Complete virtual unwrapping and reading of a rolled Herculaneum papyrus

Giorgio Angelotti[1*], Stephen Parsons[1,2], Federica Nicolardi[1,3], Youssef Nader[1], Sean Johnson[1], David Josey[1], Paul Henderson[1,4], Hendrik Schilling[1], Johannes Rudolph[1], Forrest McDonald[1], Elian Rafael Dal Prá[1], Paul Tafforeau[5], Alessandro Mirone[5], C. Seth Parker[2], Jan Paul Posma[1], Benjamin Kyles[1], Claudio Vergara[1,3], Alessia Lavorante[1,3], Rossella Villa[1,3], Maria Chiara Robustelli[1,3], Marzia D'Angelo[1,3], Gianluca Del Mastro[1,6], Michael McOsker[1,7], Kilian Fleischer[1,8], Christy Chapman[1,2], Nat Friedman[1], W. Brent Seales[1,2]

[1]Vesuvius Challenge, San Francisco, CA, USA; [2]EduceLab, University of Kentucky, Lexington, KY, USA; [3]Università degli Studi di Napoli Federico II, Napoli, Italy; [4]University of Glasgow, Glasgow, UK; [5]ESRF, Grenoble, France; [6]Università degli Studi della Campania Luigi Vanvitelli, S. Maria Capua Vetere, Italy; [7]University College London, London, UK; [8]Universität Tübingen, Tübingen, Germany.

*Corresponding author(s). E-mail(s): giorgio@scrollprize.org

Abstract

The carbonized papyrus scrolls from Herculaneum preserve the only large-scale library to survive from classical antiquity. Many scrolls have remained unread because physical unrolling risks irreversible damage. Non-invasive imaging offers a route to their texts, yet previous work on sealed Herculaneum scrolls has recovered only localized readings. Here, by combining high-resolution X-ray computed microtomography (μCT) data acquired on the BM18 beamline at the ESRF with machine-learning-assisted virtual unwrapping, we study three still-closed scrolls. For *PHerc.* 1667, a remnant of earlier attempts at physical unrolling, we achieve the complete unwrapping and reading of the preserved writing surface, making it the first Herculaneum papyrus read completely by virtual unwrapping. The text was transcribed and cross-verified by nine papyrologists. In *PHerc. Paris.* 4, the source of the Vesuvius Challenge 2023 Grand Prize text, the high-resolution scan makes ink directly visible in the tomographic volume, allowing 3D ink segmentation and comparison of recovered text with visible deposits. In *PHerc.* 139, we recover title and author-attribution evidence identifying the scroll as Philodemus, *On Gods*, Book 8. These results establish a basis for the ongoing, systematic recovery and scholarly study of texts preserved within the still-closed Herculaneum library.

The Herculaneum papyri present one of the longest-standing readability challenges in the study of the ancient world. Discovered in the eighteenth century in the Villa of the Papyri, they make up the only surviving large-scale written library from classical antiquity[1]. Carbonized by the eruption of Mount Vesuvius in AD 79, the rolls survived in a form that made reading and preservation fundamentally opposed: physical opening could reveal text, yet also irreversibly damage the material. Earlier campaigns of mechanical opening revealed extraordinary texts, especially Epicurean works, but also demonstrated the destructive cost of direct intervention. Hundreds of rolls and multi-layer roll fragments therefore remain sealed, with their contents preserved yet inaccessible[1].

Virtual unwrapping offers a non-invasive, computational approach to revealing these texts[2–11]. The recovery of two columns of Hebrew text from the severely damaged En-Gedi scroll established a general framework for this approach[2]. The En-Gedi framework operates entirely on data from X-ray computed microtomography (μCT), a form of computed tomography (CT), and combines recovery of the writing-surface geometry, rendering of the ink, digital flattening and merging of individual pieces into a coherent, readable image.

The potential impact of virtual unwrapping at Herculaneum is enormous in quantitative terms[12]. Of the 659 inventory numbers in the collection referring to unopened material, at least 388 may be sufficiently well preserved to contain substantial continuous texts, including 179 potentially complete or nearly complete scrolls[13]. Assuming an average length of 9–12 metres per scroll[14], these 179 scrolls alone could represent approximately 1.6–2.15 kilometres of papyrus, containing an estimated 21,000–36,000 columns of ancient writing. The significance is all the greater because the texts recovered and identified so far, both from physically unwrapped papyri and through partial attempts at virtual unwrapping, exclusively include works that have no parallel in any other source and were not transmitted through the medieval manuscript tradition. These — and, very likely, those still to be read — are not simply new copies of known works, but books preserving otherwise lost texts.

Extending the En-Gedi framework to the Herculaneum scrolls requires advances in both imaging and computation, with machine learning integrated into surface localization and ink detection. The extreme physical forces of the Vesuvius eruption resulted in pronounced bends, folds, breaks, and effacement of each scroll's papyrus structure, all of which pose significant challenges for the geometric modeling and flattening steps of virtual unwrapping. Additionally, the carbonized papyrus substrate on which the writing appears is radiographically similar to the carbon-based inks predominantly used across the Herculaneum collection. This leads to weak or even indistinguishable contrast between the ink strokes and the evidence of papyrus in the tomographic volume. Making matters of weak contrast even worse are the variations in ink recipes, ink contaminants, morphology of the ink deposits, and the state of preservation within and among Herculaneum scrolls. This situation means that writing may be expressed throughout the tomographic volume as a varying combination of faint metallic contrast, surface shape relief, or subtle textural signal rather than by a single, uniform signature.

Earlier studies explored several sources of signal in Herculaneum material, from phase-contrast imaging of rolled papyri to trace metallic components in some inks and fine surface deformations associated with ink on already-opened fragments[15–19]. While these and

other studies offered incremental progress toward reading the still-rolled papyri, a complete virtual unwrapping and reading remained out of reach.

We address these challenges through recent advances of the Vesuvius Challenge, an open-science innovation contest launched in 2023. The Vesuvius Challenge brought together computer science researchers, academic subject matter experts, and an online distributed community of contributors, making the challenge public through a deliberately hybrid model of competition and collaboration[20]. While prior work had established that text could be recovered from unopened or unexposed Herculaneum material[3, 21], the Vesuvius Challenge progressively scaled up the quality and quantity of readable text[22–25]. The First Letters Prize (2023) and the ensuing Grand Prize (2023) revealed multiple columns of readable text from *PHerc. Paris.* 4[25], while Vesuvius Challenge-sponsored Kaggle competitions expanded the methodological base for addressing the continuing challenges of reliable ink detection[26] and accurate surface localization[27]. Other Vesuvius Challenge milestones have included non-invasive recovery of the title of *PHerc.* 172 (Bodleian Library, Oxford)[28] and semi-automated digital unwrapping of the lower portion of that scroll[29].

Here we substantially extend the En-Gedi framework, combining high-resolution synchrotron μCT with semi-automated surface recovery and machine learning aided ink detection to recover text from three still-rolled Herculaneum scrolls. First, we report the complete virtual unwrapping of the preserved writing surface of *PHerc.* 1667 (National Library of Naples "Vittorio Emanuele III") along with a complete papyrological reading of the revealed text. Next, with *PHerc. Paris.* 4 (Académie des Inscriptions et Belles Lettres, Institut de France) we demonstrate dramatic improvements to the resolution and readability of the recovered text in comparison to the previously released Vesuvius Challenge Grand Prize region[25]. Further, the high-resolution scan makes the ink more directly visible to the human eye in the tomographic slices, enabling comparison of the model predictions with visible ink deposits. Finally, in *PHerc.* 139 (National Library of Naples "Vittorio Emanuele III"), we recover the title and author identifying the text as Philodemus, *On Gods*, Book 8. Together, these results push the virtual unwrapping of the Herculaneum scrolls beyond proof-of-concept demonstrations to a scalable research framework toward recovering the unopened library.

Technical strategy

We organize the workflow into five stages: fixturing, synchrotron scanning, virtual unwrapping, ink detection and papyrological interpretation (Fig. 1). The reconstructed CT volume contains the structure of complex carbonized papyrus surfaces on which there is writing. These papyrus surfaces are virtually unwrapped through precise localization ("segmentation") and structured mesh representations that support flattening and surface-conditioned rendering, which produces a digitally flattened 3D image of the papyrus surface with its immediate surroundings. Machine learning predicts ink from surface-aligned renderings and, for *PHerc. Paris.* 4, directly from the 3D CT volume before projection onto the flattened surface.

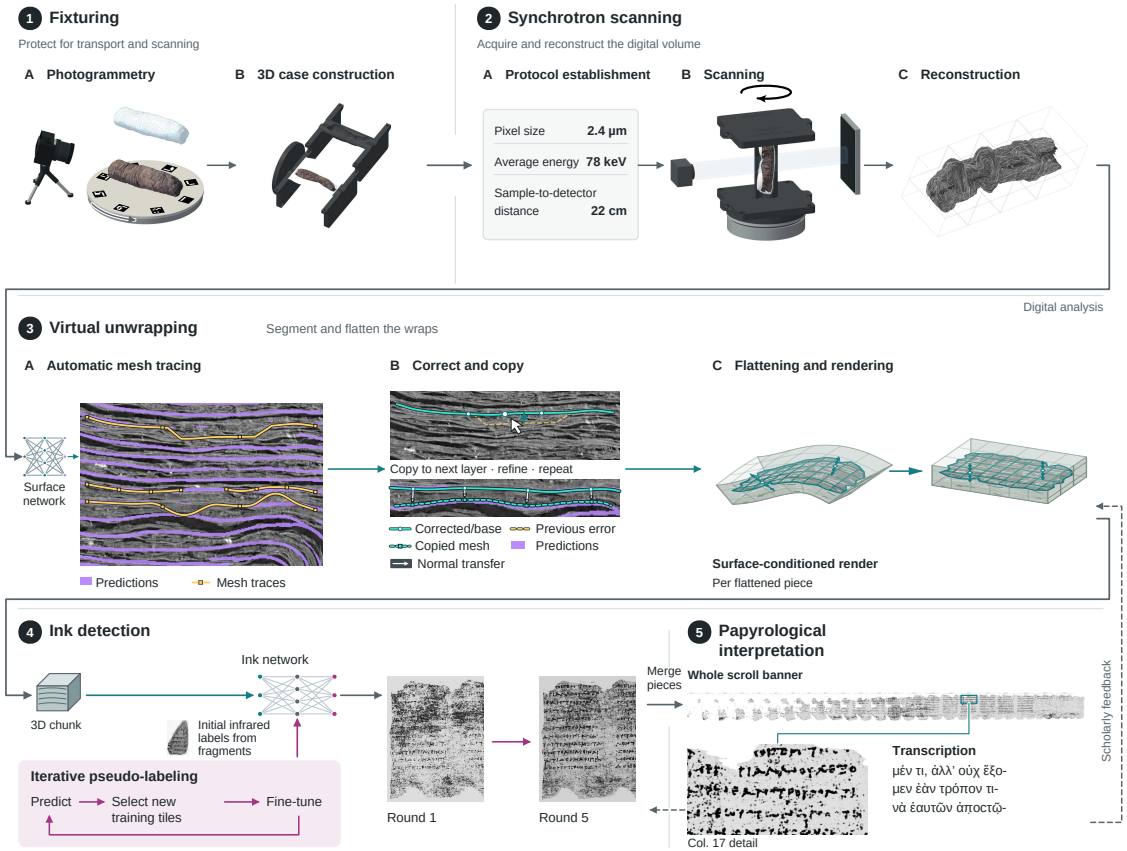


Fig. 1 | Five-stage workflow for virtual unwrapping and reading of *PHerc.* 1667. 1, Fixturing. Photogrammetry provides an exterior mesh for fitted nylon case construction; protective polytetrafluoroethylene (PTFE) sheet separates the scroll from the case. **2, Synchrotron scanning.** Protocol establishment, scanning of the cased scroll and reconstruction produce a digital CT volume. The case is shown in cutaway. **3, Virtual unwrapping.** Surface-network predictions guide automatic mesh tracing, followed by correction, copying to adjacent layers, flattening and surface-conditioned rendering of individual pieces. Representative CT and surface-prediction crops carry explanatory mesh overlays; copying and flattening are schematic. **4, Ink detection.** Local 3D chunks supply the ink network, initially trained using registered labels from infrared images of fragments. Predictions from selected tiles supply labels for subsequent fine-tuning. Round 1 and Round 5 show predictions on the same held-out column from matched models trained using the corresponding label sets. Flattened, ink-enhanced pieces are merged into a whole-scroll banner. **5, Papyrological interpretation.** A banner locator links to a detail of column 17, alongside its transcription excerpt (lines 2–4). Scholars review the image evidence and retain uncertainty in transcription. Dashed feedback paths indicate further review of geometry, rendering or ink enhancement.

Stage 1: Fixturing

Photogrammetry uses overlapping photographs to enable a 3D reconstruction of the scroll exterior, from which a fitted nylon transport and scan case is fabricated (Fig. 1, stage 1A,B; Extended Data Fig. 1a–c). Under conservator supervision, the scroll is placed in a lined case, wrapped in protective material and secured for mounting in the CT scanner (Methods).

Stage 2: Synchrotron scanning

Prior work indicated that finely detailed spatial resolution (voxel sizes in the single-digit micron range) improves the ability to capture evidence of ink[3]. But there are other factors that also play a role in acquiring tomographic scans with enough spatial resolution to reveal ink evidence. For example, variation in ink recipes and weak absorption contrast create a scanning protocol challenge. In addition, the incident energy must be significant enough to penetrate the scroll, but not so energetic that absorption variations are lost. Unwanted X-ray diffusion can occur because the scrolls themselves have regions of densely packed carbonized fibers riddled with micro-structures.

The scan protocol we use balances the need to separate papyrus layers and improve ink resolution within reasonable scan times (Fig. 1, stage 2A–C). This protocol is the result of an extensive series of imaging tests to calibrate the energy range, the effective resolution, the scanning geometry, and the level of propagation-based phase effects. The outcome of these tests is presented in Extended Data Fig. 2, while the final acquisition parameters are outlined in Methods and Extended Data Table 1. Subsequent stages operate on the reconstructed volume and its derived representations.

Stage 3: Virtual Unwrapping

Virtual unwrapping identifies evidence of the physical papyrus surface of interest (P) captured within the tomographic volume (V) and reorganizes its neighborhood into a surface-conditioned render (S), which is a 3D CT volume centered on the writing surface in flattened coordinates. To create S, a learned recto-surface segmentation model first proposes the subset of voxels in the volume V that likely contain the relevant papyrus sheet surface P. This selected surface is then modeled as an explicit quadrilateral mesh (M). The mesh is grown automatically using orientation evidence derived from the proposed voxel set, manually corrected in those places where it does not adhere to the target layer P, and then locally refined to improve agreement with P while preserving surface continuity (Fig. 1, stage 3A,B). A high-quality mesh can also be offset along its normals to a neighboring surface prediction and refined there. Finally, the mesh M is reparametrized to 2D while constraining for low in-plane metric distortion, minimizing local stretch and shear, to produce a bidirectional map between the 3D surface and its 2D counterpart. Each point in the flattened map retains corresponding 3D coordinates in the μCT volume, allowing CT values to be sampled along the local surface normal (Fig. 1, stage 3C; Fig. 2a–d).

In the surface-conditioned render $S(u, v, d)$, u, v follow the flattened mesh and d measures the offset along its local normal. This warped 3D CT volume forms a stack of aligned 2D slices, each retaining its correspondence to the original scan. Individual slices or projections provide images for inspection (Fig. 2d).

From 2D views of the surface-conditioned renders S, users may inspect the underlying papyrus layer for evidence of ink, assess the local geometric reliability of the segmentation, and distinguish surface-placement errors from those imposed by ink-enhancement or ink-interpretation processes[2–8, 30].

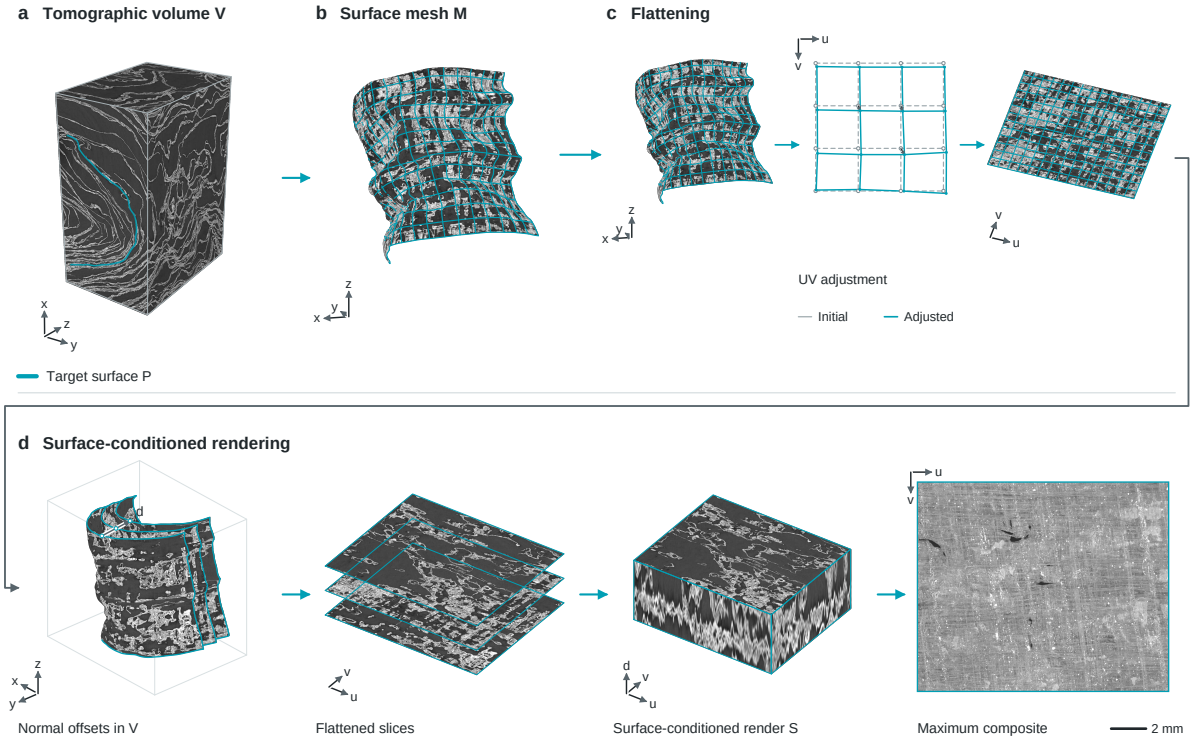


Fig. 2 | Surface geometry and rendering for virtual unwrapping of *PHerc.* 1667. **a**, Tomographic volume V with the target surface P indicated in teal. **b**, A mesh M represents the target surface. **c**, Flattening maps the surface to coordinates (u, v); the schematic grid illustrates adjustment of the initial coordinates. **d**, Sampling the CT volume at offsets d along the local surface normal produces aligned flattened slices and the surface-conditioned render S. The rightmost panel shows a maximum composite of flattened slices over the depth axis. Scale bar, 2 mm (rightmost panel).

Stage 4: Ink detection

There is wide variability in how evidence of ink appears across multiple closed scroll regions. In many cases ink is not reliably separable as a 3D object by direct volumetric segmentation within the tomographic volume. When it is, as in the case of *PHerc. Paris.* 4, our approach to ink detection can leverage the direct observability of ink traces. This scenario admits a direct 3D ink segmentation approach, allowing the identified ink in 3D to project onto virtually unwrapped (digitally flattened) surfaces (Methods). For *PHerc. Paris.* 4, inference is performed across the full reconstructed volume V in parallel with geometric tracing, and the resulting ink volume is subsequently sampled using the flattened surfaces.

More subtle cases require the ink evidence to be extracted through an enhancement process designed with supervised artificial neural models. The models take 3D chunks of S as input and predict 2D ink maps in the flattened (u, v) plane, where the supervision is defined. The models train over two related but distinct label types: direct-visibility labels

and pseudo-labels. Direct-visibility labels are used when the presence of ink can be discerned using infrared photographs of opened Herculaneum fragments that are registered to the tomographic scans of those same fragments. Pseudo-labels, on the other hand, are applied to 2D images derived from the surface-conditioned renders S of the closed scrolls, where ink presence is conservatively identified through the manual visual inspection of the surface renders and the output of previously trained ink models. Selected predictions supply pseudo-labels for successive fine-tuning rounds, progressively refining ink detection (Fig. 1, stage 4; Methods). The resulting ink-probability map enhances contrast for papyrological inspection. It is important to note, however, that our models are trained to detect *localized ink presence only* within microscopically small regions; they are *not trained* on character identities, words, transcriptions, OCR or HTR targets, or lexical labels. Together, fragment-derived supervision and restricted spatial context provide safeguards against hallucinated letterforms, directing the model toward local tomographic evidence rather than the completion of plausible words.

These two approaches to ink detection allow model predictions to be compared with ink deposits directly visible in the tomographic volume. The accuracy of ink enhancement can be tested by investigating how well the images of ink it produces align with the physically visible ink deposits rather than with artifacts introduced by rendering, surface placement, or the enhancement model.

Stage 5: Papyrological interpretation

The separately flattened and ink-enhanced pieces are merged into a single scroll view for papyrological reading (Fig. 1, stages 4–5). Papyrologists review the recovered text by inspecting the flattened renderings, ink-enhanced outputs and geometric context of the sampled surface. Accepted readings require agreement of the sampled surface with the papyrus layer, consistency of letterforms across repeated renderings or inference passes, and independent endorsement or consensus. The transcription retains underdots for uncertain letters and square brackets for restorations across physically lost substrate. Scholarly review can prompt reinspection and refinement of surface geometry, rendering or ink enhancement, following the dashed feedback paths from stage 5 to stages 3 and 4 (Fig. 1, stages 3–5).

Results

The results are organized around three questions. First, can the preserved portion of a still-rolled Herculaneum papyrus be completely unwrapped and read? The presented work on *PHerc.* 1667 addresses this question. Second, can ink recovery be checked against ink that is directly visible in the CT volume? The recent scan of *PHerc. Paris.* 4 provides this imaging-validation case. Third, can virtual unwrapping recover evidence to identify a work preserved within a sealed scroll? *PHerc.* 139 provides this title-identification case.

PHerc. 1667: complete unwrapping

PHerc. 1667 is the first Herculaneum papyrus to be completely unwrapped and read by virtual methods (Fig. 3). The surviving core is a remnant of earlier attempts at physical unrolling. The text appears to concern ethical theory and human moral perfectibility, with the mention of Aristocreon and other features suggesting a Stoic context (Methods). The sealed object is shown in Fig. 3a, while transverse and longitudinal tomographic sections reveal the dense nested architecture of the carbonized papyrus layers (Fig. 3b,c). Fig. 3d shows the corresponding 2D ink-enhanced image derived from the surface-conditioned render after virtual unwrapping: the flattened papyrus surface is displayed in three consecutive bands, with dark strokes indicating the scholar-facing ink signal used for papyrological inspection. Papyrologists localized approximately 860 cm^2 of preserved writing surface, containing the lower portion of 22 columns or column-equivalents, each comprising up to 12 lines, within the 1231 cm^2 of virtually unwrapped papyrus; black regions mark areas outside the preserved surface or physically absent or damaged substrate.

We use "complete virtual unwrapping" in a bounded geometric sense: the preserved surface is represented by a mesh whose position has been verified against the μCT volume by a human expert annotator, and is rendered in the flattened coordinate domain. We use "complete reading" in a corresponding papyrological sense: the transcription is restricted to image-supported preserved text visible in the renderings; restoration of letters corresponding to physically lost substrate is indicated by the use of square brackets in accordance with established papyrological convention; uncertain letters are marked with an underdot. The entire recovered writing surface was examined, with undecipherable passages recorded as traces and uncertainty retained in the transcription. Within this writing surface, a region of 33 cm^2 preserves ink traces that do not permit secure decipherment. The full papyrological transcription and uncertainty markings are provided in Methods.

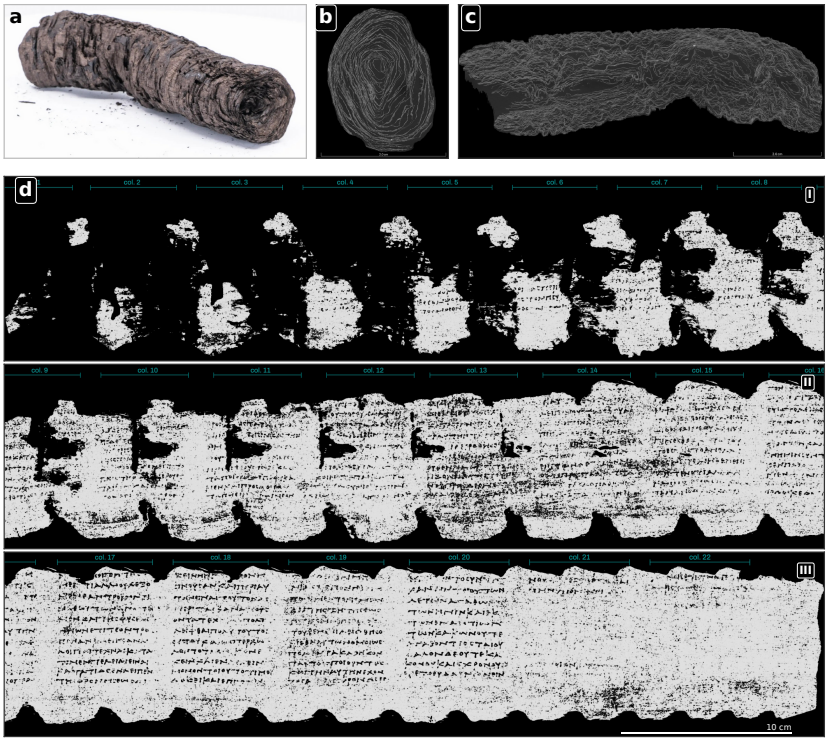


Fig. 3 | Scroll-scale virtual unwrapping of *PHerc.* 1667. a, Photograph of the sealed Herculaneum roll *PHerc.* 1667. b, Transverse tomographic section (isotropic voxel size 2.4 µm) showing the nested internal wrap structure of the carbonized papyrus. c, Longitudinal tomographic slice showing the layered structure along the scroll axis. d, 2D ink-enhanced image derived from the surface-conditioned render of the reconstructed writing surface, displayed in three consecutive bands of the flattened papyrus surface (I-III). Black background corresponds to regions outside the preserved writing surface or to physically absent or damaged substrate. Within the gray papyrus surface, dark strokes indicate the scholar-facing ink rendering used for papyrological inspection. The figure summarizes the transition from a sealed physical roll to a continuous inspectable writing surface. Scale bars as shown.

PHerc. Paris. 4: volumetric validation

PHerc. Paris. 4 tests whether surface-conditioned ink recovery corresponds to physically visible structures in the µCT volume. In the 2.4 µm high-resolution BM18 scan of *PHerc. Paris.* 4, ink-bearing strokes appear in the tomographic volume as bright, surface-associated deposits visible across several voxels (Fig. 4a,b and Extended Data Fig. 3).

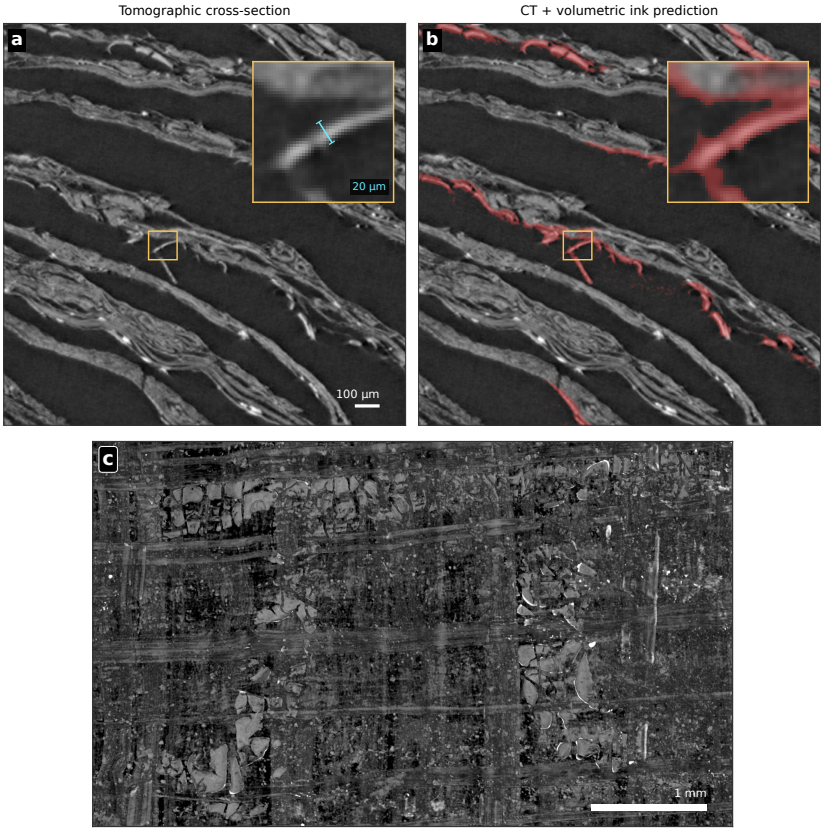


Fig. 4 | Volumetric ink segmentation in *PHerc. Paris.* 4 and projection onto the unwrapped surface. a, Native-resolution tomographic cross-section (2.4 µm isotropic voxels), with a magnified inset of a visible surface-associated deposit. The cyan ruler is a 20 µm reference marker oriented along the local in-plane normal of the deposit. b, The same CT slice overlaid in red with volumetric ink-prediction intensities; the main view and inset share the geometry and scales of a. White scale bar in a, 100 µm. c, 2D image derived from the surface-conditioned render of an ink-bearing region before projection of the volumetric segmentation, showing the underlying papyrus texture and writing traces.

These deposits are directly visible in the volume, can be segmented in 3D and then projected onto the flattened surface, where they coincide with the Vesuvius Challenge Grand Prize Banner (2023) region (Fig. 5), for which multi-column readings were previously published[23–25]. The correspondence between these deposits and the recovered text provides a directly inspectable physical reference for the reading. Agreement between the learned outputs complements this check (Methods).

PHerc. 139: identifying a sealed work

The title of ancient bookrolls was typically written twice: at the beginning, and also at the end of the text, in the innermost wraps. These are often preserved even where the outer

wraps, and with them the initial title and the opening columns were lost to the eruption, the subsequent excavation or attempts at physical opening.

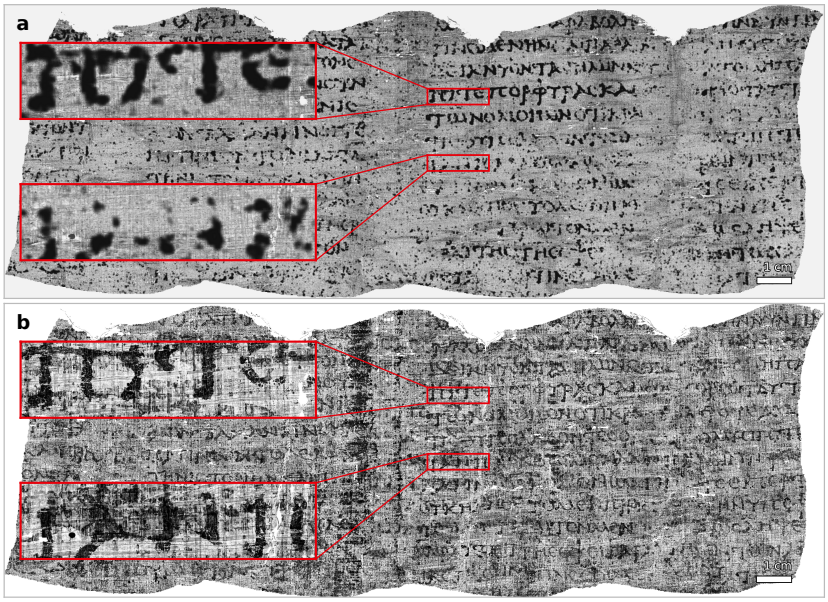


Fig. 5 | *PHerc. Paris.* 4, comparison with 2023 data. a, 2023 Grand Prize winning ink detection model[24, 25] on a virtually unwrapped portion of *PHerc. Paris.* 4 on the 7.91 µm pixel size volume of *PHerc. Paris.* 4 from the EduceLab-Scrolls dataset [22]. b, Surface rendering of the 2.4 µm data acquired at the ESRF BM18 beamline, using volumetric ink predictions to guide alpha compositing. The second confirms the reading from the 2023 Vesuvius Challenge. Zoom-in lenses highlight the difference in data quality and accurate representation of the ancient script.

PHerc. 139 provides a separate title-identification result, although its ink predictions are less sharply defined than the examples from *PHerc.* 1667 and *PHerc. Paris.* 4. The predicted strokes are broad, rounded and locally discontinuous (Fig. 6b; Extended Data Fig. 4), making their contours approximate guides to ink position rather than faithful outlines of the original letterforms. Papyrological interpretation of a title-bearing region of the reconstructed surface yields author, work-title and book-number evidence (Fig. 6). The title-bearing sequence is read as

Φι[λοδ]ήμου
Περὶ θεῶν Η̄

that is, "Philodemus, *On Gods*, Book 8". The overlined H is interpreted as the book-number indication.

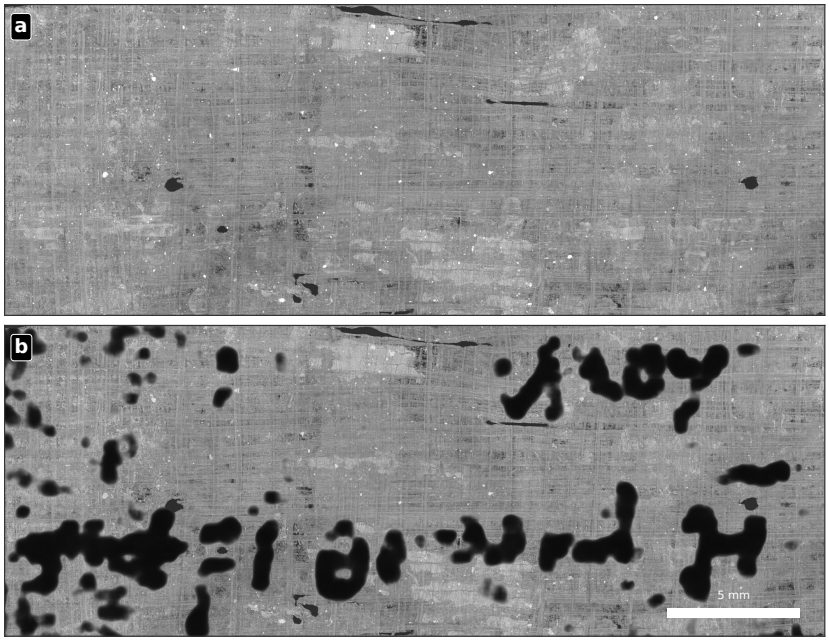


Fig. 6 | Title and author-attribution evidence in *PHerc.* 139. a, 2D image derived from the surface-conditioned render of the title-bearing region of *PHerc.* 139. b, The ink-enhanced signal shown in black against the papyrus texture. The broad, locally discontinuous predictions provide approximate stroke contours. The visible title and author-attribution sequence identifies the roll as Philodemus, *On Gods*, Book 8. Scale bar, 5 mm.

The book number is written here on the same line as the title, despite the more commonly attested practice in titles from Herculaneum, where the title and book number appear on two separate lines. A similar layout is now documented in the final title of the unopened roll *PHerc.* 172 (Φιλο̣δήμο̣υ Πε̣ρὶ | κα̣κ̣ιῶν Ᾱ̣, Philodemus, *On Vices*, Book 1)[28]. In *PHerc.* 139, the title was likely written by the same scribe who copied the text.

The presence of Philodemus' name, by far the most extensively attested Epicurean philosopher in the Herculaneum library, allows us to assign the scroll to the 1st century BC-1st century AD (with a *terminus ante quem* of 79 AD). More importantly, this discovery informs us for the first time that Philodemus' work *On Gods* comprised at least eight books. Until now, only the first book had been securely identified under this title, preserved in *PHerc.* 26 (Φιλοδήμ[ου] | Περὶ θεῶ̣ν | Ᾱ, Philodemus, *On Gods*, Book 1), together with other theological works by the same author known under different titles[31]. Building on the title identification in *PHerc.* 172, this result shows how even a short recovered title can reveal the previously unknown extent of an ancient work.

In addition to the end title, the portions of text recovered thus far preserve a number of terms and expressions that provide insight into the treatise's theological concerns: χωρὶc προνο[ί]α̣ς ("without providence"), ἀόρατα ("invisible entities"), καὶ τὸ κατὰ φύcιν ("and the

… according to nature”), καὶ πάντα χωρὶς πόνων (“and everything without troubles”), νοερόν (“intellectual”). The images which supported the reading are provided in the corresponding order in Extended Data Fig. 4.

Discussion

These results mark a step in the transition from a Herculaneum library known only through physically opened, fragmentary witnesses toward one that can also include virtually unwrapped texts with continuous lines and columns and identified titles and authors. Virtual unwrapping of the Herculaneum papyri is no longer limited to localized breakthroughs and partial-scroll discoveries[22–25]. *PHerc.* 1667 shows that scroll-scale unwrapping and reading of a still-rolled papyrus is achievable. The results from *PHerc. Paris.* 4 provide an imaging-specific payoff, because 3D segmentation of directly visible ink can ground and check surface-conditioned ink recovery. The work also opens new scholarship on previously unread texts. The title-and-attribution evidence from *PHerc.* 139, following the recovery of a title from *PHerc.* 172[28], shows that the method can identify works and authors, not only recover running prose.

This workflow does not imply that all sealed Herculaneum rolls are automatically readable. Performance depends on many factors: the general preservation state of individual scrolls; the quality of the tomographic scans; the degree of layer separation and compression; the amount of deformation suffered by the scroll; the contrast of the local ink signal; and the level of surface preservation.

Two bottlenecks remain dominant. The first is geometric. Even strong surface-prediction networks can fail in highly compressed regions, where mergers, holes and sheet switches (mistakenly leaving one layer and latching onto another, thereby breaking the continuity of any recovered writing) destabilize surface estimation. The second is radiometric. Ink evidence remains weak, varies across ink recipes and local degradation states, and is further modulated by scan conditions and diffusion of the X-rays by the samples in case of large and compacted parts. These difficulties are not independent, and their interplay helps explain the complexity of optimizing processes for readability. Better geometry produces more legible 2D images derived from surface-conditioned renders, whereas better legibility improves both expert transcription and the supervision available for ink detection.

Future gains will depend partly on improved control of effective resolution (Fig. 1, stage 2) and on a better account of how morphology, phase effects and absorption together contribute to ink visibility (stages 2 and 4). Higher-resolution and better-tuned scans (stage 2) could reduce ambiguity in both layer separation (stage 3) and ink detectability (stage 4), especially in the most compressed regions. More robust tracing networks and hybrid neural tracers that predict or refine sheet trajectories or papyrus fibers while enforcing local surface continuity and neighboring-wrap consistency should reduce the manual burden of correction and neighbor transfer (stage 3)[8]. The latter could be achieved by improving the accuracy of annotations used to train the machine learning models or relying on self-supervised methods. Global approaches that fit one explicit spiral model to the entire scroll[32], rather than growing the surface patch by patch, are a promising route to fully automated unwrapping (stage 3). Additional fragment-to-scroll supervision, together with

conservative pseudo-label expansion on intact rolls, should strengthen ink models (stage 4) across a wider range of ink recipes without sacrificing trustworthiness.

None of these remaining challenges alters the central conclusion. With the results reported here, virtual unwrapping moves beyond isolated proof-of-concept recoveries and becomes a credible scroll-scale method for restoring texts from the still-sealed Herculaneum library. Indeed, in a single integrated workflow, we show complete reading of a virtually unwrapped Herculaneum papyrus, direct tomographic validation of ink in a still-closed roll, and title-level identification of another unopened work. The question is therefore no longer whether the hidden library can be read non-invasively, but how rapidly and comprehensively its contents can be returned to scholarship. Continued advances in imaging, surface recovery and ink detection should make it possible to approach the unopened rolls not as permanently lost objects, but as a recoverable library capable of substantially enlarging the surviving record of classical thought.

References

[1] Sider, D. *The Library of the Villa dei Papiri at Herculaneum* (J. Paul Getty Museum, Los Angeles, 2005).

[2] Seales, W. B. *et al.* From damage to discovery via virtual unwrapping: reading the scroll from En-Gedi. *Sci. Adv.* **2**, e1601247 (2016).

[3] Parker, C. S. *et al.* From invisibility to readability: recovering the ink of Herculaneum. *PLoS One* **14**, e0215775 (2019).

[4] Barfod, G. H., Larsen, J. M., Lichtenberger, A. & Raja, R. Revealing text in a complexly rolled silver scroll from Jerash with computed tomography and advanced imaging software. *Sci. Rep.* **5**, 17765 (2015).

[5] Liu, C., Rosin, P. L., Lai, Y.-K. & Hu, W. Robust virtual unrolling of historical parchment XMT images. *IEEE Trans. Image Process.* **27**, 1914–1926 (2018).

[6] Rosin, P. L. *et al.* Virtual recovery of content from X-ray micro-tomography scans of damaged historic scrolls. *Sci. Rep.* **8**, 11901 (2018).

[7] Dambrogio, J. *et al.* Unlocking history through automated virtual unfolding of sealed documents imaged by X-ray microtomography. *Nat. Commun.* **12**, 1184 (2021).

[8] Kulagin, P., Polevoy, D., Chukalina, M., Nikolaev, D. & Arlazarov, V. V. Fully automatic virtual unwrapping method for documents imaged by X-ray tomography. *Document Analysis and Recognition – ICDAR 2024, Lecture Notes in Computer Science* **14806**, 233–250 (2024).

[9] Neuber, D. & Reinhart, C. Geheimnisvolle Bleischriftrolle der Mandäer: moderne Computertomographie erlaubt virtuelle Entschlüsselung eines antiken Rätsels. *ZfP-Zeitung* **132**, 39–41 (2012). In German.

[10] Lin, Y. & Seales, W. B. Opaque document imaging: building images of inaccessible texts. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV 2005)*, vol. I, pp. 662–669, IEEE (2005).

[11] Samko, O., Lai, Y.-K., Marshall, D. & Rosin, P. L. Virtual unrolling and information recovery from scanned scrolled historical documents. *Pattern Recognition* **47**, 248–259 (2014).

[12] Fleischer, K. *Die Papyri Herkulaneums im Digitalen Zeitalter: Neue Texte durch neue Techniken – eine Kurzeinführung*, Ch. 12, 72–79 (De Gruyter, Berlin and Boston, 2022). URL https://doi.org/10.1515/9783110767711-013. Wieviel Text enthalten die noch geschlossenen Rollen? Ein Quantifizierungsversuch.

[13] Nicolardi, F. Qualche considerazione sui rotoli ercolanesi ancora non svolti, in vista di un futuro svolgimento virtuale. *Atene e Roma* **18**, 121–130 (2024).

[14] Cavallo, G. *Libri scritture scribi a Ercolano: introduzione allo studio dei materiali greci* (Gaetano Macchiaroli, Napoli, 1983). Cronache Ercolanesi 13, supplemento 1.

[15] Mocella, V., Brun, E., Ferrero, C. & Delattre, D. Revealing letters in rolled Herculaneum papyri by X-ray phase-contrast imaging. *Nat. Commun.* **6**, 5895 (2015).

[16] Bukreeva, I. *et al.* Virtual unrolling and deciphering of Herculaneum papyri by X-ray phase-contrast tomography. *Sci. Rep.* **6**, 27227 (2016).

[17] Brun, E. *et al.* Revealing metallic ink in Herculaneum papyri. *Proc. Natl Acad. Sci. USA* **113**, 3751–3754 (2016).

[18] Angelotti, G., Nicolardi, F., Henderson, P. & Seales, W. B. Ink detection from surface topography of the Herculaneum papyri. *Scientific Reports* **16**, 27835 (2026).

[19] Stabile, S. *et al.* A computational platform for the virtual unfolding of Herculaneum Papyri. *Scientific Reports* **11**, 1695 (2021).

[20] Chapman, C. *"WE JUST WANNA READ THE SCROLLS!" Collaboration and competition in the Vesuvius Challenge — a study of Vesuvius Challenge Discord chats.* Theses and Dissertations–Communication. 140, University of Kentucky (2025). https://uknowledge.uky.edu/comm_etds/140.

[21] Parsons, S. *Hard-Hearted Scrolls: A Noninvasive Method for Reading the Herculaneum Papyri.* Doctoral dissertation, University of Kentucky (2023). https://uknowledge.uky.edu/cs_etds/138.

[22] Parsons, S., Parker, C. S., Chapman, C., Hayashida, M. & Seales, W. B. EduceLab-Scrolls: verifiable recovery of text from Herculaneum papyri using X-ray CT. https://arxiv.org/abs/2304.02084 (2023). Preprint at https://doi.org/10.48550/arXiv.2304.02084.

[23] Nicolardi, F. *et al.* Revealing text from a still-rolled Herculaneum papyrus scroll (PHerc. Paris. 4). *Z. Papyrol. Epigr.* **229**, 1–13 (2024).

[24] Nicolardi, F., Delattre, D., Del Mastro, G., Janko, R. & Fowler, R. The final columns of PHerc. Paris. 4 revealed through virtual unwrapping. *Cronache Ercolanesi* **54**, 9–27 (2024).

[25] Vesuvius Challenge. Vesuvius Challenge 2023 Grand Prize awarded: we can read the first scroll! https://scrollprize.org/grandprize (2024). Accessed 19 June 2026.

[26] Lourenço, A. *et al.* Vesuvius Challenge - Ink Detection. https://www.kaggle.com/competitions/vesuvius-challenge-ink-detection (2023). Kaggle.

[27] Johnson, S. *et al.* Vesuvius Challenge - Surface Detection. https://www.kaggle.com/competitions/vesuvius-challenge-surface-detection (2025). Kaggle.

[28] Vesuvius Challenge. $60,000 first title prize awarded! Title of a still-rolled Herculaneum scroll has been read for first time ever. https://scrollprize.substack.com/p/60000-first-title-prize-awarded (2025). Accessed 19 June 2026.

[29] Vesuvius Challenge. ~70% of PHerc. 172 is now digitally unwrapped. https://scrollprize.substack.com/p/70-of-pherc-172-is-now-digitally (2026). Accessed 19 June 2026.

[30] Rabinovich, M., Poranne, R., Panozzo, D. & Sorkine-Hornung, O. Scalable locally injective mappings. *ACM Trans. Graph.* **36**, 16:1–16:16 (2017).

[31] D'Angelo, M. Unpublished studies on Philodemus' On Gods, Book 1. *History of Classical Scholarship* **7**, 171–188 (2025). URL https://www.hcsjournal.org/ojs/index.php/hcs/article/view/123.

[32] Henderson, P. Virtually unrolling the Herculaneum papyri by diffeomorphic spiral fitting. *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)* 6401–6411 (2026). Preprint at https://arxiv.org/abs/2512.04927.

[33] Garrido-Jurado, S., Muñoz-Salinas, R., Madrid-Cuevas, F. J. & Marín-Jiménez, M. J. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognit.* **47**, 2280–2292 (2014).

[34] Ravi, N. *et al.* SAM 2: segment anything in images and videos. https://arxiv.org/abs/2408.00714 (2024). Preprint at https://doi.org/10.48550/arXiv.2408.00714; published at ICLR 2025.

[35] EduceLab. pgs-recon: the EduceLab photogrammetry reconstruction pipeline. GitHub, https://github.com/educelab/pgs-recon. Accessed 19 June 2026.

[36] Moulon, P., Monasse, P., Perrot, R. & Marlet, R. OpenMVG: open multiple view geometry. *Reproducible Research in Pattern Recognition – RRPR 2016, Lecture Notes in Computer*

Science **10214**, 60–74 (2017).

[37] Cernea, D. OpenMVS: open multi-view stereo reconstruction library. GitHub, https://github.com/cdcseacave/openMVS. Accessed 19 June 2026.

[38] Vijayakumar, J. *et al.* Multiresolution phase-contrast tomography on BM18, a new beamline at the European Synchrotron Radiation Facility. *Synchrotron Radiat. News* **37**, 16–25 (2024).

[39] Dejea, H. *et al.* Evolution of hierarchical phase-contrast tomography on the European Synchrotron beamlines BM05 and BM18: a whole adult human brain imaging case study. *J. Synchrotron Radiat.* **33**, 1180–1189 (2026).

[40] Mirone, A. night_rail_bm18. ESRF GitLab, https://gitlab.esrf.fr/night_rail/applications/mirone/night_rail_bm18 (2025).

[41] Paleo, P., Payno, H., Viganò, N. & Nemoz, C. Nabu: ESRF tomography processing software. Zenodo, https://doi.org/10.5281/zenodo.5916579 (2021).

[42] Paganin, D., Mayo, S. C., Gureyev, T. E., Miller, P. R. & Wilkins, S. W. Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object. *J. Microsc.* **206**, 33–40 (2002).

[43] Sheppard, A. P., Sok, R. M. & Averdunk, H. Techniques for image enhancement and segmentation of tomographic images of porous materials. *Physica A* **339**, 145–151 (2004).

[44] Angelotti, G. *et al.* Synchrotron X-ray microtomographic data of the Herculaneum papyrus PHerc 1667 at 2.399 um voxel size. European Synchrotron Radiation Facility, https://doi.org/10.15151/ESRF-DC-2488816686 (2026).

[45] Angelotti, G. *et al.* Synchrotron X-ray microtomographic data of the Herculaneum papyrus PHerc 0139 at 2.399 um voxel size. European Synchrotron Radiation Facility, https://doi.org/10.15151/ESRF-DC-2488816678 (2026).

[46] Angelotti, G. *et al.* Synchrotron X-ray microtomographic data of the Herculaneum papyrus PHerc Paris4 at 2.4 um voxel size. European Synchrotron Radiation Facility, https://doi.org/10.15151/ESRF-DC-2488890952 (2026).

[47] Angelotti, G. *et al.* Vesuvius Challenge – CT scans of Herculaneum papyri. Vesuvius Challenge, https://scrollprize.org/data (2026). Accessed 19 June 2026.

[48] Isensee, F. *et al.* nnU-Net revisited: a call for rigorous validation in 3D medical image segmentation. https://arxiv.org/abs/2404.09556 (2024). Preprint at https://doi.org/10.48550/arXiv.2404.09556; published in Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, doi:10.1007/978-3-031-72114-4_47.

[49] Kirchhoff, Y. *et al.* Skeleton Recall Loss for connectivity conserving and resource efficient segmentation of thin tubular structures. https://arxiv.org/abs/2404.03010 (2024). Preprint at https://doi.org/10.48550/arXiv.2404.03010; published in Computer Vision – ECCV 2024, doi:10.1007/978-3-031-72980-5_13.

[50] Seales, W. B., Chapman, C., Nicolardi, F. & Parker, C. S. The digital restoration of the Herculaneum papyri. *Cronache Ercolanesi* **53**, 201–211 (2023).

[51] Hara, K., Kataoka, H. & Satoh, Y. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet? *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 6546–6555 (2018).

[52] Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* **9351**, 234–241 (2015).

[53] Zou, Y., Yu, Z., Vijaya Kumar, B. V. K. & Wang, J. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. Computer Vision – ECCV 2018, Lecture Notes in Computer Science 11207, 297–313 (2018). https://doi.org/10.1007/978-3-030-01219-9_18.

[54] Zhang, J. *et al.* Detecting pest-infested forest damage through multispectral satellite imagery and improved UNet++. *Sensors* **22**, 7440 (2022).

[55] Kataoka, H., Wakamiya, T., Hara, K. & Satoh, Y. Would mega-scale datasets further enhance spatiotemporal 3D CNNs? (2020). Preprint at https://doi.org/10.48550/arXiv.2004.04968; pretrained weights at https://drive.google.com/file/d/1fCKSlakRJ54b3pEWqgBmuJi0nF7HXQc0/view.

[56] Oquab, M. *et al.* DINOv2: learning robust visual features without supervision. https://arxiv.org/abs/2304.07193 (2023). Preprint at https://doi.org/10.48550/arXiv.2304.07193; published in Transactions on Machine Learning Research (2024).

[57] Sbordone, F. Recenti tentativi di svolgimento dei papiri ercolanesi. *Cronache Ercolanesi* **1**, 23–39 (1971).

[58] Fackelmann, A. The restoration of the Herculaneum papyri and other recent finds. *Bull. Inst. Class. Stud.* **17**, 144–147 (1970).

[59] Capasso, M. & Angeli, A. Papiri aperti col metodo osloense (1983-1989): descrizione e classificazione. *Cronache Ercolanesi* **19**, 265–270 (1989).

Methods

Provenance, custody and permissions

The objects reported here are Herculaneum papyri presently held at the Biblioteca Nazionale di Napoli, Museo Archeologico Nazionale di Napoli, and the Bibliothèque de l'Institut de France, which conserves the *PHerc. Paris.* scrolls on behalf of the Académie des Inscriptions et Belles Lettres. Imaging and analysis were performed under permissions granted by the relevant authorities in compliance with the applicable laws, regulations and institutional agreements governing access, transport and study of the material. The methods below follow the five stages illustrated for *PHerc.* 1667 in Fig. 1; the volumetric ink-detection branch used for *PHerc. Paris.* 4 is described within Stage 4.

Stage 1: Fixturing

Photogrammetry.

Each sample was enclosed in a custom transport and scan case generated from a metric photogrammetric reconstruction of its exterior geometry. Photographs were acquired with a Sony α6300 camera mounted on a Joby GorillaPod tripod, using an Orangemonkie Foldio360 turntable inside an Orangemonkie Foldio3 studio with additional Halo Bar illumination. ArUco fiducial markers of known side length were fixed to the turntable base. These markers are binary square fiducials whose unique identifiers and four detected corners allow robust localization and metric scaling of an image-based reconstruction[33].

For each stable pose, the sample was photographed over a full 360° rotation at three camera heights. The object was then repositioned under conservator supervision, including inversion or alternative footholds where the morphology permitted, and the procedure was repeated until the exterior surface was covered without blind regions. Raw photographs were converted to JPEG while preserving Exif metadata. Before photogrammetric reconstruction of the sample surface, foreground masks were generated with a fine-tuned SAM 2.1 model to isolate the scroll from the studio background, turntable and fiducial markers[34]. The masked images were used for object reconstruction so that the surrounding setup did not contaminate the sample mesh. In parallel, the corresponding unmasked photographs were retained for the metric-scaling step: the pgs-recon pipeline[35] detected and triangulated the ArUco markers from the full images and applied their known side length to place the reconstructed scene at physical scale[33, 36].

Photogrammetric reconstruction was performed with pgs-recon, a Python pipeline built around OpenMVG for structure-from-motion and OpenMVS for dense reconstruction, mesh refinement and texturing[36, 37]. This produced a textured, metric-scale exterior mesh of each sample. The photogrammetry and case-generation workflow is documented in the Villa repository[1].

[1] https://github.com/ScrollPrize/villa/commit/6847063ffdb4da898ae8d1d494ebf7d71473f509

3D case construction.

The mesh was then passed to the ScrollCase workflow[2], which generated smoothed STL models for the sample surrogate and two complementary case shells. One shell incorporated the beamline tomograph mounting interface, and the two shells were designed to close with screws. Cases were fabricated by 3DnA s.r.l. from nylon powder using HP Multi Jet Fusion. During mounting, a conservator lined the cavity with a thin protective sheet, placed the sample into the corresponding shell, wrapped the sample with the protective material and closed the case. PTFE sheet was used for most objects. For *PHerc. Paris.* 4, which was prepared later, Japanese paper was used instead, and the case was polyurethane-coated to reduce surface roughness.

Stage 2: Synchrotron scanning

Protocol establishment.

Tests on different samples were used to select parameters that support virtual unwrapping (Extended Data Fig. 2; Supplementary Table 1). Earlier studies of the Herculaneum papyri emphasized propagation-based phase-contrast enhancement at effective pixel sizes of 46–47 µm[15, 16], with incident energy selection guided by the contrast-transfer function for the chosen propagation distance and detector sampling[16]. The present protocol prioritizes the separation of tightly packed papyrus layers at nearly twenty-fold finer sampling (2.4 µm). Incident energy and propagation distance were therefore selected jointly to retain sufficient phase contrast while limiting scattering-induced blurring that obscures layer boundaries. Densely packed papyrus layers lead to strong beam diffusion that makes phase contrast challenging. This phenomenon, known as decoherence, blurs the projections, and impacts the visibility of small structures. In carbonized papyri, it could be linked to the presence of graphite (a very efficient decoherer) and to the natural cellular structure of the papyrus that can act as a series of very small refractive lenses and that locally deviates the beam. The solution to counter-act decoherence effect can be to reduce the propagation distance and to increase the energy. At a fixed spatial scale in the near-field regime, both changes reduce propagation-induced phase contrast, while absorption contrast remains weak for carbon-based material at hard-X-ray energies. The aim was to keep sufficient phase-contrast effect to enhance the structures of interest, without reaching the level that would result in visible decoherence effect. The production configuration selected for complete-scroll imaging was 2.4 µm isotropic voxel size, 22 cm propagation distance and 78 keV average incident energy. Finer voxel sizes improve image quality but are currently impractical for full-scroll coverage; a targeted region of interest was nevertheless acquired at a pixel size of 1.1 µm.

Scanning.

We used a polychromatic beam from the 1.1 T side pole of the BM18 tripole wiggler. Filtration differed between acquisitions: the *PHerc.* 1667 and *PHerc.* 139 scans used 2.5 mm of sapphire and 0.20 mm of silver, whereas the *PHerc. Paris.* 4 scan used 20 mm of sapphire and 0.10 mm of gold. In the gold-filtered configuration, the thin gold filter reduces transmission above its K-edge (80 keV), narrowing the effective bandwidth.

[2] https://github.com/ScrollPrize/villa/tree/main/foundation/scrollcase

We used a scintillator-based indirect detector coupling a GAGG:Ce (Cerium-doped Gadolinium Aluminium Gallium Garnet) single crystal of 50 µm thickness with a front reflective deposit to increase its light emission efficiency. As this scintillator contains gadolinium, its X-ray stopping power increases abruptly when the energy crosses the gadolinium K-edge at 50.4 keV, making it very efficient for X-rays up to few dozens of keV above this energy. For the gold-filtered configuration, the effectively detected beam spectrum has an average of 78 keV, with enhanced detection between 50.4 keV and 80 keV. For these weakly absorbing samples, the narrowed spectrum is expected to limit beam-hardening effects. The configuration combines a filtered polychromatic beam with no focusing X-ray optics. Indeed, the only X-ray optics in the beam are sapphire and gold or silver filters, depending on the acquisition, as well as permanent polished windows (1.4 mm of diamond in the front-end and 1 mm of aluminum at the beginning of the experimental hutch).

The optic system was a tandem optic made of two photographic objectives, a Zeiss Otus 100 mm F/D 1.4 as the bottom objective, and a Nikon AF-S Nikkor 200 mm F/D 2 as the top objective. The coupling of these two objectives brings an exceptionally high numerical aperture of 0.35 that gives very high efficiency coupled with very good effective resolution. When coupled with the scintillator mentioned above, the calculated diffraction limit is of 2.72 µm for a pixel size of 2.4 µm given by the PCO edge bi-10 sCMOS camera (4.6 µm of pixel size, 85% of quantum efficiency for the green emission of GAGG:Ce).

Data were acquired on BM18[38] in helical mode[39] (Fig. 1, stage 2B,C). To extend the transverse coverage, each sample was scanned at up to four lateral positions relative to the beam and detector, each scan sampling a partially overlapping annular region around the same rotation axis; the number of positions and of projections was chosen according to the sample diameter. For some of the scans at 1.1 µm pixel size (Extended Data Table 1), the field of view was instead extended with tiled mosaics.

Tomographic reconstruction.

Volumes were reconstructed using the Night Rail software framework[39, 40], which integrates an HPC-oriented fork of Nabu[41] implementing GPU-accelerated generalized hierarchical backprojection for large helical and multi-scan BM18 acquisitions. Projections were corrected for dark current, flat field, hot pixels and detector distortion, the latter using a map calibrated independently on a slit grid and applied by bilinear resampling to both radiographs and their blending weights. Single-distance Paganin phase retrieval[42] was then applied and the resulting blur compensated with an unsharp mask[43]; the parameters used for each scan are listed in Extended Data Table 1. Intensities were then clipped to $[10^{-6}, 10]$ before logarithmic transformation.

In both the annular offset scans and the tiled mosaics, slow motion of the sample had to be compensated before the data could be combined. Two different strategies were used. When the sample and rotation-axis reference frame were translated together, their mutual geometry remained unchanged: adjacent acquisitions were aligned by correlating overlapping detector regions, relying on the dominant papyrus structures in the radiographs. Displacements were estimated every 5°, smoothly fitted versus projection angle and accumulated along the scan chain. The corrected radiographs were then blended before generalized hierarchical backprojection.

For tiled mosaics, the sample was instead displaced relative to the rotation axis, so each acquisition produced a separate subvolume. Slow motion appeared as a smooth, low-frequency deformation between subvolume coordinate systems and was represented through an elastically deformable common atlas. Corresponding control points were detected with Harris3D in overlapping regions; a CNN retained points belonging to the papyrus and rejected those from the support, protective case or background. The remaining matches determined the smooth displacement fields used to align and fuse the subvolumes.

The apparent rotation-axis position was also calibrated as a function of height, to account for slow variations of the effective center of rotation along the vertical direction, caused for example by a possible mechanical misalignment between the vertical translation axis and the rotation axis.

The 32-bit reconstructions were converted to 16-bit TIFF stacks using black and white saturation levels of 0.01%, with binned ($\times 2$, $\times 4$) and JPEG2000 copies archived at the ESRF[44–46]. For all analyses reported here, we converted the reconstructions to unsigned 8-bit OME-Zarr[47] with six resolution levels, with voxels outside the sample set to zero (about 20 TB per volume); these volumes are distributed through the Vesuvius Challenge website (see Data availability).

Stage 3: Virtual unwrapping

Recto-surface annotations and prediction.

Recto-surface prediction was formulated as supervised 3D semantic segmentation of the papyrus writing surface P in the reconstructed volume V. The target surface was the recto side of the sheet, defined operationally as the side facing the center of the scroll and corresponding, where visible, to the horizontal-fiber layer. Manual training labels were generated by voxelizing annotated surface meshes into binary volumetric masks. These labels were treated as approximate surface annotations rather than voxel-exact ground truth, because dense packing, damage, fraying and low local contrast make voxel-accurate manual labeling impractical at scale. In ambiguous or compressed regions, labels therefore marked the best effort continuous sheet surface suitable for subsequent tracing rather than requiring an unambiguous separation of recto and verso.

We trained a residual-encoder nnU-Net-style[48] 3D U-Net on CT patches normalized individually by z-score. The network produced a two-channel background/surface prediction. Examples of recto-surface predictions are shown in Extended Data Fig. 5. The training objective used Medial Surface Recall, a custom adaptation of Skeleton Recall Loss[49] for thin-structure segmentation, together with cross-entropy and soft Dice terms. This objective was chosen to emphasize recovery of continuous sheet traces and to reduce holes in the predicted thin surface. Training details are reported in Supplementary Table 2. The resulting voxelwise prediction was used only as an intermediate cue for geometric tracing: because dense predictions can contain local sheet mergers, gaps and false positives in tightly packed regions, it was converted into orientation evidence for the reconstruction of the quadrilateral mesh M described below.

Surface estimation as a quad mesh.

The mesh M that models P is assembled from smaller, locally fitted patches and stored as a tifxyz quadrilateral mesh. In this custom format, the surface is represented by a regular 2D parameter grid, and each valid grid vertex stores its 3D position in V. The three coordinate components are saved as aligned TIFF images, x.tif, y.tif, and z.tif; adjacent valid vertices define quadrilateral mesh cells.

Automatic surface estimation began from binary volumetric predictions of the recto surface. For each of the three orthogonal slice families, predicted sheet regions were reduced to centerline traces. Each trace provided local in-plane orientation evidence for the sheet, and these directions were stored in spatial bins so that nearby sheet orientations could be queried efficiently during mesh optimization.

Each base mesh was initialized from a seed point or from an already reconstructed neighboring patch. Starting from this initial mesh, the software iteratively proposed new vertices along the boundary of the 2D mesh grid and optimized their 3D positions. The optimization jointly balanced regular mesh spacing, local smoothness, agreement with nearby orientation evidence, and proximity to the binary surface predictions, allowing the mesh to follow the predicted sheet while preserving a coherent parameterization.

To form larger unwrappings, a separate surface-based tracing tool combined base meshes whose seed locations and density had been chosen to ensure substantial overlap. Starting from a seed quad, this tool optimized new mesh vertices toward the consensus surface defined by the overlapping base meshes. This enabled generation of large meshes, in some cases approaching the scale of an entire scroll, with improved quality compared to the base patches. However, mergers or holes in the binary surface predictions could still lead to holes or deviations in the large mesh, particularly where the papyrus was densely packed or anomalously curved. Areas of sufficient quality were manually extracted and used as additional mesh sources.

Where a high-quality mesh already existed for one winding of a scroll, the mesh-offset tool was also employed. It traced rays from the vertices of the base mesh along the mesh normal until each ray intersected a new surface prediction. The resulting mesh was then refined using the same prediction-based mesh optimization objectives.

For meshes from all sources, manual correction was performed on the same mesh representation using an interactive annotation interface. Annotators could add control points, move vertices to corrected positions, and use push or pull operations to displace local mesh neighborhoods along the local mesh normal. The agreement of M with P was assessed by inspecting mesh intersections overlaid on orthogonal slices of V, interactively rotating the 3D view, and visualizing a CT-textured 2D grid in real time. Regions judged geometrically consistent with a single sheet were marked with an approval mask. After edits, local reoptimization propagated the correction constraints while restoring smoothness and mesh consistency. Small holes or invalid cells in the mesh grid could be repaired by local geometric inpainting where the surrounding sheet trajectory and layer identity were unambiguous. This inpainting was used only to fill gaps in the mesh and repair geometric continuity; it did not create tomographic intensity values, ink signal, or image evidence for unsupported readings.

Flattening and rendering.

For each recovered piece, the corrected mesh M was used for final flattening and normal-direction sampling (Fig. 2b–d). Valid quadrilateral cells were converted to a triangular OBJ mesh by splitting each cell into two triangles. Per-corner texture coordinates were initialized from the original surface-grid coordinates and scaled according to the surface metadata. The mesh was parameterized with SLIM[30], initialized from these UV coordinates. The optimization used the symmetric Dirichlet distortion energy, which is usually used in the literature as an isometric parametrization objective, with the boundary constrained to its initial UV coordinates, producing a low-distortion 2D parameter domain, with flattened coordinates (u, v). The optimized UVs were shifted to non-negative coordinates and written back to OBJ.

The optimized mesh was rasterized back into tifxyz form in the flattened UV domain. For each output pixel, the corresponding 3D surface position was computed by barycentric interpolation within the containing triangle. This inverse map through M assigns each flattened coordinate (u, v) its 3D reference position in V.

For rendering, the software evaluated the flattened surface on a regular 2D canvas. At each canvas pixel, the renderer computed the corresponding 3D surface point and local surface normal. A stack of 65 samples was taken along this normal direction, centered on the reference surface. For sample index k, the normal offset was

$$d_k = (k - (N-1)/2)\Delta, \quad k = 0, \ldots, N-1,$$

with $N = 65$ and $\Delta = 1$ voxel. This gives offsets from -32 to +32 voxels; at 2.4 µm isotropic voxel sampling, the distance between the first and last samples is $64 \times 2.4 = 153.6$ µm. Tomographic intensities were sampled at each off-surface position by trilinear interpolation from the eight neighboring voxels. Samples outside the valid volume bounds were assigned zero. The resulting warped 3D CT chunk is the surface-conditioned render $S(u, v, d_k)$ (Fig. 2d). Its two in-plane coordinates follow the flattened mesh, and its depth coordinate samples the neighborhood along the local normal in V. Each fixed-depth slice is a 2D render; the central slice ($d_k = 0$) samples V on M. A selected depth range of S forms the input to the ink-detection models described below.

Stage 4: Ink detection

Fragment registration for ink labels.

The first supervised ink labels came from detached scroll fragments, whose writing was directly exposed by historical attempts at mechanical opening. Infrared photographs, acquired with a MegaVision multispectral system equipped with a 50-megapixel E7 camera back, a UV–IR apochromatic lens and four UV–IR LED light panels, were aligned to CT-based surface representations[3, 22, 50].

These fragments provide the direct-visibility labels mentioned in the Technical Strategy and serve as reference for ink location in the surface plane. Because the writing is directly exposed, infrared photography can capture it with high contrast and visibility, isolating the ink signal that is otherwise difficult to detect. A corresponding CT-based surface is generated through an orthographic projection, casting rays onto the digital volume in the

same manner that the camera images the fragment in the real world. The infrared image is then warped to align with the visible features of the digital volume. Extended Data Fig. 6 shows a registered infrared image and its CT-based surface representation.

After registration, the resulting labels were mapped into the same flattened 2D coordinate domain (u, v) associated with S used for training and inference. This permits the detector to learn from physically exposed writing while being deployed on surfaces extracted from sealed scroll interiors.

Ink detection: from fragments to sealed scrolls.

For *PHerc.* 139 and *PHerc.* 1667, each detector used local 3D chunks of S as input. During training, these chunks were aligned with 2D label regions in (u, v). Initial training used paired virtually unwrapped CT surface volumes and registered infrared ink labels from *PHerc.* 9B, *PHerc.* 343P and *PHerc.* 500P2. The infrared images accurately localize the ink signal in the surface plane, but because they capture only a top-down view of each fragment, they provide no information about the ink's position in depth. The ink detectors use 3D ResNet[51] encoders with U-Net-style[52] decoding to predict ink in the 2D coordinate domain (u, v), in which the infrared labels are defined. The decoder may operate in 2D after depth pooling or in 3D before depth reduction; the configuration used for the *PHerc.* 1667 label-coverage ablation and the model used for the *PHerc.* 139 images in Extended Data Fig. 4 are described in Supplementary Tables 3 and 4, respectively.

The ink detector sees only local 3D chunks of S, constraining the task to local CT texture and morphology and limiting the scope for predictions driven by linguistic or word-level priors. We chose an input size of 256 $\times$ 256 pixels in the surface plane, which roughly corresponds to 614 $\times$ 614 µm at 2.4 µm sampling (Extended Data Fig. 7). This input window was chosen to emphasize local CT structure and restrict the context available about the surrounding text.

The fragment-trained model detected initial ink traces on *PHerc.* 139 and *PHerc.* 1667 before any training on those scrolls, providing a pretrained backbone for scroll-specific adaptation.

Adaptation used iterative pseudo-labeling based on the model's own predictions[53] refined by an expert (Extended Data Fig. 8). Starting from an initial set of labeled tiles, we fine-tuned the model and ran inference across the surface-conditioned renders S of the scroll. Regions in which the model revealed additional candidate ink were cropped into new training tiles, with pseudo-labels derived from its predictions. Repeating this process expanded the dataset and improved the consistency and visibility of the ink signal for papyrological review. For *PHerc.* 1667, this procedure was applied at larger scale using three columns for iterative labeling and five iterations. We trained with a mixture of Dice loss and binary cross-entropy loss with label smoothing to limit overfitting to the signal and to reduce the impact of label noise[54].

Additionally, the same procedure was applied to a scanned region of interest in *PHerc.* 1667 with a pixel size of 1.1 µm, enabling the reading of previously illegible portions of the text and improving their legibility.

A label-coverage ablation (Extended Data Fig. 8) compared matched models whose ResNet3D-50 encoders were initialized from the same weights pretrained on the combined Kinetics-700 and Moments in Time video datasets[51, 55] and trained for the same number of

steps, using label sets from successive pseudo-labeling rounds. For the ablation, successive label sets were developed on segment ℓ_2. The matched models were assessed qualitatively on segment ℓ_5, which was held out throughout training and has no reference ink labels. Across the successive label sets, additional text became visible in the held-out column.

Volumetric ink segmentation in *PHerc. Paris.* 4.

Because the 2.4 µm-voxel volume V of *PHerc. Paris.* 4 preserves enough ink–substrate contrast for direct visual identification of ink, we trained a 3D residual U-Net to predict per-voxel ink probability in V from 256^3-voxel patches. A custom 3D DINOv2[56] representation was used to derive coarse 3D ink labels. Architectures and training procedures are described in Supplementary Tables 5 and 6.

Dense supervision was obtained primarily from the *PHerc. Paris.* 4-finetuned 3D DINOv2 embedding space. An expert selected a small set of visibly inked voxels in axis-aligned slices, and their DINO patch embeddings were averaged into a normalized ink prototype. Cosine similarity to this prototype produced a dense 3D ink-likeness map, which provided the main label signal.

To make these DINO-derived labels conservative, we combined two complementary signals: the DINO ink-likeness map and the output of a separately trained volumetric ink detector. This detector was trained from 2D annotations made on images derived from S, mapped into V through M and extruded along the local normal to M to give them an approximate thickness, and was therefore useful mainly as a coarse support constraint: it reduced spurious DINO responses, while the DINO map supplied the finer localization cue. In practice, the two maps were multiplied and thresholded, with low-intensity air and void voxels masked out from the second guided round onward so that labels could only occur near papyrus material.

A student detector was trained on these combined labels, then the procedure was repeated using the guided student as the coarse support model. This was intended to refine the DINO-derived supervision while preserving the DINO ink prototype as the primary localization cue. Finally, we self-distilled the DINO-guided detector by generating dense labels from two guided-detector checkpoints under test-time augmentation, retaining material and intensity gates. This final step was intended to sharpen boundaries and suppress residual artifacts. The predicted ink volume was subsequently projected onto flattened surfaces.

Stage 5: Papyrological interpretation

Validation and papyrological review.

For *PHerc.* 1667, the flattened and ink-enhanced pieces were then merged into one coherent scroll view for papyrological interpretation (Fig. 1, stages 4–5). The *PHerc.* 1667 transcription was produced and interpreted at column level by nine qualified papyrologists. Reviewers inspected the flattened renderings, ink-enhanced outputs and geometric context of the sampled surface. Accepted readings met three criteria: (i) the sampled surface was geometrically tight to the papyrus layer in the CT volume; (ii) letterforms remained consistent across repeated renderings or inference passes; and (iii) the transcription was independently endorsed or resolved by consensus.

Uncertainty in the readings was retained in the transcription. Papyrological review could prompt reinspection and refinement of surface geometry, rendering or ink enhancement (Fig. 1, dashed feedback paths from stage 5 to stages 3 and 4).

Complete *PHerc.* 1667 transcription and translation

PHerc. 1667 is today a portion of a scroll that is incomplete in both its height (8 cm out of the usual 19–24 cm of a full roll) and diameter (2 cm out of 4–6 cm for a full roll). Repeated attempts to physically open *PHerc.* 1667 were unsuccessful. Archival records document an unsuccessful attempt in the nineteenth century, after which the papyrus was classified as a *midollo*, the compact inner core of a carbonized roll from which the outer layers had been removed. Another attempt to open the papyrus was initiated in November 1969 by the conservator A. Fackelmann, but it was unsuccessful and was therefore discontinued[57, 58]. In the 1980s, the papyrus was subjected to the Oslo method, a mechanical opening technique developed for Herculaneum scrolls, which yielded ten fragments from the outer layers[59]. Although a small number of isolated letters could be identified, overlapping sheets prevented the reading of continuous text, and the roll was consequently classified as unreadable. These interventions substantially altered the physical state of the artifact, reducing the diameter of the surviving roll from 4.9 to 2 cm, and its weight from 14 g to approximately 6 g.

Following virtual unwrapping, *PHerc.* 1667 has been discovered to preserve the lower portions of the final columns of a philosophical treatise. The title, which in ancient bookrolls was typically written also at the end of the text, has not been recovered and may have been located in the now-lost upper portion of the roll. Consequently, neither the authorship nor the title of the work can be established with certainty. Despite the fragmentary state of the manuscript, the good legibility of several consecutive lines enables the identification of the principal themes addressed in the surviving portion of the text. The surviving text appears to focus primarily on ethical theory, likely in connection with the moral perfectibility of human beings, a central concern of ancient philosophy and a defining objective of Hellenistic philosophical traditions in particular. Moreover, a number of lexical and conceptual features, together with the mention of the Stoic philosopher Aristocreon, the nephew and disciple of Chrysippus, in what appears to be the final preserved column, point toward a Stoic context.

The paleographical and textual evidence—the mention of Aristocreon, whose period of activity is generally placed around the end of the 3rd century BC—supports dating *PHerc.* 1667 to the late 3rd century BC or beginning of 2nd century BC.

A comprehensive interpretation and commentary lie beyond the scope of the present paper and will be the subject of a separate papyrological study. The text presented here should be understood as a preliminary transcription of the Greek text, accompanied by an English translation, intended to support the validation of the virtual unwrapping and ink detection processes.

Coll. 1-4

traces

Col. 5

lost margin
ll. 1-2 *traces*
]. . . . τομ̣. [
]τ̣ί̣θεcθαι ψευ̣. [
]ϛ τὸ ὅμοιον. . [

« ... the similar ...»

3-4 τὸ μὴ [cυγκα|τα]τ̣ί̣θεcθαι ψευδ̣[ῶc or τὸ μὴ [cυγ|κατα]τ̣ί̣θεcθαι ψευδ̣[ῶc?

Col. 6

lost margin
l. 1 *traces*
ll. 2-4 *lost*
l. 5 *traces*
] ὁ̣ρμὰc περι[
]. c τομηcυ. [
]θεϛθαι τῶι . [

« ... impulses ... »

7-8 τὸ μὴ cυγ[κα|τατί]θεcθαι or τὸ μὴ cυγ[κατα|τί]θεcθαι?

Col. 7

lost margin
ll. 1-6 *traces*
. . . . . ψε. [. . . . .]
ϛει̣. προϛγ. [. . . . .]
. . . . . το φυϛι̣[. . . . .]
τ̣ηι̣ ἀπροπτ̣[. . .]. .

« ... deliberate(ness/ly) ... »

10 ἀπροπτ̣[ωcί]αι̣?

Col. 8

lost margin
ll. 1-5 *traces*
. ᾳ̣δε. . . [. . . . .]
l. 7 *traces*
κ. . . . η. [. . . . .]
τιν. . ε. [. . . . .]με-
ν̣ρ̣c παρα. . . [. . .]. .
ἀλλα. . . . υ. . . . .

Col. 9

866

lost margin
[. . . . .]. . . [1-2]
[. . . . .]. . ει̣εγ[1-2]
. . . . .[. .] κ̣α̣θ᾽ ὅcον
δ̣εξα̣.[. . . .].τ.c ἐ̣c̣-
5 τιν . . . .[. . . . .]. .
ὑποκατα̣.[. . . . .]
τοc τον[.]. . . . .υ
τ̣ο̣ν ἄνθ̣ρ̣[ωπον . . .]ειν
αὐτὴν ἢ cχεῖ̣ν̣ .[.]. .[. .]
10 ναι ὅτι δὲ κ̣α̣. . . .οιc

9 λαβ]εῖν Scalas

867

« … so far as … human being … this or to have/get … that/what … »

Col. 10

868

lost margin
[. . . . .]. . .[. . .]
τ̣ε̣κ̣. .[. .]. . . . . .
. . . . . .[. . .]. . . . . .
. .δετ.[. . . . .]. .
5 c̣υ̣νκριc.[. . . . .].
νῃc πο. .[. . . . . c]τά-
ναι προc̣ῆκον ἐπὶ τὸ
πᾶν̣ ἔτι .ο.[. . .]η̣ρου ἔc-
ταιδ̣έοc καὶ εν. . .ωι
10 τὸ μέγα καὶ μακρὸν

3-4 perhaps traces of a *diplē* 8]ηρου or]π̣ρου, τ̣οῦ̣ [πον]η̣ροῦ? 9 δ̣έοc or θ̣εόc? ἐν α̣ὐ̣τ̣ῶ̃ι (Hammerstaedt) or ἐν π̣ό̣γ̣ωι?

869

« … that befits/belongs on the whole still … fear will be [of? …] and … the great and long … »

Col. 11

870

lost margin
[. . . . .]κ̣α̣ι̣.[. . . .]. .
[. . . . .]. .ι̣θ. . κ̣α̣τ̣᾽ α̣ὐ̣-
τ̣ὸ̣ν̣ [. . .]α̣cτάτων . . .
. . . .[. . .].τηcιν καὶ τὴν
5 ὁρμ[ὴν . . . . .] κ̣ατὰ
τουτ̣[.]. . . . . .ι̣cηc
πρὸc [δ]ὲ τ̣ὰ̣ οὕτ̣ω̣c ἕκαc-
τα πο. . .ω̣ς̣ πεφύκα-
μεν καὶ πρὸc τὸ ταῦ-
10 τα ποιεῖν ἃ φαίνειν

3 τ̣ο̣ν̣ or τ̣ω̣ν̣ 4 ἐ̣[πιζ]ήτηcιν? 8 πονηρ̣ῶ̣ς̣? 9-10 ταῦ|τα or ταὐ|τά

867

« … according to him/it … and the impulse … For/toward each of these things in this way … we are by nature … and for/toward doing these things that … to seem … »

872 **Col. 12**

873 *lost margin*
[. . . .]ε. . .ο[. .]. κα. . .
[. .]. ου. .ο̣ς̣ ἐμ̣φαίνει.
[κοι]νό̣τατα δ᾽ ἀνθρώ̣-
[πο]ι̣ς̣ καὶ θηρσὶν τάδε
[. . . . . τ]έ̣κνα . . . . .
.[. . . . .]χ̣αι τας. α̣.
[1-2] πολὺ δὲ μάλιστα τὰ̣
κοινότατα ἕκαστα
ταῦτα συνίστησιν
`. ´ηφ. . . . .αι γὰρ ἀναγ-

2 ἐμ̣φαίνει or π̣ρ̣ο̣ς̣εμφαίνει 5 τ]έ̣κνα Rashed 10 perhaps `μ´ὴ φ

« … shows. Very common … to men and beasts these … offspring … And above all, each of the most common things constitutes these [*or* these most common things each constitute] … For, [necessity? necessary?] … »

874 **Col. 13**

875 *lost margin*
φυσικὴν ὑπ[. . . .]. . .
διὸ καὶ οὐ. . . . . .
κατὰ τὴν .π.β. . . .
ταύτην τα. .κ[.]. . .
[. . .]. α εὑρεθήσετ[αι]
καὶ οἱ βίοι προάξου̣[σιν]
οὐθέν, ἡμῶν οὔ̣τε τῆ̣ς
ἡδονῆ̣ς οὔτε τοῦ πό-
νου προσδεομένω̣ν.
ὡσαύτως δὲ καὶ τ.

1 perhaps ὑπ[όστα]σ̣ι̣ν̣ Hammerstaedt 3 perhaps ἐπ̣ι̣βο̣λ̣ή̣ν

« … natural … therefore also … according to the … this … will be found, and lives will make no progress whatsoever, as we have no need for either pleasure or pain. In the same way, also … »

876 **Col. 14**

877 *lost margin*
[. . . . .]. .α̣τα
τι αὐτο[τ]ελὲς καὶ ἐλ-
λεῖπον, οὕτως̣ βούλο̣-
μ̣α̣ι̣ λ̣έγειν α̣. . . . . .
π. . . . . π̣. . τέλεον
ἐλλείπ̣ειν τ̣[ὰ] ἀριστε̣ρ̣ά̣,
καὶ ὡσαύτ[ως] ἔ̣χειν
ἐπὶ τῶν δεξιῶν με-
ρῶν πρὸς τὰ ἀριστερά.
γίνεται δ᾽ ὁ πλεονασ-
μὸς κατὰ τὴν ὁρμὴν

1-2 perhaps κ̣α̣τά | τι 7 ἔ̣χειν Rashed

« … something complete in itself and [*or* also] lacking, in this way I want to say that … complete … the left-hand (parts) are lacking, and things stand in the same way for the right parts in relation to the left ones. There is an excess in the impulse … »

Col. 15

878

lost margin
[. . .]. .ρ̣ν̣κ̣α̣ . .[. . .]
γ. . υ̣ρ̣ν̣τ̣ος καὶ τῶν
π̣αραπ̣λ̣ησίων πάν-
των. καθ᾽ οἷον γὰρ εἶδος
ą̣ί ὁρμαὶ πεφύκασιν
γίνεσθαι, κ̣ą̣τὰ τοῦ-
τό ἐστι̣ν τὸ ἀνελλι-
πές, ὥστε μηθὲν ἐπι-
ζητε`ῖ´ν ἔτι, ἀλλὰ ἀ-
ναπ̣ε̣π̣ληρῶσθαι κα-
τὰ πᾶν ὡς ἂν κατα

1-2 κα]|τ̣ą̣λ̣ύρ̣ν̣τος? 11 perhaps κατὰ || [τὴν ἐπιστήμην Hammerstaedt

879

« … and of all similar things. For there will be what lacks nothing with respect to this kind, according to which impulses exist by nature, so as not to seek anything any longer, but to be completed in every respect as if … »

Col. 16

880

lost margin
γέ-]
γονε̣ν, δ̣ι̣[ὰ τὴν φρό]νη-
σ̣ι̣ν̣ ἐ̣γγίζουσ̣ι̣ τῆ̣ς
συντελείας. ἀ[π]ὸ̣ τού-
των δὲ μεταβą̣ί-
ν̣ουσα ἐπὶ τὸ μ̣ι̣κ̣ρ̣όν,
τὰ μὲν καθ᾽ αὑτὴν
ἐν ἡμῖν συντετέλε-
κεν πάντα οὐ δ̣ύ̣να-
μένη σ̣υμπληρῶ{ι}-
σαι τὴν φύσ̣ι̣ν, πα-
ρέδωκε δὲ ὑπὸ πρα-

12 ὑπὸ πρα||[γμάτων or ὑπὸ πρά||[ξεως

881

« … happened, by means of wisdom they approach completion. Moving from these things to the small [detail?], it [a τέχνη?] has accomplished within us all that pertains to it, even though it cannot fully complete [our?] nature, and it provided … »

882 **Col. 17**

883 *lost margin*
[. . .] . οι ζητήςο-
μέν τι, ἀλλ᾽ οὐχ ἕξο-
μεν ἐὰν τρόπον τι-
νὰ ἑαυτῶν ἀπ̣οστῷ-
μεν καὶ τῆc φύcεωc
ἡμῶν, ἔτι τε ὂν τρό-
πον ἂν ῥηθείηcαν αἱ
λοιπαὶ τέχναι κατά
τι μὲν τέλειαι εἶναι,
κατά τι δὲ ἐνλείπε`ι´ν,
τῆc φρονήcεωc δι-

« … we will inquire into something, but we will not grasp it, if in some way we depart from ourselves and from our own nature, and further, [this would be] in the same way as the remaining arts may be said to be perfect in one respect, but to be deficient in another respect, wisdom … »

884 **Col. 18**

885 *lost margin*
περὶ ἐ-]
κείνην τ̣ὴ[ν φ]ρόνη-
cιν οὖcαν, καὶ περὶ αὐ-
τὴν εἶναι, τοῦτον δὲ
περὶ τὰc βαναύcουc
ὄντα τέχν̣ᾳ̣ς ἀ̣πολε-
λεῖφθαι πολὺ τοῦ τοι-
ούτου καὶ ὥcπερ χω-
λὸν τὸ τεχνικὸν ἔ-
χειν καὶ ἐνλελειμ-
μένον, τοιοῦτό τι μοι
δοκεῖ καὶ ἐπὶ τῶν

« [If we say that] … being wisdom about that [*or* being … about that wisdom], is also about it, whereas this [person?] that deals with banausic arts is far behind from such a thing and has technical competence that is, so to speak, lame and lacking, something similar seems to me also in regard to … »

886 **Col. 19**

887 *lost margin*
[. . . .] . . [. . . . . .]
. [.] . ᾖεῖcθαι μη[δ]ε-
μιᾶc. ἐκπονηθέν-
των μέν̣τοι ἡμῶν
διὰ τῆc ζητήcεωc
ἢ τῆc μαθήcεω̣ς, κα-
τ᾽ οὐθὲν ἔτι λειφθηcό-
μεθα αὐτῶν ὁμοίωc
τέλεα τὰ καθήκον-
τ᾽ αὐτοῖc ποιοῦντεc
καὶ τὴν αὐτὴν ἔχον-
τεc αὐ̣τοῖc φρόνηcιν

« … need none. Having certainly been instructed/completed through research or learning, we will no longer be inferior to them in any respect, accomplishing the appropriate actions as they do and possessing the same wisdom as they do … »

Col. 20

888

lost margin
.[. . . .]. . . . . .
ευ. . . .ν τὸ cυντυγ-
χάνειν. τοιούτων
δὲ τῶν ἀγαθῶν ὄν-
5 των ἡμῖν, καὶ ἐκ
τῶν ἐναντίων
τῶν κακῶν οὔτε
ἀγαθόν τι ἔcται – οὐ
καλὸν δέ – οὔτε κα-
10 κόν – οὐκ αἰcχρόν – οὔ-
τε τὸ εὐδαιμονεῖν

889

« … to happen. And such being the goods for us, even from the things that are opposed to evils there will be neither anything good–nor beautiful–nor anything bad–nor ugly–, nor happiness … »

2 εὔλ̣ο̣γ̣ο̣ν? 11 perhaps τὸ εὐδαιμονεῖν || [ἐπιταθήcεται ἢ ἀνεθήcεται Hammerstaedt

Col. 21

890

lost margin
.[. . . .]. . .[. φ]ρο-
νοῦ̣ντεc μέγα καὶ
cεμνυνόμενοι καὶ
ll. 4-5 *traces*
γ̣οι. . . . . τε ἀπὸ
τοῦ̣ α̣ὐτοῦ̣ . . . . ἐ̣παι̣-
νε̣ῖ̣ν τ. . . . . γ̣
l. 9 *traces*
10 μ̣ένου οἷον κατὰ τὰ̣ς̣
ὑμνήcειc εὐ̣τ̣ον. . .

891

« … being arrogant and feeling haughty and … from the same … to praise … as according to the encomia … »

11 εὐ̣τ̣ονί̣α̣ι̣?

Col. 22

892

lost margin
. ων ε̣ἰ̣cέτι cυνα̣. . . .[. .]
. . Ἀριcτοκ̣ρέ̣ω̣ν τα̣. . . .
. . .(.) τ̣οῖc ἐχομένοιc̣ . .

893

« … still … Aristocreon [*nominative/vocative*] … to the following … »

1 cυνα̣γ or cυνα̣π̣ 2-3 perhaps τὰ̣ς̣ β̣ί|β̣λ̣ο̣υ̣ς̣? Janko 3-4 β̣ι̣|[βλίοιc]? Blank

Statistics and reproducibility

894

The complete virtual unwrapping of *PHerc.* 1667 amounted to 31 wraps and 1231 cm^2 of papyrus surface. All reported surface areas were calculated on the virtually unwrapped

895 896

meshes. The unwrapping was completed using a semi-automated wrap by wrap copy tool combined with ~25 hours per wrap of manual annotation.

The results comprise scroll-specific case studies of complete virtual unwrapping, direct volumetric ink visibility and title identification. The *PHerc.* 1667 label-coverage ablation compares matched models qualitatively on an unlabeled held-out segment. Supplementary Tables 2–6 report model configurations and, where available, the segmentation metrics and training or validation losses used to assess the models. Papyrological validation is based on review of the image evidence, geometric context and repeated renderings, with independent endorsement or consensus.

Use of AI tools

AI tools assisted with language editing, structural revision, consolidation of cited background literature, figure assembly code and explanatory vector graphics. All scientific claims, interpretations, methodological descriptions and final wording were reviewed and approved by the human authors. The human authors validated the figure assets against their identified sources. Generative image tools were not used to alter empirical CT or ink images. No AI tools were used for reading, interpreting, and translating any Greek text. No large language model is listed as an author.

Acknowledgements

We thank the Director of the National Library of Naples "Vittorio Emanuele III", Silvia Scipioni, and the curators, conservators, and staff members of the Library for access and guidance, in particular Giovanni Bova and Valentina De Martino. We acknowledge the European Synchrotron Radiation Facility (ESRF) for the provision of synchrotron radiation facilities on beamline BM18. The preliminary transcription of *PHerc.* 1667 presented in this paper was first discussed at the conference The Eternal Library (Naples, 22–26 June 2026). We are grateful to the scholars who took part in the discussion for their valuable comments and suggestions, some of which we have incorporated into the present paper. These contributions are duly acknowledged in the critical apparatus in § "Complete *PHerc.* 1667 transcription and translation" in Stage 5 ("Papyrological interpretation").

Funding

This research was supported by the Vesuvius Challenge (https://scrollprize.org), a non-profit donation-funded organization that aims to read the full collection of sealed scrolls from Herculaneum.

The Mellon Foundation sponsored the previous work by the EduceLab, University of Kentucky, that laid the groundwork for the realization of this research.

Author contributions

- G.A. led the project since June 2025, conceived the improved CT scanning protocol, contributed to photogrammetric acquisition, case design, ink segmentation models, mesh parametrization algorithms, scale-up on cloud architectures and drafted the first and the final versions of the manuscript.
- S.P. led the project from February 2024 till June 2025, pioneered ink detection pipelines and contributed to photogrammetric acquisition, design of cases and scan protocol research.
- F.N. led the papyrological efforts, selected the papyri from the National Library of Naples, and contributed to the photogrammetric acquisition and multispectral and infrared imaging.
- Y.N. contributed to the ink detection models and pseudo-labeling.
- S.J. developed the mesh tracing software and contributed to mesh tracing algorithms, ink detection models and annotation.
- D.J. contributed to the annotation efforts since May 2023, led the annotation efforts since February 2025, and developed the whole-scroll banner merging algorithm.
- P.H. supervised the machine learning and geometry research agenda and contributed to mesh fitting algorithms.
- H.S. conceived and developed the mesh tracing software and contributed to mesh tracing algorithms.
- J.R. conceived scaling on cloud architectures and contributed to data pipelines and software optimization.
- F.M. contributed to the mesh tracing software.
- E.R.D.P. contributed to fibers and ink detection models and annotation.
- P.T. and A.M. supervised, conceived and implemented the CT data acquisition and reconstruction at the BM18 beamline.
- C.S.P. pioneered CT scanning protocols and mesh tracing algorithms and contributed to the photogrammetry acquisition.
- J.P.P. led the Vesuvius Challenge online contest from January 2023 till February 2024.
- B.K. led the annotation efforts from May 2023 till February 2025.
- A.L. contributed to papyrological review with particular focus on custodial and unrolling history of *PHerc.* 1667, manuscript annotation, photogrammetric acquisition, and multispectral and infrared imaging.
- R.V. contributed to papyrological review with particular focus on script and dating of *PHerc.* 1667, manuscript annotation, photogrammetric acquisition, and multispectral and infrared imaging.
- C.V. contributed to papyrological review with particular focus on *PHerc.* 1667 coll. 1-12, manuscript annotation, photogrammetric acquisition, and multispectral and infrared imaging.
- M.C.R. contributed to papyrological review with particular focus on *PHerc.* 1667 coll. 13-22, manuscript annotation, photogrammetric acquisition, and multispectral and infrared imaging.

- M.D.A. contributed to papyrological review with particular focus on *PHerc.* 139, manuscript annotation, photogrammetric acquisition, and multispectral and infrared imaging.
- K.F. and M.McO. contributed to papyrological review and manuscript annotation.
- G.D.M. performed papyrological review.
- C.C. edited and rewrote sections of the paper, coordinated institutional relationships, permissions, conservator interactions and partner communication.
- N.F. contributed to project conception, secured resources, advised on open-science and data-release strategy.
- W.B.S. edited and rewrote sections of the paper, as well as initiated and pioneered the project and provided overall strategic guidance.

G.A., S.P., F.N., Y.N., S.J., D.J., P.H., H.S., J.R., F.M., E.R.D.P., P.T., A.M., J.P.P., C.V., M.D.A., G.D.M., W.B.S. acquired data on the BM18 beamline. G.A. and E.R.D.P. acquired data on the ID11 beamline.

Competing interests

The Authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Giorgio Angelotti. W. Brent Seales may also be contacted for material and general inquiries.

Data availability

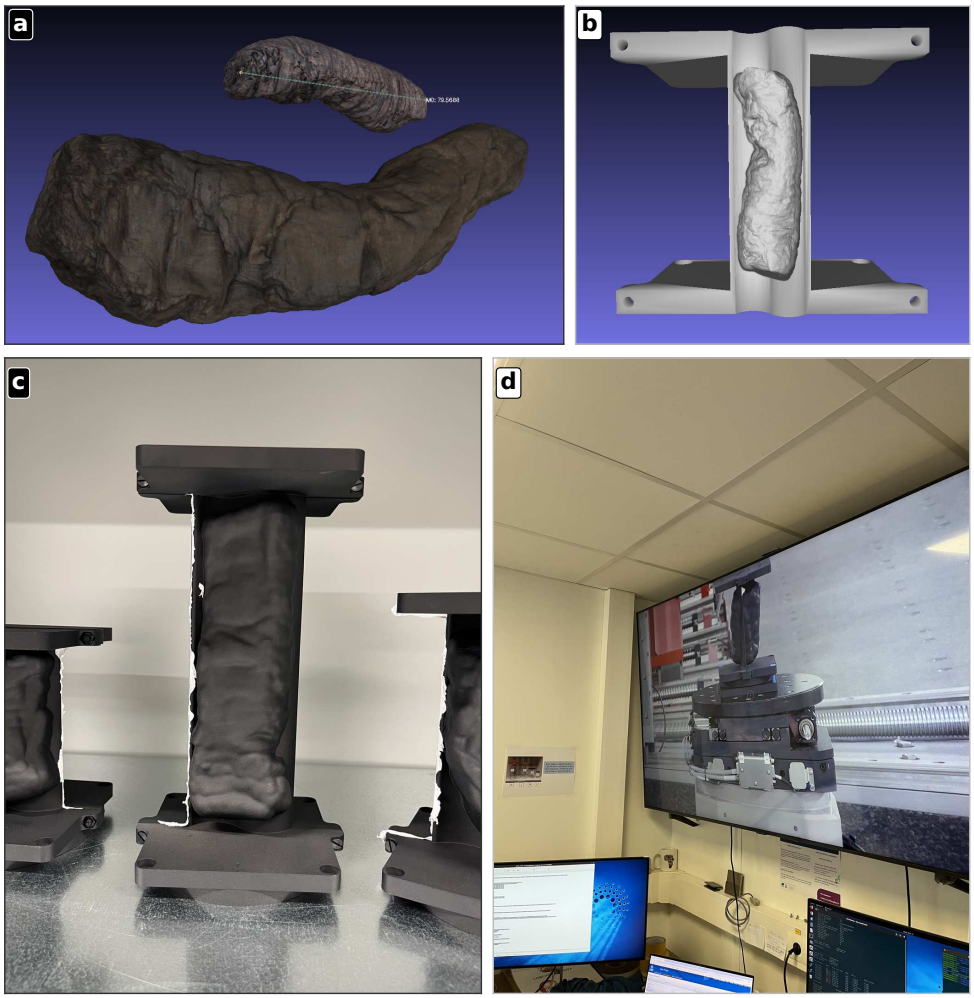
Tomographic volumes used for all analyses in this work (unsigned 8-bit, masked OME-Zarr; see Methods), reconstructed surfaces, flattened renderings, figure source data and transcription artifacts are available at https://scrollprize.org/data_browser. Data is shared under the Creative Commons 4.0 - BY NC License unless otherwise specified. JPEG2000 stack versions of the original 16-bit ESRF reconstructions, prior to 8-bit conversion, masking and OME-Zarr export, are also available on the ESRF open access database for art, history and archaeology[44–46] under the same license.

Code availability

Code sufficient to reproduce the main analyses, including mesh tracing, flattening, surface-conditioned rendering and ink-detection inference, is available at https://github.com/ScrollPrize/villa/commit/e583fb67468f483fadd73d52e06f0ab0fe5ba813.

ML model training configurations and datasets are listed in Supplementary Tables 2–6; model weights are available on the Vesuvius Challenge Hugging Face organization (https://huggingface.co/scrollprize). The *PHerc.* 1667 ink-detection ablation checkpoints are stored at https://huggingface.co/collections/scrollprize/pherc1667-ink-detection-ablation.

Extended Data

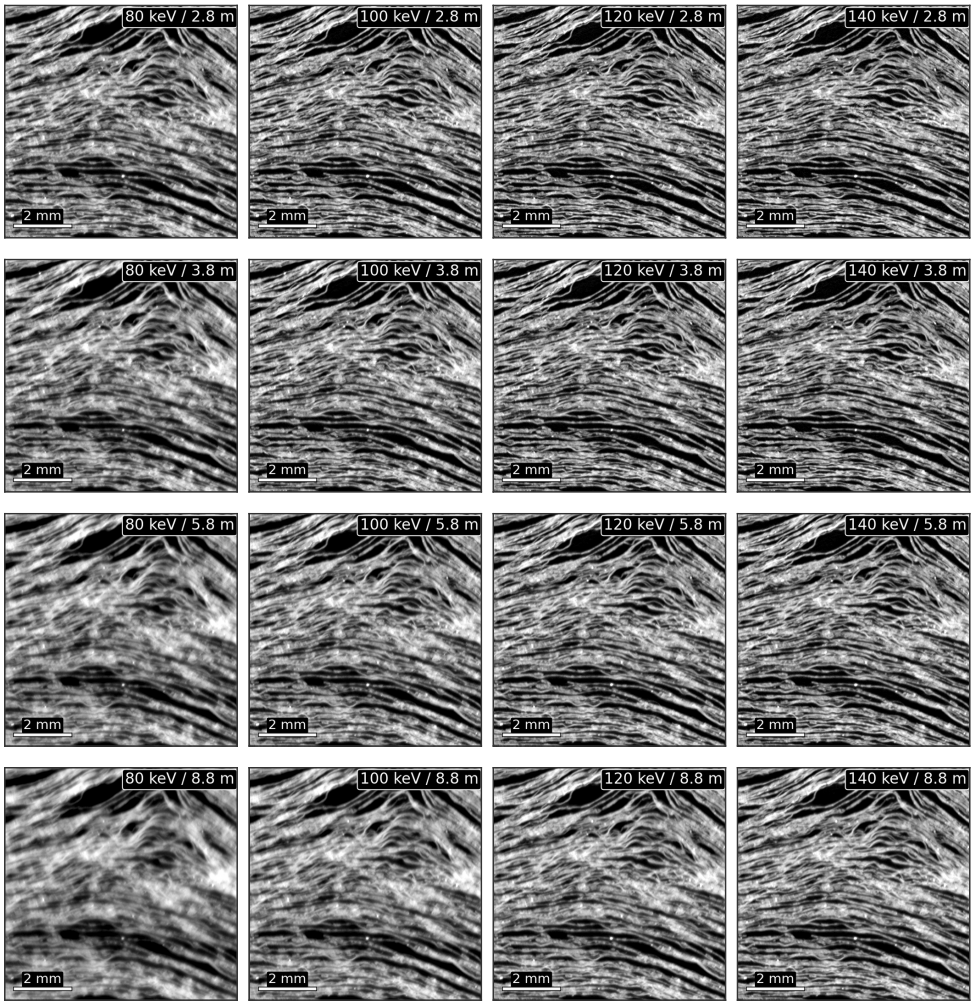


Extended Data Fig. 1 | Sample preparation and scan cases. **a**, Photogrammetric meshes of *PHerc.* 1667 (top) and *PHerc. Paris.* 4 (bottom), reconstructed from multi-view photographs; the M0 label gives the maximum chord length of *PHerc.* 1667 in millimeters. **b**, Automatically generated half-case model for *PHerc.* 1667 (non-textured render), with the photogrammetric scroll silhouette visible inside the cavity; the case is fitted to the outer morphology to minimize scan time, dose exposure and movement. **c**, Multi-jet-fusion-printed nylon cases vertically standing with scrolls inside. **d**, A cased sample mounted on the beamline rotation stage, viewed on the operator's monitor during acquisition. The scrolls shown in c and d are illustrative and are not among the objects reported in this paper; see Methods for case generation, materials and acquisition parameters.

Extended Data Table 1 | Per-scroll summary of scan parameters.

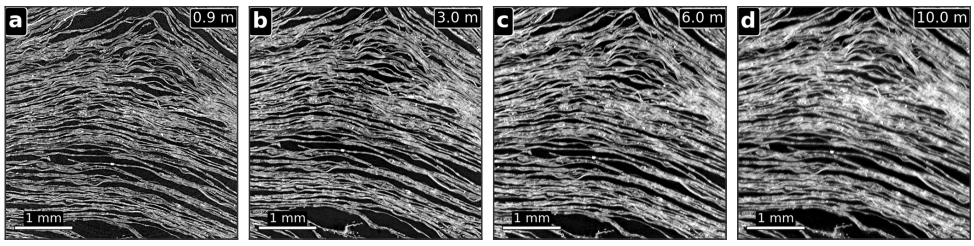
Sample ID	Role in manuscript	Beamline	Avg. energy (keV)	Voxel / pixel size (µm)	Sample-to-detector distance (m)	Phase retrieval δ/β	Unsharp coeff.	Unsharp σ (pixels)
PHerc. 9B fragment	Fragment-derived ink supervision	BM18	77	2.401	0.35	1000	4	1.2
PHerc. 343P fragment	Fragment-derived ink supervision; scan-calibration reference	BM18	77	2.403	0.22	1000	4	1.2
PHerc. 500P2 fragment	Fragment-derived ink supervision	BM18	111	2.215	0.40	1000	4	1.2
PHerc. 500P2 fragment	High-resolution fragment comparison	ID11	65	0.55	0.07	20	—	—
PHerc. Paris. 4	Volumetric ink-validation case	BM18	78	2.400	0.22	1000	4	1.2
PHerc. Paris. 3	Scanned comparison volume	BM18	78	2.400	0.22	1000	4	1.2
PHerc. 1667	Principal complete-unwrapping result	BM18	78	2.399	0.22	1000	4	1.2
PHerc. 1667 (slab)	Complement reading, detector-pixel-size comparison	BM18	59	1.129	0.22	500	4	1.2
PHerc. 1451	Scanned comparison / training volume	BM18	78	2.399	0.22	1000	4	1.2
PHerc. 332	Scanned comparison / training volume	BM18	78	2.399	0.22	1000	4	1.2
PHerc. 139	Title-identification result	BM18	78	2.399	0.22	1000	4	1.2
PHerc. 814	Scanned comparison / training volume	BM18	78	2.399	0.22	1000	4	1.2
PHerc. 841 slab	Slab scan / comparison volume	BM18	77	2.403	0.22	1000	4	1.2
PHerc. 1203 slab	Slab scan / comparison volume	BM18	77	2.403	0.22	1000	4	1.2
PHerc. 1299	Scanned comparison / training volume	BM18	78	2.399	0.22	1000	4	1.2
MANB	Scanned comparison / training volume	BM18	78	2.399	0.22	1000	4	1.2
MANBp fragment	Fragment / training volume	BM18	78	2.399	0.22	1000	4	1.2
MAN5	Scanned comparison / training volume	BM18	78	2.399	0.22	1000	4	1.2

Extended Data Fig. 2 | Imaging trade-offs in densely packed Herculaneum papyrus. Representative reconstructions showing the effects of pixel size, propagation distance, incident energy and phase retrieval on layer separability and haze-like fiber effects.

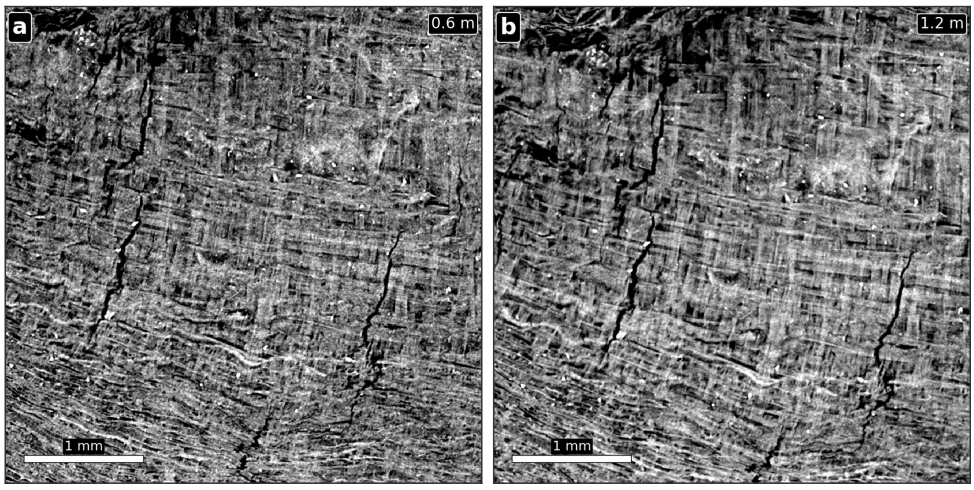


a, **Joint propagation-distance × incident-energy sweep at fixed pixel size.** Sixteen axial reconstructions of the same *PHerc.* 175A slice in a 4 × 4 matrix: rows (top → bottom) are sample-to-detector propagation distance 2.8 / 3.8 / 5.8 / 8.8 m; columns (left → right) are mean incident energy 80 / 100 / 120 / 140 keV. All cells share detector pixel size (8.009–8.013 µm effective), Paganin filter ($\delta/\beta = 1000$). Physical field of view 8 mm × 8 mm per cell; scale bar 2 mm. With pixel size held constant, reading rows top-to-bottom (fixed energy, increasing distance) shows the bright/dark phase-fringe halos broadening around every layer interface; reading columns left-to-right (fixed distance, increasing energy) shows the absorption contrast flattening as the mean

spectrum-energy rises. These tests show the decoherence effect that appears as a blurring of the slices. The most visible case is for the lower left corner, that corresponds to the lowest energy combined with the longest distance. The decoherence effect then decreases if the energy increases, and/or if the propagation distance decreases. The operationally usable corner of the matrix sits above the top-left to bottom-right diagonal, with image quality increasing with decreasing sample-to-detector propagation distance. For a pixel size of about 8 µm the empiric sweet spot for the energy is between 100 and 120 keV with propagation distances of about 3 m.

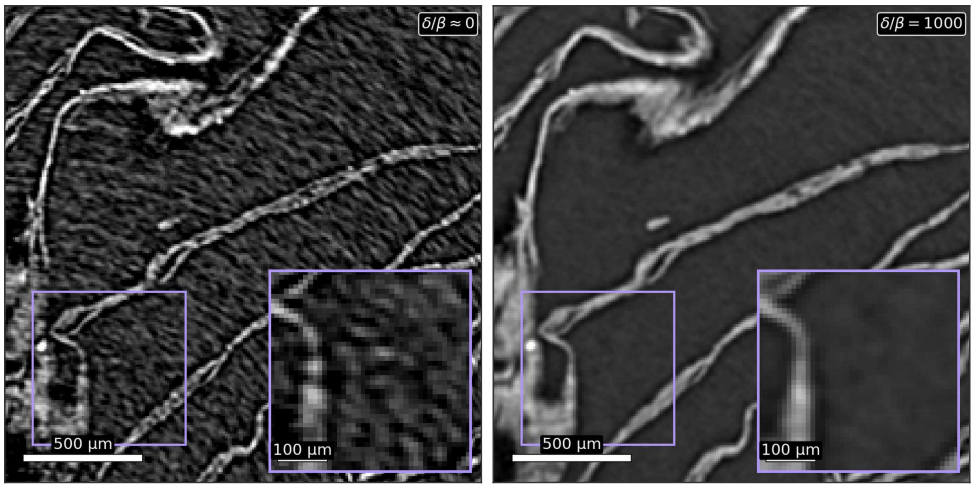


b, **Sample-to-detector propagation distance effect on retrieved phase-contrast at approximately fixed pixel size and energy.** Four reconstructions of the same *PHerc.* 175A slice at fixed mean energy (133 keV), fixed Paganin filter ($\delta/\beta = 1000$) and effective pixel size 4.1–4.3 µm as detector magnification varies with distance: 0.9 m; 3.0 m; 6.0 m; 10.0 m. Physical field of view 4 mm × 4 mm. Image quality increases with decreasing propagation distance due to the progressively increasing decoherence effect with increasing distance.

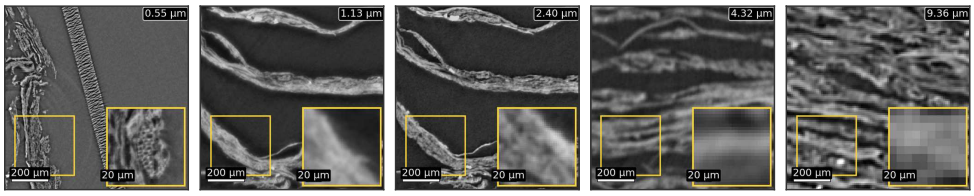


c, **Sample-to-detector propagation distance at fixed pixel size and energy.** Two reconstructions of slices from *PHerc.* 268 at fixed mean energy (133 keV), fixed Paganin filter ($\delta/\beta = 1000$), and effectively matched pixel size (4.334 µm at 0.6 m, 4.320 µm at 1.2 m): 0.6 m; 1.2 m. Physical field of view 4 mm × 4 mm; scale bar 1 mm. At 0.6 m, the crack in the slice reads as a thin dark line and the surrounding layers are well defined. Even if these slices look similar to absorption contrast, it is directly due to the phase retrieval process. In fact, even at the shortest propagation distances, the contrast is strongly dominated by phase-contrast. Absorption contrast remains very low as the

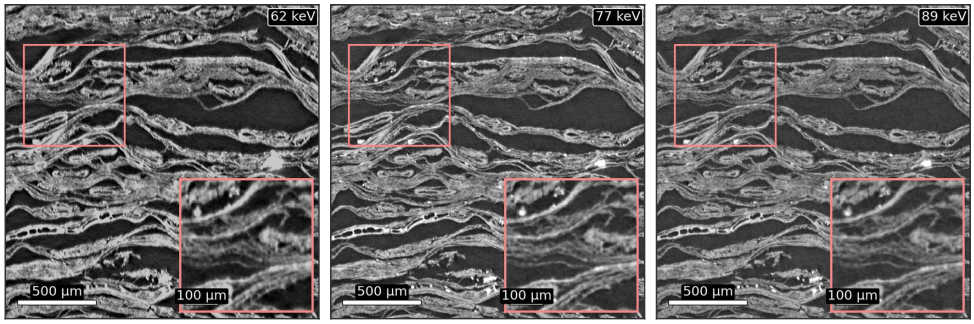
samples are mostly composed of carbon that absorbs faintly in the hard X-ray regime. The picture on the right gives an overall better visibility of the cracks and of some other structures, but at the expense of the sharpness in many places. Once again, this is due to the decoherence effect that starts to affect the resolution at 1.2 m. This comparison shows that the phase-contrast is necessary to produce usable level of contrast in these samples, but the propagation distance has to be kept low enough not to reach the decoherence level that would impact the resolution. This distance can change from one sample to another one depending on its diameter and packing density. The denser and larger it is, and the higher the decoherence effect is. These tests informed the selected propagation distance of 220 mm for the 2.4 µm pixel size setup, balancing contrast against decoherence-related blurring in large scrolls. This distance can be longer when larger pixel sizes are used (about 0.5 m at 4.6 µm and 2 m at 8 µm). *PHerc.* 268 is the only sample in the campaign with a single-sample matched sub-1 m / above-1 m pair; the full 0.9 / 3 / 6 / 10 m sweep in the previous image is on *PHerc.* 175A.



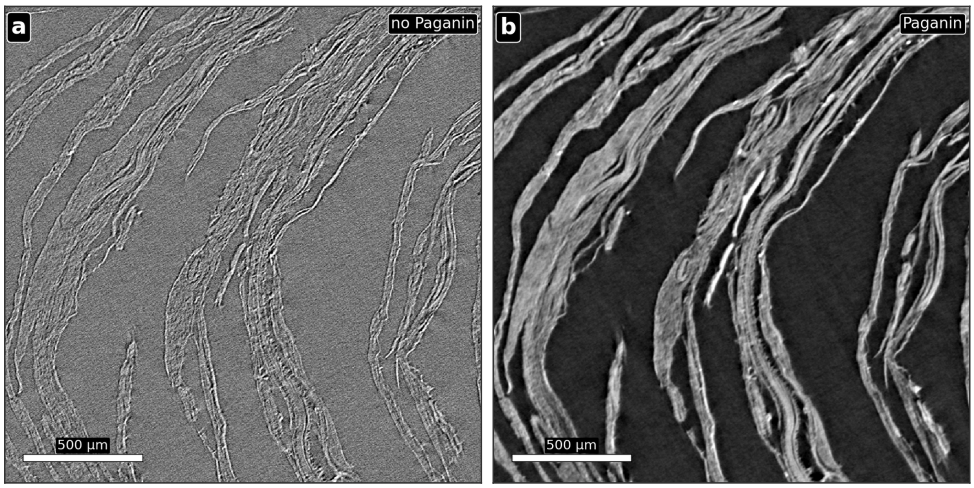
d, Companion view of the Paganin phase-retrieval trade-off on *PHerc.* 139 (9.362 µm pixel size). *PHerc.* 139, 113 keV mean, 1.2 m sample-to-detector, 9.362 µm effective pixel size, binmean2 of a 4.681 µm acquisition processed with two Paganin δ/β values illustrate the phase-retrieval trade-off. From left: δ/β ≈ 0 (no phase retrieval, equivalent to the raw filtered) back-projection and δ/β = 1000 (the production setting used for the principal 2.4 µm unwrapping and ink-detection scans). Each panel covers a 2.0 mm × 2.0 mm field of view at native pyramid level 0; the violet rectangle marks a 400 × 400 µm sub-ROI and the bottom-right inset is that ROI at 5× magnification. Scale bars are 500 µm (main) and 100 µm (zoom). At δ/β ≈ 0 fringes are visible, the image is more noisy but has also higher frequency texture signal; δ/β = 1000 (production) suppresses the fringes and obtains higher SNR and contrast at the cost of reduced apparent sharpness.



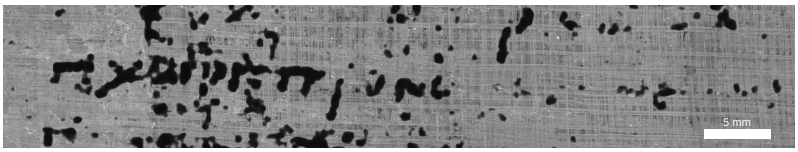
e, **Layer visibility across detector pixel sizes and acquisition conditions.** Reconstructions of densely packed Herculaneum papyrus illustrate layer visibility at five detector pixel sizes. Specimen, energy, propagation distance and phase-retrieval settings also differ, so this comparison does not isolate the effect of pixel size. From left: 0.55 µm (*PHerc.* 500P2 ID11 monochromatic small-ROI scan, 65 keV, 0.07 m propagation, $\delta/\beta = 20$), 1.129 µm (*PHerc.* 1667, BM18 59 keV polychromatic, 0.22 m propagation), 2.399 µm (*PHerc.* 1667, BM18 78 keV polychromatic, 0.22 m propagation), 4.317 µm (*PHerc.* 500P2, BM18 111 keV polychromatic, 1.2 m propagation), and 9.362 µm effective (*PHerc.* 500P2 fragment HA, BM18 113 keV polychromatic, 1.2 m, exported with binmean2 detector binning so the effective pixel size is twice). Each panel covers a 1.0 mm × 1.0 mm axial field of view at its native pyramid level 0; the yellow rectangle marks a 100 × 100 µm sub-ROI and the bottom-right inset is that ROI at 10× magnification. Scale bars at the bottom of each panel are 200 µm (main) and 20 µm (zoom inset). At 0.55 µm individual papyrus fibers are resolved (cellular cross-sections visible in the inset); at 1.129 µm fine delaminations begin to blur due to a too long propagation distance for a so small pixel size.; 2.4 µm was the selected sampling for complete imaging of large scrolls; the 4.317 µm panel is haze-limited and the 9.362 µm panel is pixel-limited, the inset shows the bare voxel grid.



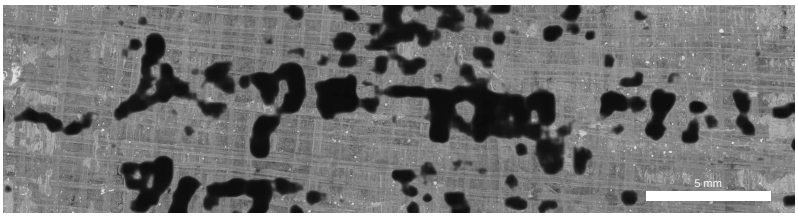
f, **Incident X-ray energy trades contrast magnitude and beam decoherence.** Three slice reconstructions of *PHerc.* 343P 2.4 µm / 0.22 m / $\delta/\beta = 1000$, 62 keV; 77 keV; 89 keV. Physical size 2.0 mm × 2.0 mm FOV. Across 62 ⟶ 89 keV, the decrease in overall contrast (faint absorption with superimposed propagation phase-contrast) is visible but gradual. Layer boundaries remain clear in every column, while layer separability increases. We found the empirical sweet spot to be at 77 keV.



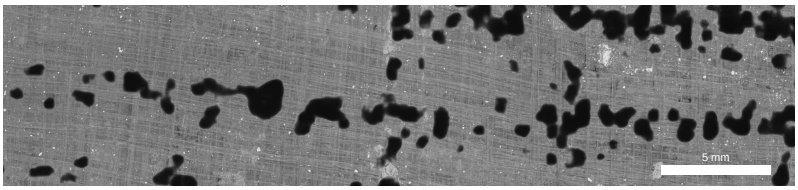
Extended Data Fig. 3 | Ink directly visible in *PHerc. Paris.* 4. Comparison of the same region of interest in the *PHerc. Paris.* 4, 78 keV, 2.4 µm pixel size, propagation distance 0.22 m **a**, No Paganin (or more accurately $\delta/\beta = 0.0001$). The slice is dominated by phase contrast fringes. Contrast is weak and washed out. **b**, Paganin ($\delta/\beta = 1000$). Contrast is strong, separation of different structures becomes clearer, ink flakes (in the center) are brighter than the surrounding papyrus.



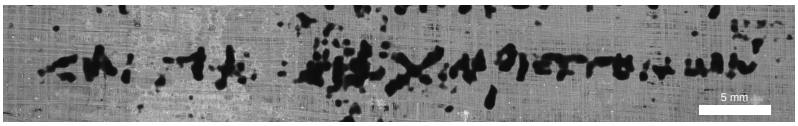
a w49: χωρὶc προνο[ί]ᾳc (“without providence”)



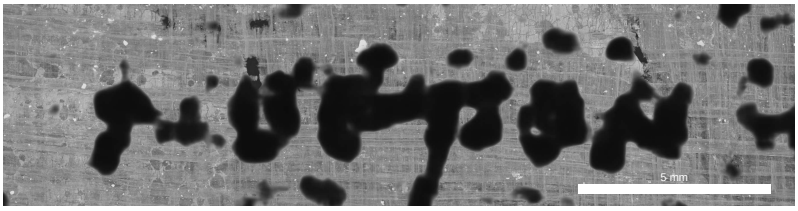
b w25: ἀόρατα (“invisible entities”)



c w34: καὶ τὸ κατᾴ φύcιν (“and the … according to nature”)

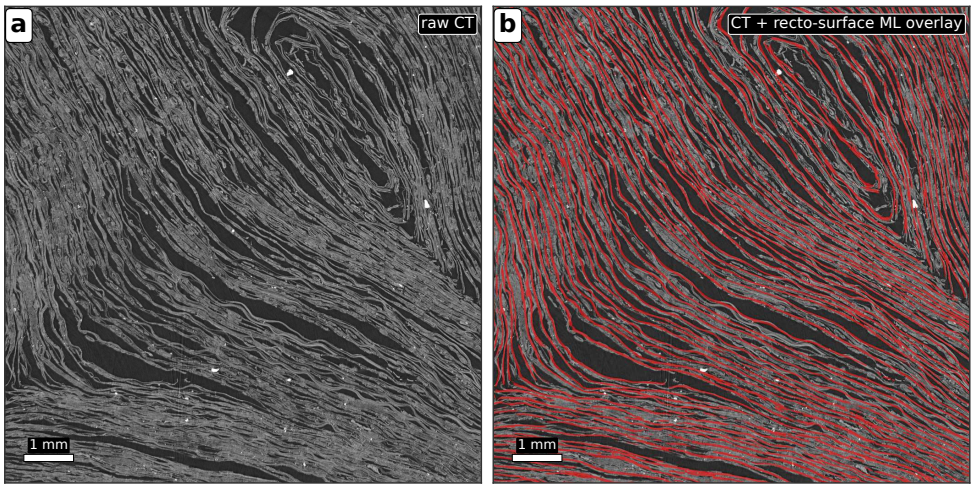


d w47: καὶ πάντᾳ χωρὶc πόνων (“and everything without troubles”)

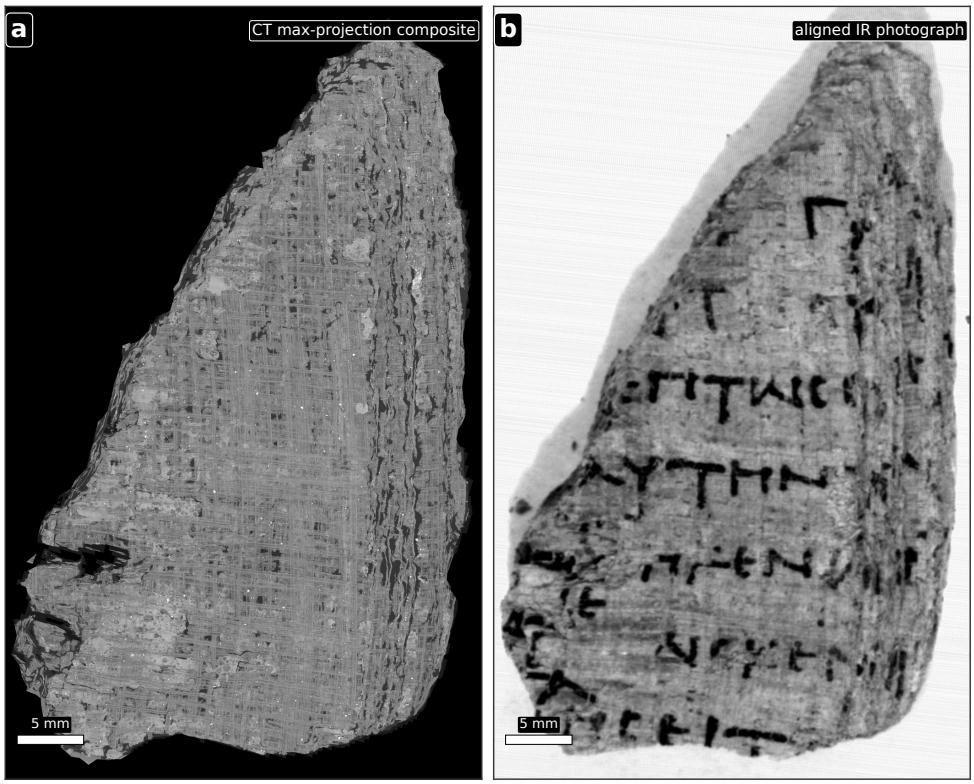


e w49: νοερόν (“intellectual”)

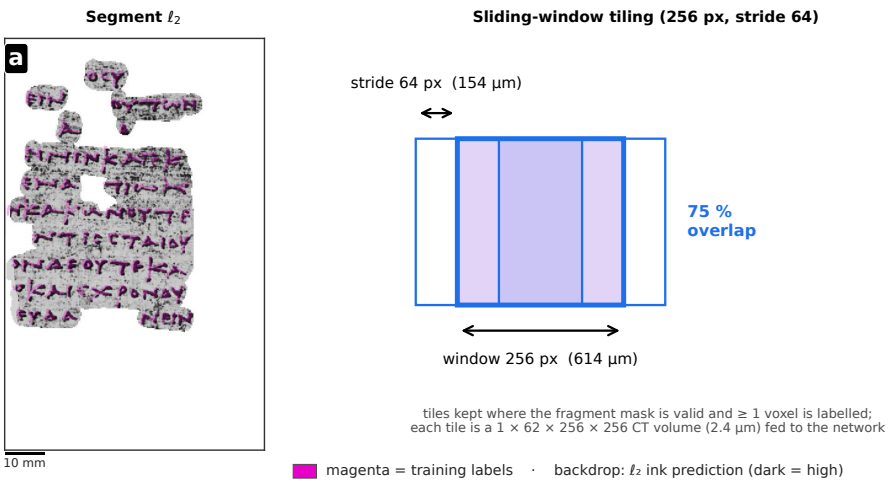
Extended Data Fig. 4 | Readings from *PHerc.* 139. Ink predictions are shown in black over maximum-projection CT composites. Scale bars, 5 mm.



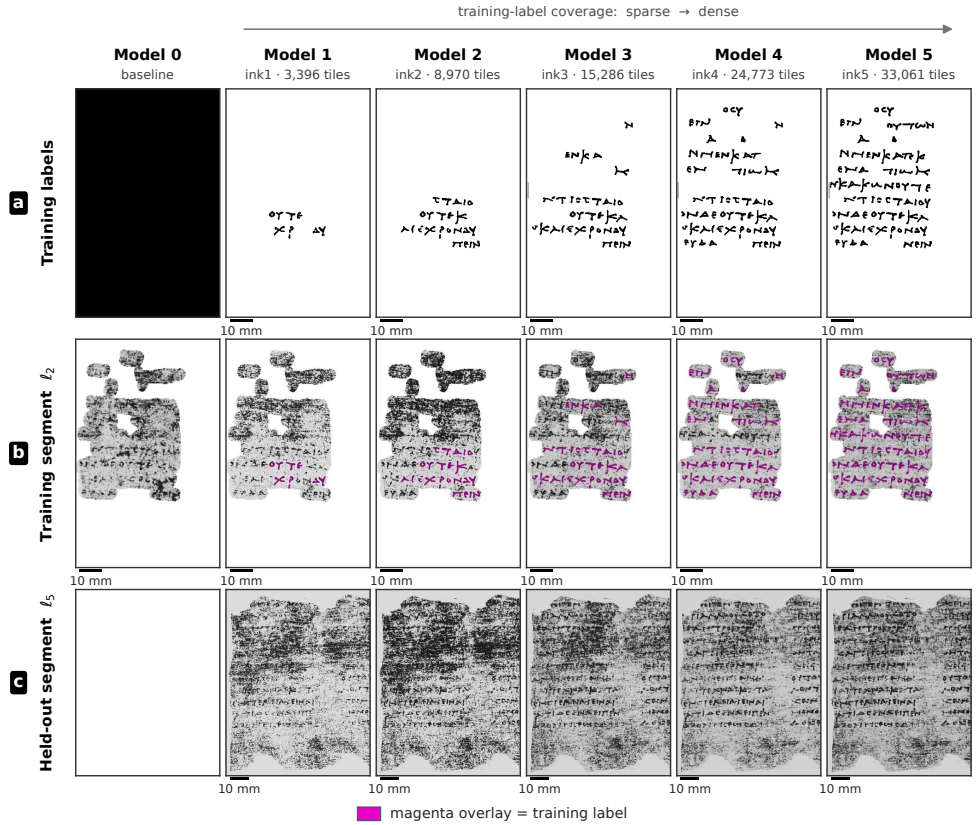
Extended Data Fig. 5 | Surface predictions. Examples of recto-surface prediction in *PHerc. Paris.* 4.



Extended Data Fig. 6 | Fragment-derived supervision. *PHerc.* 500P2 front-surface segment, a = CT max-projection composite (BM18 2.215 µm / 111 keV / 0.4 m / Paganin $\delta/\beta = 1000$, orthorectified, integrating 65 ortho-layers), b = Naples M1050 IR photograph #017 warped to the composite's coordinates via manually-placed control points. 35.8 × 58.1 mm physical FOV. Scale bars 5 mm.



Extended Data Fig. 7 | Sliding-window sampling and model input geometry. **a**, Ink predictions on *PHerc.* 1667 segment ℓ_2, with training labels in magenta. **b**, Schematic 256-pixel windows at 64-pixel stride (75% overlap along each axis). Each detector input contains 62 × 256 × 256 voxels from the surface-aligned CT stack. Scale bar, 10 mm (a).



Extended Data Fig. 8 | Pseudo-labeling ablation. Models 1–5 were trained from a common initialization pretrained on the combined Kinetics-700 and Moments in Time datasets with the same training-step count, using successively expanded label sets. Model 0 is the starting model, not trained on *PHerc.* 1667. Row a shows the training labels. Row b shows predictions on the training segment, with labels overlaid in magenta. Row c shows predictions on the held-out segment, for which reference ink labels are unavailable.

Supplementary Information

Supplementary Table 1 | Scans used for the experimental campaign

Sample ID	Figure / comparison	Beamline	Avg. energy (keV)	Voxel / pixel size (µm)	Sample-to-detector distance (m)	Phase retrieval δ/β	Reconstruction / processing note
PHerc. 175A	Energy × propagation-distance sweep	BM18	80, 100, 120, 140	8.009–8.013	2.8, 3.8, 5.8, 8.8	1000	4 × 4 matrix; 16 axial reconstructions of the same slice.
PHerc. 175A	Propagation-distance sweep at fixed energy	BM18	133	4.1–4.3	0.9, 3.0, 6.0, 10.0	1000	Four reconstructions; effective pixel size varies slightly with detector magnification.
PHerc. 175A	Paganin phase-retrieval comparison	BM18	111	2.215	0.40	0 and 1000	Same ROI reconstructed without phase retrieval and with production Paganin setting.
PHerc. 268	Matched sub-1 m / above-1 m propagation comparison	BM18	133	4.334; 4.320	0.6; 1.2	1000	Two matched reconstructions at similar effective pixel size.
PHerc. 1667	Detector-pixel-size comparison	BM18	59	1.129	0.22	500	Used as the 1.129 µm panel in the pixel-size comparison; unsharp coefficient 4 and σ = 1.2 pixels.
PHerc. 1667	Detector-pixel-size comparison / production reference	BM18	78	2.399	0.22	1000	Same as principal *PHerc.* 1667 production scan; included here only as a comparison panel.
PHerc. 500P2 fragment	Detector-pixel-size comparison	ID11	65	0.55	0.07	20	Monochromatic small-ROI scan; highest-resolution comparison panel.
PHerc. 500P2 fragment	Detector-pixel-size comparison	BM18	111	4.317	1.2	1000	Used as the 4.317 µm panel.
PHerc. 500P2 fragment HA	Detector-pixel-size comparison	BM18	113	9.362 effective	1.2	1000	Exported with binmean2 from a 4.681 µm acquisition; coarsest pixel-size panel.
PHerc. 139	Paganin phase-retrieval comparison	BM18	113	9.362 effective	1.2	≈0 and 1000	Same acquisition processed with no/near-no Paganin and with δ/β = 1000.
PHerc. 343P fragment	Incident-energy comparison	BM18	62, 77, 89	2.4	0.22	1000	Three reconstructions used to compare energy-dependent contrast and layer separability.
PHerc. 343P fragment	Propagation-distance comparison	BM18	77	2.401–2.403	0.22; 0.35	1000	Sequential acquisitions differing only in propagation distance.
PHerc. Paris. 4	No-Paganin / Paganin ink-visibility comparison	BM18	78	2.4	0.22	0.0001 and 1000	Same ROI reconstructed with near-no Paganin and with δ/β = 1000.
PHerc. Paris. 4	Production volumetric ink-validation reconstruction	BM18	78	2.4	0.22	1000	Same production volume as in Extended Data Table 1; included here because it is compared against the no-Paganin reconstruction.

Supplementary Table 2 | Training configuration of the recto-surface 3D U-Net. Values were taken from the Weights & Biases run ps256_bs2 (d5jdo9n1, 2026-03-31), archived in project vesuvius-challenge/paper.

Category	Parameter	Value
Architecture	Backbone	nnU-Net-style 3D residual-encoder U-Net
	Convolution	nn.Conv3d
	Normalization	nn.InstanceNorm3d
	Attention	Concurrent spatial/channel squeeze-and-excitation (scSE)
Input	Output channels	2 (background / surface)
	Patch size (Z × Y × X)	256 × 256 × 256 voxels
	Intensity normalization	Per-patch z-score
	Ignore label	2 (held-out / mask-out)
Training data	Volumes	*PHerc.* 0139 (three ledger pages of labels), *PHerc.* 1667, *PHerc.* 0343P, *PHerc.* 0500P2, *PHerc.* MAN Bp
	Train patches	116,531
	Validation patches	2,379
	Train / val split	0.98 / 0.02 (seed = 42)
	Min labeled-ratio gate	0.001
	Min bbox-fill gate	0.35
	Loss	Medial Surface Recall (custom Skeleton-Recall adaptation; combined with cross-entropy and soft Dice on the two-channel output, weight = 1)
Optimization	Batch size	2
	Mixed precision	AMP fp16
	Gradient clipping (L2 norm)	12
	Validation cadence	every 5 epochs (≤ 100 validation steps/epoch)
	Max epochs (scheduled)	7,500
	Final logged epoch	3,864 (run end-state: terminated)
Final validation (epoch 3,864)	Learning rate at final epoch	5.21×10^{-3}
	Mean Dice / IoU	0.596 / 0.495
	Surface (class 1) Dice / IoU	0.308 / 0.189
	Background (class 0) Dice / IoU	0.883 / 0.801

Validation Dice and IoU in Supplementary Table 2 measure agreement with the annotated recto-surface masks, before manual mesh correction.

Supplementary Table 3 | Training configuration of the ink detection model ablation

Values pulled from the wandb run vesuvius-challenge/paper/l2_ink5_l5infer (rd0qz8ps), created 2026-06-15, fine-tuned from the ResNet3D-50 release r3d50_KM_200ep.pth, pretrained on the combined Kinetics-700 and Moments in Time datasets[51, 55].

Category	Parameter	Value
Architecture	Backbone	ResNet3D-50[51] — 3D conv, BN, ReLU residual blocks
	Decoder	2D U-Net: 3 upsampling blocks of (bilinear ×2 ⟶ concat skip ⟶ 3×3 conv ⟶ BN ⟶ ReLU)
	Encoder–decoder bridge	per-stage max over the z (depth) axis ⟶ 4-D feature pyramid
	Convolution	nn.Conv3d (encoder), nn.Conv2d (decoder)
	Normalization	nn.BatchNorm3d (encoder), nn.BatchNorm2d (decoder)
	Encoder init	r3d50_KM_200ep.pth (Kinetics-700 and Moments in Time); conv1 weights summed across RGB ⟶ 1 grayscale channel
Input	Output channels	1 (sigmoid binary ink probability)
	Patch size (D × H × W)	62 × 256 × 256 voxels

Supplementary Table 3 (continued)

Category	Parameter	Value
Training data *(ink5 row)*	Intensity preprocessing	layer values clipped to [0, 200] in data.py, then Normalize(mean=0, std=1) per voxel
	Z-layer range	layers 1–62 from each segment's layer stack
	Ignore label	none (binary supervision)
	Segment	l_2 (*PHerc.* fragment)
	Label file	labels/l_2_inklabels5.png
	Train tiles	33,061 (256×256 sub-tiles at stride 64)
	Min labeled-ratio gate	label must have ≥ 1 non-zero pixel after binarization
	Min bbox-fill gate	tile fully inside the segment fragment mask
	Train / val split	100 % train; held-out segment l_5 is used for inference only (no labels)
	Inference tiles on l_5	86,765 (stride 128, 4× oversample)
	Inference tiles on l_2	22,204 (stride 128, after fragment-mask filter)
Augmentations *(training only)*	Flip	HorizontalFlip(p=0.5), VerticalFlip(p=0.5)
	Affine	ShiftScaleRotate(rot ±360°, shift ±0.15, scale ±0.10, p=0.75)
	Blur	one of {GaussianBlur, MotionBlur} (p=0.4)
	Cutout	CoarseDropout(max_holes=2, max_size=20 % of patch, p=0.5)
Optimization	Loss	0.5 × Dice + 0.5 × SoftBCE (smooth = 0.25), against label down-interpolated to 64×64
	Optimizer	AdamW
	Initial LR	2×10^{-5}
	Max LR (OneCycle)	3×10^{-4}
	LR schedule	OneCycleLR, pct_start = 0.15, final_div_factor = 10^2
	Batch size	2
	Gradient accumulation	4 (effective batch = 8)
	Mixed precision	16-mixed (PyTorch AMP)
	Gradient clipping (L2 norm)	1.0
	Max steps (scheduled)	12,396
	Steps per epoch (ink5)	4,132
	Epoch budget (ink5)	3
	Random seed	130697
	Determinism	cudnn.deterministic = True, cudnn.benchmark = False
	Compute	GPU: 1 × NVIDIA H100 80 GB
Compute	Per-run memory	≈ 12 GB GPU, ≈ 160 GB host RAM (l_2 + l_5 layer stacks held in process memory)
	Per-run wall time	≈ 2 h end-to-end (load + train + l_5 inference)
	Parallelization	5 ablation runs + 1 cross-segment baseline launched simultaneously across GPUs 0–5
Outputs	Checkpoint	checkpoints/trained/l2_ink5_last.ckpt (963 MB)
	l_2 prediction	predictions/l_2/l2_ink5_l_2_pred.{png,npy}
	l_5 prediction	predictions/l_5/l2_ink5_l_5_pred.{png,npy}
Final training loss *(ink5, step 12,396)*	Train total loss	0.5299

The held-out segment ℓ_5 in Supplementary Table 3 was used for qualitative inference only; the reported final loss is the training loss on ℓ_2.

Supplementary Table 4 | Canonical ink detector and inference for the five *PHerc.* 139 panels in Extended Data Fig. 4. Model and checkpoint metadata are from the canonical model card; inference and display settings are from the supplied panel metadata.

Category	Parameter	Value
Model	Checkpoint	`r152_3ddec_v2_l5_epoch13.ckpt`
	Architecture	ResNet3D-152 encoder, 3D decoder, learned depth-attention pooling and 2D output head
	Recipe	`new_canon_autoresearch_recipe`
	Checkpoint progress	Epoch 13; global step 113,246
Input	Source	65-layer *PHerc.* 139 surface volumes; 2.399 µm pixels
	Depth and normalization	Layers 63 down to 2 (62 layers); clip values at 200 and divide by 200
Inference	Tiling	256-pixel tiles, stride 128, Hann blending
	Execution	Half precision; no test-time augmentation
Display	CT composite	Maximum projection of 65 layers; first and 99th percentiles mapped to gray levels 91 and 194
	Overlay	Pixelwise minimum of composite and inverted ink prediction

Supplementary Table 5 | Training configuration of the 3D DINOv2 representation model (dinovol). Values were taken from the Weights & Biases pre-training run model_v2_90k_shift005_jitter105 (model_v2__shift005_jitter105, 2026-03-27), with *PHerc. Paris.* 4 finetuning from model_v2_90k_shift005_jitter105_r342500_paris4_20260416 (…__r342500__paris4__20260416, 2026-04-16). Both runs are archived in project vesuvius-challenge/paper.

Category	Parameter	Value
Architecture	Backbone	3D Vision Transformer (DINOv2 / EVA-style, "v2" config), single-channel volumetric input
	Patch embedding	nn.Conv3d, 8 × 8 × 8-voxel patches (non-overlapping)
	Embedding dimension	864
	Depth (transformer blocks)	24
	Attention heads	16
	Attention	Multi-head self-attention with mixed 3D rotary position embeddings (RoPE)
	MLP	SwiGLU, ratio 8/3 (≈ 2.667)
	Normalization	LayerNorm (pre-norm)
	Register tokens	4
RoPE settings	base 100; per-axis ("separate") normalization	train-time coord. aug. shift 0.05 / jitter 1.05 / rescale 2
	Stochastic depth (drop-path)	0.2
Input	Input channels	1 (grayscale CT)
	Global crops (Z × Y × X)	128 × 128 × 128 voxels, × 2
	Local crops	64 × 64 × 64 voxels, × 8 (scale range 0.02–0.25)
	Source sampling region	256 × 256 × 256 voxels
	Intensity normalization	Per-volume robust scaling (median / IQR)
	Voxel size	≈ 2.4 µm isotropic
Self-supervised objective	Losses	DINO (CLS-token) + iBOT (masked-patch) + KoLeo regularizer (Gram term present, weight 2, inactive)
	Centering	Sinkhorn–Knopp
	Teacher temperature	0.04 ⟶ 0.07 (warmup 25,000 steps)
	Teacher EMA momentum	0.994 ⟶ 1.0

Supplementary Table 5 (continued)

Category	Parameter	Value
Training data	Pretraining volumes (10 scrolls, ESRF ≈ 2.4 µm)	*PHerc.* 9B, 500P2, 814, 1299, 343P, 332, 139, MAN5, 1451, MANBp
	Validation volumes	*PHerc.* 332, *PHerc.* 139
	PHerc. Paris. 4 finetuning volume	*PHerc. Paris.* 4 (20260411..., 2.400 µm) — added at finetuning
Optimization	Optimizer	AdamW
	Learning rate	2.5×10^{-5} (cosine; warmup 35,000 steps → min 1×10^{-6})
	Patch-embedding LR multiplier	0.2
	Weight decay	0.04 → 0.2 (cosine)
	Gradient clipping (L2 norm)	5
	Batch size	3 (per device)
	Mixed precision	AMP
	Max iterations (scheduled)	1,000,000
	Checkpoint / validation cadence	every 2,500 / 500 steps
Pretraining end-state	Final logged step	342,558 (end-state: stopped; checkpoint_step_342500 used downstream)
	Learning rate at final step	1.95×10^{-5}
	Train / Val loss total	14.566 / 16.008
PHerc. Paris. 4 finetuning	Steps	resumed at 342,500 → final 353,686 (≈ 11,200 steps)
	Validation loss (total)	15.076

Validation losses in Supplementary Table 5 monitor the self-supervised objective on randomly sampled crops from volumes also used in training. Pretraining validation used *PHerc.* 332 and *PHerc.* 139; during *PHerc. Paris.* 4 fine-tuning, it also included *PHerc. Paris.* 4.

Supplementary Table 6 | Training configuration of the *PHerc. Paris.* 4 volumetric 3D ink-detection models. The four models share the architecture and optimization settings listed below and differ only in their supervision. Configurations were taken from the corresponding Weights & Biases runs, archived in project vesuvius-challenge/paper: feasible-voice-35 (76ezvyks, 2026-04-17), morning-rain-52 (6az0505g, 2026-04-24), divine-butterfly-57 (af3pga3d, 2026-04-24) and chocolate-wildflower-62 (4b07qv8p, 2026-04-27).

Category	Parameter	Value
Architecture	Backbone	nnU-Net-style 3D residual-encoder U-Net (vesuvius NetworkFromConfig, auto-configured)
	Convolution	nn.Conv3d
	Normalization	nn.InstanceNorm3d
	Encoder stages / feature widths	7 stages — 32, 64, 128, 256, 320, 320, 320
	Residual blocks per stage	1, 3, 4, 6, 6, 6, 6
	Decoder	Shared decoder, skip-concatenation; single ink head (1 × 1 × 1 conv)
	Parameters	1.42×10^{8} (≈ 142 M)
	Output channels	1 (ink; sigmoid + BCE)
Input	Patch size (Z × Y × X)	256 × 256 × 256 voxels
	Input channels	1
	Intensity normalization	Per-volume percentile min–max
	Mode	full-3D (surface annotations placed ± 3 voxels about the reference surface)
	Min labeled-coverage gate	0.02

Supplementary Table 6 (continued)

Category	Parameter	Value
Training data	Teacher supervision (run 76ezvyks)	Surface-conditioned 2D-ink dataset projected into 3D — 8 scrolls: *PHerc. Paris.* 4, 139, 500P2, 814, 841, 1667, MAN5, 9B
	Guidance / distillation volume (6az0505g, af3pga3d, 4b07qv8p)	*PHerc. Paris.* 4
	Seed	27
Optimization	Loss	LabelSmoothedDCAndBCELoss — soft Dice + binary cross-entropy on the 1-channel ink output (weight_dice = weight_ce = 1; Dice & BCE label smoothing 0.1; ignore-label enabled)
	Optimizer	SGD, Nesterov momentum 0.99, weight decay 3×10^{-5}
	Learning-rate schedule	Cosine with warmup (diffusers_cosine_warmup); base LR 1×10^{-2}; 5,000-step warmup
	Batch size	2
	Mixed precision	bf16
	Weight EMA	decay 0.9995 (from step 1,000); used for validation and as the frozen guidance teacher
	Max iterations (scheduled)	250,000
DINO-guided rounds	DINO model	*PHerc. Paris.* 4-finetuned dinovol (checkpoint_step_352500); 128^3 sliding windows (stride 128) $\longrightarrow 16^3$ patch tokens; Gaussian blend $\sigma = 4$; per-chunk min–max
	Reference ink embedding	Mean of 256 expert-clicked 864-D DINO patch tokens, L2-normalized
	Threshold τ	0.5 (guided rounds)
	Background mask	input intensity ≥ 50 (added from round 2 onward)
	Threshold τ	0.5
Pipeline stages & run end-state	Stage 1 · teacher (76ezvyks)	8-scroll full supervision; final step 68,000; val balanced-acc. 0.664; val loss 0.426 (EMA 0.396)
	Stage 2 · DINO-guided student (6az0505g)	init from teacher; labels = σ(teacher)·minmax(cosine$\longrightarrow$ink) > 0.5; final step 62,950; val balanced-acc. 0.715; val loss 0.317 (EMA 0.131)
	Stage 3 · + background mask (af3pga3d)	init from stage 2 (also its frozen teacher); mask intensity < 50; final step 77,000; val balanced-acc. 0.576; val loss 1.504 (EMA 1.487)
	Stage 4 · self-distillation (4b07qv8p)	init from stage 3; pseudo-labels from ensemble of stage-3 checkpoints (steps 77k + 64k) + TTA, DINO dropped, full supervision; final step 78,450; val balanced-acc. 0.603; val loss 0.352 (EMA 0.303)