




Predicting façade deterioration using machine learning approach with drone imagery and microclimate data

Downloaded from: <https://research.chalmers.se>, 2025-09-25 23:23 UTC

Citation for the original published paper (version of record):

Mandinec, J., Sasic Kalagasidis, A., Johansson, P. (2025). Predicting façade deterioration using machine learning approach with drone imagery and microclimate data. *Automation in Construction*, 178. <http://dx.doi.org/10.1016/j.autcon.2025.106443>

N.B. When citing this work, cite the original published paper.



Predicting façade deterioration using machine learning approach with drone imagery and microclimate data

Jan Mandinec^{*}, Angela Sasic Kalagasidis, Pär Johansson

Department of Architecture and Civil Engineering, Chalmers University of Technology, SE-41296 Gothenburg, Sweden

ARTICLE INFO

Keywords:

Degradation assessment
Computer vision
Degradation risk
Degradation prediction
Machine learning

ABSTRACT

Accurate prediction of façade deterioration due to microclimate effects is crucial for sustainable building management and preservation. This paper introduces a methodology for full façade risk assessment using image-based empirical evidence from real buildings. The PAIR methodology—Prepare, Analyze, Integrate, Relate—combines drone imagery, computer vision, weather data, and three-dimensional neighborhood models to create a database that organizes façade deterioration data into sections. This database supports a machine learning model to predict façade deterioration risks. A case study of 16 brick façades in Gothenburg, Sweden, demonstrated the model's strong performance ($R^2 = 0.978$, $MSE = 0.0003$) on the test sample. However, performance declined on an excluded validation façade ($R^2 = -0.467$, $MSE = 0.024$) due to limited training data and inaccuracies from prior maintenance. Despite these limitations, the methodology provides a computationally efficient alternative to full-scale hygrothermal modeling for assessing deterioration risk across entire façades.

1. Introduction

Addressing the weathering effects on building façades often requires reactive maintenance. In Sweden, about 70 % of renovations are driven by urgent needs after a component's service life ends [1]. While straightforward, this approach leads to higher costs compared to proactive maintenance, which can reduce costs by up to 18 % through timely planning and prioritization [2]. Emerging technologies, such as computer vision image analysis and drone-based façade assessments, could enable early detection of degradation, improving maintenance planning as is the case with civil engineering structures [3,4]. This paper explores leveraging of these technologies to predict façade deterioration risks more accurately. Both reactive and proactive maintenance strategies can be costly without preventive measures such as regular cleaning, inspection and maintenance, protective coating and proper drainage. Lack of such measures often leads to premature renovation needs and escalating deterioration.

Preventing moisture-induced deterioration is typically ensured by designing façades according to technical standards. In Sweden, for example, the ByggaF method [5] provides frameworks and checklists for reducing moisture risks from the planning to the operation stages. In addition to ByggaF, there are hygrothermal performance indicators that combine critical moisture saturation exceedance with dose-response

functions. This is mainly used in the design stage, allowing, for example, the risk of mold growth [6] and freeze-thaw attacks [7] to be assessed by hygrothermal simulation. Despite these methods, unexpected material degradation still occurs as current assessments do not indicate when risks will materialize, highlighting their main limitation.

1.1. Predicting moisture-induced façade degradation

In general, hygrothermal modeling and its performance indicators cannot predict the onset and evolution of degradation. Portal et al. [8] made a rare step in this direction by coupling a heat and moisture model with mechanical assessment to study the response of its lime wood supports of a historical painting to changing surrounding climate conditions. In most cases, however, dose response indicators, including freeze-thaw deterioration, are employed to estimate the risk of degradation. In theory, such indicators correlate the number of freeze-thaw cycles (doses) with the progression of damage. Laboratory studies show that increased freeze-thaw cycles lead to more severe deterioration, impacting the compressive strength and the micro/macro structure of bricks [9–11] and mortar [12,13]. However, applying the correlation to real building envelopes is challenging due to uniform laboratory conditions versus varied real-world exposures. Technical standards for testing natural stone (EN 12371:2010), calcium silicate masonry (EN

^{*} Corresponding author.

E-mail address: jan.mandinec@chalmers.se (J. Mandinec).

<https://doi.org/10.1016/j.autcon.2025.106443>

Received 14 January 2025; Received in revised form 28 July 2025; Accepted 28 July 2025

Available online 4 August 2025

0926-5805/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

772–18:2011), or clay masonry (SIS-CEN/TS 772–22:2008) all exaggerate real-world conditions by fully saturating the specimen, as noted by Mensinga [14,15], who highlights this issue and criticized the omnidirectional nature of frost exposure in the experimental setup. Additionally, he suggests that bricks deemed less resistant in labs may outperform in reality due to uncontrolled moisture saturation. In reality, even one freeze-thaw cycle can cause irreversible volumetric changes to bricks, resulting in damage, making it unclear how many cycles are needed for observable deterioration. This also applies to other indicators like the modified winter index [16] or the freeze-thaw damage risk index [17]. Therefore, hygrothermal performance indicators are used to compare different building envelope designs as shown by [7] rather than predict degradation.

Even with hygrothermal design tools predicting minimized risk at the envelope design stage, localized façade damage can still occur due to microclimate variations. Recent literature suggests assessing multiple points across the façade [18], but this is often overlooked by assuming a single set of boundary conditions for the whole façade, especially for short-wave and night-time radiation. Wind-driven rain (WDR) variation are roughly considered, with technical standard (ISO 15927-3:2009) or the SB method [19] allowing differentiation of WDR over several façade geometries. However, these semi-empirical methods have large uncertainties due to limited building configurations. These uncertainties are further exacerbated in the urban environment [20], making the use of semi-empirical methods questionable for most buildings.

Understanding the microclimate around façades can provide insights into the degradation factors and improve conservation strategies. Indeed, research is advancing in this direction. For example, Charisi and Thiis, combined ray tracing and WDR quantification based on ISO 15927-3:2009 to account for the effects of small structural elements' effect on façade surface temperatures and relative humidities [21–23]. Computational Fluid Dynamic (CFD) simulation can be applied to complex façades and urban environments. Significant efforts have been made to validate CFD models for different building configurations with field measurements [24–27]. These models have been used in urban settings [20,25] and complemented by short-wave and long-wave radiation analysis [28]. Despite these advances, however, CFD-based modeling has its limitations. In particular, such modeling is complex, computationally expensive, time consuming and requires domain expertise, presenting barriers to wider adoption.

In summary, there is a need for a methodology that overcomes the limitations of conventional microclimate and hygrothermal modeling. This could involve studying the spatial distribution of degradation on façades due to microclimate variations and relating these findings to other façades. To achieve this, it is essential to objectively quantify façade degradation. In this regard, analyzing images collected by Unmanned Aerial Vehicles (UAVs, drones) using computer vision algorithms offers a promising solution.

1.2. Use of computer vision in façades

The application of computer vision to the analyze building façade deterioration is growing, particularly for automating anomaly detection through the UAV-collected images. Significant advancements have been made especially in the field of crack detection. Initially, research involved manual inspection of automatically generated orthophotos with zoom-in capabilities [29], as well as traditional computer vision techniques such as Sobel edge detection and thresholding [30–32]. While these methods provided initial automation, they struggle to distinguish edges from actual cracks from structural patterns like frames, leading to high false positive.

Consequently, research has shifted towards deep learning techniques, extensively used in civil engineering structures [33] and now applied to façades. For example, Chen developed a two-step neural network method involving a convolutional Neural Network (CNN) for patch-level crack classification and U-net model for pixel-level crack

segmentation to address the background noise [34]. This approach requires many training images, leading researchers to use pretrained CNNs for image feature extraction, optimized for the detecting degradation patterns through transfer learning. This concept has been applied to various degradation modes, including cracking [35], spalling, peeling, biological growth or delamination [36].

Research on brick façades has increasingly turned to pretrained CNNs for anomaly detection. Katsigiannis et al. [37] benchmark the performance of six pretrained CNNs including VGG16 and MobileNetV2, for detecting cracks in brickwork using limited training data. Wang et al. [38] used the faster region-based convolutional neural network (Faster R-CNN) to spalling and efflorescence on bricks with a focus on developing real-time automated vision-based inspection system for smartphones or IP cameras.

Other studies improve the façade inspection process by mapping results to spatial representations, aiming to enhance visual documentation for inspectors. For example, Torres-González et al. [39] used a simple thresholding technique applied to HSV and CIELab colour spaces to map the presence of salts, dirt and biological growth over the orthomosaic of cladded buildings. Regarding deep learning, Chen et al. [40] re-trained the You Only Look Once (YOLOv8) algorithm to detect degraded bricks in drone imagery and mapped the detections onto orthophotos constructed from 3D point cloud models. Similarly, Yang et al. [41] leveraged a combination of K-net and UPerNet to segment cracks, aiming to map them onto a BIM model of a building reconstructed from a point cloud.

Identifying crack features helps support visual inspections by providing more detailed information about the condition. Crack segmentation as used by Yang et al. [41] was also leveraged to estimate cracks severity by measuring length, width and geometric moment. Alipour et al. [42] presented a similar example of a deep learning application where the CrackPix algorithm was developed to detect and quantify cracks in concrete structures.

Common in previous works has been to support visual inspections and reducing manual efforts. However, the data generated through such CV-based detections is rarely used beyond anomaly identification. Specifically, there is limited exploration of how CV outputs could be leveraged for predictive modeling, such as estimating future degradation or microclimatic impacts.

1.3. Aim and scope

To address previous research limitations, this paper leverages computer vision anomaly detection to establish a methodology for evaluating deteriorations risks along façades. The proposed Prepare, Analyze, Integrate, and Relate (PAIR) framework aims to estimate microclimate effects by directly predicting façade degradation. This methodology involves dividing façade orthomosaic images into sections, pairing degradation distribution and its extent with features from three-dimensional models of façades and their surrounding area, and training machine learning algorithms to predict degradation on other façades.

A key feature of the PAIR methodology is its modularity. Rather than prescribing a specific CV algorithm, it allows practitioners to select models suited to their technical capacity, data availability and target degradation types. To support this, Section 2.2 (“Analyze”) presents an overview of relevant CV algorithms, including object detection and image segmentation methods. This overview builds on the literature review in Section 1.2 and is intended to guide practitioners and researchers in selecting appropriate CV tools for the application of the methodology.

The remainder of the paper is organized as follows: Section 2 describes the methodology in detail. Section 3 illustrates the methodology with a case study of 16 brick façades in Flatås area in Gothenburg, Sweden, and describes two experimental set-ups for testing the methodology's predictions. Section 4 presents the results for the two

experiments. Finally, sections 5 and 6 contain a general discussion and conclusion.

2. Methodology

At a high level, the PAIR framework consists of four sequential phases, visualized in Fig. 1.

- 1) **Prepare:** This initial phase involves the collection of façades' orthomosaic images followed by their division into sections. This step is crucial for the subsequent façade's section-by-section analysis and predictions.
- 2) **Analyze:** At this stage, the extent of deterioration in each section is quantified using computer vision algorithms. This step is essential for establishing a baseline from which machine learning algorithms can learn to predict deterioration.
- 3) **Integrate:** The next stage builds on the degradation data by incorporating features derived from a three-dimensional model of the façade and its surroundings. These features represent the spatial characteristics and the impact of the environment on each section, thereby establishing a basis for inferring the subsequent degradation predictions.
- 4) **Relate:** In the final stage, the obtained database is employed to train machine learning algorithms. This allows to relate the degradation characteristics identified in the database to other façades, consequently making degradation predictions.

Each phase comprises several steps that are categorized into two types: prescriptive and flexible. Prescriptive steps are mandatory and must be followed to ensure the framework's integrity. For example, the framework mandates the division of façades into sections for the organization of both empirical, spatial and climate data.

On the other hand, flexible steps provide the freedom to make choices based on the specific goal to which the framework is being applied, or the tool and resources available. These steps allow for customization and adaptability, so that the framework can be tailored to the unique contexts of other geographic areas other than Gothenburg, Sweden, which was later used in the case study. For instance, in the Analyze phase, the users can select their preferred degradation metrics, such as the length of missing veneers or the number of cracks, to analyze the extent of degradation. Users are also free to choose the computer vision algorithm that best suits their needs.

This combination of prescriptive and flexible steps ensures that the framework is both robust and adaptable. This provides a structured approach while accommodating case-by-case variations and preferences. The framework's steps, both prescriptive and flexible, are illustrated in a flowchart diagram in Fig. 2. The diagram is organized within the framework phases for which the examples of outcomes are provided.

2.1. Prepare

The methodology stipulates two requirements for orthomosaic images. Firstly, the images must contain only the area of the façade that is desired to be analyzed. No requirements on the resolution or the size of the image are given. Nevertheless, the second requirement is that the degradation over the façade must be clearly recognizable. The collection of the images may thus be performed by any means that is suitable.

The most straightforward approach is to take photographs from the ground. However, this approach may necessitate more extensive pre-processing, including perspective transformation, noise reduction (i.e., sharpening), upscaling or using deep super-resolution networks [43]. Nevertheless, even with the employment of the most comprehensive pre-processing techniques, ground photography may be rendered unusable. This is particular for spacious façades where the loss of details may be significant in areas far from the camera's position, making degradation unrecognizable. In such cases, the use of Unmanned Aerial Vehicles (UAVs or drones) is more suitable. Principally, images taken by a drone require little or no pre-processing, given their ability to shoot close and in normal direction against the wall. Regardless of the method used, the photographs must be stitched together and trimmed to form an orthomosaic. This may be achieved using almost any software for graphic editing. This paper in particular employs Microsoft Image Composite Editor in a case study presented in section 3 Case study.

After obtaining an orthomosaic of a façade, each image needs to be further divided into rectangular sections of equal size. The dimension of a section is arbitrary. However, the sectioning must be aligned with a three-dimensional model (3D) model of the façade to ensure an approximate agreement between their dimensions. More details about how to derive the 3D model can be found in section 2.3 Integrate.

2.2. Analyze

At this stage, it is necessary to quantify the deterioration of each façade section. Before proceeding, however, it is necessary to specify what degradation (e.g., cracking, spalling) to look for and to establish metrics to measure the degradation. At the elementary level, the degradation may be quantified in a binary fashion, i.e., whether a section is affected or not. Alternatively, it may be quantified on a continuous spectrum. The latter is represented by the area of degradation, which is a metric used for the assessment of façades [44]. It is hereby defined as the area of a section in percentage that is degraded. Alternatively, one can employ other performance metrics like the total length of missing veneers in a section, the number of cracks, or assessing the severity of cracks using their width and orientation.

In the case of accessibility to a CV algorithm that can meet the analysis objectives, one can proceed with the degradation quantification of each façade section. Otherwise, an appropriate (computer vision) algorithm must be trained. Regardless of the chosen quantification metric, a plethora of methods are available that are capable of either detecting degradation or segmenting its area. The utilization of classical algorithms such as thresholding [45], k-mean segmentation [46] or Canny edge detection [47] are relatively straightforward. However, it is important to exercise caution when applying these algorithms in practice, given their inherent constraints. To illustrate, consider the thresholding for crack detection. This process involves transforming all pixels on a greyscale representation of an image below an arbitrary threshold to white colour and vice versa, thereby highlighting cracks. However, thresholding sensitivity to lighting conditions and façade colors may necessitate modifying the threshold for different images. Therefore, its application on a larger number of façades may be impractical.

In contrast, more advanced computer vision models may require a higher level of expertise but may offer a more robust approach to degradation analysis [48,49]. There are different computer vision models that can be categorized in object detection algorithms and




Fig. 1. High-level overview of the PAIR framework for evaluating the effect of façades microclimate by predicting degradation along façades surface

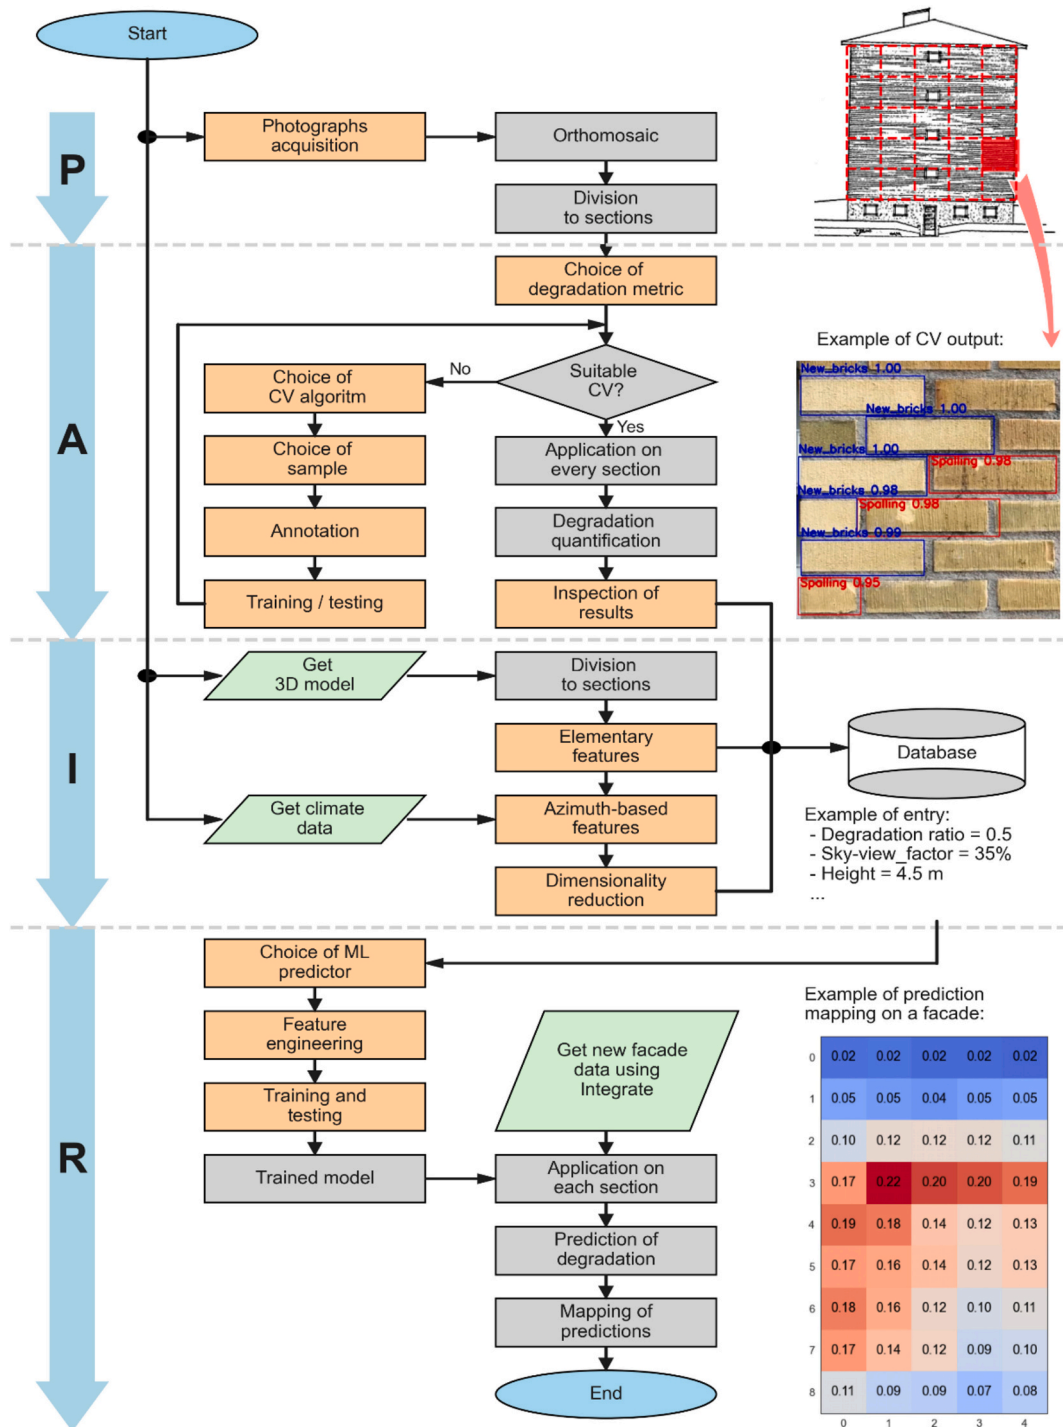


Fig. 2. Detailed overview of the PAIR framework. Flowchart steps are organized by the framework's phases, with flexible steps highlighted in orange. Images on the right illustrate the outcomes of each phase. "CV" in the Analyze phase for computer vision.

segmentation algorithms. The former can be used for detecting degraded bricks, while the latter is more suitable for identifying cracks and missing veneers between masonry. Examples of object detection algorithms are YOLO (You Only Look Once) [50], EfficientDet [51] and SSD (Single Shot MultiBox Detector) [52].

It should be noted that none of the computer vision algorithms mentioned are natively trained for detecting degradations on façades. However, it is possible to re-train them to perform this task. This was exemplified by Lee et al. [53] who retrained the Faster R-CNN [54] to detect delamination, cracks, peeled paint, and water leaks, on building

façades using 10,907 raw images. Nevertheless, such a vast number of images are unnecessary, as demonstrated by Wang et al. [38] who used only 500 images to retrain the same algorithm creating a model capable of detecting spalled and efflorescent bricks, thus illustrating the concept of transfer learning. Chen et al. [40] trained a YOLOv8-seg model using 895 high-quality images of masonry façades from Suzhou, China, mapping degradations across orthophotos constructed from 3D point cloud models.

Regarding image segmentation, one can either utilize algorithms that natively segment images, like U-net [55], FCN [56] or SegNet [57], or

combine the aforementioned object detection algorithms with other algorithms like vision transformers. YOLO-based algorithms are a popular choice among researchers in this regard. Cao [58] combined the YOLOv7 framework with BlendMask-based segmentation [59] to create YOLOM—a model that facilitates façade inspection tasks. Similarly, Fu and Angkawisittpan [60] leveraged YOLOv5 with Swin transformer [61] to classify and segment 5 types of degradation (cracks, delamination, plant growth, alkalization and staining) using 150 images for each type. Outside of YOLO family, Yang et al. [41] leveraged a combination of K-net [62] and UPerNet [63] to segment cracks, using 400 training images.

In addition to significantly reducing the number of required training images, transfer learning also adapts the used algorithm to the specifics of the analyzed façades—the same type of damage may appear differently on different bricks—potentially providing higher accuracy. Both the annotation and the training process can differ depending on the algorithm. For instance, each algorithm may require different hyperparameters to be tuned for optimal performance. Once the algorithm is tuned and applied to all façade sections, it is recommended to inspect the outcomes of the degradation analysis. This should be done before moving into the Integrate phase of the methodology.

This paper, respectively the case study showcasing the methodology, leverages the YOLOv4 framework to detect bricks subjected to spalling due to weathering. The framework was chosen based on the following criteria, which may be adopted by other assessors as well:

- 1) **Degradation metric definition:** To quantify the extent of degradation on each façade section, a variant of the area of degradation metric—degradation ratio—is chosen, considering only the presence of degradation but not, for example, its depth or severity. This requires the computer vision algorithm to count the total number of bricks and the number of degraded bricks in each section, prioritizing an object detection approach.
- 2) **Image characteristics:** An algorithm should be able to distinguish colors and textures on brick façades. In this regard, brick walls are generally highly variable. In consequence, an algorithm with strong generalization (CNN-based model) is preferred over simpler techniques like thresholding or contour analysis.
- 3) **Data availability:** The number of real images available for training is typically limited. Thus, the concept of transfer learning may be utilized by prioritizing algorithms that can be fine-tuned for the specifics of the case study.
- 4) **Performance and Robustness:** Existing and validated algorithms on similar tasks can be used with benefits. For this reason, an in-house model based on YOLOv4, capable of detecting spalled bricks, is selected and further fine-tuned to the case study using image samples from all analyzed façades.

Information on the training process of the algorithm and results of its application are given in [sections 3 Case study](#) and [3.1 Degradation analysis](#).

2.3. Integrate

The database consisting of the state of degradation for each section is further enhanced by the incorporation of environmental and spatial information. To achieve this integration, a 3D model of the façade and its neighborhood is required. The façade in the 3D neighborhood model must be segmented into sections that correspond to those defined in the Prepare phase. The environmental and spatial data is then defined for the center of each section. The collection of environmental information focuses on factors that are known to contribute to the occurrence and progression of degradation. These mainly include wind-driven rain, night-sky radiation cooling and solar radiation. The collection of the environmental and spatial factors specific to each façade section are hereby divided into elementary and azimuth-based features. The elementary features may include spatial factors such as the height from

the ground and the distance of the section's center from the edge of the façade, two of which are the influencing factors to the rain catch ratio. Additionally, the sky-view factor, which contributes to night-sky radiation cooling, is to be collected. Further data collection focuses on aggregating weather loads exposure on each façade section. This can be achieved by, but is not limited to, transforming weather factors such as mean wind speeds or cumulative wind-driven rain intensity data to e.g. 36 directions separated by 10° azimuths, thereby creating azimuth-based features.

To provide information on façade microclimate, azimuth-based features can be transferred further to account for neighboring buildings. This is done sequentially. Initially, distances from the center of each façade section to nearest object are measured in the directions of the azimuths. Secondly, the unobstructed wind and wind-driven rain impingement is assumed from azimuths in which the distances are 100 m or longer. Lastly, azimuth-based features are transferred using linear interpolation. This transformation process is illustrated in [Fig. 3](#).

In addition to the environmental and spatial data, general information about the façades can be collected, including the year of construction, year of renovation, or information about the composition of the façades. However, it is important to note that the absence of certain general characteristics like a year of construction may not be essential in the context of degradation predictions as other factors like the wind-drive rain exposure may prove to be sufficient or more influential. Moreover, the relative importance of a feature may vary depending on the machine learning algorithm employed.

The Integrate process may lead to the creation of many features and thus to a large feature space. Such database set-up may lead to a phenomenon known as the curse of dimensionality, a term coined by Richard Bellman [64]. As the number of dimensions increases, the volume of the feature space expands exponentially and so does the amount of data needed for effectively covering the space. Put differently, data points may become sparse, making underlying patterns in data undetectable. Therefore, it is recommended to reduce the feature space before advancing further.

Two common approaches for dimensionality reduction are Principal Component Analysis (PCA) [65] and binning, but other approaches may also be adopted. PCA involves the transfer of data onto a new coordinate system in the direction of the data's greatest variation. This method is invariant to the feature's type, and therefore, it can be applied to a part or the entire database at once. Binning, on the other hand, aggregates features of the same type onto bins. In the context of the methodology, it may be applied to azimuth-based features, whereby incoming wind-driven rain from two or more directions is combined.

The Integrate process adopted in the case study, which later demonstrates the methodology, produced three elementary features for each section of each façade, i.e. sky view factor, height above ground level and distance of the section center from the nearest edge of the façade. The wind-driven rain in each section was calculated for each 10° azimuth, based on the multiplication of the horizontal rain intensity with the mean wind speed coming from the corresponding direction and with the cosine transformation of the wind direction normal to the façade surface. The obstructions were accounted for as described above. Both binning and the PCA were adopted to reduce the azimuth-based WDR features, forming 9 binned WDR features, and 2 principal components. In addition, the sum and mean values for each façade sections were calculated from WDR from all directions. Finally, a façade constant feature, which categorizes the façades based on their material composition, is formed. More information on the specific implementation of the Integrate process is given in [section 3.2 Description of features](#).

2.4. Relate

At this final phase, features are utilized in machine learning algorithms to perform degradation predictions on façade sections.

The selection of relevant features is a crucial step in the construction

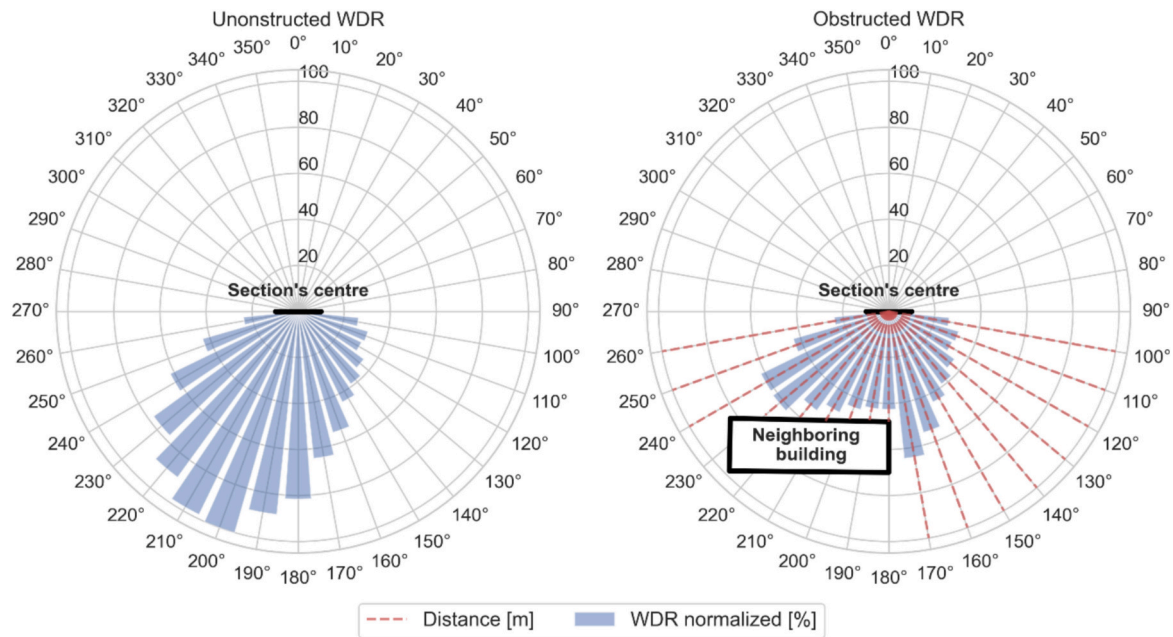


Fig. 3. Accounting for obstruction from buildings and other nearby objects in azimuth-based features. Left: unobstructed wind-driven rain at the center of a façade section. Right: obstructed wind-driven rain. Obstructions are accounted for by measuring distances from the section center to the nearest building at every 10° azimuth (red dashed lines). These distances are then interpolated considering a maximum of 100 m. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

of machine learning predictors. It serves to reduce the issue of overfitting, reduce computational costs and, most importantly, enhance the predictive performance of the model [66]. Common feature selection methods include filter, wrapper and embedded methods. In practical terms, the choice of the feature selection method depends on multiple aspects such as the dimensionality of data or the computational resources. Later in this paper a method for feature selection embedded in a random forest regression algorithm is employed in the case study. The selection of this ensemble algorithm is supported by the results of an initial ablation study conducted during the preparation process of this study, as well as its capacity to interpret feature importances, resist collinearity and overfitting, and model complex non-linear relationship.

To ensure that the selected features in a model generalize well to new unseen data, strategies such as hyperparameter tuning and/or cross-fold validation can be employed. Both strategies are generally used to reduce overfitting, which happens when the model learns the noise instead of the underlying patterns. The case study presented later in this paper combines both strategies with the objective of minimizing the risk of overfitting to the greatest extent possible. In practical terms, cross-fold validation involves splitting the dataset into K equally sized folds. $K-1$ folds are used for training the model, while the remaining folds are used to validate the performance. This process is repeated until all folds have been used for validation separately. The model's final performance is then assessed by averaging a performance metric (e.g., mean squared error (MSE), mean absolute error (MAE), coefficient of determination (r^2)) for all validations. All of this may be considered as one iteration within the process of searching for optimal hyperparameters of machine learning algorithms, specifically the random forest regressor used later in the case study. Typically, such a search is performed either by exhaustively trying all possible combinations of parameters specified in a grid to see which combination yields the best results, or by trying combinations of parameters randomly sampled from the grid. In the context of the case study described here, cross-fold validation is performed using r^2 as the performance metric combined with exhaustive grid search for hyperparameter tuning. More information regarding the specific use of the feature selection and the optimization is given in 3.3 Experimental set-up.

3. Case study

The methodology is further exemplified and tested by investigating a case study conducted on a sample of façades located in the Flatås area in Gothenburg, Sweden. The objective of the case study is to predict the extent of degradation, i.e., degradation ratio, across a façade. The sample consists of 16 façades with a common outer layer of yellow brick and rendered lightweight concrete basement walls beneath the brickwork. The façades in the sample all show signs of microclimate induced deterioration. Furthermore, they are all from a similar building stock in a single climatic zone, assuming the façades to follow a similar deterioration pattern. Strict inclusion criteria were applied to ensure that only façades predominantly exhibiting microclimate-related weathering—excluding those affected by factors such as water leakage or foundation movement—were selected, thereby providing a controlled environment for method evaluation; however, the limited availability of such suitable cases constrained the dataset size. Ten of these façades are constructed as non-ventilated brick façades. According to a technical documentation for one of the buildings, these façades are made of yellow hollow bricks, attached to lightweight concrete blocks. The remaining façades in the sample have a ventilated cavity between the outer layer of brick and the rest of the construction. On closer inspection, the ventilated façades can be further divided into three subcategories, each with a different type of yellow brick. The buildings were all constructed between 1954 and 1965. The selection of the façades was based on the presence of observed degradation. Additionally, tenants' privacy was considered by preferring parts of façades with no or small window openings. In total nine of the 16 façades that were investigated face south, two façades face north and west, respectively, and three façades face east. The position of each façade within the neighborhood is shown in Fig. 4.

In the preparatory phase, a drone survey was conducted for each façade using the DJI Mavic 3, which is equipped with 4/3 CMOS Hasselblad camera. The drone was operated manually, with photographs taken at a distance ranging 3–6 m from the façade surface. Furthermore, the photographs were taken on an overcast day to prevent the formation of shadows on the façades. The obtained photographs were then stitched



Fig. 4. Satellite image of the neighborhood with the façades studied in the case study highlighted. Façades 1–9 face south, 10–11 face north, 12–13 face west, and 14–16 face east. The satellite image was produced using Google Earth © 2018.

together using the Microsoft Image Composite Editor, resulting in orthomosaics of the façades. The division of each orthomosaic into sections was carried out in conjunction with a three-dimensional model of the neighborhood, obtained from the Digital Twin Cities Centre (DTCC) model [67], forming sections with an area of approximately 1 m². This step was performed using the Grasshopper and Ladybug plugins for Rhinoceros 7, specifically the LB ViewPercent function. Prior to

the division, however, the DTCC's model was manually modified to partially compensate for its limitations. In particular, roof overhangs were incorporated into the model. Moreover, basement walls were integrated into the model, with their respective heights determined on-site. The areas of façades subjected to further analysis were identified based on the delineation of the basement walls, the presence of larger openings, and the coverage of the examined façades by bushes in a few




Fig. 5. Slice of the 3D model of the neighborhood, showing the examined areas on the south-facing façades 1, 3, and 7, and the west-facing façade 12. Examined areas are highlighted in gold. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

cases. The areas of interest over the façades are illustrated in Fig. 5, which shows a slice of the 3D model of the neighborhood.

It should be noted that the division process did not result in squared sections in all façades due to their dimensions, the dimensional uncertainties associated with the height of the buildings in the DTCC model, and the way the function performs the division. Furthermore, the dimensions of the façades are not identical. The smallest examined part of a façade is 3 m wide and 8 m high whereas the façade with the largest area is 13.15 m wide and 11.5 m high. As a result, the number of sections in façades within the dataset varies from 24 to 132. A list of all the façades used in the case study, including information on orientation, dimensions, number of sections and whether the façade is ventilated or not, can be found in the Supplementary Data section, appended at the end of this paper.

3.1. Degradation analysis

To pinpoint the impact of microclimate on a façade based on observed degradation, it is necessary to consider previous renovations. Naturally, the newly renovated façade will display no degradations but may display the signs of the replaced brickwork, indicating that those bricks were previously damaged. This became clear upon examining the façades, where some bricks exhibited a noticeably brighter colour. Other areas of brickwork were also visually different from the rest. In such, however, the determination of whether they have been replaced was often uncertain. Therefore, prior to the degradation analysis, a survey was conducted among building owners to gather information on the history of façade renovations. Despite efforts to identify the locations of replaced brickwork and the dates of restoration, most attempts were unsuccessful. At best, only the date of the last replacement was determined. Consequently, the following analysis will focus solely on bricks that exhibit signs of physical degradation (such as spalling) and those that have been clearly replaced.

To assess the microclimate impact on façades, a ratio of the number of degraded bricks (including those that have been replaced) to the total number of bricks is employed to quantify the extent of degradation in each façade section. The analysis itself is conducted using a BuildSense model, which is an in-house object detection model based on YOLOv4, which can detect spalled and efflorescent bricks, but also identify missing veneers in a stone masonry wall [68]. Although newer computer vision algorithms like YOLOv11 or DETR may provide better accuracy and computational efficiency, the YOLOv4-based model was preferred due to its convenience. Specifically, the YOLOv4 architecture was configured with an input resolution of 416×416 pixels, batch size of 4, subdivisions of 4, and a learning rate of 0.0001. The training policy included a burn-in period of 400 iterations and learning rate step-downs at 1600 and 1800 iterations. Additional augmentations such as mosaic and HSV transformations (saturation, exposure, hue) were enabled via the training configuration. The model was retrained three times to produce three different models, one to detect spalled bricks, the second to detect replaced bricks and the third to count the total number of bricks in a section of the façade. The training process was conducted using a sample of 5 to 8 images from each façade in the dataset, hence training on 90 images in total. These were manually annotated using Label Studio [69] and YOLO format, with annotations also made for parts of bricks located on sections' edges. Image resolution varied across the dataset depending on drone positioning, with photos captured from distances ranging between 3 and 6 m from the façade. No manual image resizing or augmentation was applied; the YOLO training pipeline handled all preprocessing steps internally.

The images were divided into two sets, with 80 % allocated to the training set and 20 % to the testing set. Max batches, i.e., the maximum number of iterations the algorithm goes through, was set to 2000. The training took approximately 10 min on Intel(R) Core(TM) i9-10,900 CPU @ 2.80GHz with GPU NVIDIA GeForce RTX 2070 SUPER with 8GB VRAM and 64GB system RAM for each model. The achieved average

precisions were 99 %, 87 % and 90 % for detection of brick, replaced brick and for spalled bricks, respectively. For inference, a confidence threshold of 0.9 and an IoU (Intersection over Union) threshold of 0.4 were applied for non-maximum suppression (NMS). The models were then applied to all façades, iterating through all the sections, and the resultant class label, bounding box coordinates and confidence score were stored in a database with the inferred number of bricks, degraded brick and replaced bricks. This process is illustrated in Fig. 6.

Subsequently, a manual inspection of the analysis outputs was conducted to ensure the quality of the degradation data. This was done by mapping all detections over the orthomosaics. It was observed that in cases of two façades, the model for the detection of replaced bricks resulted in a large number of false positives, located mostly under roof overhangs, which could be a result of imperfect lighting conditions when the photographs were taken. These instances were erased from the database to ensure the quality of the target variable in the following predictions. The results from the analysis, i.e., the target variable for building machine learning predictors, are given in Fig. 7.

As visible in Fig. 7, the degradation distributions are, in general, highly skewed towards zero. Out of a total of 1171 façade sections, 672 (57.4 %) were classified as completely undamaged. Many other sections exhibit only a few detections. Degradation ratios below 0.05 were recorded in 909 sections (i.e., 77.6 %). Such observations are chiefly prevalent in north, west and east façades with only 26 sections exceeding the threshold. Degradation ratios of such outliers are below 0.2 except for one section in façade 16 where 24.3 % of bricks were classified as degraded/replaced. In contrast, the sections of the south-facing façades exhibit markedly higher levels of degradation. While the third quartile (Q3), which indicate the upper end of central 50 % of the points in a distribution, of the north, west and east-facing façades never exceeded the degradation ratio of 0.025, the third quartile of the south-facing façades exceeded this ratio on seven occasions, reaching a value of 0.05 or above. Furthermore, the upper bound whiskers, representing 1.5 times the interquartile range (IRQ) from Q3, are typically higher in south-facing façades. These results suggest that the most severe conditions in Flatås originate from the south. Fig. 8 shows the positioning and prevalence of the degradation ratio on façades oriented towards the south.

As illustrated in Fig. 8, the degradation is more pronounced towards the edge of the façades, indicated by the thick continuous lines. This phenomenon can be observed in all façades except for façade 9. Here, the bricks are mainly degraded towards the dashed line, after which the façade continues. Similarly, the lowest portions of the brickwork were also affected in most other façades. It is noteworthy that only façades 3, 6 and 8 have their bottom edge positioned within a half of a meter from the ground, which could indicate the presence of capillary suction. However, the lowest areas of façade 3 seem to be relatively unaffected by this compared to façades 6 and 8. The lowest edges of the other façades are positioned well above two meters from the ground, which suggests that the observed degradation may be influenced by rainwater runoff. The part of the façade underneath the brickwork is a render on lightweight concrete blocks which allows water to run off the façade. Furthermore, the façades tend to exhibit deterioration either around the middle or the upper third section of the examined area. This pattern is the most pronounced for façades 1 and 2. The former façade displays a distinct horizontal stripe of degradation spanning across the façade just above its midpoint, while the area above this stripe shows minimal or no degradation, probably due to the roof overhang. In contrast, the degraded areas on the latter façade do not exhibit such a distinct pattern but rather display degradation in clusters. These are again positioned above the midpoint of the façade. Interestingly, the most pronounced cluster is positioned much closer to the upper edge of the façade than in façade 1, despite the two façades having identical roof overhangs.

Notwithstanding the observed similarities in the façades' degradation patterns, it is evident that each façade can be regarded as uniquely damaged. For instance, façades 2 and 5 are almost identical yet the level

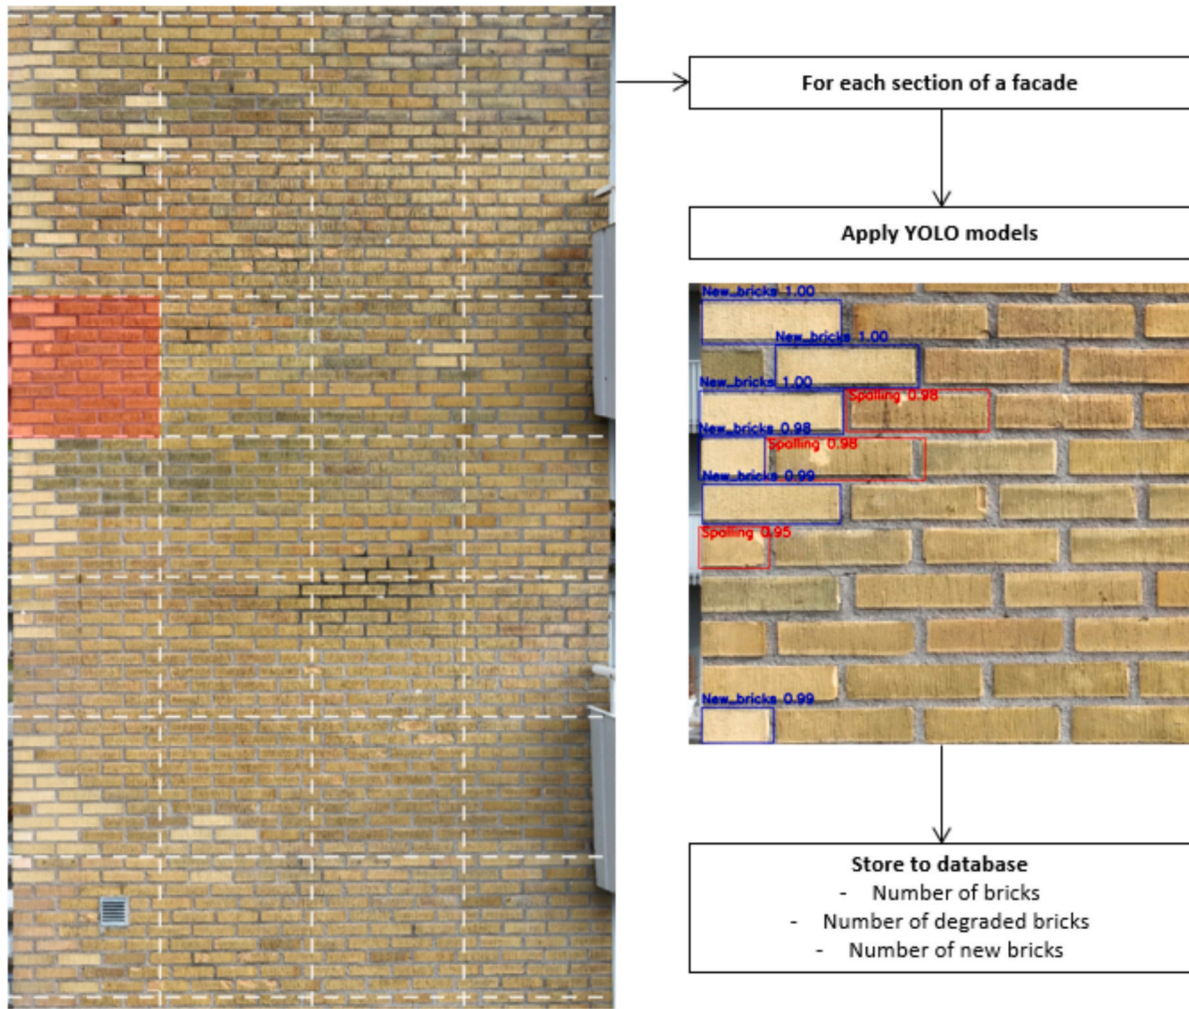


Fig. 6. Application of YOLOv4 model retrained to detect spalled and replaced bricks. The model is iteratively applied to all sections of each façade to determine degradation ratio of per section. Left: division of a façade into sections. Right: output of the YOLOv4 model. The example is presented on Façade 8.

of their deterioration is distinct. Furthermore, other façades may also be regarded as dissimilar to one another. A comparison of façade 1 or 4 to the others is illustrative of this point.

3.2. Description of features

The database, which now comprises degradation ratios for each façade section, is further complemented by descriptive features. The elementary features, namely height from the ground, sky-view factor and distance from the closest edge of the façade to section center, are derived directly from the modified DTCC's 3D model of the neighborhood. These are further complemented by a façade type, which is in the subsequent analysis referred to as the "Façade constant". This categorical variable serves to distinguish non-ventilated and the three types of ventilated façades.

The four elementary features are further complemented by an azimuth-based wind-driven rain (WDR) that considers obstruction from neighboring objects. The generation of such is further instituted by the past modelled hourly weather data for Gothenburg, Sweden, under the validated climate scenario A1B. The projection contains weather data for the period between 1960 and 2023 covering the exposure for most of the service life for the façades in the dataset. The wind and rainfall data in the winter months, which are used further for the generation of features, are summarized in Fig. 9.

As can be observed, the intensity of the exposure is amplified on the

south and south-west sides. The predominant wind direction is between 200° and 210° azimuths, while the greatest rainfall intensity occurs between 200° and 240° azimuths. The calculation of WDR-based azimuth features on each façade section is conducted utilizing rainfall intensity data in the winter months. The cosine rule, which is a fundamental principle in any commonly known semi-empirical WDR method, is employed to project precipitation on the façade surfaces. The precipitation projections are further compounded to 36 directions separated by 10° azimuths and multiplied by mean wind speeds from corresponding direction, which are accounted for neighboring obstructions as described in section 2.3. The resulting 36 WDR features are further reduced using binning and principal component analysis. The former aggregates WDR to bins set apart at 40° intervals, while the latter reduces the 36 dimensions to 2 principal components. Finally, the feature space is completed by aggregating the mean and sum of WDR. The total number of features for building a machine learning predictor of degradation is thus 17. Spearman's correlations for each of these 17 features are shown in Fig. 10. These are presented for individual façades, for façades within their respective orientations and for the entire dataset.

The features display notable variability in their correlations with the degradation ratios across different façades. To illustrate, the sky-view factor in south-oriented façades ranges from a strong and positive correlation ($r = 0.79$) in façade 1 to a negative and negligible correlation ($r = -0.13$) in façade 7. A similar variability is evident across other




Fig. 7. Degradation assessment of all façades using the YOLOv4 model to detect spalled and replaced bricks. Each point represents the degradation ratio estimated in one façade section. Boxplots indicate the degradation distribution across each façade by showing its median value and interquartile range.

features. Furthermore, features demonstrate moderately strong correlations at best when façades are sampled to their respective orientation groups. However, correlations for the entire dataset showed only weak relationships between the features and the degradation ratios.

3.3. Experimental set-up

The objective of the case study is to construct a machine learning predictor that can predict degradation ratios across façades. The performance of the predictor is to be evaluated using two holdout samples. The first sample is produced by excluding data for one entire façade,

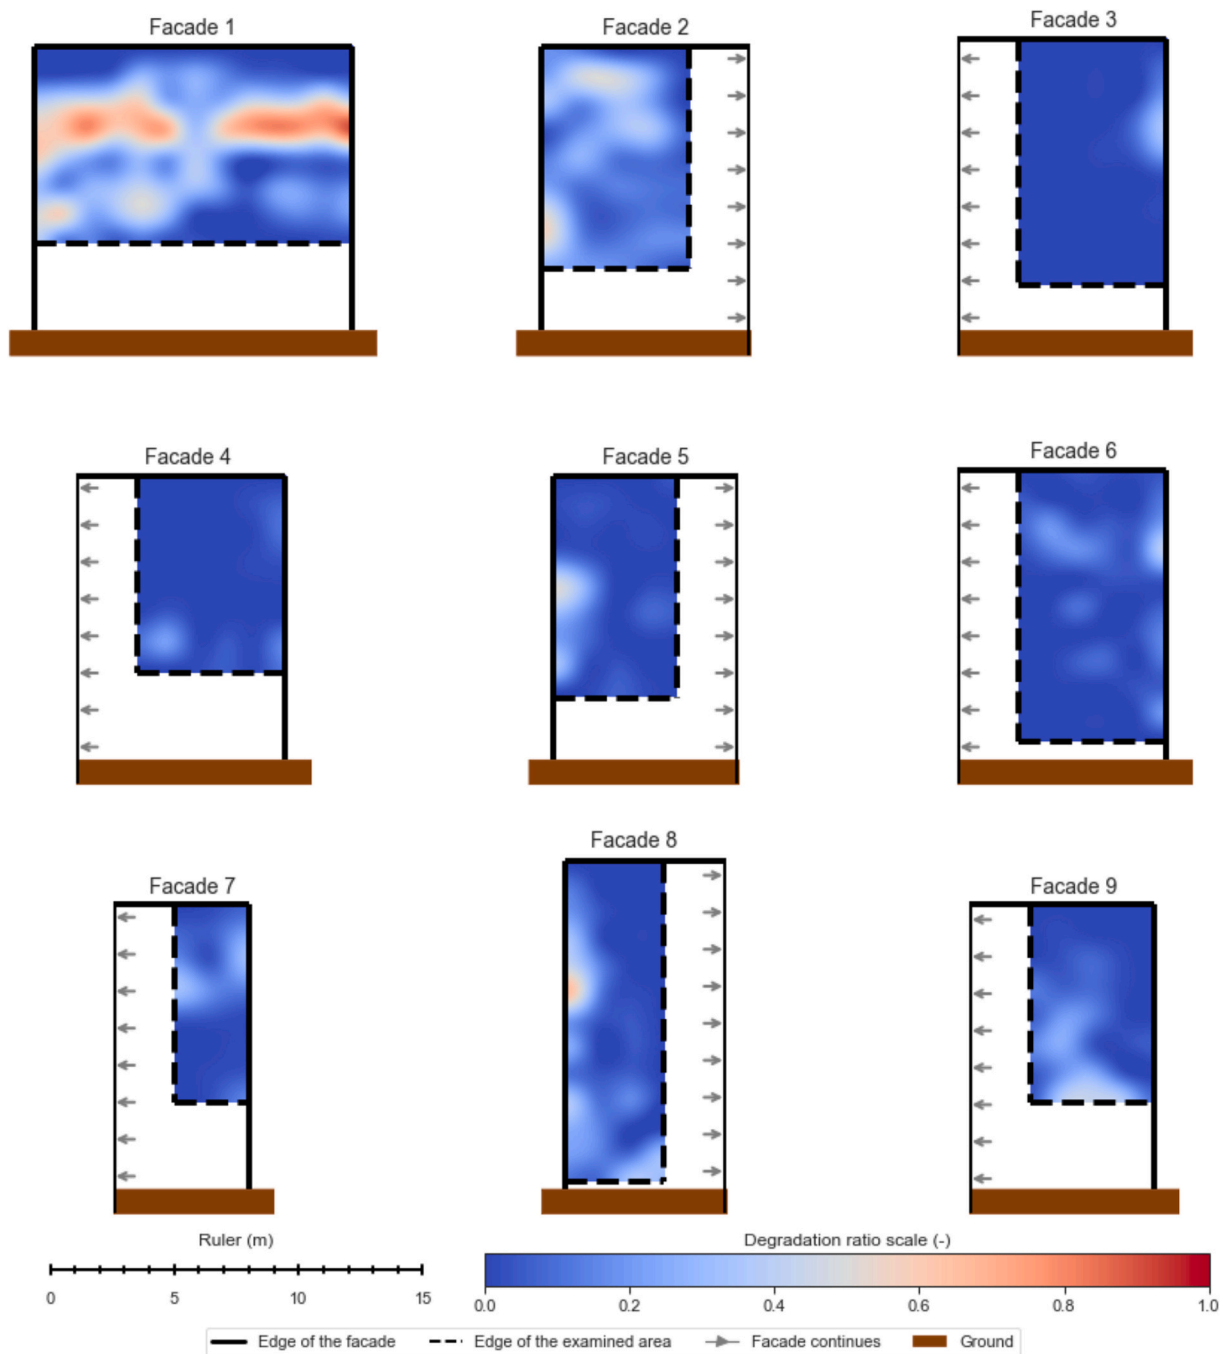


Fig. 8. Positioning and prevalence of the degradation ratio on south-facing façades. Heatmaps were generated by upscaling degradation data arrays (sections) by a factor of 100 using spline interpolation. The façades are depicted to their actual scale.

which is hereinafter referred to as the validation façade sample. The second sample is generated through a process in which the data are augmented by adding random Gaussian noise to the degradation ratios and adjusting augmented values that fell below zero to zero, randomly shuffled and then divided according to the 80:20 convention. The smaller portion of the data is the holdout sample hereinafter called the test sample. The latter constitutes the training data for the algorithm.

The division of the cases study façades in different samples aims to test two different aspects of the model's performance. The test sample is used to inquire which features and how well those features may inform the degradation prediction across façades. It assesses the model's generalization to unseen data from the same degradation distribution as the training data. Conversely, the validation façade tests the model's

ability to generalize to unseen data which may be from a different degradation distribution, as hinted at in Section 3.1, where all façades were regarded as exhibiting unique damage patterns across their surfaces.

For validation, Façade 2 is hereafter selected as the validation façade. As Fig. 7 and Fig. 8 show, it displays the second highest degree of degradation after façade 1. Unlike Façade 1, however, which exhibits degradation levels that significantly exceed those observed in other façades, and which may be regarded as an outlier whose extent of degradation cannot be drawn from the data of the remaining façades, high degradation ratios in Façade 2 might be predicted based on Façade 1. This makes the prediction of degradation levels across Façade 2 a challenging but feasible undertaking. Fig. 11 shows a photograph of




Fig. 9. Wind and rain data for Gothenburg, Sweden, based on the A1B projection for the period 1960 and 2023. Left: wind-rose diagram. Right: heatmap showing the cumulative sum of horizontal rain intensity binned in 40-degree intervals. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Façade 2.

The remaining data, i.e., all data without Façade 2, are augmented by a factor of 10 as explained at the beginning of this section, resulting in a training and testing dataset, comprising 9008 and 2252 datapoints, respectively. The validation façade encompasses data for 45 sections. All portions of the dataset were standardized i.e., each feature was scaled by the removal of its mean value and subsequent division by its standard deviation. This is done to ensure uniformity in the contribution of the features to the prediction model.

The prediction model was constructed using a random forest regression algorithm as its core component. Random forest was selected based on the results of an initial ablation study that included alternative models such as XGBoost, Support Vector Regression (SVR), and a basic feed-forward neural network with random forest achieving the highest r^2 and lowest Mean Square Error (MSE). Random forest was ultimately chosen for its robustness to collinearity and overfitting, ability to capture complex non-linear relationships, and its capacity to provide feature importance insights, aiding in understanding key predictive variables. This last quality is used to select features from the feature space. This is done by constructing a random forest regressor comprising 100 decision trees, which are used to estimate the relative importance of each feature. Subsequently, only features with an importance score above a threshold, derived from the mean importance of all features, are kept. The automatic feature selection, hereinafter referred to as Experimental set-up 1, is augmented with the manual selection of features. This is primarily to ascertain the impact of varying feature selections on the predictions of the validation façade, which, as previously suggested, exhibit different degradation patterns compared to the rest. This is further referred to as Experimental set-up 2.

To reduce the possibility of overfitting, hyperparameter tuning conjunct with cross-fold validation is adopted. The former was carried

out using a grid search approach, which optimizes the model's performance by exploring a range of hyperparameter combinations. Cross-fold validation, with k-folds equal to 5, was employed throughout the process to ensure the model's robustness across different subsets of the training data. The r^2 scoring method was employed to train the model, with the objective of emphasizing the model's ability to explain variance in the target variable.

The hyperparameter tuning process considers five distinct hyperparameters, which are deemed the most important when constructing the random forest regressor. The number of trees, as the name implies, controls the maximum number of trees that can be included in the forest. Its default value is 100 trees. Additionally, the tuning process entails the examination of 50 and 200 trees. The second hyperparameter is the maximum depth which defines the maximum number of computational nodes for each tree. The default value 'None' for this parameter, which allows trees to expand until all leaves contain less than a minimum number of samples that are required to split a tree node (i.e., the third hyperparameter that is considered), is complemented by values 10, 20 and 30. For the third hyperparameter, the values of 2 (the default value), 5 and 10 are used in the optimization process. The fourth hyperparameter is the minimum number of samples that is required to form a leaf node. Its default value is complemented by values of 2 and 4. Finally, the maximum number of features to consider when searching for the best split is considered in the tuning by considering all n features, the square root of n features, and the binary logarithm of n features.

4. Results

The case study example of the methodology so far has progressed through three of the four phases, generating degradation data and descriptive features for each façade section. These are used in the final

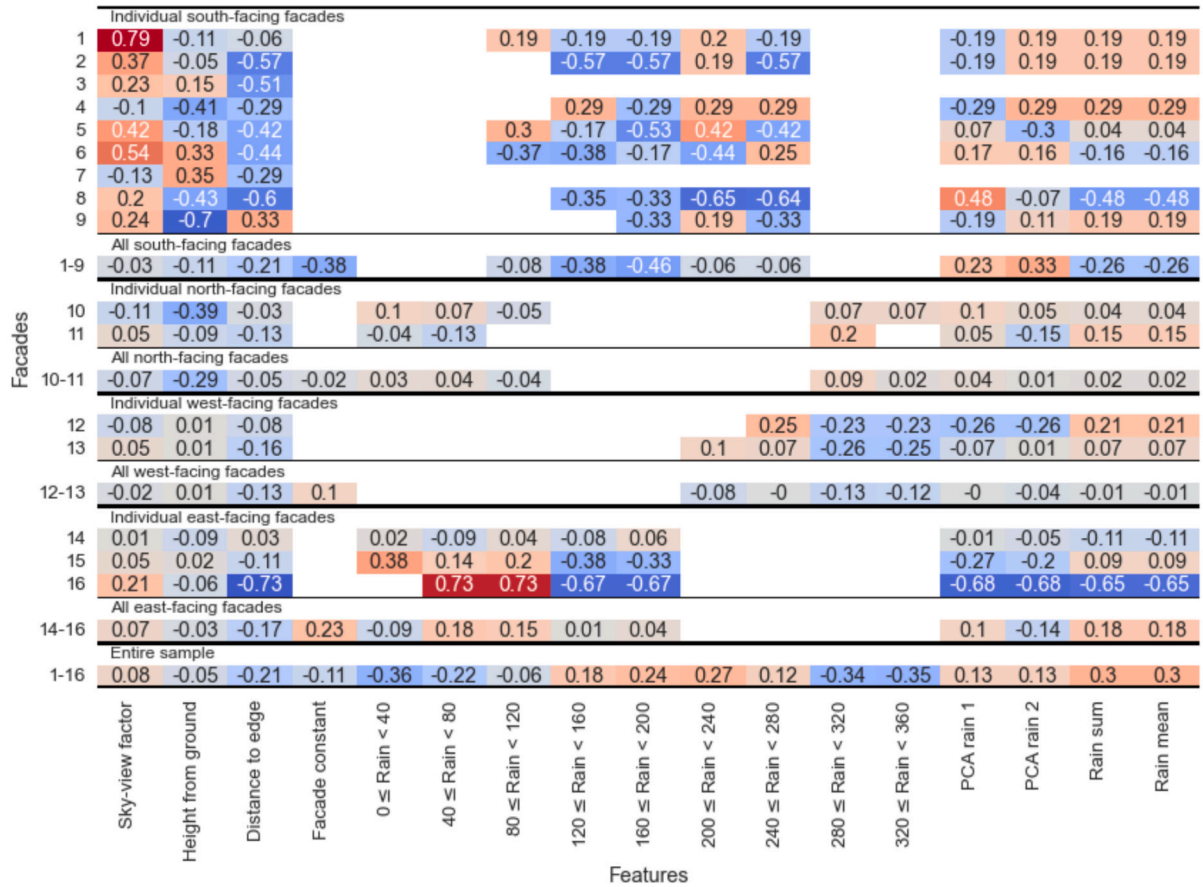


Fig. 10. Spearman correlation of features with the degradation ratio across façades. Note, in the instances where there were gaps, the correlation coefficient was not defined due to the presence of constant inputs.

Relate phase of the methodology demonstrating the predictive capabilities of the methodology. The outcomes in this section are presented in the context of the two experimental set-ups, which have been defined in [section 3.3](#) Experimental set-up.

4.1. Experimental set-up 1 – Automatic feature selection

The relative importance of features, the selection threshold, and features selected for constructing the degradation predictor are shown in [Table 1](#). The feature selection threshold (0.059) was derived by calculating the mean of the relative importance of all features.

Five features comprising the three elementary ones i.e., sky-view factor, height from the ground and the distance from the closest corner, and the two principal components derived from rain azimuth-based rain features, are included in the following analysis. The training process resulted in a set of hyperparameters which are shown in [Table 2](#) alongside with r^2 and MSE scores on the hold out samples.

The model's predictions on the test sample resulted in both high variance ($r^2 = 0.978$) and low error ($MSE = 0.0002$), indicating a good predictive capability. In contrast, the metrics exhibit a lack of predictive power the validation façade ($r^2 = -1.152$; $MSE = 0.0332$). These results are illustrated in [Fig. 12](#), which depicts a comparison between the actual degradation ratios in sections and the predicted ratios for both holdout samples.

In the case of the test hold-out sample, essentially all predictions are in proximity to the red dashed line, which represents the ideal best possible fit between the actual and predicted degradation ratios. Nevertheless, predictions over the validation façade are generally low, with the highest degradation estimate 0.12. Predictions over most of the sections, however, are close to 0. The disparity over the validation

façade is further illustrated in [Fig. 13](#).

The actual spatial distribution of the degradation ratios over the validation façade, as determined through computer vision analysis, demonstrates higher ratios towards the edge of the façade on the left and above its center. The highest ratio (0.54) is located near the bottom left corner. The ratios get progressively lower further away from the façade's edge, except for the area above the center i.e., rows 1–3. In this region, the degradation ratios initially increase, but then, at the right edge of the brick portion of the façade, undergo a sudden decline towards zero. In contrast, the predicted spatial distribution exhibits no such pattern. The sections above the center, i.e., rows 0–3, exhibit minimal variation. The only discernible pattern is observed from row 4 onward. In this area, the left edge of the façade once again exhibits the highest degradation ratios with the maximum of 0.12. The ratios then decline progressively but are elevated at the far-right sections.

4.2. Experimental set-up 2 – Manual feature selection

This experimental set-up involves testing different feature combinations to evaluate their impact on the predictions over the validation façade. Three elementary features from Set-up 1, i.e., the sky-view factor, height from the ground and the distance from the closer edge, are kept constant throughout the analysis. The focus, however, was on the azimuth-based features, which were modified in each iteration. To illustrate the impact of these features on the predictive performance, three examples are provided: (1) the feature $0 \leq \text{Rain} < 40$, (2) $200 \leq \text{Rain} < 240$, and (3) combination of two azimuth-based features i.e., $0 \leq \text{Rain} < 40$ and $160 \leq \text{Rain} < 200$. To further analyze the features' impact on predictions, the experiment is at the end repeated with the inclusion of the "Façade constant" in the feature space. Although this feature



Fig. 11. Photograph of Façade 2 that is hereafter regarded as the validation façade for testing the ability of the predictor to generalize to unseen data, which may be drawn from different degradation distribution. Right: detailed view of the area highlighted with a red mask in the photograph, showing bricks undergoing spalling (lighter-colored areas inside bricks). Bricks with a lighter overall appearance, mostly located on the left side, are considered as replaced bricks in further analysis. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1

Assessment of preliminary feature importance as estimated by the Random Forest regression algorithm. A selection threshold of 0.059 results in the inclusion of the top five features.

	Feature	Importance
1	Sky-view factor	0.311
2	PCA rain 2	0.219
3	PCA rain 1	0.155
4	Height from ground	0.124
5	Distance to closest edge	0.060
6	$160 \leq \text{Rain} < 200$	0.030
7	$200 \leq \text{Rain} < 240$	0.029
8	$240 \leq \text{Rain} < 280$	0.018
9	$0 \leq \text{Rain} < 40$	0.015
10	$120 \leq \text{Rain} < 160$	0.014
11	Rain sum	0.009
12	Rain mean	0.008
13	$80 \leq \text{Rain} < 120$	0.004
14	Façade constant	0.003
15	$40 \leq \text{Rain} < 80$	0.001
16	$280 \leq \text{Rain} < 320$	0.000
17	$320 \leq \text{Rain} < 360$	0.000

Table 2

Model's estimated training parameters. Included are r^2 and MSE metrics achieved on the test sample and the validation façade.

Technique	Hyperparameters	Optimal Values	Test sample r^2	Test sample MSE	Val. façade r^2	Val. façade MSE
Random Forest Regression	Number of trees	200	0.978	0.0002	−1.152	0.0332
	Maximum depth of each tree	N/A				
	Minimum samples required to split a node	2				
	Minimum samples required at each leaf node	1				
	The number of features to consider	sqrt				

showed low estimated importance in Set-up 1 (Table 1), its inclusion is justified as the presence of ventilated cavity is known to have significant impact on the hygrothermal performance of façades and thus the occurrence of degradation. Moreover, its importance within the random forest regression might be amplified when combined with feature subsets. Regardless of the example, the hyperparameters were optimized and the algorithm trained as described in section 3.3 Experimental set-up. Table 3 shows r^2 and MSE score for both hold-out samples for the three examples without the use of 'Façade constant' in the feature space.

The model, as presented in Table 3, demonstrated comparable performance across all instances in the test hold-out sample. Additionally, both the r^2 and MSE scores were almost identical to those recorded in the Experimental set-up 1, where the features were automatically selected by the model. However, the metrics for the validation façade exhibit insufficient performance. This is specifically reflected in the negative r^2 scores, which were achieved in all instances. The performance of all three examples is graphically shown in Fig. 14.

In accordance with the established metrics, the impact of selecting different features on the model's performance is negligible in the test hold-out sample. However, the utilization of different features alters predictions over the validation façade. The model, in general, tends to

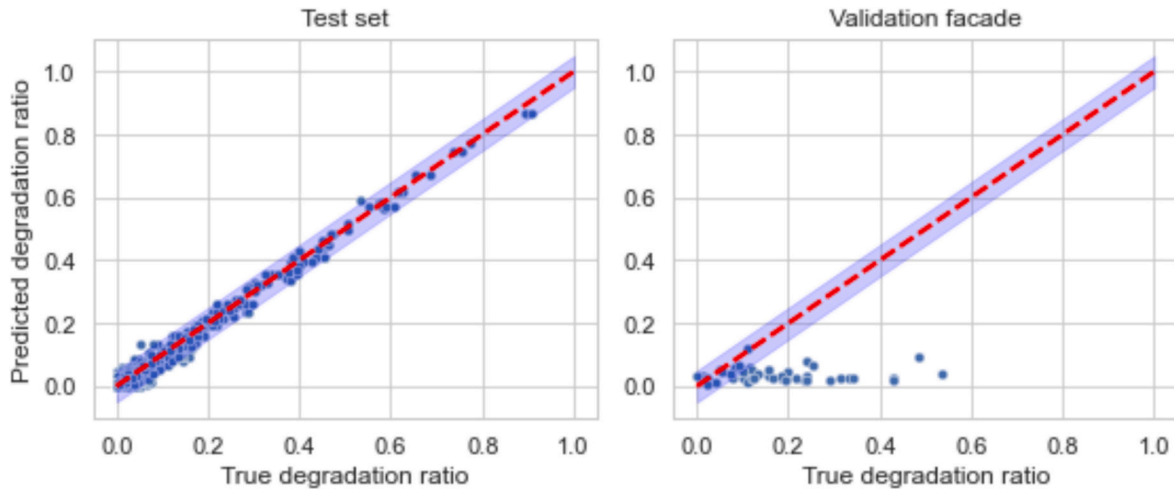


Fig. 12. Application of the Random Forest regressor on the test set (left) and on the validation façade (right), utilizing features selected in Table 1 - illustrating the comparison between predicted and actual degraded/undegraded ratio in each section of the façade. The red dashed line represents the theoretical best possible agreement between the predicted and actual values. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 13. Spatial comparison of actual degradation based on computer vision analysis (left) and predicted degradation ratios (right) over the validation façade.

predict low degradation ratios. The estimated maximum is 0.3 (feature subset with $200 \leq \text{Rain} < 240$, in the middle). Inherently, the predictions at the lower end of the true degradation spectrum demonstrate a higher degree of precision. This is particularly evident in the feature subsets on the left and right, specifically $0 \leq \text{Rain} < 40$ and $160 \leq \text{Rain} < 200$, where most predictions fall within a 5 % deviation from the optimal fit. Nevertheless, as the true degradation ratios approach a value of 0.2, the prevalence of underestimated predictions increases across the examples. The results for the Validation façade are further complemented in Fig. 15, which depicts the spatial distribution of the predictions.

The predictions based on different subsets of features exhibited distinct spatial patterns over the validation façade. The use of the range $0 \leq \text{Rain} < 40$ (on the left) produced a pattern that most closely resembles the true degradation distribution (Fig. 13). However, the spatial comparison deviates particularly in the upper parts of the façade, i.e., rows 0–2, which exhibited a high prevalence of degradation with the maximum of 0.43. Conversely, the predictions are underestimated with the recorded maximum of 0.11. In general, the degradation is more prevalent on the left side of the façade and in row 3 where the predicted maximum (0.29) occurs. This pattern is also recognizable in the case to the right, which employs two azimuth-based features, i.e., $0 \leq \text{Rain} < 40$,

Table 3

Model performance for three different subsets of azimuth-based features. Each subset includes three elementary features, namely the sky-view factor, height from the ground and distance from the closer edge.

Subset with following azimuth-based features:	Test sample r2	Test sample MSE	Validation façade r2	Validation façade MSE
$0 \leq \text{Rain} < 40$	0.978	0.0003	-0.467	0.024
$200 \leq \text{Rain} < 240$	0.975	0.0003	-0.948	0.031
$0 \leq \text{Rain} < 40$, $160 \leq \text{Rain} < 200$	0.978	0.0002	-0.246	0.019

$160 \leq \text{Rain} < 200$. In this instance, however, the degradation ratios steadily increase from the bottom right corner towards the left corner of row 3, with a maximum value of 0.25 occurring on column 1. This progression is also evident in the initial case, although the transition is less pronounced. The final case (in the middle), which employs the feature $160 \leq \text{Rain} < 200$, has larger differences. The predicted pattern resembles the letter “H”, with a more pronounced right side. In comparison to all the other cases, the model in this instance predicted the highest degradation ratios (0.3) on the right side of the examined wall i.e., in the furthest distance from the edge of the façade.

To illustrate the influence of the “Façade constant” on the predictions over the validation façade, Fig. 16 shows the spatial distribution of the

degradation over the façade for the three examples.

In general, the addition of the “Façade constant” in the feature space kept its excellent r2 scores for the test hold-out sample but inherited and further amplified the inaccuracies from the original results (Fig. 15). Specifically, their r2 scores has worsened across cases with -0.585, -1.091 and -0.317 for the cases $0 \leq \text{Rain} < 40$, $200 \leq \text{Rain} < 240$, and $0 \leq \text{Rain} < 40$ and $160 \leq \text{Rain} < 200$, respectively. The drop in performance is especially noticeable in the case on the left ($0 \leq \text{Rain} < 40$) and in the middle ($200 \leq \text{Rain} < 240$). The use of the $0 \leq \text{Rain} < 40$ (depicted on the left) produced a pattern over the façade that, on its left side, resembles the true degradation distribution (Fig. 13). This resemblance, however, falls apart particularly on the right side of the façade where the predictions at the edge of the examined area (column 4) were significantly overestimated. There in row 3, which by itself forms a stripe of sections with elevated predicted degradations, the maximum degradation of 0.42 was predicted. Another notable deviation is in the upper parts of the façade, i.e., rows 0–2, which exhibited a high prevalence of degradation with the maximum of 0.43. Conversely, the predictions are underestimated with a recorded maximum of 0.20. A similar pattern can be seen in the middle case ($200 \leq \text{Rain} < 240$) with the difference that the predicted degradation ratios are generally lower compared to the previous case. This is invalidated only in column 4 where some predictions are higher with the maximum of 0.38 at rows 4 and 5. To some extent the pattern is observable also in the case to the right, which employs two azimuth-based features, i.e., $0 \leq \text{Rain} < 40$ and $160 \leq \text{Rain}$

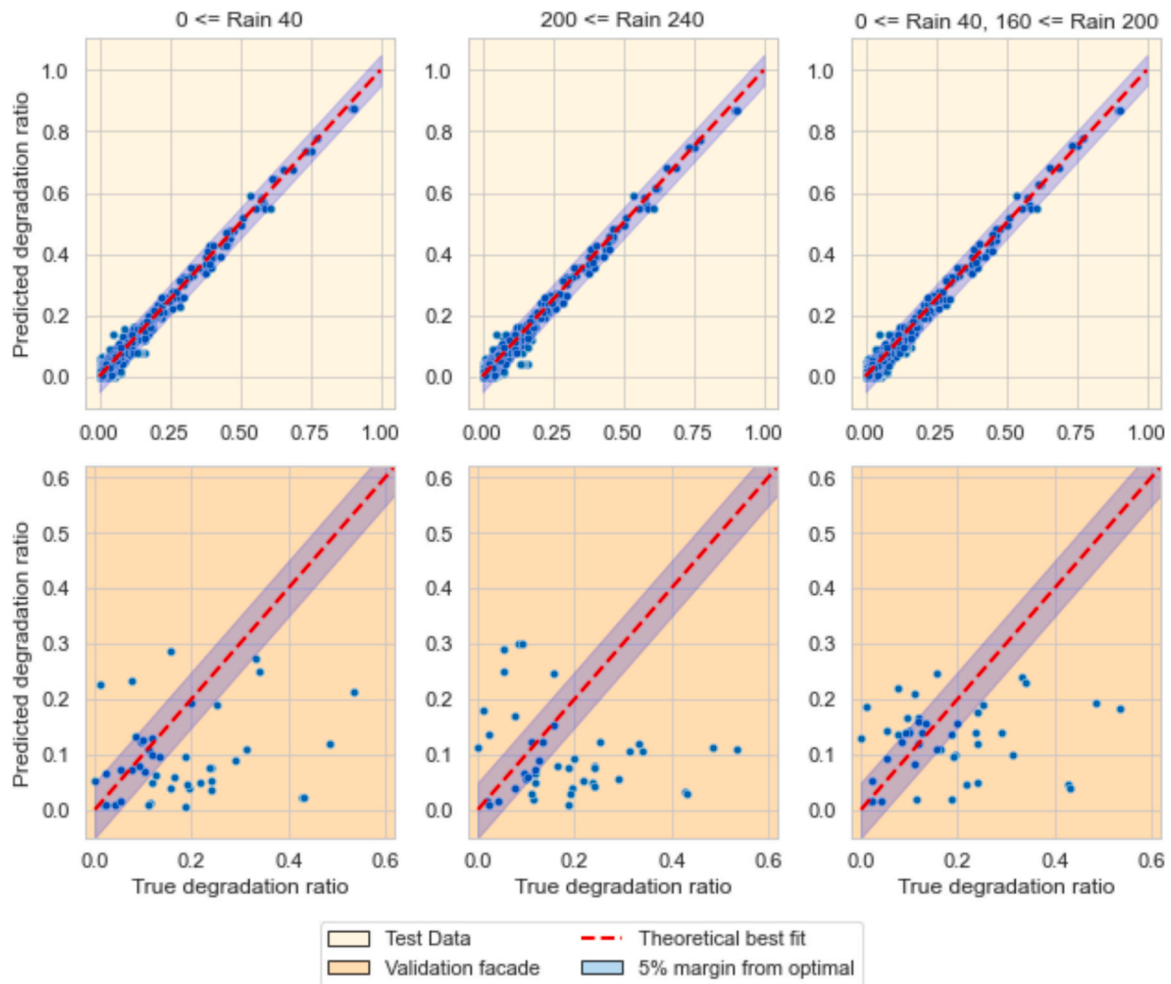


Fig. 14. Application of the Random Forest regressor on the test set (first row) and on the validation façade (second row), utilizing three different subsets of azimuth-based features: $0 \leq \text{Rain} < 40$ (left), $200 \leq \text{Rain} < 240$ (middle) and $0 \leq \text{Rain} < 40$, $160 \leq \text{Rain} < 200$ (right). Each subset is accompanied by three elementary features: sky-view factor, height from the ground and distance from the closer edge. The red dashed line represents the theoretically best possible agreement between predicted and actual values. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 15. Spatial distribution of degradation predictions for three subsets of azimuth-based features: $0 \leq \text{Rain} < 40$ (left) $200 \leq \text{Rain} < 240$ (middle) and $0 \leq \text{Rain} < 40, 160 \leq \text{Rain} < 200$ (right). Each subset is accompanied by three elementary features: sky-view factor, height from the ground, and distance from the closer edge.




Fig. 16. Additional influence of the "Façade constant" feature on the spatial distribution of degradation predictions for three subsets of azimuth-based features: $0 \leq \text{Rain} < 40$ (left), $200 \leq \text{Rain} < 240$ (middle), and $0 \leq \text{Rain} < 40, 160 \leq \text{Rain} < 200$ (right). These are accompanied in all cases by three elementary features: sky-view factor, height from the ground, and distance from the closer edge.

<200. In this instance, however, like the original case in Fig. 15, the degradation ratios increase steadily from the lower right corner to the left corner of row 3, with a maximum value of 0.24 occurring in column 1.

5. Discussion

The two experimental set-ups provided key insights. Firstly, the methodology effectively identifies areas under significant environmental stress, likely to degrade. This is supported by the test hold-out sample results, where the random forest regression model estimated most degradation ratios with a margin of error below 5 %. However, applying the model to the validation façade highlights a crucial caveat for practical implementation. The test hold-out sample was from the same distribution as the training data, while the validation façade, which was set aside for later testing, came from a different distribution. Thus, inaccurate predictions in the validation likely stem from distribution shift, as often noted in the literature.

In essence, the distribution shift results in the model's insufficient generalization to the unseen data. To illustrate this one might consider the case of a self-driving car that has been trained and successfully tested in the streets of Gothenburg, Sweden, and then deployed in unfamiliar and narrower streets of Thessaloniki, Greece. This is analogous to the situation with the validation façade. The model becomes familiar with, i. e., overfit to, the training and test sets, which demonstrate certain characteristics. However, the model underperforms in the validation façade, exhibiting degradation patterns and characteristics that are likely distinct from those observed in other façades. This coincides with the observation made in section 3.1 Degradation analysis where the degradation patterns across each façade in the dataset were regarded as unique upon observation of Fig. 7 and Fig. 8. In other words, the model learned patterns from the training set that are probably not present in the validation façade. Consequently, the selection of distinct features also resulted in disparate degradation predictions for the validation façade. This was, among others, illustrated with the feature "Façade constant", which categorizes façades based on whether they have a ventilated cavity or not. Although this feature is considered significant from the building physics perspective, it had a negative impact on the prediction of degradation over the validation façade. In conclusion, a feature may prove highly predictive in most façades but irrelevant or even misleading when applied to a façade with its distinctive degradation patterns.

Addressing the distribution shift is essential for the further development of the methodology. In this context, the shift is primarily due to the inadequate assessment of the façades' degradation, which does not fully incorporate their maintenance history. The majority of the buildings in the dataset were erected at the beginning of the 1960s, with one exception, which was constructed in 1954. Despite the long existence of the buildings, the efforts to ascertain the history of façade renovation yielded inadequate information. The most dated renovation effort recorded in the survey was undertaken in 2011. In such cases, bricks of markedly brighter colour indicate the recent replacement. However, a closer examination of a few façades in the database reveals other areas of brickwork, which appearance also deviates from the rest. In such situations it is unclear whether these bricks were replaced, or whether their distinct visual characteristics are caused by other factors such as the deviation of material properties coming from a different batch of bricks. Therefore, for practical reasons, the degradation analysis considered only damaged bricks and bricks that have markedly brighter colour. However, this approach probably imposed a significant limitation as the outcomes of the degradation analysis may have been more similar across the façades.

It is also necessary to consider other factors that may contribute to the distribution shift. It can be argued that the differences in environmental exposure across the façades have been considered by including the most influential factors; however, it is possible that others could be

added in order to capture the exposure more comprehensively. For instance, including additional factors in the feature space such as sun exposure and shading, or the inclusion of other objects than buildings, such as trees, which can affect frost exposure in early spring, could have led to a more accurate inference of façades' degradation. Similarly, the more comprehensive description of façade material might have helped to explain the disparities between the degradation patterns. However, given the excellent performance of the models on the test-holdout sample, incorporating maintenance history in the degradation analysis is deemed more important.

To address the issue of the distribution shift, two distinct approaches are worth considering. The first approach involves expanding the training dataset. The current database comprises 16 façades from one neighborhood, which were expected to show similar degradation patterns as they come from similar building stock in a single climate zone. Arguably, many more façades can be added to the dataset, as there are numerous similar buildings throughout Sweden. It may be the case that capturing more façades that exhibit similar degradation patterns to the validation façade could diminish the gap caused by the distribution shift. However, the incorporation of new façades should be well considered as façades may be damaged for reasons other than micro-climate weathering (e.g., water leakage, moving foundations), breaking the integrity of the data. Moreover, it is likely that the inclusion of several more façades will not narrow the distribution gap, but the inclusion of hundreds of brick façades, which exceeds the scope of this paper, may achieve this objective.

The second approach focuses on enhancing the degradation analysis itself. Currently based on supervised computer vision models trained on manually annotated data, this method may struggle to identify unannotated but anomalous areas. This limitation becomes more pronounced without access to historical drone surveys, which would otherwise provide valuable temporal context. An unsupervised machine learning method, such as k-means clustering applied features extracted from deep vision networks, may offer a more flexible solution. It could group bricks into clusters—identifying those damaged, replaced a decade ago, or earlier—without labels. The number of clusters could be user-defined or inferred by the algorithm. This approach could provide a more nuanced understanding of the brickwork's condition over time, potentially revealing subtle patterns of deterioration and repair. Coupling it with non-destructive testing could further validate the findings, ultimately helping to resolve the distribution shift through a richer, more independent analysis.

The narrowing of the gap caused by the distribution shift would make the methodology more suitable for assessing which parts of façades are under failure conditions and likely to degrade. It could also indicate which sections might sustain greater damage. However, for a nuanced understanding of when the degradation onsets, considering the composition of the façade and its material properties is essential. For instance, façades that allow for drying or brickwork that allows for a higher level of moisture before saturation may be exposed to later onset of freeze-thaw degradation. Moreover, material properties can vary within a single façade, so the degradation ratio of each section could scale the results derived from Monte Carlo-based hygrothermal simulation. This could provide insight into how many freeze-thaw cycles a façade must undergo before degrading. Consequently, the predicted degradation ratio of each section could help determine when and to what extent degradation occurs.

Although the framework is showcased only on brick façades, its prescriptive and flexible steps ensure theoretical robustness. However, ensuring practical robustness, in addition to addressing the distribution shift, would require the use of computer vision models that are trained to detect degradation of various types on different façade materials, similar to those CV models presented in [40] or [61]. Moreover, as different materials are expected to have different hygrothermal and thus different degradation mechanisms and patterns, scalability of the PAIR framework would require material- and climate-zone specific calibration.

Vandemeulebroucke et al. [70], for instance, emphasized this point when they analyzed results of 34,560 hygrothermal simulations, showing that, in some scenarios, façade degradation risks may primarily be driven by location, while in other cases, it may be driven by brick type. Ultimately, the framework, respectively its ML-based degradation predictor, could be used to predict the location of degradation on façades in entire neighborhoods across countries and climate zones.

6. Conclusion

This paper developed a methodology for the assessment of the risk of degradation along façade surfaces. The methodology is designed to address the primary limitations of conventional microclimate and hygrothermal modeling by estimating the effects of microclimate over façades through direct prediction of façade degradation. This is made possible by the four-step methodology, which comprises the following stages: Prepare, Analyze, Integrate, and Relate (PAIR). These stages combine data from drone imagery of façades and their analysis by computer vision with weather data and data mined from a 3D model of the façade surroundings. The result is a machine learning predictor for façade surface degradation prediction.

The methodology was exemplified in a case study comprising 16 brick façades from the Flatås area in Gothenburg, Sweden. These were subjected to a drone survey, which yielded high-resolution images of their surfaces. Subsequently, the drone imagery was analyzed, detecting damaged and replaced bricks during previous renovations. This was achieved by retraining an in-house computer vision model (BuildSense) based on You Only Look Once (YOLOv4) computer vision framework. The degradation data was conjugated with data from DTCC model of the neighborhood. The performance of the final model was evaluated using two distinct hold-out samples of the data, namely the validation façade hold-out sample and the test hold-out sample. The former was constructed by setting aside data for one façade at the outset of the analysis. The choice of the validation façade was made upon examining the degradation patterns of the façades in the dataset, choosing a façade with the second highest degree of degradation with the incentive that its highly and less degraded sections could be predicted based on other façades in the dataset. The latter was constructed by shuffling the remaining data and randomly selecting 20 % of the data. Consequently, the predictor exhibited excellent performance in the case of the test hold-out sample. In contrast, the model demonstrated inadequate predictive capabilities when evaluated on the validation façade. These findings were consistent across both experimental set-ups tested, which differed in their feature selection methods. Additionally, the predictions on the validation façade have exhibited high sensitivity to the selected features. This behavior was attributed to the distribution shift, which was likely caused by the validation façade exhibiting distinctive degradation patterns that were not adequately represented in the training data.

Although YOLOv4 was selected for its ready availability within the BuildSense system, recent object detection models such as YOLOv11, and transformer-based approaches like DETR have demonstrated improvements in both accuracy and efficiency. Future work may involve evaluating these newer models to enhance detection robustness, reduce manual verification, and increase scalability in broader applications. Additionally, unsupervised clustering techniques could be explored to better distinguish between degraded and replaced bricks. Alternatively, degradation analysis could be enhanced or replaced using point cloud data, which is not available in the current study. Nevertheless, state-of-the-arts methods like Vision Transformer segmentation or PointNet++ could further enhance the accuracy and robustness of the analysis. In addition, point cloud-based 3D reconstruction techniques may eliminate the need for the 3D model of the neighborhood, potentially making the application of the PAIR methodology more accurate and streamlined. The degradation assessment could also be enhanced by incorporating data from other sensors, such as infrared cameras or ground-penetrating

radar—technologies that avoid practical challenges associated with both invasive (e.g., moisture probes, thermocouples) and environmental (e.g., surface temperature, ambient humidity, wind exposure) continuous measurements. Cost, unexpected data loss due to sensor malfunction, and the need to periodically change batteries are some of the challenges that would limit the incorporation of numerous façades into the dataset, hindering the resolution of the distribution shift.

Finally, the methodology has the potential to offer a computationally inexpensive risk of degradation assessment for façades in whole neighborhoods. However, its practical implementation is currently impeded by the distribution shift. Future research should focus on resolving this issue. This may be overcome by enhancing computer vision analysis in conjunction with other non-destructive testing methods. This would give a comprehensive understanding of façade's maintenance records and expanding the number of façades in the database.

Declaration of generative AI in scientific writing

During the preparation of this work the author(s) occasionally used DeepL in order to polish the writing. ChatGPT was used to shorten the abstract and parts of the introduction. After using these tools, the author (s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

CRediT authorship contribution statement

Jan Mandinec: Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Data curation, Conceptualization. **Angela Sasic Kalagasidis:** Writing – review & editing, Supervision, Resources, Conceptualization. **Pär Johansson:** Writing – review & editing, Supervision, Resources, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors are grateful for funding from Formas, the Swedish research council for sustainable development (grant 2019-01402).

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.autcon.2025.106443>.

Data availability

Raw data, i.e., drone images of the façades, are available in [71]. Other processed data will be made available upon reasonable request.

References

- [1] L. Thuvander, P. Femenías, K. Mjörnell, P. Meiling, Unveiling the process of sustainable renovation, *Sustainability* 4 (6) (Jun. 2012) 1188–1213, <https://doi.org/10.3390/su4061188>.
- [2] R. Ruparathna, K. Hewage, R. Sadiq, Multi-period maintenance planning for public buildings: a risk based approach for climate conscious operation, *J. Clean. Prod.* 170 (Jan. 2018) 1338–1353, <https://doi.org/10.1016/j.jclepro.2017.09.178>.
- [3] D. Ai, G. Jiang, S.-K. Lam, P. He, C. Li, Computer vision framework for crack detection of civil infrastructure—a review, *Eng. Appl. Artif. Intell.* 117 (105478) (Jan. 2023), <https://doi.org/10.1016/j.engappai.2022.105478>.
- [4] Q. Chen, J. Cao, S. Zhu, Data-driven monitoring and predictive maintenance for engineering structures: technologies, implementation challenges, and future directions, *IEEE Internet Things J.* 10 (16) (Aug. 2023) 14527–14551, <https://doi.org/10.1109/JIOT.2023.3272535>.

- [5] Fuktcentrum, Industry standard Byggaf-method for moisture safety of construction process, Accessed: Jan. 16, 2023. [Online]. Available: https://www.fuktcentrum.lth.se/verktyg_och_hjaelpmedel/fuktsaekert_byggande/byggaf_metoden/.
- [6] A. Hukka, H.A. Viitanen, A mathematical model of mould growth on wooden material, *Wood Sci. Technol.* 33 (6) (Dec. 1999) 475–485, <https://doi.org/10.1007/s002260050131>.
- [7] E. Vereecken, L. Van Gelder, H. Janssen, S. Roels, Interior insulation for wall retrofitting – a probabilistic analysis of energy savings and hygrothermal risks, *Energ. Buildings* 89 (Feb. 2015) 231–244, <https://doi.org/10.1016/j.enbuild.2014.12.031>.
- [8] N.W. Portal, A.W.M. van Schijndel, A.S. Kalagasidis, The multiphysics modeling of heat and moisture induced stress and strain of historic building materials and artefacts, *Build. Simul.* 7 (Jun. 2014) 217–227, <https://doi.org/10.1007/s12273-013-0153-4>.
- [9] M. Uranjek, V. Bokan-Bosiljkov, Influence of freeze–thaw cycles on mechanical properties of historical brick masonry, *Constr. Build. Mater.* 84 (Jun. 2015) 416–428, <https://doi.org/10.1016/j.conbuildmat.2015.03.077>.
- [10] I. Netinger Grubeša, M. Vračević, J. Ranogajec, S. Vučić, Influence of pore-size distribution on the resistance of clay brick to freeze–thaw cycles, *Materials* 13 (10) (Jan. 2020) 2364, <https://doi.org/10.3390/ma13102364>. Art. no.
- [11] J. Yue, C. Ma, L. Zhao, Q. Kong, X. Xu, Z. Wang, et al., Study on deterioration of gray brick with different moisture contents under freeze–thaw environment, *Materials* 15 (5) (Jan. 2022) 1819, <https://doi.org/10.3390/ma15051819>. Art. no.
- [12] B. Funk, D. Göhler, B. Sachsenhauser, M. Stintz, B. Stahlmecke, B.A. Johnson, et al., Impact of freeze–thaw weathering on integrity, internal structure and particle release from micro- and nanostructured cement composites, *Environ. Sci. Nano* 6 (5) (May 2019) 1443–1456, <https://doi.org/10.1039/C8EN01397G>.
- [13] A. Török, B. Szemerey-Kiss, Freeze-thaw durability of repair mortars and porous limestone: compatibility issues, *Prog. Earth Planet. Sci.* 6 (1) (Jun. 2019) 42, <https://doi.org/10.1186/s40645-019-0282-1>. Art. no.
- [14] P. Mensinga, J. Straube, C. Schumacher, Assessing the freeze-thaw resistance of clay brick for interior insulation retrofit projects, in: Presented at the Thermal Performance of the Exterior Envelopes of Whole Buildings - 11th International Conference, Clearwater, Florida, USA, 2010 [Online]. Available: https://web.ornl.gov/sci/buildings/conf-archive/2010%20B11%20papers/203_Mensinga.pdf.
- [15] P. Mensinga, Determining the Critical Degree of Saturation of Brick Using Frost Dilatometry, Ph.D. Thesis, University of Waterloo, Waterloo, Ontario, Canada, 2009. Accessed: Jun. 01, 2022. [Online]. Available: <https://uwspace.uwaterloo.ca/handle/10012/4638>.
- [16] J. Koci, J. Maděra, M. Keppert, R. Černý, Damage functions for the cold regions and their applications in hygrothermal simulations of different types of building structures, *Cold Reg. Sci. Technol.* 135 (Mar. 2017) 1–7, <https://doi.org/10.1016/j.coldregions.2016.12.004>.
- [17] X. Zhou, D. Derome, J. Carmeliet, Hygrothermal modeling and evaluation of freeze-thaw damage risk of masonry walls retrofitted with internal insulation, *Build. Environ.* 125 (Nov. 2017) 285–298, <https://doi.org/10.1016/j.buildenv.2017.08.001>.
- [18] X. Zhou, D. Derome, A. Kubilay, J. Carmeliet, Comprehensive study of moisture risk on building facades based on spatial distribution of wetting and drying, in: Presented at the Thermal Performance of the Exterior Envelopes of Whole Buildings XV International Conference, Clearwater, Florida, USA, 2022 [Online]. Available: <https://www.ashrae.org/File%20Library/Conferences/Specialty%20Conferences/Buildings%20XV%20-%20Papers/C029.pdf>.
- [19] J. F. Straube and E. F. P. Burnett, “Simplified prediction of driving rain deposition,” Presented at the International Building Physics Conference, Eindhoven, September 12–21, 2000, pp. 375–382. [Online]. Available: <https://users.encs.concordia.ca/~raojw/crd/reference/reference001203.html>.
- [20] X. Zhou, A. Kubilay, D. Derome, J. Carmeliet, Comparison of wind-driven rain load on building facades in the urban environment and open field: A case study on two buildings in Zurich, Switzerland, *Build. Environ.* 233 (Apr. 2023) 110038, <https://doi.org/10.1016/j.buildenv.2023.110038>. Art. no.
- [21] S. Charisi, T.K. Thiis, P. Stefansson, I. Burud, Prediction model of microclimatic surface conditions on building facades, *Build. Environ.* 128 (Jan. 2018) 46–54, <https://doi.org/10.1016/j.buildenv.2017.11.017>.
- [22] S. Tsoka, T.K. Thiis, Calculation of the driving rain wall factor using ray tracing, *J. Wind Eng. Ind. Aerodyn.* 179 (Aug. 2018) 190–199, <https://doi.org/10.1016/j.jweia.2018.06.008>.
- [23] T.K. Thiis, I. Burud, A. Flø, D. Kraniotis, S. Charisi, P. Stefansson, Monitoring and simulation of diurnal surface conditions of a wooden façade, *Procedia Environ. Sci.* 38 (Jan. 2017) 331–339, <https://doi.org/10.1016/j.proenv.2017.03.088>.
- [24] B. Blocken, J. Carmeliet, Validation of CFD simulations of wind-driven rain on a low-rise building facade, *Build. Environ.* 42 (7) (Jul. 2007) 2530–2548, <https://doi.org/10.1016/j.buildenv.2006.07.032>.
- [25] A. Kubilay, D. Derome, B. Blocken, J. Carmeliet, Numerical simulations of wind-driven rain on an array of low-rise cubic buildings and validation by field measurements, *Build. Environ.* 81 (Nov. 2014) 283–295, <https://doi.org/10.1016/j.buildenv.2014.07.008>.
- [26] A. Khalilzadeh, H. Ge, H.D. Ng, Effect of turbulence modeling schemes on wind-driven rain deposition on a mid-rise building: CFD modeling and validation, *J. Wind Eng. Ind. Aerodyn.* 184 (Jan. 2019) 362–377, <https://doi.org/10.1016/j.jweia.2018.11.012>.
- [27] A. Kubilay, D. Derome, B. Blocken, J. Carmeliet, Wind-driven rain on two parallel wind buildings: field measurements and CFD simulations, *J. Wind Eng. Ind. Aerodyn.* 146 (Nov. 2015) 11–28, <https://doi.org/10.1016/j.jweia.2015.07.006>.
- [28] A. Kubilay, D. Derome, J. Carmeliet, Coupling of physical phenomena in urban microclimate: a model integrating air flow, wind-driven rain, radiation and transport in building materials, *Urban Clim.* 24 (Jun. 2018) 398–418, <https://doi.org/10.1016/j.uclim.2017.04.012>.
- [29] J. Choi, C.M. Yeum, S.J. Dyke, M.R. Jahanshahi, Computer-aided approach for rapid post-event visual evaluation of a building façade, *Sensors* 18 (9) (Sep. 2018) 3017, <https://doi.org/10.3390/s18093017>. Art. no.
- [30] S. Sankarasrinivasan, E. Balasubramanian, K. Karthik, U. Chandrasekar, R. Gupta, Health monitoring of civil structures with integrated UAV and image processing system, *Procedia Comput. Sci.* 54 (Jan. 2015) 508–515, <https://doi.org/10.1016/j.procs.2015.06.058>.
- [31] C. Eschmann, C. Kuo, C. Kuo, C. Boller, Unmanned aircraft systems for remote building inspection and monitoring, in: Presented at the 6th European Workshop on Structural Health Monitoring, 2012. Accessed: Oct. 01, 2024. [Online]. Available: <https://www.semanticscholar.org/paper/Unmanned-Aircraft-System-s-for-Remote-Building-and-Eschmann-Kuo/25af5ace3b8e925045a0f13587340c1fc6f14dcd>.
- [32] F.C. Pereira, C.E. Pereira, Embedded image processing Systems for Automatic Recognition of cracks using UAVs, *IFAC-Pap.* 48 (10) (Jan. 2015) 16–21, <https://doi.org/10.1016/j.ifacol.2015.08.101>.
- [33] C.-Z. Dong, F.N. Catbas, A review of computer vision-based structural health monitoring at local and global levels, *Struct. Health Monit.* 20 (2) (Mar. 2021) 692–743, <https://doi.org/10.1177/1475921720935585>.
- [34] K. Chen, G. Reichard, X. Xu, A. Akanmu, Automated crack segmentation in close-range building façade inspection images using deep learning techniques, *J. Build. Eng.* 43 (Nov. 2021) 102913, <https://doi.org/10.1016/j.jobte.2021.102913>. Art. no.
- [35] Y. Chen, Z. Zhu, Z. Lin, Y. Zhou, Building surface crack detection using deep learning technology, *Buildings* 13 (7) (Jul. 2023) 1814, <https://doi.org/10.3390/buildings13071814>. Art. no.
- [36] J. Guo, Q. Wang, Y. Li, P. Liu, Façade defects classification from imbalanced dataset using meta learning-based convolutional neural network, *Comput. Aided Civ. Inf. Eng.* 35 (12) (2020) 1403–1418, <https://doi.org/10.1111/mice.12578>.
- [37] S. Katsigiannis, S. Seyedzadeh, A. Agapiou, N. Ramzan, Deep learning for crack detection on masonry façades using limited data and transfer learning, *J. Build. Eng.* 76 (Oct. 2023) 107105, <https://doi.org/10.1016/j.jobte.2023.107105>. Art. no.
- [38] N. Wang, X. Zhao, P. Zhao, Y. Zhang, Z. Zou, J. Ou, Automatic damage detection of historic masonry buildings based on mobile deep learning, *Autom. Constr.* 103 (2019) 53–66, <https://doi.org/10.1016/j.autcon.2019.03.003>.
- [39] M. Torres-González, J. Valença, B.O. Santos, A. Silva, M.P. Mendes, StainView: a fast and reliable method for mapping stains in facades using image classification in HSV and CIELab colour space, *Remote Sens* 15 (11) (Jan. 2023) 2895, <https://doi.org/10.3390/rs15112895>. Art. no.
- [40] X. Chen, J. He, S. Wang, Deep learning-driven pathology detection and analysis in historic masonry buildings of Suzhou, *Npj Herit. Sci.* 13 (May 2025) 197, <https://doi.org/10.1038/s40494-025-01783-y>. Art. no.
- [41] L. Yang, K. Liu, R. Ou, P. Qian, Y. Wu, Z. Tian, et al., Surface defect-extended BIM generation leveraging UAV images and deep learning, *Sensors* 24 (13) (Jan. 2024) 4151, <https://doi.org/10.3390/s24134151>. Art. no.
- [42] M. Alipour, D.K. Harris, G.R. Miller, Robust pixel-level crack detection using deep fully convolutional neural networks, *J. Comput. Civ. Eng.* 33 (6) (Nov. 2019), [https://doi.org/10.1061/\(ASCE\)CP.1943-5487.0000854](https://doi.org/10.1061/(ASCE)CP.1943-5487.0000854).
- [43] B. Lim, S. Son, H. Kim, S. Nah, K.M. Lee, Enhanced deep residual networks for single image super-resolution, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Jul. 2017, pp. 1132–1140, <https://doi.org/10.1109/CVPRW.2017.151>.
- [44] P.L. Gaspar, J. de Brito, Quantifying environmental effects on cement-rendered facades: a comparison between different degradation indicators, *Build. Environ.* 43 (11) (Nov. 2008) 1818–1828, <https://doi.org/10.1016/j.buildenv.2007.10.022>.
- [45] Z. Niu, H. Li, Research and analysis of threshold segmentation algorithms in image processing, *J. Phys. Conf. Ser.* 1237 (2) (Jun. 2019) 022122, <https://doi.org/10.1088/1742-6596/1237/2/022122>.
- [46] S.M. AqilBurney, H. Tariq, K-means cluster analysis for image segmentation, *Int. J. Comput. Appl.* 96 (4) (Jun. 2014) 1–8, <https://doi.org/10.5120/16779-6360>.
- [47] J. Canny, A computational approach to edge detection, in: M.A. Fischler, O. Firschein (Eds.), *Readings in Computer Vision*, Morgan Kaufmann, San Francisco (CA), 1987, pp. 184–203, <https://doi.org/10.1016/B978-0-08-051581-6.50024-6>.
- [48] A. Plaksyvyi, M. Skubiewska-Paszkowska, P. Powroznik, A comparative analysis of image segmentation using classical and deep learning approach, *Adv. Sci. Technol. Res. J.* 17 (6) (Dec. 2023) 127–139, <https://doi.org/10.12913/22998624/172771>.
- [49] N. O'Mahony, S. Campbell, A. Carvalho, S. Harapanahalli, G.V. Hernandez, L. Krpalkova, et al., Deep learning vs. traditional computer vision, in: *Advances in Computer Vision*, K. Arai, S. Kapoor (Eds.), in: *Advances in Intelligent Systems and Computing* 943, Springer International Publishing, Las Vegas, USA, 2020, pp. 128–144, https://doi.org/10.1007/978-3-030-17795-9_10.
- [50] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: unified, real-time object detection, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Jun. 2016, pp. 779–788, <https://doi.org/10.1109/CVPR.2016.91>.
- [51] M. Tan, R. Pang, Q.V. Le, EfficientDet: scalable and efficient object detection, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Jun. 2020, pp. 10778–10787, <https://doi.org/10.1109/CVPR42600.2020.01079>.
- [52] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, et al., SSD: single shot multibox detector, in: *Computer Vision–ECCV 2016*, Amsterdam, Netherlands, 2016, pp. 21–37, https://doi.org/10.1007/978-3-319-46448-0_2.
- [53] K. Lee, G. Hong, L. Sael, S. Lee, H.Y. Kim, MultiDefectNet: multi-class defect detection of building façade based on deep convolutional neural network,

- Sustainability 12 (22) (Jan. 2020) 22, <https://doi.org/10.3390/su12229785>. Art. no.
- [54] S. Ren, K. He, R. Girshick, J. Sun, Faster R-CNN: Towards real-time object detection with region proposal networks, in: *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, in NIPS'15, MIT Press, Cambridge, MA, USA, Dec. 2015, pp. 91–99, <https://doi.org/10.1109/TPAMI.2016.2577031>.
- [55] O. Ronneberger, P. Fischer, T. Brox, U-net: convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Munich, Germany, 2015, pp. 234–241, https://doi.org/10.1007/978-3-319-24574-4_28.
- [56] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: Presented at the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE Computer Society, Jun. 2015, pp. 3431–3440, <https://doi.org/10.1109/CVPR.2015.7298965>.
- [57] V. Badrinarayanan, A. Kendall, R. Cipolla, SegNet: a deep convolutional encoder-decoder architecture for image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (12) (Dec. 2017), <https://doi.org/10.1109/TPAMI.2016.2644615>.
- [58] M.-T. Cao, Drone-assisted segmentation of tile peeling on building façades using a deep learning model, *J. Build. Eng.* 80 (Dec. 2023) 108063, <https://doi.org/10.1016/j.jobbe.2023.108063>. Art. no.
- [59] H. Chen, K. Sun, Z. Tian, C. Shen, Y. Huang, Y. Yan, BlendMask: Top-down meets bottom-up for instance segmentation, in: Presented at the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Jun. 2020, pp. 8570–8578, <https://doi.org/10.1109/CVPR42600.2020.00860>.
- [60] X. Fu, N. Angkawisittpan, Detecting surface defects of heritage buildings based on deep learning, *J. Intell. Syst.* 33 (1) (Jan. 2024) 20230048, <https://doi.org/10.1515/jisys-2023-0048>. Art. no.
- [61] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, et al., Swin transformer: Hierarchical vision transformer using shifted windows, in: Presented at the 2021 IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, Canada, Oct. 2021, pp. 9992–10002, <https://doi.org/10.1109/ICCV48922.2021.00986>.
- [62] W. Zhang, J. Pang, K. Chen, C.C. Loy, K-net: towards unified image segmentation, in: *Advances in Neural Information Processing Systems*, Curran Associates, Inc, 2021, pp. 10326–10338. Accessed: Jun. 27, 2025. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/hash/55a7cf9c71f1c9c495413f934dd1a158-Abstract.html.
- [63] T. Xiao, Y. Liu, B. Zhou, Y. Jiang, J. Sun, Unified perceptual parsing for scene understanding, in: V. Ferrari, M. Hebert, C. Sminchisescu, Y. Weiss (Eds.), *Computer Vision – ECCV 2018*, ECCV 2018, Springer International Publishing, 2018, pp. 432–448, https://doi.org/10.1007/978-3-030-01228-1_26. Lecture Notes in Computer Science.
- [64] R. Bellman, Dynamic programming, *Science* 153 (3731) (Jul. 1966) 34–37, <https://doi.org/10.1126/science.153.3731.34>.
- [65] K. Pearson, On lines and planes of closest fit to systems of points in space, *Lond. Edinb. Dublin Philos. Mag. J. Sci.* 2 (11) (Nov. 1901) 559–572, <https://doi.org/10.1080/14786440109462720>.
- [66] U. Stańczyk and L. C. Jain, Eds., *Feature selection for data and pattern recognition*, Vol. 584. in *Studies in Computational Intelligence*, Vol. 584. Berlin, Heidelberg: Springer, 2015. doi:<https://doi.org/10.1007/978-3-662-45620-0>.
- [67] V. Naserentin, A. Logg, D. Wästberg, DTCC builder: a mesh generator for automatic, efficient, and robust mesh generation for large-scale city modeling and simulation, *J. Open Source Softw.* 8 (86) (2023), <https://doi.org/10.21105/joss.04928>.
- [68] J. Mandinec, P. Johansson, Towards an automatized and objective assessment of data from visual inspections of building envelopes, *Acta Polytech. CTU Proc.* 38 (Dec. 2022) 57–64, <https://doi.org/10.14311/APP.2022.38.0057>.
- [69] M. Tkachenko, M. Malyuk, A. Holmanyuk, N. Liubimov, LABEL STUDIO: DATA LABELING SOFTWARE [Online]. Available: <https://github.com/HumanSignal/lab-el-studio>, 2020.
- [70] I. Vandemeulebroucke, S. Caluwaerts, N. Van Den Bossche, Geographical variability in hygrothermal simulation results of historical building envelopes with interior insulation in Belgium, *J. Build. Eng.* 98 (2024) 111229, <https://doi.org/10.1016/j.jobbe.2024.111229>. Art. no.
- [71] J. Mandinec, Drone images of 16 brick façades in Flatås, Gothenburg, Dataset, Jan. 08, 2025, <https://doi.org/10.17632/r2yz9b8cfs.1>.