

<https://doi.org/10.1038/s40494-026-03023-3>

Multimodal synchrotron X-ray analysis of palm-leaf manuscripts: structural and compositional signatures of botanical genus, manufacture, and provenance



Laura Gallardo¹✉, Mark Busch^{1,2}, Anastasia Poliakova^{3,4}, Giovanni Ciotti^{3,5} & Patrick Huber^{1,2}

Palm-leaf manuscripts are cultural artifacts from South and Southeast Asia, yet diverse preparation methods and degradation histories complicate material characterization and comparison. We use a multimodal synchrotron approach combining small- and wide-angle X-ray scattering with X-ray fluorescence to extract structural and compositional descriptors from unprocessed leaves, writing supports, and manuscripts. Supervised learning evaluates differences related to genus, processing stage, and geographical origin. The processing stage is classified reliably, whereas genus separation is clearer in processed materials than in unprocessed leaves. Origin classification shows partial separability, with distinct regional profiles and overlaps reflecting shared preparation traditions, class imbalance, and botanical and degradation-related factors. Classification relies on multivariate combinations of crystallinity, fibril orientation, and elemental descriptors. This scalable, non-destructive framework supports comparison of collections and investigation of material similarities among traditions. The classifications represent comparative material profiles rather than definitive provenance assignments and must be combined with codicological and historical evidence.

Palm-leaf manuscripts (PLM) are written artifacts that constitute an integral part of the cultural and material heritage of South and Southeast Asia. This writing support was used for centuries to preserve religious, belletristic, legal, and scientific knowledge. Although paper later replaced palm leaf as the dominant writing support, manuscript production survives in some religious and craft contexts. These surviving practices provide valuable information about the tradition, although historical records documenting them are scarce^{1,2}.

Palm leaves are lignocellulosic composites containing cellulose microfibrils embedded in hemicellulose and lignin, with surface waxes and variation related to botanical identity and leaf maturity. The principal genera used for manuscript production are *Corypha* L. and *Borassus* L.³, both of which belong to the family Arecaceae, subfamily Coryphoideae. Preparation commonly involves cutting, boiling, drying, polishing, and the application of surface treatments¹, which can alter cellulose crystallinity, porosity, composition, mechanical behavior, and stability⁴. In inscribed

manuscripts, incision and the application of carbon-based ink introduce further structural and compositional modifications⁵.

The diversity in the preparation methods of palm leaves² is expected to impact the physical and chemical properties of the manuscripts. Extensive efforts are focused at the moment to document these practices^{6–13}, as well as to classify the plants used for manuscript preparation and conservation¹. However, the available evidence remains incomplete. Historical accounts of production are scarce, and detailed technical descriptions are largely limited to sources from the twentieth century onward, leaving the earlier development and regional diversification of these practices poorly understood¹. Moreover, modern accounts often describe only the general sequence of operations and omit parameters such as treatment duration, temperature, vessel material, ingredient quantities, or the composition of boiling and preservation mixtures^{2,14}.

Relating the physicochemical properties of a manuscript to a specific production method is further complicated by changes occurring after

¹Hamburg University of Technology, Institute for Materials and X-ray Physics, Hamburg, Germany. ²Centre for X-ray and Nano Science CXNS, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany. ³Centre for the Study of Manuscript Cultures CSMC, Universität Hamburg, Warburgstraße 26, Hamburg, Germany. ⁴Georg-August-Universität Göttingen, Albrecht-von-Haller-Institute for Plant Sciences, Department for Palynology and Climate Dynamics, Göttingen, Germany. ⁵Department of History and Cultures, University of Bologna, Bologna, Italy. ✉e-mail: laura.gallardo@tuhh.de

manufacture. Material properties are affected by natural ageing, environmental exposure, handling, and conservation treatments. Multiscale studies of historical manuscripts have identified changes in anatomy, chemical composition, and mechanical properties, including cell-wall damage, microbial erosion, degradation of polysaccharides, and loss of tensile strength¹⁵. Preservation practices may introduce further modifications through the application of oils, botanical extracts, arsenic-containing compounds, or modern consolidants¹⁶. Consequently, the measured properties of a manuscript may reflect several overlapping influences rather than one identifiable production recipe.

A range of analytical approaches has been applied to investigate these different aspects of PLM materiality. Microscopic, spectroscopic, diffraction, elemental, and mechanical methods have been used to characterize palm-leaf anatomy, deterioration, composition, and writing materials^{15,17–21}. These studies provide detailed information on individual material components and degradation mechanisms, but comparatively few combine structural and compositional information across multiple length scales.

Spectroscopic and compositional measurements have also been used to differentiate palm-leaf materials. FTIR, vibrational microspectroscopy, XRF, and other techniques can capture variation related to botanical origin, preparation, preservation, and environmental exposure^{22–24}. They can often be applied non-invasively, although detailed measurements of large collections may remain time-consuming. FTIR and related vibrational methods are also primarily sensitive to near-surface functional groups and are rarely combined with structural measurements that probe several length scales.

Because the contributions of botanical origin, manufacture, ageing, and conservation are difficult to separate using individual parameters, multivariate and statistical methods have increasingly been applied to palm-leaf materials. Principal component analysis and other multivariate approaches have been used to identify patterns in spectroscopic, elemental, or morphological datasets and to explore distinctions among manuscript groups^{6,22,23}. These studies show that material groups can be differentiated even when the contributions of individual underlying mechanisms cannot be resolved. However, analyses are often based on unsupervised dimensionality reduction and do not necessarily evaluate the predictive performance or robustness of different classification models. These limitations motivate an approach that combines complementary structural and compositional measurements with a supervised framework for evaluating classification performance.

Synchrotron X-ray techniques offer several advantages for the study of heritage materials because they enable rapid, spatially resolved, and non-destructive measurements. Raster scanning enables evaluation of local heterogeneity and derivation of representative values by averaging over selected material regions. Compared with predominantly surface-sensitive spectroscopic methods, small- and wide-angle X-ray scattering provide complementary information on the internal organization of the substrate. WAXS probes cellulose crystallinity and crystallite dimensions, whereas SAXS is sensitive to larger-scale fibrillar organization and anisotropy. When combined with XRF, structural and elemental information can be collected simultaneously from the same region. Although X-ray scattering and fluorescence have been applied to a range of cultural-heritage materials, structural studies of lignocellulosic objects frequently require microsampling^{25–27}, and comparatively few studies combine SAXS, WAXS, and XRF entirely non-destructively^{28–31}.

Here, we combine simultaneous SAXS, WAXS, and XRF measurements with supervised machine learning to compare unprocessed palm leaves, modern writing supports, and historical manuscripts. Physically interpretable descriptors of cellulose crystallinity, crystallite dimensions, amorphous scattering, fibrillar orientation, and elemental composition were extracted from raster-scan measurements. We evaluate the extent to which these descriptors distinguish botanical genus, processing stage, and geographical origin, and use feature importance and stratified analyses to identify the material properties underlying these classifications. The primary objective is not definitive botanical or provenance assignment, but the

identification of structural and compositional similarities among manuscript groups and their possible relationships to material selection, preparation, ageing, and conservation history.

Methods

The dataset comprised selected folios from the CSMC collection and private collections, together with manuscript fragments, dried leaves, and partially or fully prepared writing supports obtained from institutions, collectors, production centers, archives, and botanical gardens. These include the State and University Library (Carl von Ossietzky Staats- und Universitätsbibliothek Hamburg); the manuscripts collection of the Ecole Française D'extreme-Orient (EFEO) in Puducherry, India; the collection of the Department of Indology, French Institute (IFP) in Puducherry, India; the Kerala Manuscript Preservation Center and the Vedic Research Center (VRC), Vadakke Madham Brahmaswam (VMB) in Thrissur (Keralam, India). Some manuscript fragments had detached naturally from degraded objects and were made available for scientific analysis. Sample metadata, including the documented origin, botanical identification, processing stage, estimated manuscript age (when applicable), and holding institution, are provided in Supplementary Data 1. A list of measurements and experimental details is provided in Supplementary Data 2. Palm leaf and palm-leaf manuscript samples were collected with permission for scientific study. To the best of our knowledge, all necessary permissions for their analysis were obtained at the time of collection.

Because many objects have incomplete collection histories, these metadata have different levels of certainty. Labels were assigned from catalog records, donor documentation, fieldwork information, botanical examination, and expert assessment, where available. They should therefore be understood as the best available classifications rather than independently verified ground truth. In particular, the historical geographical origin, exact preparation procedure, age, and degradation history are unknown or only approximately known for part of the dataset.

Synchrotron measurements

All measurements were performed non-destructively. Synchrotron X-ray measurements were performed at two facilities: the P62 beamline at PETRA III, DESY³², during two beamtimes^{33,34}, and the BM02 beamline at the ESRF during one beamtime, where only manuscript fragments and reference materials were measured³⁵. At both beamlines, small-angle X-ray scattering (SAXS), wide-angle X-ray scattering (WAXS), and X-ray fluorescence (XRF) were recorded simultaneously, as illustrated in Fig. 1.

Because the measurements were acquired using different experimental configurations and detector systems, the comparability of the extracted structural and compositional descriptors was evaluated before combining the datasets. This included a direct comparison of measurements taken on the same reference material, along with additional statistical tests based on the classification framework described below. The distribution of sample classes across campaigns and the beamtime-classification results are presented in Supplementary Section S1 and Supplementary Figs. 1 and 2. Direct comparisons of matched samples measured at BM02 and P62 are provided in Supplementary Section S2 and Supplementary Fig. 3.

All samples were measured over raster-scan areas of a few millimeters, with the scan dimensions determined mainly by the size and geometry of the available object or fragment. When sufficient material was available, measurement regions were selected to capture spatial variability within the palm-leaf material, including areas near inscriptions, leaf edges, binding holes, and different anatomical regions. For small fragments, the choice of position was restricted by the available sample area. A standardized position could therefore not be used for all objects.

In the statistical analysis, each measurement corresponds to a single raster-scan region from which a single set of SAXS, WAXS, and XRF descriptors was extracted. Each scan contained multiple spatial points, and the descriptors were calculated by averaging over the relevant writing-support pixels rather than from individual points. Inscriptions were identified using elemental maps (see “Methods” section) and excluded because




Fig. 1 | Multimodal synchrotron-radiation X-ray analytics of a palm-leaf manuscript. Setup combining X-ray fluorescence (XRF) (a), wide-angle X-ray scattering (WAXS) (b) and small-angle X-ray scattering (SAXS) (c) as arranged at the beamlines P62 and BM02 of the PETRA III synchrotron and ESRF, respectively. The incident synchrotron X-ray beam goes in the z -direction through the manuscript, positioned at an angle α with respect to the normal plane of the beam (xy). The manuscript folio scatters X-rays detected by the WAXS detector near the sample and by the SAXS detector, away from the sample, to capture smaller scattering angles. Black lines were added in the detector images to indicate the integration from the

beam center towards growing q . An inset is included for the SAXS detector to illustrate the direction of the angle χ , used in the azimuthal integration of the intensity. Incoming X-rays also trigger X-ray fluorescence in the sample. The XRF detector captures the spectral fingerprints (elemental lines) that were later identified and fitted using pyMCA. The sample is moved in the dashed lines directions in the sketch, allowing the performance of raster scans; the scan region in this measurement is included on the right. Properties can be mapped in the scan region; for example, the iron K-line intensity distribution is shown.




Fig. 2 | SAXS azimuthal intensity distribution analysis for manuscript 10_15, folio 7, measured at P62. a Normalized azimuthal intensity distribution across the analyzed q -band. In light blue, the azimuthal intensity distributions (normalized to 1) for q in the range 0.14 nm^{-1} to 0.22 nm^{-1} . In black, the average intensity distribution, displaying maxima at approximately 0° and $\pm 180^\circ$, and $\pm 90^\circ$, one of them coinciding with a blind spot of the detector. b Probability density distribution function $P(\chi)$ (light blue) as calculated from the average intensity distributions over


the writing support. As peaks are symmetrical (two orientations for the same direction), the interval is limited to 180° indicated in panel (a) by two gray vertical dashed lines. $P(\chi)$ is then fitted using four different Voigt contributions: halo and core in the longitudinal direction (LD), represented in dark and light brown, respectively; and core and halo contributions in the orthogonal direction (OD), represented by dark and light green, respectively. Global fit is represented by a black dashed line.

they may alter the local structure and elemental composition. They typically showed increased iron and reduced calcium relative to the surrounding support. Calcium maps were also used to exclude points corresponding to the sample holder or air. This masking reduces the contribution of the incised regions but cannot remove diffuse ink residues or substances distributed over the wider leaf surface.

Separate scan regions, including measurements from different folios of the same manuscript, were retained as individual observations rather than combined into an object-level average. Some objects are therefore represented by more than one measurement, and location-dependent variability remains part of the dataset.

The wavevector q -ranges used for both SAXS and WAXS detectors differed from beamtime to beamtime, as well as the angle α between the samples and the incident beam. Two beamtimes were realized at $\alpha = 0^\circ$, with the exception of a few samples used for the study of the general structure of the leaf, which were measured in a range from 0° to 50° . In the remaining beamtime, measurements were taken for every sample at 32° and 45° . The influence of sample rotation on the extracted WAXS descriptors is evaluated in Supplementary Section S3 and Supplementary Fig. 4. While the effects of rotation are in general not very relevant, deviations tend to arise at larger rotation angles; hence, for those samples measured at different rotation angles, the lower rotation was kept, and

Fig. 3 | Radially integrated WAXS intensity averaged over the scan region (light blue) for sample ms41, measured at P62. The three cellulose reflections, (110), (200), (004), observable in this q -range are indicated by vertical black dashed lines. They were fitted using Voigt functions. An amorphous background contribution was also considered. Global fit is represented in red.



the rest were discarded for the training. Data were converted and calibrated using standard beamline-specific workflows. Detector calibration included geometric calibration and conversion to scattering vector q using reference standards. Detector artifacts, including masked regions (e.g., beamstop and bad pixels), were excluded prior to integration.

For all scans, WAXS detector images were radially integrated to obtain one-dimensional intensity profiles $I(q)$. SAXS detector images were azimuthally integrated for orientation analysis.

Raster scans of empty sample holders were acquired under identical conditions, and the resulting background signal, after smoothing, was subtracted from the data. Where applicable, intensities were normalized to account for variations in acquisition conditions.

XRF spectra were fitted at each scan point using pyMCA. Subsequent preprocessing in Python included masking of holder/air points, identification of inscription regions from elemental maps, and averaging over the writing-support region prior to feature extraction.

All extracted parameters were collected in a master table together with metadata, including sample name, beamtime, rotation angle α , and processing stage (unprocessed, writing support, manuscript). Additional metadata (origin, genus, species, age) was only partially available, as the history of some of the measured manuscripts is unknown (see Supplementary Data 3).

From the SAXS patterns, anisotropy was quantified by azimuthal (sector χ) integration, see Fig. 1c. The intensity is a function of the position in the manuscripts determined by the coordinates x and y , as illustrated in Fig. 1. The intensity was integrated in 1° azimuthal sector and analyzed at constant q . To improve signal-to-noise ratio, intensities were averaged over a low- q band $\Delta q = q_{\max} - q_{\min}$, where $q_{\min} = 0.14 \text{ nm}^{-1}$ and $q_{\max} = 0.22 \text{ nm}^{-1}$:

$$\bar{I}(x, y, \chi) = \frac{1}{\Delta q} \sum_{q=q_{\min}}^{q_{\max}} I(x, y, \chi, q) \quad (1)$$

Averaging over Δq yields stable peak shapes while preserving the anisotropy signal, see Fig. 2a.

Because inscriptions may alter local fiber orientation, they were masked before computing the scan-averaged orientation distribution. The resulting writing-support intensity was:

$$I_{ws}(\chi) = \frac{1}{n_{ws}} \sum_{x, y \in ws} \bar{I}(x, y, \chi), \quad (2)$$

being n_{ws} the number of measurement points. The probability density distribution function in an angular sector χ_i was defined as:

$$P(\chi_i) = \frac{I_{ws}(\chi_i)}{\sum_j I_{ws}(\chi_j)}, \quad (3)$$

where j takes values in the selected 180° range.

Using $P(\chi)$ instead of raw intensity makes the fits scale-invariant and facilitates comparison between scans and beamtimes.

The peaks in $P(\chi)$ were fitted using Voigt functions (Fig. 2b). For each principal direction, longitudinal (LD) and orthogonal (OD), a two-component model consisting of a narrower core contribution and a broader halo contribution, together with a constant baseline, was applied to all measurements.

The extracted features include the total full width at half maximum in each direction (FWHM^{LD} , FWHM^{OD}), the widths of the halo and core contributions ($\text{FWHM}_{\text{halo}}$, $\text{FWHM}_{\text{core}}$), the relative area of the core contribution with respect to the halo ($f_{\text{core}}^{\text{LD}}$, $f_{\text{core}}^{\text{OD}}$), and the peak-to-baseline ratios (I/b).

The core and halo terms provide an empirical description of narrower and broader orientational contributions within each peak. They may reflect differences in the degree of organization among the structural components of the leaf, but they cannot be directly assigned to specific tissues based on the present measurements. Accordingly, f_{core} is interpreted as the relative contribution of the narrower fitted component rather than as a direct measure of tissue thickness. For the statistics, the difference between the total FWHM of the peaks was also considered $\Delta\text{FWHM} = |\text{FWHM}^{\text{OD}} - \text{FWHM}^{\text{LD}}|$.

The position χ of the intensity maximum is related to the orientation direction in real space μ , depending on the incidence angle α of the beam on the sample:

$$\tan \mu = \frac{\tan \chi}{\cos \alpha}. \quad (4)$$

The relative orientation ($\Delta\mu$) was also extracted as a descriptor and considered in the statistics.


WAXS diffractograms were obtained by radial integration of 2D diffraction images and then averaged over the sample scan region, excluding the inscriptions when present. The resulting curves were fitted using Voigt functions for cellulose Bragg reflections (110), (200), and (004) together with a broad amorphous contribution and a piecewise linear background, see Fig. 3.

While a quantification of amorphous material is also not possible, we defined an amorphous contribution A to enable comparison between samples:

$$A = \frac{\int_{10}^{20} I_{\text{am}}(q) dq}{\int_{10}^{22} I_{\text{corr}}^+(q) dq}, \quad (5)$$

where $I_{\text{am}}(q)$ is the fitted amorphous contribution and $I_{\text{corr}}^+(q)$ is the positive part of the background-corrected WAXS intensity within the range $10 - 22 \text{ nm}^{-1}$. The descriptor, therefore, represents the area of the fitted amorphous peak relative to the total positive scattering intensity. It should be interpreted as a relative measure of the prominence of the amorphous scattering contribution rather than as an absolute amorphous fraction.

Fig. 4 | XRF spectrum averaged over the scan area for sample Lookwa (black line). Vertical dashed gray lines are placed at relevant peaks, and the text next to them shows the elemental lines contributing to the peaks, with the first contribution in bold letters being the strongest. The spectrum was fitted using pyMCA³⁶, which calculates elemental distribution over the scan region based on the fits. The bright green line represents the continuum fit of the background, and the red dashed line represents the total fit of the spectrum, including elemental contributions.



Crystallinity indices were computed for each reflection using the areas method:

$$C_{hkl} = \frac{S_{hkl}}{S_{hkl} + S_{am}} \quad (6)$$

where S_{hkl} is the area of the crystalline peak, and S_{am} is the amorphous/background contribution in the corresponding region. Crystallite sizes were estimated using the Scherrer equation:

$$D_{hkl} = \frac{K\lambda}{\text{FWHM} \cos \theta} \quad (7)$$

with $K = 0.9$ for rod-like crystallites and θ being the diffraction angle. Because the measured peak width also contains a contribution from instrumental broadening, the calculated D_{hkl} values may underestimate the true crystallite size and are therefore interpreted as approximate lower bounds. Accordingly, they are used primarily as relative descriptors for comparison between samples rather than as absolute crystallite-size measurements.

XRF spectra were analyzed using pyMCA³⁶. After identifying all elemental lines in the scan (see Fig. 4), pyMCA performs per-point fitting and returns concentration-like intensity maps. For each element e , mean intensities were computed in the inscription region I_e^{ink} and in the rest of the manuscript region I_e^{support} . Since absolute concentrations require detailed knowledge of the sample matrix and thickness, we work with relative, composition-based descriptors. Hence, centered log-ratios (clr) were computed as:

$$\text{clr}_e^{\text{ink/support}} = \ln \frac{I_e^{\text{ink/support}}}{g(I^{\text{ink/support}})} \quad (8)$$

where $g(I^{\text{ink/support}})$ is the geometric mean of detected elements in the ink/support region.

Machine learning-based statistical analysis

Machine-learning (ML) classification models provide a practical way to test whether measured physical descriptors ("features") contain enough information to discriminate between sample categories ("labels"), such as origin, genus, or processing stage. Since different ML algorithms impose different assumptions about the structure of the decision boundary (linear vs. non-linear, parametric vs. non-parametric, ensemble-based vs. kernel-based), comparing multiple models is informative in itself: strong performance across several model families suggests that the separation is robust, whereas

performance restricted to a single model type can indicate a more specific structure in the data (e.g., non-linear separations, feature interactions, or strong collinearity). Comparing different models also allows us to select the most appropriate model for our problem.

A data frame was built in Python with all the measurements' metadata together with the physical parameters extracted from SAXS, WAXS and XRF. Missing feature values were imputed using the feature-wise mean calculated independently within each training fold. Across the complete feature table, missingness ranged from approximately 4–5% for the elemental clr descriptors to 12–15% for the SAXS and WAXS descriptors. Missingness was not distributed uniformly among sample groups; detailed proportions by feature family are provided in Supplementary Section S6 and Supplementary Table S1, while the complete pattern of available and missing values is provided in Supplementary Data 3. The data frame can be filtered for the analysis of specific groups of samples, for example, only manuscripts or only from one beamtime. For each analysis, a target property (e.g., origin, stage or genus) is selected, and the dataset is divided into measurements with known and unknown labels. Labels for genus, origin, and processing stage were assigned from the best available catalog information, donor documentation, botanical analysis, fieldwork records, and expert assessment. Measurements without an assignable label were excluded from the corresponding supervised task. However, the reliability of the available labels is not uniform, and some included metadata, particularly geographical origin, manuscript age, and historical processing, remain approximate. Possible label uncertainty, therefore, constitutes an additional source of error that is not represented by the cross-validation scores.

Known labels are encoded into integer classes. A minimum of nine measurements per class was used as a pragmatic inclusion threshold. This retained the smallest origin group while excluding classes with still fewer observations, for which stratified cross-validation would be highly unstable. The threshold should not be interpreted as an optimal minimum sample size, and results for classes close to this limit are considered exploratory. For each classification task, the number of samples per class is summarized in Table 1. Origin classification was performed on processed samples only.

Six different algorithms were considered for the classification of the samples: Random Forest (RF), Support Vector Machine (SVM), Logistic Regression (LR), Principal Component Analysis followed by Logistic Regression (PCA+LR), Gradient Boosting (GB) and Gaussian Process Classifier (GP). These models were imported from the scikit-learn 1.8.0 module in python³⁷. The main properties of each model are summarized in Table 2.

No external hyperparameter optimization was performed. Fixed model settings were used for all classification tasks: RF used 200 trees and a maximum depth of 10; SVM used an RBF kernel with $C = 1$ and

Table 1 | Number of measurements per class used for the main classification tasks

Classification task	Class	Number of measurements
Genus	<i>Corypha</i>	100
	<i>Borassus</i>	83
Processing stage	Manuscript	154
	Unprocessed	43
	Writing support	41
Origin	India	78
	Indonesia	62
	Thailand	21
	Sri Lanka	14
	Myanmar	12

Origin classification was performed using processed measurements only.

gamma="scale"; LR used $C = 1$ and the `lbfgs` solver; PCA retained 95% of the explained variance; and GB used 300 estimators, a learning rate of 0.05, and a maximum depth of 3. The GPC used a constant-times-RBF kernel with an added white-noise component, whose kernel parameters were optimized internally during each model fit.

To compare the algorithms, stratified 5-fold cross-validation was applied to the labeled dataset for each classification task. The measurements were divided into five folds while preserving the class proportions as closely as possible. In each iteration, four folds were used for model training, with the remaining fold used for validation, ensuring that each labeled measurement served once as validation data. No separate held-out test set was used; all reported performance values were obtained from cross-validation of the complete labeled dataset. For algorithms sensitive to feature scaling, standardization was performed independently within each training fold to avoid information leakage:

$$z = \frac{X - \mu_{\text{train}}}{\sigma_{\text{train}}}, \quad (9)$$

where X is the original feature value, and μ_{train} and σ_{train} are the mean and standard deviation estimated exclusively from the training portion of the corresponding fold. The resulting transformation was then applied to the validation portion of that fold.

Cross-validation was performed at the measurement level. Because separate scan regions or folios from the same object were retained as individual measurements, observations from one object could occur in both the training and validation folds. The reported scores, therefore, quantify classification of measurement-level descriptors and may overestimate performance for entirely unseen objects.

Model performance was quantified using the mean cross-validated balanced accuracy (CV-mean), with fold-to-fold variability reported as the standard deviation (CV std). Balanced accuracy was chosen because the class sizes were imbalanced, and this metric averages recall across classes, thereby reducing bias toward the majority class. For each target property, the best-performing model was selected as the one with the highest CV-mean.

To assess whether the observed performance reflected informative structure rather than trivial class imbalance, the best-performing model was compared with two dummy classifiers implemented in `scikit-learn`³⁷: `most_frequent`, which always predicts the majority class of the training fold, and `stratified`, which predicts randomly according to the class proportions of the training fold. Both dummy classifiers were evaluated using the same cross-validation splits as the trained models. In addition, permutation feature importance was computed by randomly shuffling the class labels while keeping the features unchanged, and the full cross-validation evaluation was repeated 50 times. A model was considered

informative when its performance exceeded both the strongest dummy baseline and the 95th percentile of the permutation-null distribution.

To visualize classification behavior across classes, confusion matrices were computed from out-of-fold (OOF) cross-validation predictions. These matrices summarize the correspondence between true and predicted classes and were reported in normalized form, with each row representing the fraction of measurements in a true class assigned to each predicted class.

In parallel with multivariate classification, the distributions of individual features across classes were examined using box plots. For each feature, the statistical test was selected based on tests of normality and homogeneity of variance. When these assumptions were satisfied, differences among classes were evaluated using one-way analysis of variance followed by Tukey's honestly significant difference test for pairwise comparisons. Otherwise, the Kruskal–Wallis test was used, followed by Dunn's pairwise tests with Holm adjustment for multiple comparisons. Global p -values obtained across the evaluated features were additionally adjusted using the Benjamini–Hochberg false-discovery-rate procedure. Pairwise differences were considered statistically significant when $p_{\text{adj}} \lesssim 0.05$.

Some algorithms, including logistic regression, random forests, and gradient boosting, provide native measures of feature importance based on model coefficients or impurity-based splits learned during training. For these models, both native and permutation-based feature importances were evaluated, as they do not necessarily rank variables in the same way. Permutation feature importance was calculated by randomly shuffling a feature's values across samples while keeping all other features unchanged and measuring the resulting decrease in model performance. Features showing high native importance but low permutation importance may contribute to the definition of decision boundaries without being strongly discriminative on their own, or may encode information already captured by correlated variables. Because correlated descriptors can share predictive information and influence feature-importance rankings, relationships among descriptor families and individual features were additionally evaluated using Spearman correlations (Supplementary Section S5 and Supplementary Figs. 6 and 7).

After model comparison and validation, the best-performing model for each property was retrained using 100% of the available labeled data after applying the minimum class-size filter. The trained classifier was then used to estimate the similarity of unlabeled measurements to each class. When supported by the classifier, class membership probabilities were extracted using `predict_proba`. The highest class probability was used as a confidence estimate, while the full probability vector was used to visualize how similar each unlabeled measurement was to each class.

This design provides a framework for assessing whether the extracted physical descriptors contain sufficient information to discriminate between sample classes, while controlling for class imbalance and spurious correlations.

Even when automatic labeling is not the primary goal, these probability outputs provide a representation of sample similarity. Samples assigned to a class with high confidence can be interpreted as strongly aligned with that class's feature profile. In contrast, low-confidence or ambiguous samples may indicate a lack of similar samples in the training set or greater sample complexity, for example, due to poor preservation or unusual preparation. In this way, the supervised models serve both as predictors and as structured similarity probes to support interpretation and guide targeted follow-up analyses.

Results

We evaluate the performance of the models described in Table 2 according to three kinds of classification properties: according to origin, genus and stage. Origin refers to the modern country associated with the manuscript. Genus refers to the botanical genera *Borassus* and *Corypha*, and stage refers to the processing of the sample. We distinguish unprocessed samples (raw, dry leaves), writing support (modern reference samples, with no writing or ink), and manuscripts, which include measurements on fragments as well as on folios belonging to manuscripts. Class distributions for each classification task are reported in Table 1 and are taken into account in the interpretation

Table 2 | Summary of the classification algorithms evaluated

Model	Linear	Scaling needed	Outputs probabilities	Typical strengths/weaknesses in this context
Random Forest (RF)	Non-linear	No	Yes (empirical; may be imperfectly calibrated)	Robust for heterogeneous tabular features; captures interactions automatically. Can be less interpretable and may bias importance toward some feature types.
Support Vector Machine (SVM)	Non-linear	Yes	Yes (when enabled; can be less stable)	Learns complex boundaries; often strong with limited samples. Sensitive to hyperparameters and scaling; training cost increases with dataset size.
Logistic Regression (LR)	Linear	Yes	Yes	Interpretable and stable; good when classes are close to linearly separable in feature space. Can underfit if separations are strongly non-linear.
PCA + Logistic Regression (PCA+LR)	Linear in PC space	Yes	Yes	Reduces dimensionality and collinearity; can improve stability when features are correlated. Interpretation shifts from physical features to principal components.
Gradient Boosting (GB)	Non-linear	No	Yes	Strong performance on structured/tabular data; captures subtle non-linearities. More sensitive to hyperparameters and can overfit if not controlled.
Gaussian Process Classifier (GPC)	Non-linear (kernel)	Yes	Yes	Probabilistic model with principled uncertainty; flexible decision functions. Computationally expensive; may be unstable with many samples/features.

Scaling refers to the standardization of features (zero mean, unit variance) performed within the training folds to avoid information leakage.

of the results, particularly for origin classification, where class imbalance is more pronounced, see Table 1. Because cross-validation was not grouped by object, the reported performance should be interpreted as measurement-level rather than independent object-level classification.

The groups under the classification properties are characterized by different physical features that may or may not have more or less complex correlations, hence the best model to describe classification is not necessarily the same in every case, obtaining indeed comparable cv-mean for Gradient Boosting (GB) and Random Forest (RF) for the genus classification, and RF for classifications according to stage and origin, see cv-mean values in Table 3. Figure 5 illustrates the difference in the performance of the algorithms. In the graph, we see that the differences between the algorithms for genus classifications tend to be small and within the standard deviation. GB and RF yield the same cv-mean, but GB was selected because it provides more balanced OOF predictions across the genera. In the case of the stage and origin classifications, we observe a difference between the top two performing models: tree models GB and RF, versus the rest of the models. This suggests that non-linear decision boundaries and feature interactions may be particularly relevant for these tasks.

To make sure the best performing models for each group are not doing a trivial classification based on class imbalance, we compare the cv-mean values with the performance obtained for the dummy classifiers and by the permutation null distribution. The values obtained are presented in Table 3. For all three classification tasks, the trained models passed the predefined informativeness criterion: their performance exceeded both dummy baselines and the 95th percentile of the corresponding permutation-null distribution.

Genus differentiation: intrinsic properties vs. processing effects

Two palm genera are mainly used for the production of palm-leaf manuscripts: *Borassus* and *Corypha*. From *Corypha*, the species *Corypha umbraculifera* L. (talipot palm) is the most commonly used, although *Corypha utan* Lam. is also common in some areas, and some cases have been recorded for *Corypha taliera*³⁸. From the *Borassus* genus, the species *Borassus flabellifer* L. (palmyra palm) is employed^{1,39}. The leaves of both genera are differentiated among other features by their venation, *Corypha* species exhibiting continuously distributed transverse veins and minor veins within the mesophyll for extra mechanical support, as leaves are larger. The venation in *Borassus* is brick-like. Additionally, *Corypha* leaves are dorsiventral (sides of the leaves are different), while *Borassus* leaves are generally isolateral (both sides of the leaves look the same). The size and type of the stomata (pores for gas exchange) are also a discriminating feature, being kidney-shaped and double the size in *Borassus*, and tetracytic in *Corypha*³⁸.

Even in processed stages, the palm genus used for the manufacture can usually be determined by the naked eye based on the thickness, stiffness, and venation architecture that is normally distinguishable on the foliage. However, the classification is not trivial when only fragments are available or when the degradation stage is advanced, and venation cannot be distinguished anymore. In these cases, microscopic inspection can be used to examine the stomata³⁸. Nevertheless, a single measurement at the synchrotron may provide information simultaneously on genus and provenance. As different areas use different palm species, genus or species determination already offers clues to geographical origin or dissemination routes³⁸.

In Fig. 6, the confusion matrix is built using the out-of-fold (cross-validated) predictions of each class: *Borassus* and *Corypha*. The results illustrate the categorization performance of the algorithm depending on the specific class. The numbers indicate the fraction of samples classified into the predicted category.

Both genera show a high fraction of samples correctly classified into their categories. While the aim was to achieve classification based on intrinsic genus-specific physico-chemical properties, the contribution of processing-related effects cannot be excluded and is examined in more detail below.



Table 3 | Model performance and baseline tests for each target property group.

Property	Best model	CV bal. acc.	Dummy MF	Dummy strat.	Perm. null	Perm. 95th
Genus	GB	0.75 ± 0.09	0.500	0.468	0.500 ± 0.051	0.583
Stage	RF	0.82 ± 0.04	0.333	0.339	0.335 ± 0.018	0.366
Origin	RF	0.68 ± 0.05	0.200	0.193	0.199 ± 0.024	0.237

Reported scores are balanced accuracy (mean ± standard deviation across cross-validation folds). Dummy baselines use `DummyClassifier` with `most_frequent` and `stratified` strategies. The permutation null distribution is obtained by shuffling labels and recomputing cross-validation scores (50 repeats).

Fig. 5 | Comparison of the mean cross-validated (CV) balanced accuracy achieved by the classification algorithms for the different target properties. The cv-mean is displayed above each bar.

Standard deviation is represented by horizontal black dashed lines. On top, the measurements considered for every classification are displayed: All samples for species and stage classification, and only processed samples for origin classification. The following color code is used for the bars: dark green, random forest; orange, support vector machine (SVM); purple, logistic regression; pink, principal component analysis plus logistic regression (PCA+LR); light green, gradient boosting; yellow, Gaussian process.

**Fig. 6 | Confusion matrix calculated from cross-validated out-of-fold (OOF) predictions of genus using the gradient boosting (GB) algorithm.** Rows represent the true genus and columns the predicted genus (*Borassus* or *Corypha*). Each cell shows the fraction of predictions within the corresponding true class, with darker blue indicating a higher fraction.

In Fig. 7, the top important features according to the permutation feature importance (a) and the assigned weights (b) are visualized. In Fig. 7b the features that were not ranked as important in the permutation feature importance have bars in a lighter orange color. Both rankings place the crystallinity index C_{200} and the elemental-centered log ratio clr_K as the most important features for classification. In the native importance chart, these features are followed by other elemental features: clr_{Ca} and clr_{Zn} , while in the permutation chart, the cellulose crystallite size in the (004) direction, D_{004} occupies third place. The elemental clr_{Cu} is also graded with high importance in both charts. Orientation parameters also play a role, but their importance is comparatively lower.

Box plots of some of the most important features are shown in Fig. 8, revealing overlapping distributions among genera that suggest that the model is relying on complex feature interactions rather than direct differences. C_{200} is the only feature displaying markedly different distributions between the genera, being slightly higher and broader in *Corypha*. The remaining features show largely overlapping distributions. Differences between processed and unprocessed groups can be seen in C_{200} and clr of chromium, copper and calcium. *Corypha* and *Borassus* both belong to the palm family, so it is not surprising that they show similar composition and structural properties, and a classification based on structural features is only achieved based on the combination of these. The classification is probably not achieved solely by intrinsic taxonomic markers; the effects of processing history and subsequent alterations are likely genus-specific and may also play a role in the classification.

Among the elemental descriptors, potassium and calcium may partly reflect the natural composition of the leaf substrate. However, none of the

Fig. 7 | Feature importance for genus classification. In (a) feature importance is calculated using permutation importance for the trained gradient-boosting (GB) classifier (green) and in (b) the classifier's native feature-importance measure (orange). Both methods identify clr_K and C_{200} as the two most influential descriptors. In (b), lighter orange indicates features assigned low importance by the permutation method. Here, clr_X denotes the centered log-ratio-transformed intensity of element X ; C_{hkl} and D_{hkl} denote, respectively, the cellulose crystallinity index and crystallite size associated with the (hkl) Bragg reflection; $\text{FWHM}_{\text{core}}^{\text{OD}}$ denotes the full width at half maximum of the core component in the orthogonal direction (OD); and I_b^{OD} denotes the intensity relative to the fitted baseline, b , in the OD.





Fig. 8 | Box plots of selected important features driving genus classification. The distributions are shown separately for *Borassus* (brown boxes) and *Corypha* (green boxes) genera and for unprocessed and processed groups. Here, clr_X denotes the

centered log-ratio-transformed intensity of element X ; C_{hkl} and D_{hkl} denote, respectively, the cellulose crystallinity index and crystallite size associated with the (hkl) Bragg reflection.

measured elemental signals can be assigned unambiguously to a single source. Although inscription regions were masked, carbon-based ink was spread across the folio and residues may remain outside the incised lines. The XRF descriptors may therefore contain contributions from the original leaf, preparation residues, ink, handling, environmental contamination, and later conservation treatments. This is particularly relevant for elements such as iron, manganese, calcium, and copper, while chromium, zinc, titanium, nickel, and arsenic are also difficult to interpret unequivocally. The box plots in Fig. 8 further show that several elemental descriptors vary with both genus and processing stage. Their importance in the classifier should therefore be interpreted as reflecting differences in the overall material history rather than as direct markers of botanical origin or a specific preparation step.

In an attempt to disentangle these contributions, the training and evaluation workflow described for genus classification was tested by training on different subgroups of samples based on their processing stage. One subset includes unprocessed samples (dry leaves) and the other contains processed samples (writing support and manuscripts). A comparison summary between the trainings is shown in Fig. 9. First, the training was repeated in exclusively processed samples, yielding similar results, achieving 0.78 cv-mean for the Gaussian Process classifier, similar top permutation-important features, and a similar confusion matrix, indicating that this sample group is informative enough for sample classification and even yields lower confusion coefficients *Borassus*-*Corypha*.

Then the workflow was repeated using only the unprocessed sample subset, achieving the best cv-mean of 0.65 for the Random Forest classifier,

sitting only slightly over the permutation 95th percentile mark. The associated confusion matrix (Fig. 9c) reveals strongly asymmetric behavior, with *Corypha* being classified much more successfully than *Borassus*. Therefore, the trained classifier is weakly informative but comparatively less robust, especially for *Borassus*. The permutation importance plots (Fig. 9d-f) further reveal that while full set and processed subset trainings rely on crystallinity-related, elemental, and orientational descriptors, the unprocessed subset training is driven entirely by C_{200} . This is consistent with the box plots in Fig. 8, where C_{200} shows partial genus-dependent separation in the unprocessed group but with substantial overlap between *Borassus* and *Corypha*, which would lead to high confusion coefficients.

The strong discrepancy between processed and unprocessed subsets indicates that the classification is not primarily driven by intrinsic genus-specific differences. Instead, the results suggest that processing and degradation introduce changes from which the algorithms learn. The full-dataset classification should therefore not be interpreted as botanical identification independent of the material history.

The full-sample model was retained because it provided the most comprehensive classification framework, maintained high cross-validated performance, and remained applicable to samples from all processing stages. After model evaluation, the selected classifier was refitted using all available labeled measurements and then applied to samples with missing genus labels, as shown in Fig. 10.

For objects measured more than once, the predicted class probabilities were averaged across measurements, and the corresponding standard




Fig. 9 | Genus-classification results for the full dataset and the processed and unprocessed subsets. **a** Maximum cross-validated (CV) balanced accuracy obtained for each dataset using Gradient Boosting (GB) for the full dataset (light green), Gaussian Process (GP) classification for the processed subset (yellow), and Random Forest (RF) for the unprocessed subset (dark green). **b, c** show the confusion matrices (CM) obtained from out-of-fold predictions for the processed and unprocessed subsets, respectively. **d–f** show the corresponding permutation feature importances

(PI) for the full, processed, and unprocessed datasets. Here, clr_X denotes the centered log-ratio-transformed intensity of element X ; C_{hkl} and D_{hkl} denote, respectively, the cellulose crystallinity index and crystallite size associated with the (hkl) Bragg reflection; $\text{FWHM}_{\text{core/halo}}^{\text{OD/LD}}$ denotes the full width at half maximum of the core or halo component in the orthogonal direction (OD) or longitudinal direction (LD); $I_b^{\text{OD/LD}}$ denotes the intensity relative to the fitted baseline, b , in the OD or LD; and $f_{\text{core}}^{\text{OD/LD}}$ denotes the relative area of the core contribution with respect to the halo.

deviations were used to represent within-object variability. Multiple measurements were included intentionally to sample different regions of the same object. In manuscripts, these regions could correspond to different folios or to distinct areas of a folio, such as text regions, edges, or locations near the binding hole. In the case of fragments, repeated measurements were sometimes acquired in comparable areas to assess consistency between beamlines or measurement campaigns.


Because the extracted descriptors can vary spatially within an object, identical predictions across repeated measurements are not necessarily expected. The averaged probabilities and their standard deviations nevertheless show that the variability is generally moderate and that repeated measurements usually lead to the same final genus assignment, differing mainly in confidence. This variability may reflect local differences in preservation, handling, or material accumulation. Edges, binding holes, first folios, and inscribed regions may be more strongly affected by use, exposure, inking, or conservation treatments, although inscription-related contributions were reduced as far as possible through the masking procedure described in the Methods.

Structural and compositional features across processing stages

A classification by processing stage of the sample is conducted to learn about the evolution of structural and compositional properties during the leaf's life. We compare, therefore, unprocessed leaves, reference writing supports (processed, modern), and manuscripts (processed and possibly affected by preservation strategies and degradation).

The confusion matrix for the palm leaf stage classification, obtained using the best-performing algorithm (RF), is shown in Fig. 11. It shows that most manuscripts are correctly categorized as manuscripts, with a fraction of 0.97. Unprocessed samples follow with 0.74, and last are the writing support samples with 0.73 correct classifications. Manuscripts are the most complex class evaluated, as they include measurements performed on a very heterogeneous group of palm leaves treated with different processes and of different ages, in contrast to unprocessed and writing-support samples, which are all modern. Of the writing supports, 15% were classified as manuscripts. This may reflect structural and compositional similarities associated with their shared status as processed leaves, although it does not imply that the same preparation methods were used. A further 12% were

Fig. 10 | Average predicted genus probabilities for unlabeled samples. On top of every sample bar the number of measurements considered for the calculation of the average and standard deviation (horizontal dashed lines). Bar section in brown corresponds to *Borassus*-predicted probability and bar section in green to *Corypha*.



classified as unprocessed, possibly because some writing supports underwent comparatively mild treatments.

Figure 12 shows the feature importance as calculated both using the permutation feature importance (a) and the native importance (b) of the RF algorithm. Only 3 features have a relevant effect on the permutation feature importance, and these are also among the top weighted features by the algorithm (native). The native important features of lighter colors have very small permutation importance, which may indicate that they are relevant to build decision boundaries or in combination with other features.

We see that the model is sensitive to centered log-ratio clr elemental concentrations of copper and nickel, as well as crystallinity index in the (200) crystallographic plane C_{200} . Native importance also grades with high values other elemental clr such as chromium, titanium, potassium or chlorine, and the amorphous contribution A . Box plots for feature distributions according to stage for selected top features are shown in Fig. 13.

Among the structural descriptors, C_{200} shows the clearest progression across processing stages. Unprocessed leaves exhibit the broadest distribution, writing supports display lower and more narrowly distributed values, and manuscripts are shifted further toward lower values. This pattern indicates that C_{200} is influenced not only by material preparation, but also by later ageing and degradation. C_{200} should therefore be interpreted as a combined descriptor of botanical origin, processing history, and subsequent material alteration.

The amorphous contribution A also changes with processing stage, but follows a non-monotonic trend. Values increase from unprocessed leaves to writing supports and then decrease again in manuscripts. This behavior suggests that A does not represent a simple progression toward greater or lower disorder, but instead reflects the combined influence of treatment-induced structural changes and longer-term ageing processes, which may have opposite effects.

The elemental variables require more cautious interpretation. clr_{Ni} shows a pronounced stage dependence, with lower values in manuscripts than in unprocessed leaves and writing supports, but this trend cannot be assigned unambiguously to a specific preparation or degradation mechanism. Similar uncertainty applies to other highly ranked elemental descriptors, such as clr_{Cu} , clr_{As} , and clr_{Cr} . These variables clearly contribute to classification, but their chemical origin may involve overlapping effects of intrinsic composition, preparation practices, environmental exposure, and conservation history.

Copper remains particularly relevant because clr_{Cu} generally increases in processed material, see Fig. 8. One possible source is the use of copper vessels during boiling, which has been reported in palm-leaf preparation and may facilitate metal transfer to the leaf substrate⁴⁰. Metal uptake can also depend on vessel composition and solution pH⁴¹. However, the XRF measurements do not allow copper originating from cookware, ink residues, conservation treatments, environmental contamination, or the intrinsic leaf composition to be distinguished. The observed increase, therefore, indicates

that copper may be introduced or redistributed during manuscript production or later material history, but its specific source cannot be identified from the present data.

Material similarities and differences across geographic origins

While genus and processing stage are absolute categories, origin categorization is a more relative concept that is based on modern geographical divisions and can reflect on different traditions shared across modern borders. In the same way, a single country like India may include many different traditions that we summarize in a single label. This is therefore an intrinsic limitation of our classification model. The objective of this analysis is not to establish a practical provenance-assignment tool, but to examine similarities and differences among the material profiles associated with the available origin labels, ideally based on the impact on the material of the different preparation methods and conservation strategies followed by different traditions. At the same time, some contribution from ageing-related processes can be expected, as suggested by the stage-dependent differences shown in Fig. 13. For this training, we reduced the measurement set by masking out unprocessed samples. This was decided because (i) Classification of samples according to genus was already difficult on unprocessed samples, implying that structural and compositional information may be less clear due to possible contaminants or lower contrast due to the higher amount of amorphous substances in the leaves. (ii) Most importantly, because we want the learning curve to be based on manuscript traditions rather than possible taxonomic or compositional differences between the same species coming from different regions. Still, genus-dependent processing traditions and material selection remain indirectly encoded in the dataset and may contribute to the observed separations. Stratified genus- and origin-classification performance is shown in Supplementary Fig. 8, while the corresponding genus-separated feature distributions are shown in Supplementary Fig. 9. These analyses show that some origin-related differences remain within a single genus, but several genus–origin combinations are absent or represented by very few measurements, preventing complete separation of the two effects. The following interpretation distinguishes between patterns directly observed in the measured descriptors and possible historical or technological explanations, which cannot be resolved independently within the present dataset.

Figure 14 shows the confusion matrix built from the OOF predictions. The results show that the descriptors contain origin-related information, but the degree of separation differs among countries. For Myanmar and Sri Lanka, the low recall may reflect genuine overlap with other material groups, limited available measurements or both. Their class-specific confusion patterns should therefore be regarded as exploratory. Repeating the origin classification with class-weighted models did not improve the recall of the smallest classes, indicating that their lower performance cannot be attributed solely to the absence of class weighting; the corresponding robustness analysis is presented in Supplementary Section S4 and Supplementary Fig. 5.

Within the present dataset, India, Indonesia, and Thailand show the highest class-specific recall values (>0.80), whereas Sri Lanka and Myanmar are correctly classified in only 0.43 and 0.42 of cases, respectively. Thai samples are correctly classified in 81% of the cases, but 19% of the samples were also classified as Indian. Myanmar is correctly classified in 42% of the cases, and tends to be confused with the rest of the countries except Sri Lanka. Resemblances between the classes may originate from similarities in the raw material (plant species and their specific taxonomy) as well as from similarities in the production methods of the manuscripts. These results are compatible with a recent review that analyzes published preparation recipes, focusing on the plants employed during the processing and later preservation strategies¹. Here, they report similar plant use between Indonesia and India, which could explain the India-Indonesia and Indonesia-India confusion coefficients. The impact of plant use extends beyond compositional differences, as substances introduced by these plants may determine material chemical response during boiling and favor specific degradation pathways. Based on plant use, Poliakova et al. also report high similarity coefficients between Tamil Nadu (India) and Sri Lanka¹, where they found that plants used in Sri Lanka for manuscript production are identical to those of Tamil Nadu, with the exception of two plant species used for preservation exclusively in the latter region. Since most of the measured and

evaluated Indian samples in this work are from Tamil Nadu and Kerala, this could explain the high confusion coefficient between India and Sri Lanka that we also obtained. Since Sri Lanka is close to Tamil Nadu and Kerala, it is also plausible that some of their manuscript-preparation recipes and preservation strategies are similar, which could also contribute to the high confusion coefficient. While descriptions of these processes are not detailed, Subramaniam reports two different recipes for tender and mature leaves in Tamil Nadu⁴², both of them involving boiling in water. Different sources describe different Sri Lankan recipes demonstrating variability^{2,43,44}, although all of them agree that leaves need to be boiled. It is plausible that some of these recipes result in manuscripts more similar to Indian, while some other recipes may include specific Sri Lankan traditions or plant use, leaving a distinctive print that allows differentiation from Indian manuscripts, thereby being correctly classified by our algorithm (43% of cases).

Myanmar's confusion coefficients are difficult to interpret, as the literature on the manuscript production in this area is very scarce. Similarities to Indian and Sri Lankan preparation recipes are acknowledged². While the Myanmar-India coefficient is indeed 0.25, the Myanmar-Sri Lanka coefficient is zero. Since Myanmar is the smallest class, it is plausible that our sample set does not include specific examples reflecting high similarity to Sri Lankan manuscripts. It could also be the case that these similarities are eclipsed by the stronger similarities to India and Indonesia. The distribution of out-of-fold predictions over different origins could be an indication of high diversity in the production methods, but it could also mean that the production methods result in average properties of all structural and compositional parameters, thus complicating the categorization of these manuscripts.

Figure 15 shows the top important features as calculated using the permutation method (a), and as established by the native weights imposed by the random forest algorithm (b). The features driving classification are compositional and crystalline structure-related, while fiber orientation-associated features do not appear among the top graded features. All features in the permutation importance chart also appear in the native importance chart in a different order. However, the native importance chart contains additional features in light orange that are not present in the permutation importance chart. These may contribute to the construction of internal decision boundaries. The top common important features in both panels include the crystallinity index of the Bragg (200) reflection, as well as compositional centered-log ratios clr of zinc, manganese, potassium and calcium. Figure 16 shows box plots for some of the top important features, where clear distinctions are observed between regions.

The clr elemental distributions of the displayed elements are, in general, similar across manuscripts of different origins and probably acted in the classifier to build decision boundaries or in combination with other features. Exceptions are manganese and potassium clr concentrations, which show distinctive ranges in the box plots according to the origin of the samples. For manganese, the differences are especially important between India/Sri Lanka/Myanmar and Indonesia/Thailand. The clr concentration of potassium distributions for different origins results in high differentiation of





Fig. 11 | Confusion matrix calculated from out-of-fold (OOF) cross-validated predictions of stage with the Random Forest (RF) algorithm. True stage class shows fraction of predicted class: unprocessed, writing support, and manuscript, with darker blue indicating a higher fraction.

Fig. 12 | Feature importance for processing stage classification. In (a) feature importance was calculated using permutation feature importance (green) and in (b) the native importance (orange) on the best trained model: Random Forest. Features in light orange in the native importance chart were graded with low importance in the permutation feature importance. Here, clr_X denotes the centered log-ratio-transformed intensity of element X ; C_{hkl} and D_{hkl} denote, respectively, the cellulose crystallinity index and crystallite size associated with the (hkl) Bragg reflection; A denotes the amorphous contribution.






Fig. 13 | Box plots of the most important features driving processing-stage classification: the cellulose crystallinity index associated with the (200) Bragg reflection (C_{200}), the amorphous contribution A , and the centered log-ratio-transformed intensity of nickel (clr_{Ni}). Adjusted global p -values (p_{adj}) shown above

the box plots were obtained using the Kruskal–Wallis test, with Benjamini–Hochberg adjustment across the evaluated features. Boxes representing unprocessed samples are shown in green, writing supports in orange, and manuscripts in purple.




Fig. 14 | Confusion matrix calculated from out-of-fold (OOF) predictions during cross-validation of origin classification, using the Random Forest (RF) algorithm. True stage class shows a fraction of the predicted class: India, Indonesia, Myanmar, Sri Lanka and Thailand. Each cell shows the fraction of predictions, with darker blue indicating a higher fraction.

Indonesian samples from the other classes due to its lower values, but all distributions are actually quite different and a very low adjusted coefficient p_{adj} was obtained, denoting high statistical reliability of the data. Elemental differences between samples of different origin can be due to both the origin of the palm and the composition of the soil, as well as the elements introduced over time.

Several facts may have an influence on the amorphous and crystalline contents of the manuscripts. Preparation methods have a direct impact on cellulose crystalline structure but also an indirect influence, as the resulting material may be more or less prone to certain degradation mechanisms. These features were already investigated in previous sections (see Figs. 8 and 13), where we observed that the crystallinity index moves in a smaller range for *Borassus*, meaning that differences between regions may be partially determined by the raw material species used in every region.

The crystallinity index C_{200} is especially distinctive for Thai samples, that show very high values in Fig. 16. This is consistent with the more extended use of *Corypha* in this area, which shows higher crystallinity indices than *Borassus*, see Fig. 8. However, the distribution reveals an average value of ≈ 0.025 for *Corypha*, and the average crystallinity index for Thai samples would sit right at the top of the *Corypha* box over ≈ 0.05 , hence additional factors may play a role.

While the specific preparation of the measured Thai manuscript samples is unknown, some Thai preparation recipes are reported to be very unique, using a kiln to bake the leaves instead of the usual boiling method and resulting in the exudation of a black substance by the leaves¹⁴. The chemical nature of this substance is, to our knowledge, not identified, but we hypothesize that it may correspond to a non-cellulosic amorphous substance, which could influence the apparent crystallinity. Additionally, the leaves are exposed to fire after baking. Both treatments could contribute to structural rearrangement during drying, thereby increasing the measured crystallinity index, although this mechanism is not directly tested in the present work and needs further investigation.

The amorphous contribution A (Fig. 16) also stays low for Thai samples. This descriptor is not an amorphous-to-crystalline ratio, but is calculated from the fitted area of the broad amorphous scattering contribution relative to the total positive scattering intensity within the selected q -range. The amorphous contribution A is also comparatively low for the Thai group. Because this descriptor represents the relative prominence of a fitted broad scattering contribution rather than a chemically identified amorphous fraction, its low value cannot be linked directly to a potential exudation of material during kiln treatment. Botanical composition, sample age, preparation, and preservation history may all contribute to the observed distribution. A also presents high values for Myanmar and Sri Lanka. This feature may be key for the cases in which Sri Lankan and Myanmar manuscripts are correctly classified.

As discussed in the previous section, the smaller difference between unprocessed leaves and writing supports suggests that manuscript ageing may have a larger impact on crystallinity than the processing methods themselves, which may contribute to the high crystallinity observed in the Thai group, as many of the measured samples correspond to modern writing supports. Nevertheless, preparation methods are expected to influence the degradation pathways experienced by manuscripts over time. The selection of leaves at different maturity stages may also contribute to differences in amorphous and crystalline content, thereby influencing provenance determination. The crystalline and amorphous contributions, which strongly influence origin classification, are therefore complex and multifactorial, likely reflecting a combination of processing history, material selection, and degradation effects that cannot be resolved individually within the present dataset.

Ageing and degradation could not be included as independent variables in the statistical analysis. The age of many manuscripts is unknown, or only approximately assigned, and chronological age does not necessarily represent their degradation state, which also depends on environmental exposure, handling, storage, and conservation history. Because no consistent degradation indicator was available across the full dataset, these effects could not be corrected for or evaluated as a separate classification target.

Fig. 15 | Feature importance for classification according to origin. **a** Permutation (green) and **b** native (orange) feature importance of the random forest algorithm, trained to distinguish between samples of different origins. Feature bars are light orange in the native importance chart **(b)**, and those graded with low importance in the permutation feature importance. Here, clr_X denotes the centered log-ratio-transformed intensity of element X ; C_{hkl} and D_{hkl} denote, respectively, the cellulose crystallinity index and crystallite size associated with the (hkl) Bragg reflection; A denotes the amorphous contribution.





Fig. 16 | Box plots for the top important feature driving origin classification with the random forest algorithm. Boxes have assigned colors: India, dark green; Indonesia, orange; Thailand, purple; Sri Lanka, pink; Myanmar, light green.

Adjusted pairwise p -values (p_{adj}) shown above the box plots were obtained using the Kruskal–Wallis test, with Benjamini–Hochberg adjustments across the evaluated features.

Despite this complexity, the Random Forest origin classifier performs well for some countries but remains unreliable for the smallest groups, especially Myanmar and Sri Lanka. Their confusion with other classes may reflect genuine similarities in material composition or preparation practices. Such similarities are partly supported by the work of Poliakova et al.¹, and plausible historical or technological explanations are proposed above. However, these interpretations remain hypotheses that cannot be confirmed with the present data. Moreover, the limited number of measurements available for these groups is likely to affect their classification performance, although its contribution cannot be separated from genuine material similarities.

The classification results are further constrained by the dataset's observational nature. The measured groups differ in botanical material, processing history, age, degradation, storage, handling, and conservation, while the associated metadata vary in completeness and certainty. Cross-validation evaluates reproducibility relative to the assigned labels but cannot establish that these labels are historically exact or separate the contributions of correlated variables. The reported performances and feature associations

should therefore be interpreted as patterns within the present reference dataset rather than as universally transferable causal relationships.

The model should therefore be used mainly to identify material similarities between groups rather than to assign provenance with certainty. Larger and more balanced datasets will be needed to determine whether the observed confusion patterns represent shared manuscript traditions or insufficient class representation.

Discussion

In this work, we presented a non-invasive multimodal synchrotron approach for the characterization of palm leaves and palm-leaf manuscripts. Physical descriptors derived from SAXS, WAXS, and XRF measurements were combined with supervised machine learning to investigate differences among samples. Genus-related distinctions were most clearly expressed in processed materials and reflected a combination of intrinsic botanical properties, processing-induced modifications, and subsequent material alteration. The analysis highlights that both structural markers (crystallinity- and fibril-related parameters) and compositional markers (in

particular selected elemental signatures) contribute to the classification and provide interpretable links to material properties.

Beyond predictive performance, the workflow enables a comparative assessment of manuscript traditions by quantifying similarities among labeled samples. Within the present dataset, the results reveal regional similarities, including overlap between Sri Lankan and South Indian material profiles. Thai samples show a comparatively distinct and consistent feature profile. This separation may be related to reported preparation practices, such as kiln treatment, although the underlying mechanisms remain to be verified directly.

The Myanmar group shows a more heterogeneous classification pattern. This may reflect genuine variability in preparation practices, but the limited number of measurements prevents a firm interpretation. More generally, when applied to manuscripts of unknown origin or production history, the classifier provides a probability-based similarity profile across the learned classes. Such outputs should not be interpreted as definitive provenance assignments, but can support further analysis when combined with codicological, historical, and technological evidence.

These findings demonstrate the potential of non-destructive synchrotron methods for comparing heritage materials and for identifying material patterns potentially related to production, use, ageing, and conservation history.

The present dataset is diverse but does not capture the full variability of palm-leaf manuscript traditions, and some geographic classes remain underrepresented. This limits the generalizability of the results and may contribute to differences in classification performance among groups. Future work should therefore expand the reference dataset by including objects from additional collections and manuscript traditions. Further measurements on controlled reference materials are also needed to clarify how specific treatments influence the structural and compositional evolution of palm leaves over time.

Data availability

The raw, corrected, and integrated data of the synchrotron X-ray measurements at the ESRF are available in the ESRF data repository after the embargo period (April 2027)³⁵. DESY datasets, including raw, corrected, and integrated data, as well as background-subtracted data, are listed in the DESY data repositories and are available upon request^{33,34}. The DESY repositories also contain photos of the manuscript folios and the rough measurement area considered in each scan. Sample metadata and measurement information are provided in Supplementary Data 1 and 2, and the master table containing the extracted descriptors used in the statistical analysis is provided in Supplementary Data 3. Additional details on the fits performed and the machine learning outputs are available in TORE¹⁵.

Code availability

Custom Python scripts and Jupyter notebooks used for data processing and analysis are available in ref. 45.

Received: 12 May 2026; Accepted: 16 September 2026;

Published online: 06 October 2026

References

1. Poliakova, A., Ciotti, G., Perumal, P. & Balachandran, N. Plants used for palm-leaf manuscript production and conservation in South and Southeast Asia: a review of the literature. *Bot. Lett.* **173**, 333–358, <https://doi.org/10.1080/23818107.2026.2634848> (2026).
2. Wiland, J. et al. A literature review of palm leaf manuscript conservation-Part 1: a historic overview, leaf preparation, materials and media, palm leaf manuscripts at the British Library and the common types of damage. *J. Inst. Conserv.* **45**, 236–259 (2022).
3. Freeman, R. Turning over old leaves: palm leaves used in South Asian manuscripts. *Book Pap. Group Annu* **24**, 99–102 (2005).
4. Marrot, L., Lefeuvre, A., Pontoire, B., Bourmaud, A. & Baley, C. Analysis of the hemp fiber mechanical properties and their scattering (Fedora 17). *Ind. Crops Prod.* **51**, 317–327 (2013).
5. Sharma, D., Singh, M., Krist, G. & Velayudhan, N. M. Structural characterisation of 18th century Indian Palm leaf manuscripts of India. *Int J. Conserv Sci.* **9**, 257–264 (2018).
6. Poliakova, A., Ciotti, G., Helman-Wazny, A. & Fromm, J. Phytolith assemblages from palm leaves and palm-leaf manuscripts: what is the difference and what it could mean?. *Front. Plant Sci.* **15**, 1482790 (2024).
7. Poliakova, A., Ciotti, G. & Blinnikov, M. Historical use of grasses in south Indian palm-leaf manuscript production reconstructed from phytolith analysis. *Rev. Palaeobot. Palynol.* **352**, 105610 (2026).
8. Helman-Wazny, A. & Grabowsky, V. Postcards from Laos. <https://www.csmc.uni-hamburg.de/about/blog/2023-03-02-laos.html> (2023).
9. Helman-Wazny, A. & Grabowsky, V. Postcards from Thailand. <https://www.csmc.uni-hamburg.de/about/blog/2023-03-13-thailand.html> (2023).
10. Helman-Wazny, A. & Grabowsky, V. A Fading Tradition, Part 1: Illuminating the Living Tradition of Lao Manuscript Culture. <https://www.csmc.uni-hamburg.de/news/2025-07-09-a-fading-tradition-laos.html> (2025).
11. Grabowsky, V. Tai manuscripts in the Dhamma script domain: surveying, preservation and documentation, part 1. *Manuscr. Cult.* **1**, 16–23 (2009).
12. Acri, A. *Palm-Leaf Manuscripts in Today's Bali*. <https://www.insideindonesia.org/editions/edition-121-jul-sep-2015/palm-leaf-manuscripts-in-today-s-bali>. Edn 121 (2015).
13. Lessler, F. Preserving Bali's Cultural & Literary History through the Palm Leaf Project. <https://blog.archive.org/2019/10/09/preserving-balis-cultural-literary-history-through-the-palm-leaf-project/> (2019).
14. Agrawal, O. P. *Conservation of Manuscripts and Paintings of South-East Asia* (Butterworths, London, 1984).
15. Chu, S., Lin, L. & Tian, X. Evaluation of the deterioration state of historical palm leaf manuscripts from Burma. *Forests* **14**, 1775 (2023).
16. Zhou, Y., Wang, K. & Han, L. Synergistic ATRP-IL approach for improving the mechanical properties of fragile palm leaf manuscripts. *npj Herit. Sci.* **13**, 510 (2025).
17. Schulz, A., Lautner, S., Fromm, J. & Fischer, M. Not stealing from the treasure chest (or just a bit): analyses on plant derived writing supports and non-invasive DNA sampling. *PLoS ONE* **13**, e0198513 (2018).
18. Sharma, D., Singh, M. R. & Dighe, B. Chromatographic study on traditional natural preservatives used for palm leaf manuscripts in India. *Restaurator* **39**, 249–264 (2018).
19. Singh, M. R. & Sharma, D. Investigation of pigments on an Indian palm leaf manuscript (18th–19th century) by SEM-EDX and other techniques. *Restaurator* **41**, 49–65 (2020).
20. Sharma, D., Singh, M., Krist, G. & Nair, M. V. Pigment analysis of palm leaf manuscripts of India. *Curr. Sci.* **118**, 285–292 (2020).
21. Delhey, M., Kindzorra, E., Hahn, O. & Rabin, I. Material analysis of Sanskrit palm-leaf manuscripts preserved in Nepal. *J. Int. Assoc. Buddh. Stud.* **36/37**, 119–152 (2015).
22. Voges, L. F., Horn, N., Ciotti, G. & Seifert, S. Differentiation of species and provenance of palm-leaf manuscripts using Fourier transform infrared spectroscopy and chemometrics. *Chemosensors* **13**, 196 (2025).
23. Sathiyamani, S. et al. Material characterisation of 19–20th century manuscripts from Northern Thailand. *Restaurator* **45**, 117–140 (2024).
24. Ciotti, G. & Bosch, S. *The Glory of Vishnu: Tracing the Origins of Palm-leaf Manuscripts*. <https://www.csmc.uni-hamburg.de/hamburgs-schriftschatze/english/the-glory-of-vishnu.html> (2023).

25. Guo, J. et al. Molecular and crystal structures of cellulose in severely deteriorated archaeological wood. *Cellulose* **29**, 9549–9568 (2022).
26. Svedström, K., Bjurhager, I., Kallonen, A., Peura, M. & Serimaa, R. Structure of oak wood from the Swedish warship Vasa revealed by X-ray scattering and microtomography. *Holzforschung* **66**, 355–363 (2012).
27. Wess, T. J. et al. The use of small-angle X-ray diffraction studies for the analysis of structural features in archaeological samples. *Archaeometry* **43**, 117–129 (2001).
28. Haas, S. & Helman-Ważny, A. The application of small- and wide-angle x-ray scattering to the study of the paper support and pigment distribution. In *The Zhangzhung Nyenyü 'Tsakalis': a Cross-Disciplinary Analysis*. 95–113 (De Gruyter, 2025).
29. Bergmann, U. & Knox, K. T. Pseudo-color enhanced x-ray fluorescence imaging of the Archimedes Palimpsest. In *Document Recognition and Retrieval XVI*, vol. 7247 of *Proceedings of SPIE-IS&T Electronic Imaging* (eds. Berkner, K. & Likforman-Sulem, L.) 724702 (SPIE, 2009).
30. Pouyet, E. et al. Revealing the biography of a hidden medieval manuscript using synchrotron and conventional imaging techniques. *Anal. Chim. Acta* **982**, 20–30 (2017).
31. Mocella, V., Brun, E., Ferrero, C. & Delattre, D. Revealing letters in rolled herculeanum papyri by x-ray phase-contrast imaging. *Nat. Commun.* **6**, 5895 (2015).
32. Haas, S., Sun, X., Conceição, A. L. C., Horbach, J. & Pfeffer, S. The new small-angle x-ray scattering beamline for materials research at Petra III: Saxsma beamline p62. *J. Synchrotron Radiat.* **30**, 1156–1167 (2023).
33. Huber, P. Palm leaves as writing support: compositional and multiscale structural analysis to trace the history of Asian manuscripts. <https://doi.org/10.60883/E007EE79FB> (2026).
34. Huber, P. Understanding the production of palm leaf manuscripts: additional measurements for data completeness. <https://doi.org/10.60883/410671131A> (2026).
35. Gallardo Domínguez, L. & Busch, M. *Palm Leaves as Writing Support: Compositional and Multiscale Structural Analysis to Trace the History of Asian Manuscripts*. <https://doi.org/10.1515/ESRF-ES-1520005738> (ESRF, 2027).
36. Solé, V. A., Papillon, E., Cotte, M., Walter, P. & Susini, J. A multiplatform code for the analysis of energy-dispersive x-ray fluorescence spectra. *Spectrochim. Acta Part B* **62**, 63–68 (2007).
37. Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
38. Chen, Q. et al. Plant selection in ancient Tibetan palm-leaf manuscripts: a novel approach to rapid species identification. *npj Herit. Sci.* **14**, 116 (2026).
39. Sah, A. Palm Leaf manuscripts of the world: material, technology and conservation. *Stud. Conserv.* **47**, 15–24 (2002).
40. Padmakumar, P. K., Sreekumar, V. B., Rangan, V. V. & Renuka, C. Palm Leaves as writing material: history and methods of processing in Kerala. *Palms* **47**, 125–129 (2003).
41. Ali Sultan, S. A. et al. Assessing leaching of potentially hazardous elements from cookware during cooking: a serious public health concern. *Toxics* **11**, 640 (2023).
42. Subramaniam, P. Palm-leaf manuscripts in Tamil. In *Palm-Leaf and Other Manuscripts in Indian Languages. Proceedings of the National Seminar* (eds. Hikosaka, S. & Samuel, G. J.) 169–193 (Institute of Asian Studies, Pondicherry University, Pondicherry, 1995).
43. Alahakoon, C. N. K. The division of labour in the production of Sri Lankan palm leaf manuscripts. *Tech. Rep.* **2**, <https://www.jstor.org/stable/43855214>. (2012).
44. Nishanthi, M. & Wijayasundara, N. D. Preservation and conservation of palm leaf manuscripts at the Library of University of Sri Jayewardenepura. *Vidyodaya J. Humanit. Soc. Sci.* **7**, 81–96, <https://doi.org/10.31357/fhss/vjhss.v07i02.06> (2022).
45. Gallardo, L., Busch, M. & Huber, P. Synchrotron SAXS, WAXS and XRF-derived structural and compositional descriptors of palm leaves and palm-leaf manuscripts and statistical analysis. <https://doi.org/10.15480/882.17456> (2026).

Acknowledgements

We acknowledge fruitful discussions with Malgorzata Grzelec (Center for the Study of Manuscript Cultures, CSMC, University of Hamburg) and Hossein Ostovar (Institute of Chemical Reaction Engineering, Hamburg University of Technology) for his assistance with machine learning methods. The authors gratefully acknowledge the institutions that provided access to manuscripts and contributed fragments for analysis. In particular, we thank the CSMC and its library, as well as the African Asian Institute (AAI) of the University of Hamburg, whose manuscripts are on permanent loan at the CSMC. We also thank the individuals who generously provided access to objects for measurement: Prof. Ewa Wilden, Prof. Volker Grabowsky, Prof. Jörg Fromm, Dr. Chieng Li, Alex Malik, and Malgorzata Grzelec. We thank Dr. Alan Paton (Head of Science Collections), Dr. William J. Baker (Herbarium Curator and a Senior Research Leader) and MSc Saba Rokni (curator-botanist) of the Royal Botanic Gardens (Kew, UK) for providing us with herbarium reference material from RBG Kew. We thank Dr. Irina Wandrey and Dr. Katrin Janz-Wenig for access to the manuscripts stored in the State and University Library (Carl von Ossietzky Staats- und Universitätsbibliothek Hamburg), University of Hamburg, Germany. Prof. Dominic Goodall is acknowledged for access to the manuscript collection of the Ecole Française d'Extrême-Orient (EFEO) in Puducherry, India. Dr. Margherita Trento and Dr. Deviprasad Mishra are acknowledged for access to the manuscript collection of the Department of Indology, French Institute (IFP) in Puducherry, India. Prof. Hugo David, Dr. Olga Nowicka and Dr. Murali Krishnan are acknowledged for help in collecting fallen off micro-fragments of the PLMs stored and digitized at the Kerala Manuscript Preservation Center and the Vedic Research Center (VRC), Vadakke Madham Brahmaswam (VMB) in Thrissur (India) and for research collaboration with the Project “Digital Preservation of Kerala Archives” (DiPiKA, CSMC Hamburg-VRC/VMB, Thrissur). Samples were obtained in the frameworks of the “Palm-Leaf Manuscript Profiling Initiative” (PLMPI) and the project “Paleogenomic Studies of Written Artifacts of Different Origin.” We thank Prof. Agnieszka Helman-Ważny for providing leaves in different stages of preparation collected during her fieldwork. We thank the synchrotron facilities that granted beamtime for this research, the European Synchrotron Radiation Facility (ESRF) and Deutsches Elektronen-Synchrotron (DESY), as well as the beamlines BM02 and P62. We are grateful to the beamline scientists Dr. Isabelle Morfin and Dr. Gilbert Chachine (ESRF) and Dr. Sylvio Haas (DESY) for their support during the experiments.

Author contributions

L.G. conceived the study, performed the experiments, carried out data analysis, and wrote the manuscript. M.B. contributed to the design of the experiments and setups, performed the experiments and assisted with data preprocessing. A.P. obtained botanical material from herbarium collections and documented it in the Supplementary Data 3. Microscopically and/or genomically identified palm species for all samples were indicated. G.C. contributed to the assessment of manuscript origin based on codicological data and to the historical and manuscript-related interpretation. P.H. supervised the project and contributed to data interpretation. All authors reviewed and approved the final manuscript.

Funding

L.G. received support from the Deutsche Forschungsgemeinschaft (DFG) through the Excellence Strategy of the Federal Government and the federal states under EXC 2176/2 “Understanding Written Artifacts” (UWA). Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at

<https://doi.org/10.1038/s40494-026-03023-3>.

Correspondence and requests for materials should be addressed to Laura Gallardo.

Reprints and permissions information is available at <http://www.nature.com/reprints>

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

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

© The Author(s) 2026