

RESEARCH

Open Access



Prediction of the indoor climate in cultural heritage buildings through machine learning: first results from two field tests

Christian Boesgaard¹, Birgit Vinther Hansen², Ulla Bøgvad Kejser^{2*}, Søren Højlund Møllerup¹, Morten Ryhl-Svendsen³ and Noah Torp-Smith¹

Abstract

Control of temperature and relative humidity in storage areas and exhibitions is crucial for long-term preservation of cultural heritage objects. This paper explores the possibilities for developing a proactive system, based on a machine-learning model (XGBoost), for predicting the occurrence of unwanted indoor environmental conditions: either a too high or a too low relative humidity, within the forthcoming 24 h. The features used in the model were hourly indoor and outdoor climate recordings, and it was applied to two indoor heritage environments; a storage facility and a church building. The test accuracy (f1-score) of the model was good (0.93 for high RH; 0.93 for low RH) when applied to the storage building, but only 0.78; 0.62 (high RH; low RH) for the church building test. Challenges encountered include difficulties in obtaining good historical climate data sets for training and testing the model, and the dependency of external IT systems, which, if they fail, inactivates the model without a warning. Several issues call for more research: A desirable improvement of the model would be predictions for periods longer than 24 h ahead, still maintaining a high test accuracy. Further perspectives of using machine learning for indoor environmental forecasting could be for indoor air pollution, or energy consumption due to climate control.

Keywords: Relative humidity, Indoor environment, Prediction, Climate forecast, Environmental control, Storage, Church building, Time-series analysis

Introduction

Heritage institutions, such as galleries, libraries, archives, and museums (GLAM) strive to safeguard and prolong the useful lifetime of the cultural artefacts in their collections. Typically, the main part of the collections is in storage, and heritage preservation in large depends on the environmental control in those facilities. As early as 1979, it was stated by UNESCO that “probably more harm has been done to museum collections through improper storage than by any other means” [1].

Thus, inappropriate temperature and humidity levels speed up the deterioration of physical materials and may cause irreparable damage to cultural artefacts. Chemical decay of organic objects is accelerated by warm and humid conditions, while a too dry atmosphere may cause mechanical damage such as shrinkage, warping, or cracks. Climate also influences the living conditions of insects, pests and microorganisms, and high humidity may accelerate such attacks. Different materials have different optimal storage conditions; however, most collections are well preserved at a cool temperature (below 20 °C) and a moderate (around 40–60%) relative humidity (RH). Additional damaging factors of the environment are light and UV, which fades colour and decomposes material’s surfaces, and air pollution, which initiates several deterioration processes, ranging

*Correspondence: ubk@kb.dk

² Royal Danish Library, Preservation, Søren Kierkegaards Plads 1, 1221 Copenhagen K, Denmark

Full list of author information is available at the end of the article

from metal and stone corrosion to discolouration and brittleness of organic materials such as paper and textile [2]. In summary, management of heritage storage as well as exhibition environments is a fundamental collection care task, typically realised by more or less advanced heating, ventilation, and air conditioning (HVAC) systems.

To control climate set points and optimise the environment in buildings used for cultural heritage preservation, managers need to assess and understand the environmental conditions. Therefore, it is best practice to regularly monitor the indoor environment, of which temperature and relative humidity are the main factors. There are international standards and guidelines directed at different types of heritage collections [3–5].

Traditionally, indoor air quality is monitored by measuring important features of the environment, such as temperature, relative humidity, sometimes air pollutants and light levels, and manually analysing the collected data (climate curves, exceedance of set limits, observations of minima and maxima, etc.) to discover problematic behaviours. Modern HVAC systems can automatically monitor and control the environment and based on pre-setting, e.g., of humidity, they can provide alerts in case the specifications are exceeded.

Thus, facility managers usually control the building environment by analysing data in a retrospective way. They monitor the environment and act when they get indications of unwanted environmental conditions, such as too high or too low levels of relative humidity or highly fluctuating humidity levels. In other words, mitigation actions are triggered by environmental events that have already occurred and they may therefore be initiated too late, when risk is imminent or damage to cultural heritage has already taken place. Ideally, managers should be able to execute actions and adjust the environment before any adverse conditions occur, thereby preventing the risk of harming the cultural heritage collections.

In many fields, scientists are increasingly using machine learning (ML) to improve data analysis. ML is a subfield of computer science and artificial intelligence, based on mathematics and statistics [6]. It encompasses widely used technologies, such as systems that recommend goods based on data about users' past preferences and behaviour, or systems that support weather forecasts based on analysis of previously measured meteorological data. Due to the ability to solve complex non-linear problems, ML techniques are often used in prediction models.

The term machine learning stems from the way the computer programs, also known as algorithms, are developed. Instead of directly programming the

algorithms to do a task, they automatically adjust the way they carry out a task by processing data and finding patterns within it. In this sense, we say that the algorithms learn from data. This analogy between how humans learn and how the algorithms are developed is widely used in ML jargon. For example, it is common to refer to data as 'experience' or 'observations' that algorithms are 'trained' on to solve a problem. ML tasks can further be divided according to whether they perform a classification or a regression, whether they are based on decision trees, clustering, deep learning, or the like.

In this study, we made use of Time series machine learning models, which are based on continuous or repeated measurement data and used this historical data, e.g., from a thermo-hygrometer, to predict future values (in this case temperature and relative humidity).

Numerous studies within the cultural heritage field use ML, typically for tasks around automatic text recognition, annotation and classification of images, and recommendations based on user preferences. In the field of conservation science and heritage preservation studies are more limited. Typically, they focus on identification and classification of materials e.g., pigments [7], or structures e.g., wood types [8]. Other studies use ML to monitor cultural heritage collections or sites for abnormalities. Thus, Zou et al. used ML (deep learning) on image data to support inspection of historical buildings in the Forbidden City in China and locate missing or impaired heritage components [9], while Kejser et al. used ML for classifying the acidity of historic paper samples [10]. Pei et al. [11] used machine learning to predict household mite infestation based on indoor climate conditions and found that the extreme gradient boosting (XGBoost) model was the most suitable method when compared to other methods such as logistic regression and support vector machine (SVM).

For non-cultural heritage buildings, ML algorithms have been used in modern 'smart buildings', often with the aim of energy savings. The development here is mainly on forecasting indoor temperature, and in that regard balancing ventilation and heating/cooling units for as high a thermal comfort as possible while consuming a minimum of energy [12, 13]. Fan et al. [14] used deep learning based methods to predict the cooling load of a building 24 h ahead.

However, the special demand for prediction of indoor humidity and air pollution levels, which is central for GLAM institutions, is less explored. Pernia et al. [15] used ML (k-means clustering) to support the analysis of the indoor air quality of a Belgian church and identify periods of elevated risks for heritage conservation. The authors concluded that the ML method found patterns




Fig. 1 Storage Hall P at the National Museum of Denmark, Ørholm facility—outdoor (left) and indoor (right). The indoor climate sensors are located in the middle of each of the three sections of the building, near the mezzanine

in the data that could help select the best mitigation action and alert guardians about potential environmental risks. Based on an analogy between a human being and a building, La Russa [16] envisions that new methodologies, including wireless sensor networks, artificial intelligence, machine learning and visual programming language, can be used to collect and analyse building data and propose solutions based on decision-making models. This concept was further developed by La Russa & Santagati [17] in relation to museum collections stored in historic buildings, with the goal to improve the conservation of the collections as well as the architecture.

The aim of the present study was to use ML for early warning of unwanted environmental conditions by processing environmental data as it is recorded daily, and by combining indoor and outdoor environmental datasets. Our hypothesis was that based on patterns within the data, probabilities for upcoming environmental events or periods of incorrect levels can be forecasted so that preventive actions can be initiated before it affects the preservation of cultural heritage collections. We used a modern museum storage hall, and a historic church building, as examples for testing a ML model, built on past indoor climate data collected at the two sites. The objective of this work was also to give conservators and conservation scientists an idea of how machine learning can improve data analysis, and advance cultural heritage preservation.

Methods

Test sites

Two indoor heritage sites were selected for the prediction of incidents where the relative humidity varied beyond

conservation set points. Both were located in Zealand, Denmark; one being a museum storage hall, and the other a rural church. Denmark has a temperate climate (Köppen classification Dfb), where the annual average temperature is 8 °C and the relative humidity is 82% RH.

The storage hall is a 1500 m² (11,000 m³) purpose-built museum storage facility from 1990 belonging to the National Museum of Denmark (Fig. 1). It is located in the suburban area Ørholm north of Copenhagen (coordinates: 55.800855, 12.506972). The facility houses a mixed collection of larger cultural history objects, such as vehicles, boats, furniture, and sculptures, made of wood, metal and stone. The hall, which consists of three interconnected sections, is climatically controlled by semi-passive means, i.e., it has a thermally well-insulated building envelope with a low natural ventilation rate at 0.1 per hour, which minimises the influence of weather, aided by a mechanical humidity control system. For conservation reasons the indoor climate is set to be controlled within the interval of 40–60% RH, while temperature is allowed to vary with the seasons. However, the relative humidity in summer and autumn often exceeds the limit of 60%. The storage building was previously described in detail by Padfield [18], and Ryhl-Svendsen et al. [19].

The other site was a mediaeval church located in the village of Annisse (coordinates: 55.982406, 12.170868) (Fig. 2). The original church building dates from the twelfth century, with vaults added around 1400, and a new roof from 1967. It is a typical Danish country-side church with massive stone, brick and lime mortar walls and vaults, of about 120 m² (500 m³). The air exchange rate was not measured, but based on previous experience



Fig. 2 Annisse Church—outdoor (left) and indoor (right). The indoor climate sensor is located in the nave at the pulpit

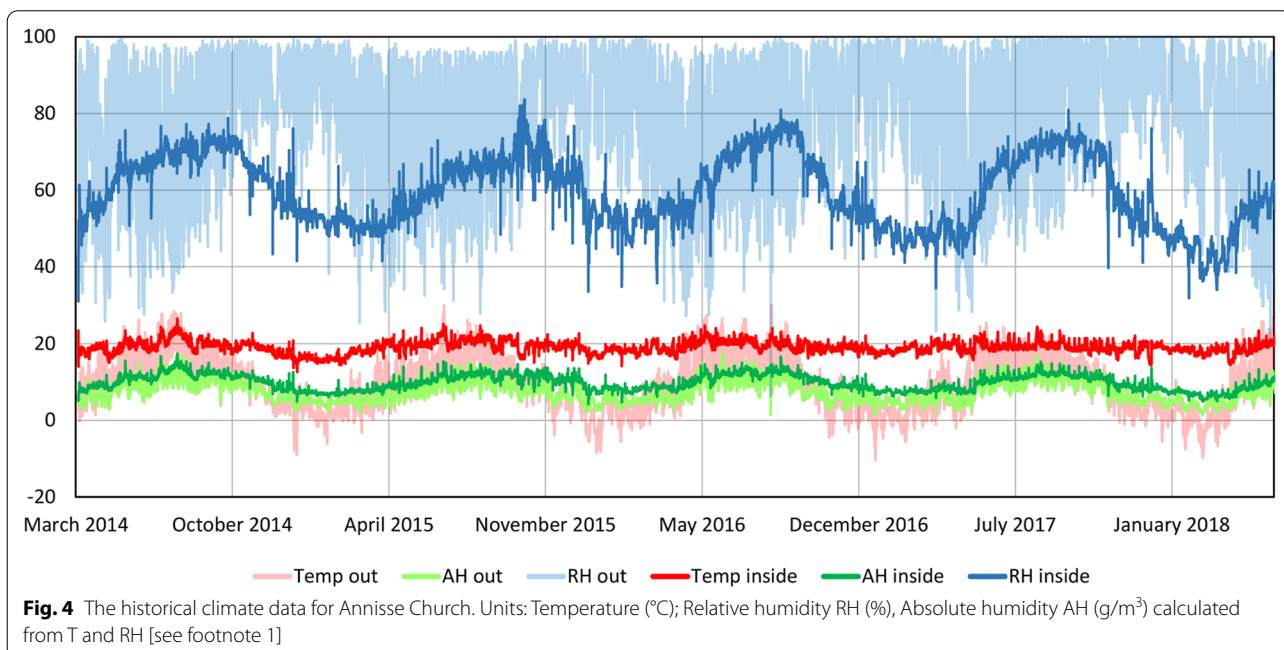

it is estimated to be less than 0.5 per hour. The vaults have been thermally insulated in the attic by a perlite-based mortar, and the church room is heated by electric elements to maintain about 20 °C in the cold seasons [20]. Although the heating of the church indirectly controls the relative humidity, there is no strict climate control. We chose an upper and lower relative humidity limit of 75% and 45% RH, respectively, as the conservation limits for the Annisse Church ML model. Acknowledging the concept of historical climate defined in EN 15757 [21] the limits are typical for Danish country-side churches without strict climate control. RH levels above 45% prevents desiccation of materials and levels below 75% RH prevents mould growth. A climate within this interval is not uncommon in such historical buildings with little climate control, on the other hand, exceeding beyond this is still expected several times a year.

Historical data sets

The temperature and relative humidity of the Ørholm storage hall have been monitored since it was put in use in 1990, however by different monitoring set-ups and instruments over time. We chose to focus on the data set recorded by the current building management control (BMS) sensors, which were installed in 2008, in order to avoid instrumental bias from recordings by different sensors over time. At the time of training our ML algorithm in 2021, the data set consisted of about 12.5 years of continuous temperature and relative humidity recording, with measurements time-stamped at 1-h intervals (almost 110,000 readings). The historical climate data for the Ørholm storage facility used

for building and testing the machine learning model is shown in Fig. 3. The store is not heated, so the temperature is allowed to drift slowly over the seasons, within the interval 10–23 °C (average 16 °C). The relative humidity is controlled mechanically to remain within 40–60% RH, which it was during the monitoring period for more than 90% of the time (average 51% RH). However, exceedances occur, and extreme single episodes over the 12.5 year monitoring period were 30 and 72% RH.

The indoor climate of Annisse Church is normally not monitored, however, a recent data set from a 4.2 year monitoring period (March 2014–May 2018) was available from a previous research project [20]. During this period temperature and relative humidity were recorded inside the nave at 1-h intervals (about 36,800 readings). Although there was almost a four-year gap up to the current application of our ML algorithm at the site, there had been no alterations to the building or the heating regime, so we considered the data set still representative of the present indoor climate. The church is heated for comfort, and for 90% of the time the temperature was between 16 and 22 °C (average 19 °C), while the average relative humidity was 60% RH. Relative humidity varies considerably between seasons; lowest in late winter and highest in late summer/early fall. However, for more than 90% of the time it remains within the chosen conservation limits of 45–75% RH. The annual magnitude was about 35–80% RH, with extreme single episodes over the 4.2 year monitoring period of 31 and 84% RH. The extreme low RH episodes occurred during winter, when outside dry air enters the heated church reducing RH further; while the high



RH episodes occurred late summer when very humid outdoor air (possibly due to rain), enters the colder church, cooling the air and thus increasing RH further. The Annisse Church indoor climate data sets used for building and training the ML model is shown in Fig. 4. The dataset contains measured temperature and relative humidity, and the calculated absolute humidity, for indoors and outdoors.

Outdoor climate data were acquired for the same periods as the indoor data sets, from nearby weather stations of the national Danish Meteorological Institute (DMI). Data was collected in the same format as for the indoor data sets; time-stamped air temperature and relative humidity records at 1-h intervals. For the Ørholm storage building the nearest DMI weather station was Jægersborg (4 km), and for Annisse Church

it was Sjølsmark (19 km). Even though small short-term differences between the outdoor measurements at the weather stations and the sites may occur, for example due to a very local rainstorm, these are considered insignificant on a 24-h and longer basis. This is also supported by the fact that the climate data from the two weather stations we used are largely uniform, even though there is also a certain distance between them (approx. 15 km).

Climate monitoring

At the Ørholm site data was acquired through the BMS connected climate sensors, used for the humidity control system (Hygromaxx S Transmitters, Novasina, Lacerna, Switzerland). Climate readings were accessible at a central BMS server, which could be accessed in real-time via an application programming interface (API). Historical data was archived by The National Museum in a library at eclimatenotebook.com (Image Permanence Institute, Rochester, USA).

In Annisse the historical data was recorded by battery-driven climate sensors with data loggers (TinyTag 2 Plus, Gemini Dataloggers, Chichester, UK), and at present by a narrow-band IoT (SIM card) connected climate sensor, which is accessed in real-time by API (Roomalyzer, IoT fabrikken Aps, Roskilde, Denmark).

All sensors are electronic, including thermistor temperature sensors and polymer-based chips humidity sensors. HygroMaxx S has an accuracy of ± 0.5 °C and $\pm 1.5\%$ RH, and TinyTag2 Plus an accuracy of ± 0.5 °C and $\pm 3.0\%$ RH. The loggers were placed in the middle of the rooms shielded from heat radiation and ventilation and away from surfaces with a temperature different from the air (EN 16242) [22].

The DMI weather station data (historical and real-time) was accessed through the DMI Open Data web portal (<https://www.dmi.dk/frie-data/>).

All the historical climate data from both sites, as well as the climate data harvested from DMI were time stamped at 1-h intervals. For consistency this measurement frequency was maintained in the current ongoing climate monitoring.

Machine learning model

A prototype model was developed for predicting the indoor air quality (in this case the relative humidity). The model was based on data from the Ørholm facility:

- Outdoor humidity, hourly observations;
- Outdoor temperature, hourly observations;
- Indoor humidity, hourly observations; and
- Indoor temperature, hourly observations.

Since there are three climate sensors in the Ørholm facility, the indoor observations were converted to their medians to eliminate any outliers from drifting sensors. In order to possibly enhance the precision of the model the calculated absolute humidity was also added to the data (for outdoor and indoor median value).¹ The resulting data set for Ørholm consisted of 105,797 observations from 2009-01-01 to 2021-03-02.

The aim of the model was to predict whether the relative humidity would be outside the acceptable interval (40–60% RH) during the following 24 h from a given point in time. The 24 h forecast window was selected as the minimum amount of time to be useful in practice, and provide sufficient time for a facility manager to take action (e.g., engage extra humidity control equipment), while still being short enough that it could be expected to be an achievable goal modelling-wise.

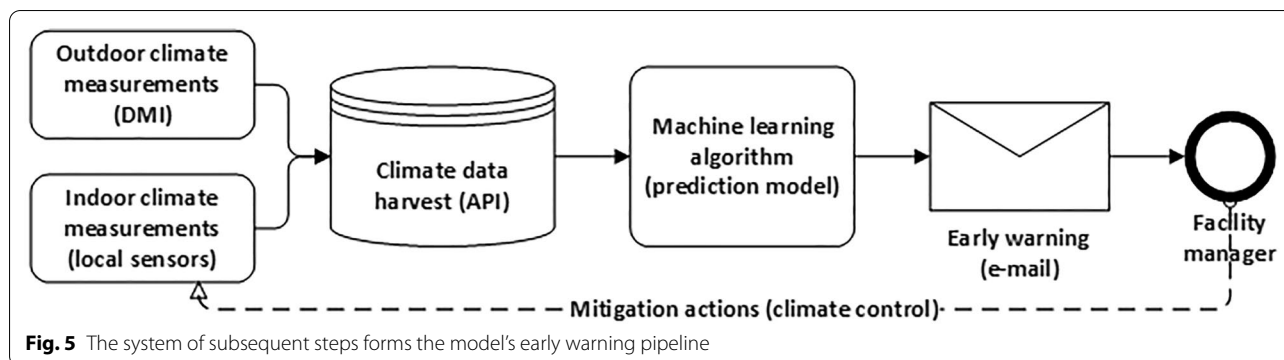
As the primary data are time series (data points ordered in time), we used time series analysis to extract statistical information [23]. We started by doing time series specific exploratory data analysis and combined this with domain knowledge to assess what was possible. The Python programming language package tsfresh [24] was used to calculate and test a large number of time series characteristics against the prediction targets, based on the available data. The dataset was extended with the following features found significant by tsfresh:

- 24 h mean of the indoor median humidity;
- 24 h exponential moving average of the indoor median humidity;
- 24 h median of the indoor median humidity; and
- 24 h minimum and maximum of the indoor median humidity.

The features produced by tsfresh are based on the provided data, which means that no new data is added, but the features provided by tsfresh can enhance the signal and reduce noise by focusing on different aspects of the original data.

We only selected some of the features produced by tsfresh, namely the simplest features with the best signal so as to produce an explainable model that still performed well. Based on this dataset the model was built. As it turned out that the outdoor climate data did not improve the model performance in the Ørholm case it was excluded. We worked with two standard ways of building models for time series analysis, namely

¹ The absolute humidity of air (AH) was calculated from the measured temperature and relative humidity by using the formulae [31]: $AH = (1322.9 * (RH/100) * EXP(T/(T + 238.3)) * 17.2694)/(T + 273.16)$, where T is the air temperature in degree C; RH is the relative humidity of air in percent; and AH is the absolute humidity of air in g/m³.



RandomForest [25] and XGBoost [26]. RandomForest is relatively easy to comprehend, compared to XGBoost, but sometimes the latter yields slightly better results.

It is standard practice to build ML models in the following manner: First, “hold off” some data from the model building, in order to be able to evaluate the model's performance—a standard split is to use 80% of the data set for training, and the remaining 20% for evaluation (this is done to ensure the model is able to perform on unseen data, corresponding to how the model would be used for prediction in a system). The newest (last) 20% of data is used for testing to ensure the model is trained on historical data only. Then feed the training data into the model, and train it on that. This results in a model, which then can be evaluated using the evaluation data (or test data, which it is often called in literature). For each of the template models (RandomForest/XGBoost), there are some “hyperparameters” that should be decided upon. We used a tool, GridSearchCV from sklearn [27], to test a number of combinations of these hyperparameters, and reported the combination that gave the best model performance. In the same way, a model (XGBoost) was trained on historic climate data from Annisse Church. Results of the models' performance are reported in the “[Results and discussion](#)” Section below.

Early warning pipeline

The ML model engages in a system consisting of a number of subsequent steps, from which the prediction based on climate monitoring may issue an early warning on an upcoming incorrect humidity. This is done in a pipeline of tools which are initiated to ensure that the system actively monitors and harvest the climate data from the sites (Fig. 5). There are a number of practical decisions involved here, but in short; the model is stored and is running the predictions on an internal server (at DBC Digital's infrastructure).

Climate data is fetched daily from the local servers or data cloud solutions at the museum, church or

meteorological institute according to the following steps (exemplified below by the Ørholm site at The National Museum of Denmark):

- At a specified time in the morning, a job is initiated automatically, to start the process.
- A program logs on to the required virtual private network (VPN) to fetch daily measurements from the Ørholm facilities, stored on the National Museum of Denmark's BMS. If, for some reason, this procedure fails, a notification is sent to the relevant developers who then need to act on this, because this is critical to the system.
- Another program is started, which initially loads the model (XGBoost), and feeds it the previous day's measurements from Ørholm. The result of this is a prediction of whether there will be too high (or low) humidity in the coming 24-h period.
- If an extreme humidity level is predicted, an email is sent to the relevant facility manager who can then act proactively. If no extreme levels are predicted, a notification about this is sent to a developer, who is thus informed that the pipeline has been executed successfully.

Results and discussion

ML model testing and evaluation

The models were evaluated on the ability to predict for the coming 24 h if the indoor relative humidity would be too high or too low (humidity episodes), as compared with the defined thresholds. The evaluation factors for the prediction were the precision (the fraction of relevant episodes among the retrieved episodes), the recall (the fraction of relevant episodes that were retrieved), and the f1-score, which expresses the test accuracy in a single metric (the harmonic mean of the precision and recall). The data used for testing the models is described in the Section “[Historical data sets](#)”.

Table 1 Evaluation for the model based on RandomForest, Ørholm storage hall

Prediction for next 24 h	Precision	Recall	F1-score
Humidity too high (> 60% RH)	0.84	0.80	0.82
Humidity too low (< 40% RH)	0.87	0.90	0.88

The model predicted episodes of indoor relative humidity outside 40–60% RH

Table 2 Evaluation for the model based on XGBoost, including AH, Ørholm storage hall

Prediction for next 24 h	Precision	Recall	F1-score
Humidity too high (> 60% RH)	0.93	0.92	0.93
Humidity too low (< 40% RH)	0.92	0.93	0.93

The model predicted episodes of indoor relative humidity outside 40–60% RH

As mentioned above (Machine learning model) a model was first built based on RandomForest from sklearn, followed by one based on XGBoost. For completeness, in the example of Ørholm (including AH) we show the evaluation numbers for both, to demonstrate that in our case XGBoost performed better, and by how much (see Tables 1 and 2). Note that since the outdoor climate data did not influence the model's performance it was omitted in the Ørholm case. As can be seen, the numbers for XGBoost are better for the Recall, meaning that this model more often detects a situation where humidity will be too high or low in the coming 24 h period.

Table 3 shows the results of XGBoost excluding AH. Comparing the results in Tables 2 and 3 reveals that the model's overall accuracy is not improved by AH.

For the example of Annisse Church two XGBoost models were tested; one based on all six input parameters (indoor and outdoor temperature, relative and absolute humidity: Table 4), and another one only on humidity parameters (excluding temperature: Table 5).

XGBoost results in a lower performance in the Annisse Church case compared to Ørholm (see Tables 2 and 4). This may be due to the larger volume of data from the Ørholm facility compared to Annisse Church. However, there can be other environmental related reasons that can explain the better performance such as the building's construction, the surrounding conditions, and the different conservation limits.

In the case of the Annisse Church we have not yet had the chance to evaluate thoroughly the significance of including/excluding the outdoor temperature and RH on the performance of the model, but excluding outdoor temperature seems to improve the overall accuracy.

Table 3 Evaluation for the model based on XGBoost, excluding AH, Ørholm storage hall

Prediction for next 24 h	Precision	Recall	F1-score
Humidity too high (> 60% RH)	0.95	0.91	0.93
Humidity too low (< 40% RH)	0.96	0.95	0.95

The model predicted episodes of indoor relative humidity outside 40–60% RH

Table 4 XGBoost model evaluation for Annisse Church, for prediction of episodes of indoor relative humidity outside 45–75% RH

Prediction for next 24 h	Precision	Recall	F1-score
Humidity too high (> 75% RH)	0.74	0.83	0.78
Humidity too low (< 45% RH)	0.96	0.29	0.45

This model was based on all input parameters (temperature, relative humidity, absolute humidity, indoors and outdoors)

Evaluation of the early warning system

Regarding the model used for the Ørholm data, the model was built and applied to real-time data monitoring starting from November 4th, 2021, and for Annisse Church the model was built and applied on April 17th, 2022. Since then, the models have retrieved data once a day, performed the analysis, and, if the relative humidity has been predicted to exceed the set limits within the next 24 h, a warning message has been sent by email.

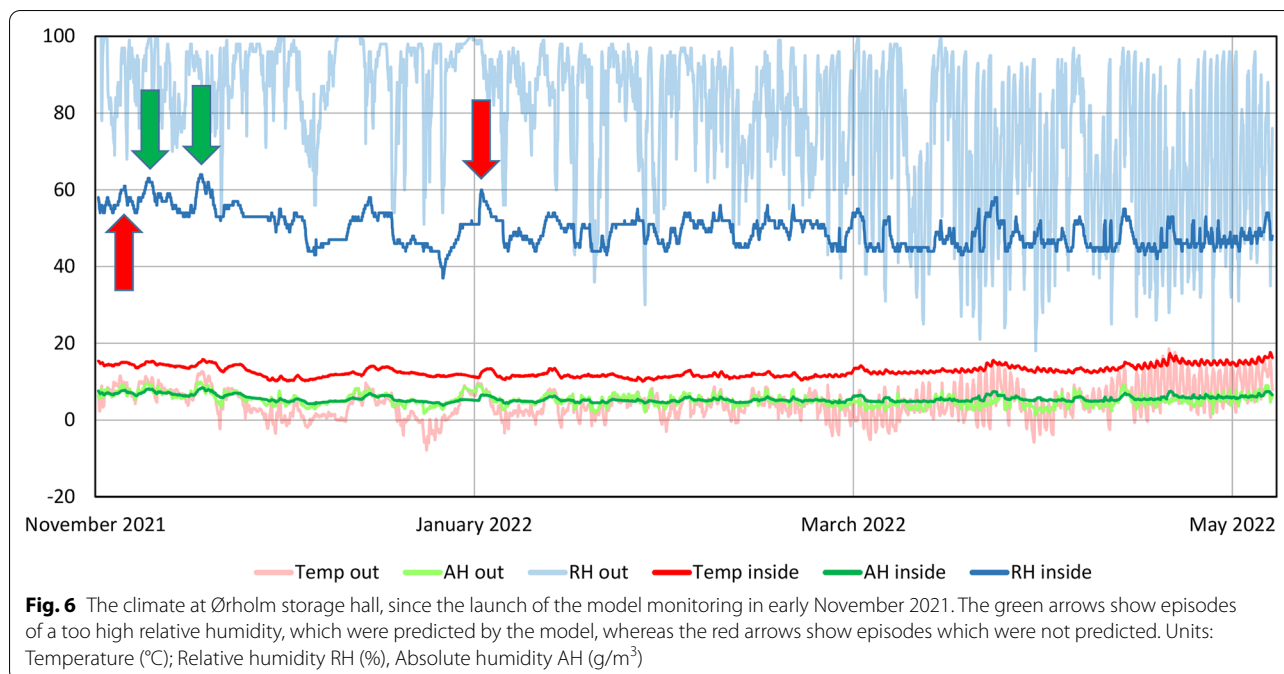
For Ørholm two warnings have been issued, on November 11th and November 19th (Fig. 6). From Annisse Church, no warnings have been issued yet (as of June 2022).

When analysing the indoor climate of the Ørholm building since the model was applied, we found that besides the two issued warnings in November 2021, a third episode also in November was not identified, nor was an episode in January 2022 (although very brief, and just above 60% RH). This is shown in Fig. 6, with green arrows pointing out retrieved episodes, and red arrows pointing to the missed ones. This 50/50 success rate is poor compared to the rather good f1-score of 0.92. However, as the statistical basis is very small (four episodes)

Table 5 XGBoost model evaluation for Annisse Church, for prediction of episodes of indoor relative humidity outside 45–75% RH

Prediction for next 24 h	Precision	Recall	F1-score
Humidity too high (> 75% RH)	0.73	0.83	0.78
Humidity too low (< 45% RH)	0.98	0.45	0.62

In this model temperature was excluded, so the results are based only on relative humidity, absolute humidity, indoors and outdoors



it is too early to conclude anything from this yet. It will be interesting to follow the predictions of the model over the next years' time. The system could be modified to use a lower threshold for warnings. This would be expected to improve the recall, but also result in false warnings (lower precision).

Based on the historical climate data from the Annisse Church and the short period the early warning system has been running (from Spring to early Summer), it was expected that no alarms would be sent out. The relative humidity will typically only become too high during the fall (Fig. 4), which will be tested as we continue to run the model for Annisse for at least another year.

Two immediate problems arose during the test runs. The warning system for Ørholm failed at some point due to an update of an adjacent system which was not taken into account in the email notification routine. This failure might explain why no email was sent out for the too high relative humidity episode in January 2022 (Fig. 6) and if not found out, could be misinterpreted as having no current risk. The current problem was solved, and the system works, however, it emphasized that errors related to software communicating automatically is a potential risk.

The other issue is that the model is currently set up to harvest, analyse and send out warnings every morning for the following 24 h. Therefore, any alarming humidity levels that the model predicts may be given less than 24 h

ahead and could even be almost simultaneously with the unwanted rise in humidity, as seen in Fig. 6.

Challenges and future perspectives

It is desirable that a future version of the model is able to predict a longer time period ahead; e.g., 2 days, or more. This requires, however, good historical data sets containing many episodes of the type of unwanted humidity event which the model can learn to predict from (e.g., too high relative humidity). The Annisse Church case is a good example of this. It had an f1-score of 0.78 for predictions variations outside 45–75%RH; an interval chosen for conservation reasons. However, if for the sake of example, this interval was narrowed down to 50–70% RH, the f1-score became much better; namely 0.95. This is obviously due to the many more episodes of relative humidity outside the 50–70% RH band, which optimises the training of the model.

Domain knowledge about the input parameters is an important factor in choosing the basis for building a ML model. Our model is based on the most common climate parameters, which routinely are measured and therefore readily available in most heritage institutions; temperature and relative humidity. But in addition to this, the absolute humidity was added to the model [see footnote 1], as the algorithm might not by itself discover this rather complex relation, when analysing patterns in the indoor climate variations. This was a decision we

took as conservation professionals, knowing that taking the absolute humidity of air into account might reinforce the validity of the analysis. However, as described in the “Results and discussion” Section the addition of AH did not improve the performance of the model.

During the project it became quite clear that, although, as already mentioned, climatic parameters are quite commonly measured indoors, it can be hard to find historical data sets for long and persistent periods, which are required for building good quality models. Data formats, measurement resolution, and lack of data storage routines may be a challenge [28]. Outdoor historical and real-time weather data may be accessible by local weather stations, or even by national meteorological services, but that is not a matter of course. On a future perspective it could be interesting to include forward-looking weather prognosis as a parameter for indoor climate prediction, given the prognosis could be expressed in a numerical way.

Another challenge is to find good data on other parameters, for example, on energy consumption, or air pollution. At the beginning of this project, we had a vision of including several air quality parameters to the model. However, as no historical data could be identified for the buildings in question, we are only now in the beginning of establishing a proper data set, after one first year of continuous air pollution monitoring (PM_{2.5} measured outside and inside the Ørholm building). How this develops will be reported in a future publication.

From early on an “use best technology available” approach provided the standard for environmental control in heritage institutions [29], however, today’s guidelines are to a higher degree based on a cost–benefit approach balancing preservation and the resources involved [30]. The current climate and energy crisis has increasingly put emphasis on sustainability in preservation, and managers are also aware of the importance of measuring the facility’s energy consumption. Another perspective of using ML for indoor environmental conditions is the potential to predict the energy consumption for climate control in advance. There is a need to investigate these issues further, which is yet another area where ML programmers, museum facility managers, and conservators could unite on developing energy-saving and conservationally feasible indoor climate control solutions.

Conclusions

In this research, we investigated the usability of supervised machine learning methods for predicting the occurrence of harmful environmental conditions inside

buildings used for preservation of cultural heritage. More specifically, we conducted a case study based on historical environmental data from two heritage facilities and tested the ability of two different machine-learning algorithms, namely RandomForest classifier (RFC) and extreme gradient boosting algorithm (XGBoost) to forecast incidents of too low or too high levels of relative humidity inside the facility.

The model had a quite good f1-score of 0.93/0.95 (RH too high/RH too low) in predicting the humidity in the humidity-regulated storage facility (Ørholm), but only 0.78/0.62 in the temperature regulated church (Annisse).

Machine learning technologies can reveal new knowledge and insights from data collected over decades in heritage institutions, and we believe that the technology can be used to develop effective data driven control of the indoor environment. In the described cases it revealed for the Ørholm facility that the outdoor weather conditions in general did not have much influence on predicting the indoor relative humidity level, and in the Annisse case it seems that the outdoor temperature did not. This indicates that the significance of the different parameters depends on the conditions at the actual location.

The main limitation of the study is that we have only been able to train the prototype on data sets from two cases where the indoor climate is relatively well regulated. This is due to a lack of long-period data sets in good quality, including a sufficient number of unwanted incidents for the model to train on.

Future work includes applying the prototype with the use of weather forecasts to spaces with more instant impact from the climate outside the building and less climatically regulated historical buildings. Likewise, we plan to investigate the use of an adjusted prototype to predict other environmental risks such as inappropriate temperature or pollutant levels, the effect of initiated climate control methods, and the energy consumption involved.

Abbreviations

AH: Absolute humidity; AI: Artificial Intelligence; API: Application programming interface; BMS: Building management system; CNN: Convolutional neural network; DMI: Danish Meteorological Institute; GLAM: Galleries, Libraries, Archives and Museums; HVAC: Heating, ventilation and air conditioning; IT: Information technology; ML: Machine learning; PM: Particulate matter; SVM: Support vector machine; VPN: Virtual private network; XGBoost: EXTreme Gradient Boosting.

Acknowledgements

The authors thank Lars Aasbjerg Jensen, National Museum of Denmark, for providing climate data from the Ørholm facility and all his help throughout the project; Poul Klensz Larsen, National Museum of Denmark, for providing

climate data from Annisse Church; and Alice Ryhl, DTU Compute, for application of the model to Annisse Church climate data.

Author contributions

MRSV, BVH and UBK were responsible for conceptualisation, methodology, formal analysis, validation, writing, reviewing and editing of the paper. CBO, SHM, and NTS enabled data harvest, developed the machine learning model and executed the data analysis. MRSV organised sensor installation, instrument calibration, and data collection at the sites. UBK was responsible for project management. All authors contributed to the final manuscript. All authors read and approved the final manuscript.

Funding

The study was funded by the Danish Ministry of Culture (Project No. FPK.2020-0009). The funding body had no influence on the design or interpretation of the study.

Availability of data and materials

The model and datasets generated and analysed during the current study are available in the Zenodo repository: <https://doi.org/10.5281/zenodo.6589119>. The data is also available at Github: <https://github.com/DBCDK/orholm-prediction-demo/tree/v.1.0.2>.

Declarations

Competing interests

The authors declare that they have no competing interests.

Author details

¹DBC Digital A/S, Tempovej 7, 2750 Ballerup, Denmark. ²Royal Danish Library, Preservation, Søren Kierkegaards Plads 1, 1221 Copenhagen K, Denmark. ³Royal Danish Academy, Institute of Conservation, Esplanaden 34, 1263 Copenhagen K, Denmark.

Received: 1 June 2022 Accepted: 28 September 2022

Published online: 28 October 2022

References

- Johnson EV, Horgan JC. Museum collection storage. Paris: UNESCO; 1979.
- CCI, Agents of deterioration. Canadian Conservation Institute (Government of Canada). 2017. <https://www.canada.ca/en/conservation-institute/services/agents-deterioration.html>. Accessed 31 May 2022.
- ISO 11799:2015. Information and documentation—document storage requirements for archive and library materials. Geneva: International Standard Organization; 2015.
- EN 16893:2018. Conservation of cultural heritage—specifications for location, construction and modification of buildings or rooms intended for the storage or use of heritage collections. Brussels: European Standard Organization CEN; 2018.
- ASHRAE. Museums, galleries, archives and libraries, in ASHRAE application handbook, american society of heating, refrigerating and air-conditioning engineers (ASHRAE), Atlanta, 24.1–24.46; 2019.
- Mitchell TM. Machine learning, McGraw-Hill Science. ISBN: 0071154671. 1997. <http://www.cs.cmu.edu/afs/cs.cmu.edu/user/mitchell/ftp/mlbook.html>. Accessed 31 May 2022.
- Sevtlidis V, Pavlidis G. Effective Raman spectra identification with tree-based methods. *J Cult Herit*. 2019;37:121–8. <https://doi.org/10.1016/j.culher.2018.10.016>.
- Kobayashi K, Hwang S-W, Okochi T, Lee W-H, Sugiyama J. Non-destructive method for wood identification using conventional X-ray computed tomography data. *J Cult Herit*. 2019;38:88–93. <https://doi.org/10.1016/j.culher.2019.02.001>.
- Zou Z, Xuefeng Z, Peng Z, Fei Q, Niannian W. CNN-based statistics and location estimation of missing components in routine inspection of historic buildings. *J Cult Herit*. 2019;38:221–30.
- Kejser UB, Vinther Hansen B, Ryhl-Svendsen M, Boesgaard C, Møllerup S. Teaching machines to think like conservators—machine learning as a tool for predicting the stability of paper based archive and library collections. *Transcending Boundaries: Integrated Approaches to Conservation: ICOM-CC 19th Triennial Conference Preprints*, Beijing, 17–21 May 2021. 1 s. <https://www.icom-cc-publications-online.org/4417/Teaching-machines-to-think-like-conservators-Machine-learning-as-a-tool-for-predicting-the-stability-of-paper-based-archive-and-library-collections>. Accessed 31 May 2022.
- Pei J, Gong J, Wang Z. Risk prediction of household mite infestation based on machine learning. *Build Environ*. 2020;183: 107154. <https://doi.org/10.1016/j.buildenv.2020.107154>.
- Alawadi S, Mera D, Fernández-Delgado M, et al. A comparison of machine learning algorithms for forecasting indoor temperature in smart buildings. *Energy Syst*. 2020. <https://doi.org/10.1007/s12667-020-00376-x>.
- Chen C-C, Lee D. Artificial intelligence-assisted heating ventilation and air conditioning control and the unmet demand for sensors: part 1. Problem formulation and the hypothesis. *Sensors*. 2019;19:1131. <https://doi.org/10.3390/s19051131>.
- Fan C, Xiao F, Zhao Y. A short-term building cooling load prediction method using deep learning algorithms. *Appl Energy*. 2017;195:222–33. <https://doi.org/10.1016/j.apenergy.2017.03.064>.
- Leyva Pernia PD, Demeyer S, Schalm O, Anaf W. A data mining approach for indoor air assessment, an alternative tool for cultural heritage conservation. *IOP Conf Ser Mater Sci Eng*. 2018;364:012045. <https://doi.org/10.1088/1757-899X/364/1/012045/meta>.
- La Russa F. HS—BIM: historical sentient—building information model. 5. 17–27. 2019. https://www.researchgate.net/publication/343979778_HS_-_BIM_Historical_Sentient_-_Building_Information_Model. Accessed 31 May 2022.
- La Russa FM, Santagati C. An AI-based DSS for preventive conservation of museum collections in historic buildings. *J Archaeol Sci Rep*. 2021;35: 102735. <https://doi.org/10.1016/j.jasrep.2020.102735>.
- Padfield T. Low-energy climate control in museum stores. A postscript. *ICOM-CC Triennial Conference*, Edinburgh; 1996. p. 68–71.
- Ryhl-Svendsen M, Jensen LA, Bøhm B, Larsen PK. Low-energy museum storage buildings: climate, energy consumption, and air quality, UMTS research project 2007–2011: final data report. Department of Conservation, National Museum of Denmark, Lyngby; 2012. p. 121.
- Larsen PK. Climatic protection of historical vaults with lime—perlite mortar. *Stud Conserv*. 2020. <https://doi.org/10.1080/00393630.2020.1758499>.
- EN 15757:2010. Conservation of cultural property—specifications for temperature and relative humidity to limit climate-induced mechanical damage in organic hygroscopic materials. Brussels: European Standard Organization CEN; 2010.
- EN 16242:2012. Conservation of cultural heritage—procedures and instruments for measuring humidity in the air and moisture exchanges between air and cultural property. Brussels: European Standard Organization CEN; 2012.
- Nielsen A. Practical time series analysis, prediction with statistics and machine learning, O'Reilly Media; 2019.
- Christ M, Braun N, Neuffer J, Kempa-Liehr AW. Time Series Feature extraction on basis of scalable hypothesis tests (tsfresh—a Python package). *Neurocomputing*. 2018;307:72–7. <https://doi.org/10.1016/j.neucom.2018.03.067>.
- Wikipedia, RandomForest. https://en.wikipedia.org/wiki/Random_forest. Accessed 31 May 2022.
- Wikipedia, XGBoost. <https://en.wikipedia.org/wiki/XGBoost>. Accessed 31 May 2022.
- SciKit Learn, GridSearchCV. https://scikit-learn.org/stable/modules/generated/sklearn.model_selection.GridSearchCV.html. Accessed 31 May 2022.
- Padfield T. Why keep climate records—and how to keep them. *Museum Microclimate*, Contributions to the Copenhagen Conference, 19–23 November 2007. The National Museum of Denmark; 2007. p. 157–163.
- Thomson G. The museum environment. 2nd ed. London: Butterworth-Heinemann; 1986.
- ICOM-CC. Environmental guidelines, ICOM-CC and IIC Declaration, International Council of Museums—Committee for Conservation. 2014. <https://www.icom-cc.org/en/environmental-guidelines-icom-cc-and-iic-declaration>. Accessed 31 May 2022.

31. Padfield T. Calculator for atmospheric moisture, Conservation Physics Website. 2009. <https://www.conservationphysics.org/atmcalc/atmocalc.php>. Accessed 31 May 2022.

Publisher's Note

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

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
