting over a test set. However, the setup used in the current work is quite different. Indeed, there is just a 4–5 years series, i.e., repeated measurements of each time window. This is a very limiting factor, because it means no significant pieces of the data may be extracted for comparing the reconstruction of a training sample with the corresponding sample of a test set. Another option could be to compare the reconstruction of a sample with a nearby one used as a test sample, under the assumption that they are similar. In practice, however, this last assumption is quite arguable, let alone the fact that this would prevent the use of a conspicuous part of the available data. Moreover, the goal is to start from the original, spoiled, time series and to generate another series similar to the previous one but without any anomalous peak. The goodness of the reconstruction can be



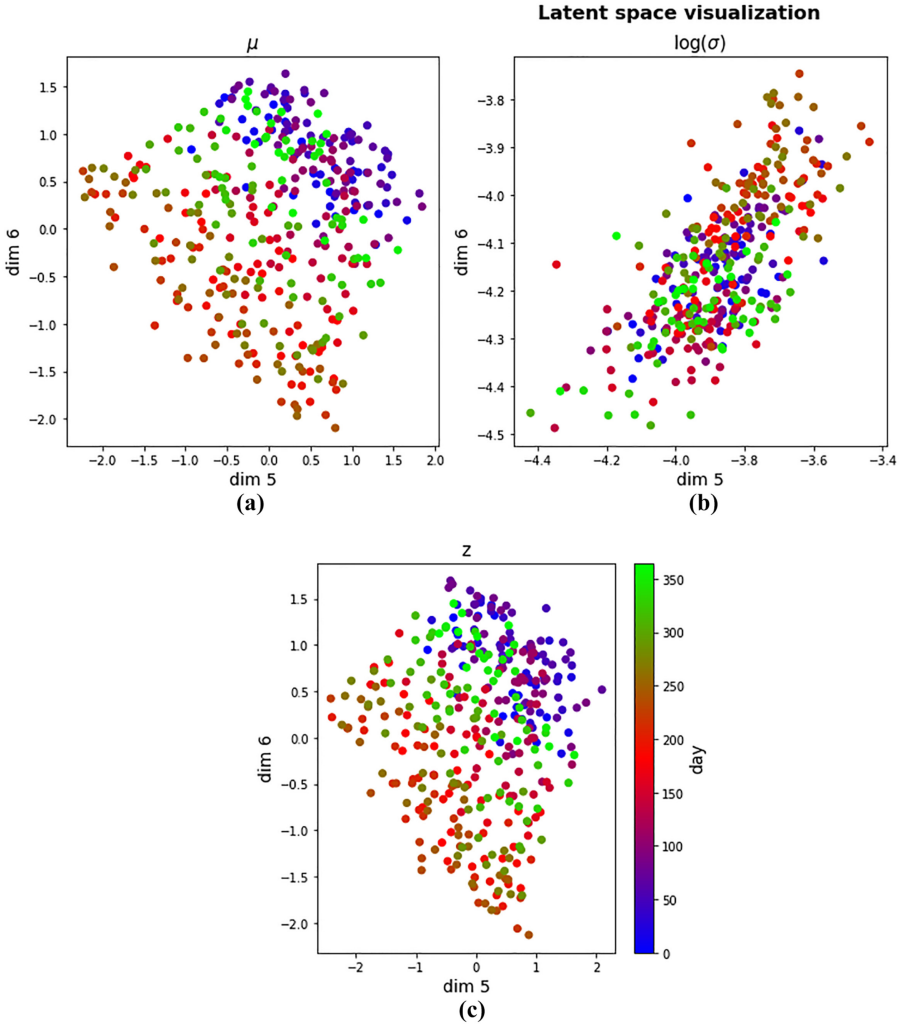
**Figure 7.** Example of the reconstructed time series by VAE 1 (red curves) superimposed to the original signal (dashed blue curves) for: (a) internal humidity, (b) internal temperature, (c) external humidity and (d) external temperature



**Figure 8.** (a) Reconstruction of the entire internal temperature time series of Ringebu Church by VAE1 (red line) and original dataset (blue line) with zooms on: summer time interval (b), artificially heated time interval (c), winter time interval with a high level of corruption caused by artificial heating and churchgoers



**Figure 9.** (a) Reconstruction of the entire internal temperature time series of Ringebu Church by VAE1 (red line), original dataset (purple dashed line) and original dataset with linear interpolation between cut intervals, which correspond to the actual input of the VAE (blue line). (b–d) Zoom of year periods with different missing data interval length



**Figure 10.** Scatterplots of 2 dimensions (13 vs 9) of: (a) means and logarithm of the standard deviations (b) of the multivariate Gaussian generated by every data given in input to the VAE; (c) the latent space variable  $z$  extracted from the multivariate Gaussians (c)

therefore assessed by the capability of the VAE to either generate a resembling time series *and* to cut off the anomalies. These arguments drove the decision to use the entire dataset of the station as a training set, and to look at the reconstruction obtained by feeding the VAE with the original Ringebu series. Then, relying on the score defined in Section 3.2 and also on a qualitative assessment, the results have been evaluated.

Moreover, the approach followed using the station dataset for the training phase is the best possible to the authors' knowledge, but a note of caution has to be taken since it is just a *plausible* reconstruction of the possible series that would have been expected in case many years of records had been available.

With respect to other classical VAE architectures, the one presented in this work outperforms the others mainly thanks to two key aspects. First, the introduction of the LSTM layer, which allows the encoder to detect important features of the time series and to keep memory of them for a long period of time. In addition, it contains a larger number of trainable

---

parameters and makes the training a little slower, but its contribution to the final results is fundamental. The second beneficial aspect is the implementation of convolutional layers with kernel sizes equal to the daily and half-day periodicities, which were evident in the dataset.

Finally, the results have shown that the VAE is able to reconstruct the time series even in the absence of some fractions of input data.

## 5. Conclusions

In this work an attempt to implement an unsupervised machine learning method for the problem of reconstructing the internal climate time series of a heritage building, a stave church in this case, has been proposed. As previously underlined, there are some criticalities within this approach, which are due to the poor amount of available data.

Nevertheless, it has been successfully managed to reconstruct, for the vast majority of the year period, a realistic time series of the natural internal climate of the church, providing a method that is able to reliably fill missing data of many hours' time span, up to some days.

There is surely room for improvements and further studies. For example, the current work has been focused on denoising the original time series; for this kind of task, possibly even a simpler auto encoder architecture could have worked as well (although a VAE is typically more versatile). In other words, the generating properties of the VAE have not been exploited thoroughly, because the latent space of such a complex dataset is very difficult to explore, and it is hard even to define what is hoped to be thrown out from it. However, future perspectives could deal with:

- (1) Finding a way to understand the geometry of the latent space, to generate completely new samples;
- (2) Studying and implementing conditional variational auto encoders, which allow generating a complete time series from the input of just some control parameters (in this case the sine and cosine of day and year, for instance). Using them, it could be even possible to generate new data by just specifying the needed time window, provided that the architecture is able to learn the seasonality of one sample.

This way, optimized tools would be available and could be used in several different fields. For example, it can be used to forecast financial time series and business metrics such as stock market price, sales. It could be useful for filling gaps in data from obsolete weather stations and, in general, it could be implemented in any signal processing analysis and control system. In conclusion, these approaches would be of extreme usefulness in studying and predicting the indoor microclimate of historical buildings; as a matter of fact, they would lead towards smarter choices for the management of the indoor microclimate, in the framework of preservation and conservation of the historical and cultural heritage, by limiting, when possible, time and money consuming on-the-field monitoring campaigns.

## References

- American Society of Heating Refrigerating and Air-Conditioning Engineers (2011), "ASHRAE handbook - HVAC applications, www.Ansi.Org American society of heating, refrigerating and air-conditioning Engineers, Inc.", available at: <http://www.ashrae.org>.
- Arslan, N. and Sekertekin, A. (2019), "Application of long short-term memory neural network model for the reconstruction of MODIS land surface temperature images", *Journal of Atmospheric and Solar-Terrestrial Physics*, Elsevier, Vol. 194 June, p. 105100.
- Bertolin, C., Camuffo, D. and Bighignoli, I. (2015), "Past reconstruction and future forecast of domains of indoor relative humidity fluctuations calculated according to EN 15757:2010", *Energy and Buildings*, Elsevier B.V., Vol. 102, pp. 197-206.



- Camuffo, D. (2019), *Microclimate for Cultural Heritage*, Elsevier Science.
- Camuffo, D., Bertolin, C., Bonazzi, A., Campana, F. and Merlo, C. (2014), "Past, present and future effects of climate change on a wooden inlay bookcase cabinet: a new methodology inspired by the novel European Standard EN 15757:2010", *Journal of Cultural Heritage*, Elsevier Masson SAS, Vol. 15 No. 1, pp. 26-35.
- Ciulla, G., Lo Brano, V. and Orioli, A. (2010), "A criterion for the assessment of the reliability of ASHRAE conduction transfer function coefficients", *Energy and Buildings*, Vol. 42 No. 9, pp. 1426-1436.
- de Wit, M. (2006), "HAMBBase: heat, Air and moisture model for building and systems evaluation", *Bouwstenen*, Vol. 100, ISBN: 90-6814-601-7.
- Frasca, F., Siani, A.M., Casale, G.R., Pedone, M., Bratasz Strojcki, M. and Mleczkowska, A. (2017), "Assessment of indoor climate of Mogiła Abbey in Kraków (Poland) and the application of the analogues method to predict microclimate indoor conditions", *Environmental Science and Pollution Research*, Vol. 24 No. 16, pp. 13895-13907.
- Giaconia, C. and Orioli, A. (2000), "On the reliability of ASHRAE conduction transfer function coefficients of walls", *Applied Thermal Engineering*, Vol. 20 No. 1, pp. 21-47.
- Goodfellow, I., Bengio, Y. and Courville, A. (2016), *Deep Learning*, MIT Press, Massachusetts.
- Han, J., Lee, G.H., Park, S., Lee, J. and Choi, J.K. (2021), "A multivariate time series prediction-based adaptive data transmission period control algorithm for IoT networks", *IEEE Internet of Things Journal*, IEEE, Vol. 4662, p. 1.
- Kramer, R., van Schijndel, J. and Schellen, H. (2013), "Inverse modeling of simplified hygrothermal building models to predict and characterize indoor climates", *Building and Environment*, Elsevier, Vol. 68, pp. 87-99.
- Liguori, A., Markovic, R., Dam, T.T.H., Frisch, J., van Treeck, C. and Causone, F. (2021), "Indoor environment data time-series reconstruction using autoencoder neural networks", *Building and Environment*, Elsevier, Vol. 191 December 2020, p. 107623.
- Mehta, P., Bukov, M., Wang, C.H., Day, A.G.R., Richardson, C., Fisher, C.K. and Schwab, D.J. (2019), "A high-bias, low-variance introduction to Machine Learning for physicists", *Physics Reports*, Vol. 810, pp. 1-124.
- Olstad, T.M., Ørnhoi, A.A., Jernæs, N.K., de Ferri, L., Freeman, A. and Bertolin, C. (2020), "Preservation of distemper painting: indoor monitoring tools for risk assessment and decision making in kvernes stave church", MDPI AG, *Climate*, Vol. 8 No. 2, doi: [10.3390/cli8020033](https://doi.org/10.3390/cli8020033).
- Varas-Muriel, M.J., Martínez-Garrido, M.I. and Fort, R. (2014), "Monitoring the thermal-hygrometric conditions induced by traditional heating systems in a historic Spanish church (12th-16th C)", *Energy and Buildings*, Elsevier B.V, Vol. 75, pp. 119-132.
- Wen, T. and Keyes, R. (2019), "Time series anomaly detection using convolutional neural networks and transfer learning", available at: <http://arxiv.org/abs/1905.13628>.

### Corresponding author

America Califano can be contacted at: [america.califano@unipd.it](mailto:america.califano@unipd.it)

---

For instructions on how to order reprints of this article, please visit our website:

[www.emeraldgroupublishing.com/licensing/reprints.htm](http://www.emeraldgroupublishing.com/licensing/reprints.htm)

Or contact us for further details: [permissions@emeraldinsight.com](mailto:permissions@emeraldinsight.com)