

Complete virtual unwrapping and reading of a rolled Herculaneum papyrus

Giorgio Angelotti^{1*}, Stephen Parsons^{1,2}, Federica Nicolardi^{1,3}, Youssef Nader¹, Sean Johnson¹, David Josey¹, Paul Henderson^{1,4}, Hendrik Schilling¹, Johannes Rudolph¹, Forrest McDonald¹, Elian Rafael Dal Prá¹, Paul Tafforeau⁵, Alessandro Mirone⁵, C. Seth Parker², Jan Paul Posma¹, Benjamin Kyles¹, Claudio Vergara^{1,3}, Alessia Lavorante^{1,3}, Rossella Villa^{1,6}, Maria Chiara Robustelli^{1,3}, Marzia D'Angelo^{1,3}, Gianluca Del Mastro^{1,7}, Michael McOsker^{1,8}, Kilian Fleischer^{1,9}, Christy Chapman^{1,2}, Nat Friedman¹, W. Brent Seales^{1,2}

¹VESUVIUS CHALLENGE, SAN FRANCISCO, CA, USA; ²EDUCELAB, UNIVERSITY OF KENTUCKY, LEXINGTON, KY, USA; ³UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II, NAPOLI, ITALY; ⁴UNIVERSITY OF GLASGOW, GLASGOW, UK; ⁵ESRF, GRENOBLE, FRANCE; ⁶UNIVERSITÀ DI SALERNO, SALERNO, ITALY; ⁷UNIVERSITÀ DEGLI STUDI DELLA CAMPANIA LUIGI VANVITELLI, S. MARIA CAPUA VETERE, ITALY; ⁸UNIVERSITY COLLEGE LONDON, LONDON, UK; ⁹UNIVERSITÄT TÜBINGEN, TÜBINGEN, GERMANY.

*CORRESPONDING AUTHOR(S). E-MAIL(S): GIORGIO@SCROLLPRIZE.ORG

ABSTRACT

The carbonized papyri from Herculaneum preserve the only large-scale library to survive from classical antiquity, but many unopened rolls remain unread because physical opening risks irreversible damage [1]. X-ray computed microtomography (μ CT) and virtual unwrapping offer a non-invasive route to their texts, yet previous work on sealed Herculaneum scrolls has recovered only localized readings or limited surface regions. Here, using high-resolution phase-contrast μ CT acquired on the BM18 beamline at the European Synchrotron Radiation Facility (ESRF), together with improved computational unrolling and machine learning, we achieve the complete virtual unwrapping and reading of PHerc. 1667 under explicit coverage and papyrological-review criteria. This makes PHerc. 1667 the first Herculaneum papyrus to be fully digitally unrolled and read for extended scholarly study without physical opening. In PHerc. Paris 4, the optimized scan protocol makes ink directly visible in the tomographic volume, allowing three-dimensional ink segmentation and independent validation of surface-conditioned ink recovery. In PHerc. 139, we recover title and author-attribution evidence identifying the scroll as *Philodemus, On Gods, Book 8*. These results move virtual unwrapping of the Herculaneum scrolls beyond isolated demonstrations towards a scalable framework for systematic recovery of the still-unopened library.

The Herculaneum papyri present one of the longest-standing access problems in the study of the ancient world. Discovered in the eighteenth century in the Villa of the Papyri, they are the only surviving large-scale 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].

A non-invasive route to these texts is provided by X-ray μ CT combined with virtual unwrapping, in which the internal sheet geometry is reconstructed, flattened into a readable domain and examined for traces of writing[2–8]. Herculaneum material is difficult for X-ray μ CT because the carbonized papyrus support is radiographically similar to many carbon-based inks, yielding weak stroke–substrate contrast. The problem is compounded by variation in ink recipe, contaminants, deposit morphology and preservation, which means that writing may be expressed as faint metallic contrast, surface relief or subtle textural signal rather than by a single uniform signature. Earlier phase-contrast μ CT studies reported candidate text signals in sealed Herculaneum material, but they did not establish a reproducible route from whole closed samples to *scroll-scale* scholarly reading[9, 10]. Subsequent work showed that useful signals may arise from metallic components in some inks, from hyperspectral contrast in opened material, or from fine surface deformations associated with ink on fragments[11–13]. What was missing was an end-to-end route from whole sample to readable surface: a coupled method that integrates optimized synchrotron imaging, geometric reconstruction, ink recovery and papyrological review.

Many of the advances required to address this problem were developed through the Vesuvius Challenge, an open-science effort launched in 2023 that brought together our research team, academic collaborators and an online distributed community of contributors. That programme established that text could be recovered from unopened or still-rolled Herculaneum material and progressively extended the readable extent of individual rolls[14–17]. The Vesuvius Challenge also made the problem public through a deliberately hybrid model of competition and collaboration. The 2023 Grand Prize established multi-column reading from PHerc. Paris 4[17], while open Kaggle competitions on ink detection[18] and surface detection[19] expanded the methodological base for two bottlenecks: ink recovery and surface extraction. Subsequent Vesuvius Challenge milestones included non-invasive recovery of the title of PHerc. 172[20] and semi-automated digital unwrapping of the lower approximately 70% of that scroll[21].

These advances form the immediate foundation for the present work: not another local detection of possible ink, but a scroll-scale workflow for turning closed or still-rolled Herculaneum material into inspectable scholarly surfaces.

Here we show that such a workflow can recover ancient text at scale, subject to the coverage and validation criteria specified below. By combining optimized high-resolution synchrotron μ CT with moderate phase contrast, advanced large-volume tomographic reconstruction[22], semi-automated virtual unwrapping and deep neural networks, we recover text from three closed or still-rolled Herculaneum rolls in complementary evidential

regimes. We report complete virtual unwrapping of the preserved writing surface of PHerc. 1667 and complete papyrological reading of the text revealed throughout the preserved roll portion catalogued under this inventory number. We also recover title and author-attribution evidence in PHerc. 139, identifying the roll as *Philodemus, On Gods, Book 8*. In PHerc. Paris 4, the optimized scan provides a distinct evidential mode: ink is directly visible in the tomographic volume and can be segmented in three dimensions rather than inferred only from a surface-conditioned model. Together, these results move virtual unwrapping of the Herculaneum scrolls from proof-of-concept demonstrations towards a scalable research framework for recovering the unopened library.

TECHNICAL STRATEGY

The workflow couples optimized tomography, surface reconstruction, surface-conditioned rendering and ink recovery. At the voxel sizes required for intact Herculaneum rolls, image quality is not determined by nominal sampling alone: densely packed carbonized fibers can reduce effective layer separability, while strong phase effects and weak absorption contrast can either obscure or reveal writing depending on sample geometry and ink recipe[14]. The optimized scan regime used here was selected to balance layer separation, ink visibility and manageable scan time; acquisition parameters are given in [Methods](#) and [Extended Data Table 1](#). This optimized scanning approach is the result of an extensive series of imaging tests to optimize the energy range, the resolution level, the scanning geometry, and the level of phase-contrast. Results of these tests are presented in [Extended Data](#).

Virtual unwrapping converts the tomographic volume into an inspectable writing surface. A learned recto-surface segmentation model provides volumetric cues, but the final surface is represented as an explicit quadrilateral mesh. The mesh is extracted, manually corrected where necessary, locally re-optimized and parameterized into a low-distortion near isometric two-dimensional domain. This explicit mesh representation prevents the reading from being tied to a single rendered image. Any point on the flattened surface has corresponding three-dimensional coordinates in the CT volume, allowing reviewers to inspect the underlying papyrus layer, assess local geometric reliability and distinguish surface-placement errors from ink-interpretation errors[2–8, 23].

Ink recovery uses two related but distinct regimes. In most sealed-scroll regions, ink remains too subtle for direct volumetric segmentation and is enhanced on surface-conditioned renderings using models trained from fragment-derived supervision and conservative sealed-scroll pseudo-labels. These models are used as visibility amplifiers for expert inspection, not as autonomous reading systems: they are not trained on character identities, words, transcriptions, OCR targets or lexical labels. PHerc. Paris 4 is treated separately because the optimized scan makes ink-bearing deposits directly visible in the volume, allowing three-dimensional segmentation of ink and projection of that segmentation onto the flattened surface[3, 14]. This volumetric segmentation provides an independent test of whether surface-conditioned ink recovery corresponds to physically visible, surface-associated deposits; the result is presented below as an imaging-validation case.

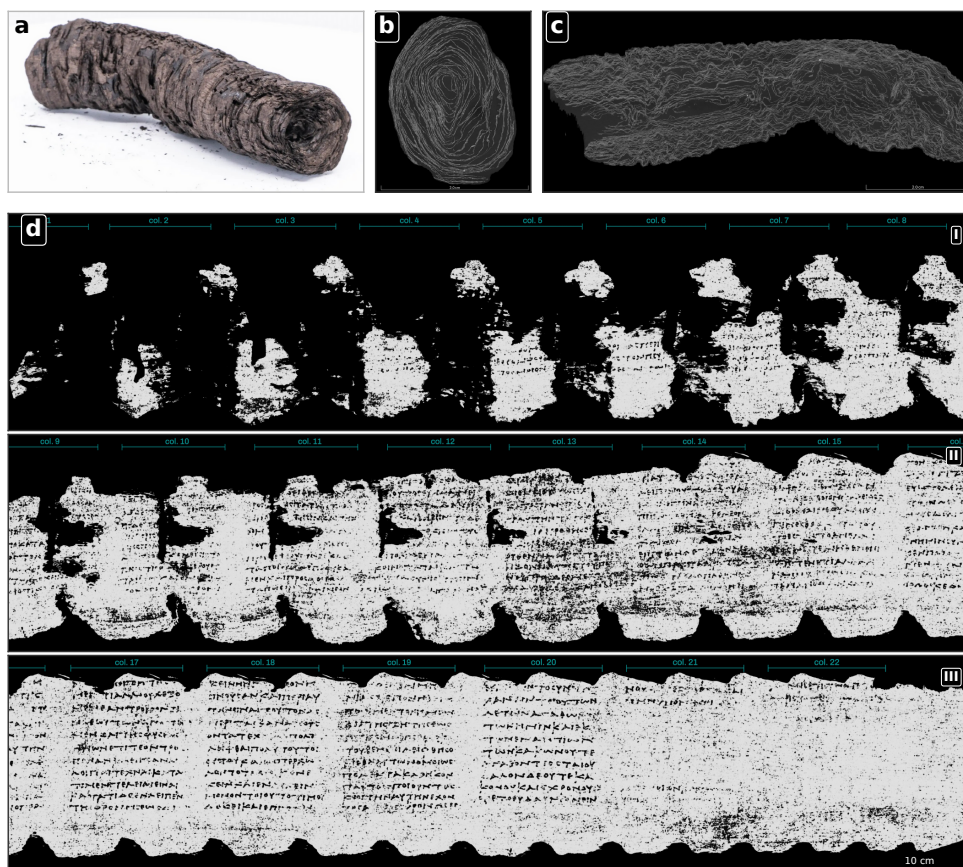


Fig. 1 | Scroll-scale virtual unwrapping of PHerc. 1667. a, Photograph of the sealed Herculaeum roll PHerc. 1667. b, Transverse tomographic section (isotropic voxel size $2.4\ \mu\text{m}$) showing the nested internal wrap structure of the carbonized papyrus. c, Longitudinal tomographic slice showing the layered structure along the scroll axis. d, Surface-conditioned rendering of the reconstructed writing surface after virtual unwrapping and ink enhancement, displayed in three consecutive bands of the flattened parameter domain (I-III). Black background corresponds to regions outside the preserved writing surface or to physically absent or damaged substrate. Within the grey 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.

RESULTS

The Results are organized around three questions. First, can an intact Herculaeum roll be completely unwrapped and read? The presented work on PHerc. 1667 addresses this question. Second, can ink recovery be independently 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 paratextual evidence sufficient to identify a sealed work? PHerc. 139 provides this title-identification case.

COMPLETE UNWRAPPING AND READING OF AN UNOPENED PAPYRUS: PHERC. 1667

PHerc. 1667 provides the principal complete-unwrapping result of this work (Fig. 1). The sealed object is shown in Fig. 1a, while transverse and longitudinal tomographic sections reveal the dense nested architecture of the carbonized papyrus layers (Fig. 1b,c). Fig. 1d shows the corresponding surface-conditioned rendering after virtual unwrapping and ink enhancement: the reconstructed writing surface is displayed in three consecutive bands of the flattened parameter domain, with dark strokes indicating the scholar-facing ink signal used for papyrological inspection. The reconstructed surface contains 22 columns or column-equivalents over approximately 860 cm² of preserved writing surface; black regions mark areas outside the preserved surface or physically absent or damaged substrate.

This is a complete reconstruction and ink-enhanced rendering of the preserved writing surface assigned to PHerc. 1667, rather than a selected local region. We use “complete virtual unwrapping” in a bounded geometric sense: the preserved surface included in the claim is represented by an approved mesh, can be inspected in the CT volume 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 complete reading covers all image-supported preserved text, with untranscribed material restricted to the surviving portions of the first three text columns, amounting to 33 cm², where the fragmentary state of the physical object does not permit a secure interpretation of the visible ink traces. The full papyrological transcription and uncertainty markings are provided in [Methods](#).

VOLUMETRIC VALIDATION OF INK RECOVERY: PHERC. PARIS 4

PHerc. Paris 4 tests whether surface-conditioned ink recovery corresponds to physically visible structures in the CT volume. In the optimized high-resolution BM18 scan of PHerc. Paris 4, thanks to phase retrieval, ink-bearing strokes appear as bright, surface-associated deposits with apparent cross-sectional thicknesses of approximately 10–20 μm (Fig. 2 and [Extended Data Fig. 3](#)).

These deposits are directly visible in the volume, can be segmented in three dimensions and then projected onto the flattened surface, where they coincide with the Vesuvius Challenge Grand Prize Banner region (Fig. 3), the first multi-line reading from a still-rolled Herculaneum papyrus[15–17]. This provides an independent physical validation case showing that, under at least one optimized scan regime and ink preservation state, the surface-conditioned ink signal corresponds to directly visible surface-associated deposits.

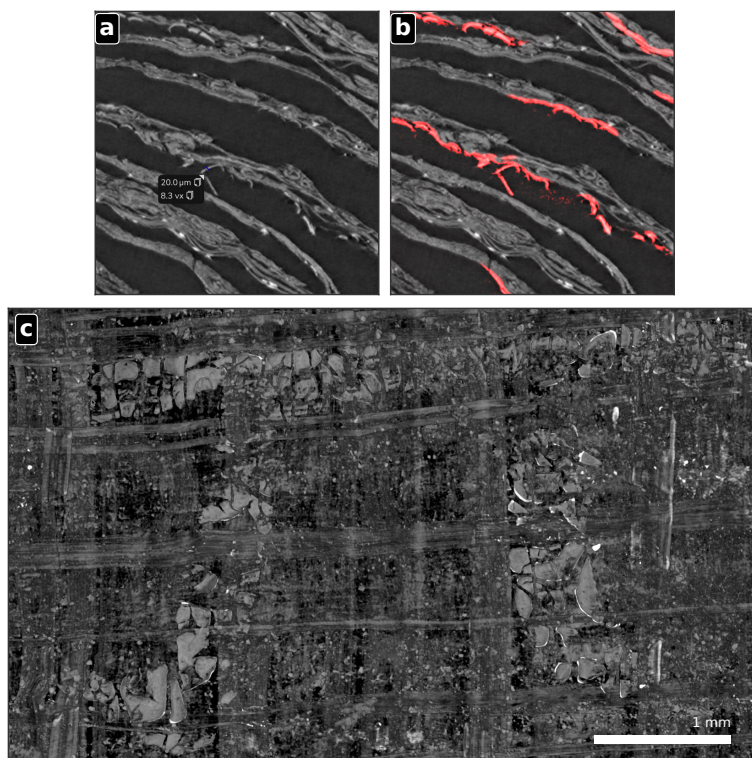


Fig. 2 | Volumetric ink segmentation in PHerc. Paris 4 and projection onto the unwrapped surface. a, Phase-retrieved tomographic cross-section using $2.4\ \mu\text{m}$ isotropic voxel size through a visible ink-bearing surface deposit, with a representative apparent thickness measurement. b, The same cross-section with the segmented ink mask overlaid in red. c, Flattened surface-conditioned rendering of an ink-bearing region before projection of the volumetric segmentation, showing the underlying papyrus texture and writing traces.

IDENTIFICATION OF A SEALED WORK FROM TITLE EVIDENCE: PHerc. 139

PHerc. 139 provides a separate title-identification result. In a title-bearing region of the reconstructed surface, we recover author, work-title and book-number evidence (Fig. 4). The title-bearing sequence is read as

Φι[λοδ]ήμου
περὶ θεῶν Ḥ

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

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

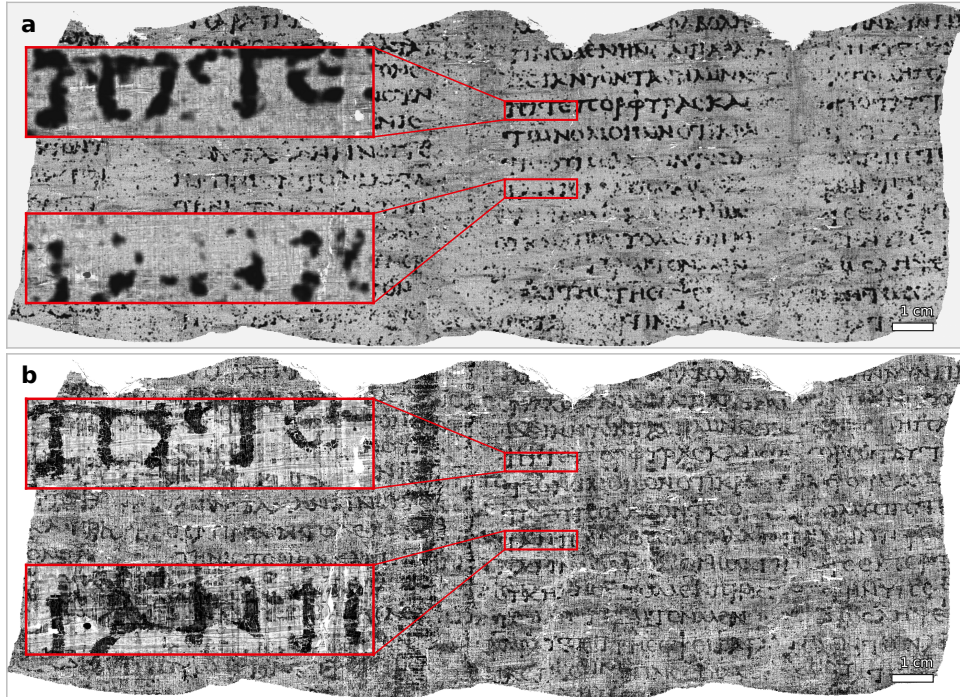


Fig. 3 | PHerC. Paris 4, comparison with 2023 data. a, 2023 Grand Prize winning Ink Detection Model [16, 17] on a virtually unwrapped portion of PHerC. Paris 4 on the 7.91 μm pixel size volume of PHerC. Paris 4 from the Educelab-Scrolls dataset [3] b, 3D ink segmentation enhanced alpha composite rendering against the BM18 2.4 μm data. The second confirms the reading from the 2023 Vesuvius Challenge. Zoom-in lenses highlight the difference in data quality.

roll PHerC. 172 (Φιλοδήμου Περι | καχιῶν \bar{A} , Philodemus, *On vices* book 1)[20]. The title was 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 (I BCE), 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 was known, preserved in PHerC. 26 (Φιλοδήμου | Περι θεῶν | \bar{A} , Philodemus, *On gods* book 1), together with other theological works by the same author known under different titles.

The PHerC. 139 result is important because title-bearing regions can identify sealed books even when continuous prose reading remains incomplete. The workflow therefore recovers not only passages of text, but also library metadata: author, work title and book number.

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 theological concerns: $\chi\omega\rho\iota\sigma$

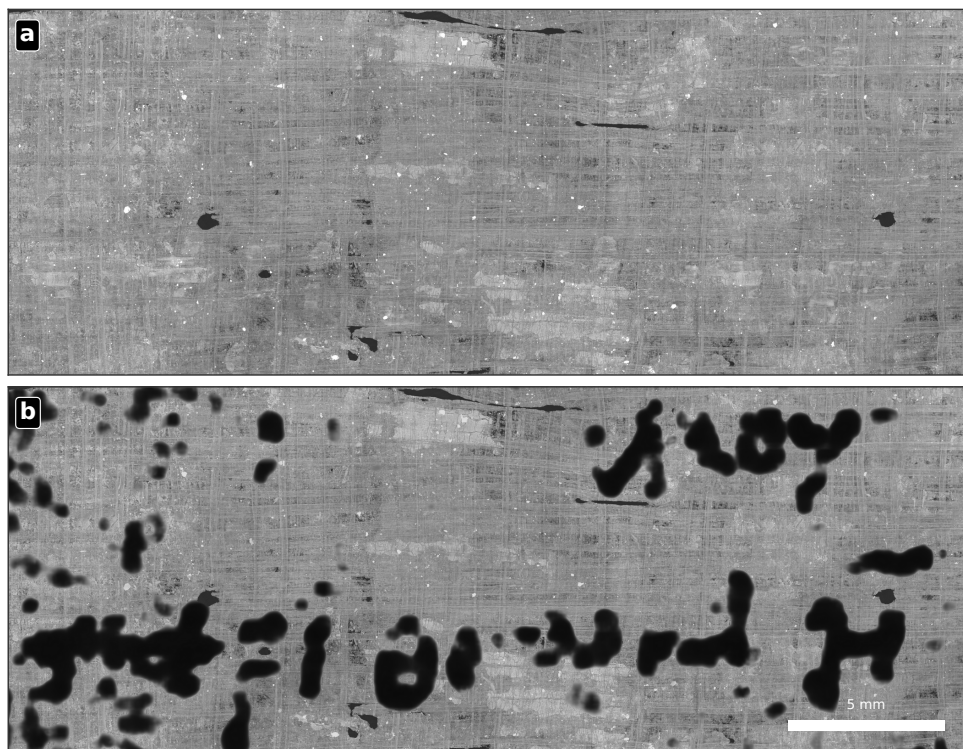


Fig. 4 | Title and author-attribution evidence in PHerc. 139. a, Surface-conditioned rendering of the title-bearing region of PHerc. 139. b, The ink-enhanced signal shown in black against the papyrus texture. The visible title and author-attribution sequence identifies the roll as *Philodemus, On Gods, Book 8*. Scale bar, 5 mm.

προνο[τ]ας (“without providence”), ἀόρατα (“invisible entities”), καὶ τὸ κατὰ φύσιν (“according to nature”), καὶ πάντα χωρὶς πόνων (“and everything free from troubles”), νοερόν (“intellectual”). The images which supported the reading are provided in the corresponding order in the [Supplementary Information](#).

DISCUSSION

These results mark a transition in the virtual recovery of the Herculaneum library. What had previously been demonstrated in local regions, selected banners or partial scroll surfaces is shown here to operate across the preserved writing surface of an intact roll, while also yielding independent volumetric validation and title-level identification in separate samples. They move the programme from localized breakthroughs and partial-scroll demonstrations towards an end-to-end workflow for complete unwrapping and reading. The results from PHerc. 1667 provide the strongest test of that pipeline so far, showing that scroll-scale virtual

unwrapping and complete reading of a still-sealed scroll are achievable. The recovered title-and-attribution evidence from PHerc. 139, after the recovery of the title-and-attribution from PHerc. 172[21], adds a second kind of payoff: the method can identify works and authors, not only recover running prose. The PHerc. Paris 4 scan adds a third, imaging-specific payoff, because directly visible three-dimensional ink segmentation can ground and audit surface-conditioned ink recovery.

This workflow does not imply that all sealed Herculaneum rolls are automatically readable. Performance depends on scroll general preservation state, scan quality, layer separation, deformation, local ink contrast and 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 destabilize surface estimation. The second is radiometric. Ink segmentation 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. Better geometry produces more legible surface-conditioned views, whereas better legibility improves both expert transcription and the supervision available for ink detection. For this reason, progress depends on treating scanning, learned surface and ink models, geometric reconstruction and scholar-facing image formation as a coupled system.

Future gains will depend partly on improved control of effective resolution and on a better account of how morphology, phase effects and absorption together contribute to ink visibility. The present results are best understood as operating in a mixed-signal regime, rather than as reading topography alone[13]. Higher-resolution and better-tuned scans should reduce ambiguity in both layer separation and ink detectability, especially in the most compressed regions. More robust tracing networks and hybrid neural tracers that predict or refine sheet trajectories while enforcing local surface continuity and neighbouring-wrap consistency should reduce the manual burden of correction and neighbour transfer[8, 24]. The latter could be achieved by improving the accuracy of annotations used to train the machine learning models or relying on self-supervised methods. Additional fragment-to-scroll supervision, together with conservative pseudo-label expansion on intact rolls, should strengthen ink models across a wider range of ink recipes without sacrificing trustworthiness.

None of these remaining challenges alters the main conclusion. The key transition marked by the present work is therefore from exceptional local recovery to systematic scroll-scale recovery. The remaining challenge is not whether sealed Herculaneum texts can be read non-invasively, but how broadly, robustly and efficiently the workflow can be extended across the still-unopened library.

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] 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).
- [10] Bukreeva, I. *et al.* Virtual unrolling and deciphering of Herculaneum papyri by X-ray phase-contrast tomography. *Sci. Rep.* **6**, 27227 (2016).
- [11] Brun, E. *et al.* Revealing metallic ink in Herculaneum papyri. *Proc. Natl Acad. Sci. USA* **113**, 3751–3754 (2016).
- [12] Tournié, A. *et al.* Ancient Greek text concealed on the back of unrolled papyrus revealed through shortwave-infrared hyperspectral imaging. *Sci. Adv.* **5**, eaav8936 (2019).
- [13] Angelotti, G., Nicolardi, F., Henderson, P. & Seales, W. B. Ink detection from surface topography of the Herculaneum papyri. *Sci. Rep.* 58467 (2026).
- [14] 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>.
- [15] Nicolardi, F. *et al.* Revealing text from a still-rolled Herculaneum papyrus scroll (PHerc. Paris. 4). *Z. Papyrol. Epigr.* **229**, 1–13 (2024).

- [16] 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).
- [17] Vesuvius Challenge. Vesuvius Challenge 2023 Grand Prize awarded: we can read the first scroll! <https://scrollprize.org/grandprize> (2024). Accessed 19 June 2026.
- [18] Vesuvius Challenge. Vesuvius Challenge – Ink Detection. Kaggle, <https://www.kaggle.com/competitions/vesuvius-challenge-ink-detection> (2023). Accessed 19 June 2026.
- [19] Vesuvius Challenge. Vesuvius Challenge – Surface Detection. Kaggle, <https://www.kaggle.com/competitions/vesuvius-challenge-surface-detection> (2025). Accessed 19 June 2026.
- [20] 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.
- [21] 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.
- [22] Mirone, A. *et al.* Eikonal phase retrieval: unleashing the potential of fourth-generation sources for enhanced propagation-based tomography on biological samples. *J. Synchrotron Radiat.* **32**, 1291–1301 (2025).
- [23] Rabinovich, M., Poranne, R., Panozzo, D. & Sorkine-Hornung, O. Scalable locally injective mappings. *ACM Trans. Graph.* **36**, 16:1–16:16 (2017).
- [24] Quab, 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).
- [25] 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).
- [26] Cernea, D. OpenMVS: open multi-view stereo reconstruction library. GitHub, <https://github.com/cdcseacave/openMVS>. Accessed 19 June 2026.
- [27] 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).
- [28] 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.

- [29] EduceLab. pgs-recon: the EduceLab photogrammetry reconstruction pipeline. GitHub, <https://github.com/educelab/pgs-recon>. Accessed 19 June 2026.
- [30] Paganin, D. M., Gureyev, T. E., Pavlov, K. M., Lewis, R. A. & Kitchen, M. J. Phase retrieval using coherent imaging systems with linear transfer functions. *Opt. Commun.* **234**, 87–105 (2004).
- [31] Angelotti, G. *et al.* Vesuvius Challenge – CT scans of Herculaneum papyri. Vesuvius Challenge, <https://scrollprize.org/data> (2026). Accessed 19 June 2026.
- [32] 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.
- [33] 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.
- [34] Sbordone, F. Recenti tentativi di svolgimento dei papiri ercolanesi. *Cronache Ercolanesi* **1**, 23–39 (1971).
- [35] Capasso, M. & Angeli, A. Papiri aperti col metodo osloense (1983-1989): descrizione e classificazione. *Cronache Ercolanesi* **19**, 265–270 (1989).
- [36] 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).

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 Institut de France. 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.

SAMPLE PREPARATION

Each sample was enclosed in a custom transport and scan case generated from a metric photogrammetric reconstruction of its exterior geometry. Photogrammetry reconstructs a three-dimensional object from overlapping two-dimensional photographs by estimating

camera poses, matching image features and recovering surface geometry through structure-from-motion and multi-view-stereo reconstruction[25, 26]. Photographs were acquired with a Sony α 6300 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[27].

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[28]. 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 detected and triangulated the ArUco markers from the full images and applied their known side length to place the reconstructed scene at physical scale[25, 27, 29].

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[25, 26, 29]. This produced a textured, metric-scale exterior mesh of each sample. The mesh was then passed to the ScrollCase workflow¹, 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.

TOMOGRAPHIC SCANNING

Many tests have been performed on different samples in order to investigate the best combination of parameters that would give the best readability of the scrolls (see Extended Data). The optimized imaging regime described in the main text was selected following these tests in order to balance multiple competing contrast mechanisms. 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

¹<https://github.com/ScrollPrize/villa/tree/main/foundation/scrollcase>

small refractive lenses that locally deviates the beam. The solution to counter-act decoherence effect can be to reduce the propagation distance, and to increase the energy. Both of these actions decrease the effect of phase contrast, the absorption contrast being very low for carbon-based material in hard X-rays. 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 optimal configuration selected for the complete scrolls imaging was 2.4 μm isotropic voxel size, 22 cm propagation distance and 78 keV average incident energy. Finer voxel sizes lead to higher image quality, but scanning full samples is currently impractical.

We used a polychromatic beam (BM18 1.1T side pole from tripole wiggler filtered with 20mm of sapphire and 0.1mm of gold. Using a thin gold filter makes it possible to cut the energies above the K-edge of gold in the spectrum (80 keV), and then to narrow the effective bandwidth.

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. As a result, the effectively detected beam spectrum has an average of 78 keV, with an enhanced detection between 50.4 keV and 80 keV. Considering the low absorption of the samples, this spectrum is narrow enough not to have any detectable beam-hardening effect. It produces data very close to monochromatic beam conditions, while allowing much faster scan, perfectly stable beam and nearly perfect preservation of the native source coherence. Indeed, the only X-ray optics in the beam are mirror polished sapphire and gold filters, as well as permanent polished windows (1.4mm of diamond in the front-end and 1mm of aluminum at the beginning of the experimental hutch).

The optic system was a tandem optic made of two photographic objectives, a Zeiss Otus 100mm 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).

The BM18 helical tomography data were acquired as an up-to-four-position, laterally offset helical acquisition. The exact number of lateral positions as well as the number of projections were chosen according to the diameter of the virtual cylinder containing the specimen.

TOMOGRAPHIC RECONSTRUCTION

The rotation axis was fixed in the sample reference frame at the acquisition level; residual intra-scan deformations of the sample holder were corrected afterwards by the overlap-alignment procedure described below. The different data sets were obtained by translating the whole sample-rotation-axis assembly laterally with respect to the beam/detector coordinate system. Thus, the same fixed sample axis was placed at different lateral positions

in the detector field of view. Each scan therefore sampled a partially overlapping annular region of the specimen. This annular tiling strategy extended the transverse coverage of the reconstructed volume while keeping the specimen and its rotation axis mutually fixed throughout the nominal acquisition geometry.

DETECTOR PRE-PROCESSING

Each projection was corrected for dark current and normalized by the corresponding flat-field signal. Flat-field images acquired during the scan were averaged before normalization. Spurious hot pixels were corrected using a local median criterion: pixels exceeding the median of their 3×3 neighbourhood by more than 0.1 were replaced by that median value.

DETECTOR DISTORTION CORRECTION

A pre-calibrated two-dimensional detector-distortion correction was applied to all radiographs before helical gridding and reconstruction. The correction map was obtained independently from a slit-grid calibration. Detected slit-crossing coordinates were fitted with bivariate polynomial models to produce dense source-coordinate maps,

$$\text{coords_source}_x(z, x), \quad \text{coords_source}_z(z, x),$$

which define, for each pixel in the rectified detector plane, the corresponding sub-pixel position in the raw detector image. Corrected radiographs were therefore obtained by backward bilinear resampling,

$$I_{\text{corrected}}(z, x) = I_{\text{raw}}(\text{coords_source}_z(z, x), \text{coords_source}_x(z, x)).$$

The same coordinate transformation was applied to the associated weight maps, ensuring that radiographic intensities and blending weights were expressed in the same rectified detector coordinate system.

PHASE RETRIEVAL AND CONTRAST RESTORATION

After flat-field correction and detector-distortion correction, single-distance Paganin phase retrieval[30] was applied with $\delta/\beta = 1000$.

The retrieved images were sharpened using an unsharp-mask operation, in order to compensate for the blurring induced by the Paganin phase retrieval,

$$I_{\text{unsharp}} = (1 + c) I - c G_{\sigma}(I),$$

with $c = 4.0$ and $\sigma = 1.2$ pixels.

The images were then clipped to the interval $[10^{-6}, 10]$ before applying the logarithmic transform.

MULTI-SCAN WEIGHTING AND RADIOGRAPH FUSION

The laterally offset radiographic data sets were combined at the projection/sinogram level before volume reconstruction. Smooth error-function apodisation windows were used to

provide gradual transitions at the detector borders and in the regions where neighbouring annular acquisitions overlapped.

CENTRE-OF-ROTATION CALIBRATION

The centre of rotation was calibrated before the final helical reconstruction. This calibration did not only provide a single scalar value, but a z-dependent estimate of the apparent rotation-axis position. This correction accounts for slow variations of the effective centre of rotation along the vertical direction, due for example to a possible mechanical misalignment between the z-translation axis and the rotation axis. The z-dependent centre-of-rotation correction was used during helical reconstruction.

This z-dependent centre-of-rotation correction is distinct from the overlap-based motion correction described below. The latter estimates residual relative shifts between neighbouring laterally offset scans from their overlap regions, as a function of projection angle, and applies the fitted corrections to the framewise horizontal and vertical translations stored in the ETF geometry.

ANGULAR-DEPENDENT OVERLAP ALIGNMENT

In addition to the centre-of-rotation calibration, residual relative motion between the laterally offset helical scans may occur because of slight deformations of the sample holder, consisting of a nylon cage and a Teflon membrane envelope, during the scan. These movements of few dozens of micrometers could be due on one hand, to the effect of the dose accumulation and on the other hand, to the change of humidity level between the sample storage room (whose humidity is controlled at 50% +/-5%) and in the experimental hutch, where only the maximum hygrometry is controlled at 55% (air can be dryer than in the storage room). These movements were corrected using an overlap-correlation procedure. This correction was applied before the final concatenated helical reconstruction and acts directly on the framewise scan geometry.

Overlap-correlation maps were first computed for each adjacent pair of scans with an angular sampling step of 5 degrees. For each adjacent scan pair and each sampled angle, narrow overlapping detector stripes were extracted from the two scans at the corresponding helical height. The two overlap images were first brought to their nominal relative horizontal position using the motor-derived translations. The signals were then weighted by their validity maps and band-pass conditioned. This conditioning consisted of a weighted low-pass normalization followed by subtraction of a weighted broader Gaussian baseline, enhancing structural features while suppressing illumination differences and slowly varying background contributions.

For each sampled angle, the relative displacement between the two overlap images was estimated by minimizing a weighted absolute-difference loss over a two-dimensional grid of offsets. The raw per-angle displacement minima were then fitted independently for each adjacent scan pair. The fitted corrections were applied cumulatively along the scan chain.

These corrections compensate for sample motion, residual mechanical errors and geometrical inconsistencies between the laterally offset acquisitions before the radiographs are merged into the final helical reconstruction.

RECONSTRUCTION

The final volume was reconstructed using a GPU-accelerated generalized hierarchical back-projection algorithm adapted to the helical acquisition geometry. Reconstruction was performed on the concatenated and overlap-corrected helical data set.

The initial 32 bits reconstructed stacks were converted first in 16 bits tif stack using saturation levels in black and white at 0.01%. These data were transferred out of the ESRF for further processing and analysis as described below. As the datasets are very large (from dozens of Tb up to more than 100 Tb for the largest one), binning versions by factor 2 and 4 were calculated, as well as compressed 16 bits version in jpeg2000 format with a lossy compression of 10. This last conversion aimed principally for long term storage of the reconstructed data at the ESRF, as well as to make them possible to download through the new ESRF online database for art, history and archaeology (<https://cultural-heritage.esrf.fr/tomo>).

Finally, reconstructed volumes are converted to uint8 OME-Zarr[31] with 6 multi-scales. Every voxel out of the samples has been masked to zero, and hence full zeroes chunks are not stored at all to save space. The average size of such a reconstructed volume is 20 TB.

RECTO-SURFACE ANNOTATIONS AND PREDICTION

Recto-surface prediction was formulated as supervised three-dimensional semantic segmentation of the papyrus writing surface in the reconstructed volume. The target surface was the recto side of the sheet, defined operationally as the side facing the scroll umbilicus 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 labelling 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[32] three-dimensional U-Net on z-score-normalized CT patches. The network produced a two-channel background/surface prediction. The training objective used Medial Surface Recall, a custom adaptation of Skeleton Recall Loss[33] 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 the Supplementary Information.

The resulting voxelwise prediction was used only as an intermediate cue for geometric tracing. The final surface representation was not the segmentation volume itself, because dense predictions can contain local sheet mergers, gaps and false positives in tightly packed regions; instead, the prediction was converted into orientation evidence for the quad-mesh reconstruction described below.

SURFACE ESTIMATION AS A QUAD MESH

The reconstructed sheet surface is assembled from locally, smaller fitted patches, and stored as a tifxyz quadrilateral mesh. In this custom format, the surface is represented by a regular

two-dimensional parameter grid, and each valid grid vertex stores its three-dimensional position in the tomographic volume. 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 centreline 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.

A surface was initialized from a seed point or from an already reconstructed neighbouring patch. Starting from this initial mesh, the software iteratively proposed new vertices along the boundary of the two-dimensional mesh grid and optimized their three-dimensional 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. The optimization failed in regions where the topology of the predictions did not match that of the intended papyrus surface, for example in regions with artifacts such as mergers or holes. These failure cases usually corresponded to regions where the papyrus was densely packed or exhibited anomalous curvature.

Initial unwrappings for new scans were created with a surface-based tracing tool that used a large dataset of smaller meshes generated by the seed- and prediction-based mesh-growing procedure described above. Seed locations and seed density were chosen to ensure substantial overlap between base meshes. The tool then generated a new mesh from a seed quad by iteratively proposing and refining vertices, optimizing each vertex toward the consensus surface defined by 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, some holes and deviations remained where the binary surface predictions contained errors. 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 neighbourhoods along the local mesh normal. Surface placement was assessed by inspecting mesh intersections overlaid on orthogonal CT slices, interactively rotating the three-dimensional view, and visualizing a CT-textured two-dimensional grid in real time. Regions judged geometrically consistent with a single sheet were marked with an approval mask. After edits, local re-optimization 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

The corrected tifxyz surface was used as the input for final flattening and rendering. 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[23], 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 two-dimensional parameter domain. 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 three-dimensional surface position was computed by barycentric interpolation within the containing triangle. In the workflow, UV-domain rasterization resamples points in parametrization space, allowing the conversion from the flattened OBJ to a tifxyz surface.

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

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

with $N = 65$ and $\Delta = 1$ voxel. This gives offsets from -32 to +32 voxels; at $2.4 \mu\text{m}$ isotropic voxel sampling, the distance between the first and last samples is $64 \times 2.4 = 153.6 \mu\text{m}$. Tomographic intensities were sampled at each off-surface position by trilinear interpolation from the eight neighbouring voxels. Samples outside the valid volume bounds were assigned zero.

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, 14].

These fragments serve as a digital ground truth for the ink signal. 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.

After registration, the resulting labels were mapped into the same surface-aligned space 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

We start by training models on the fragment-aligned infrared images, on a dataset composed of PHerc. 9B, PHerc. 343P and PHerc. 500P2. These 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. We therefore train a U-Net architecture with a 3D ResNet encoder and a 2D decoder: the encoder features are pooled along the depth dimension before decoding, which sidesteps the need for depth-wise localization and lets the model predict ink directly in the 2D plane in which the infrared labels are defined.

The ink detector operates on local surface-aligned patches and was not trained on character identities, words, transcriptions, OCR targets or lexical labels. This constrains the task to local CT texture and morphology and removes any risk that predictions would be driven by linguistic or word-level priors. We choose an input size of a 256 pixel window, which roughly corresponds to 614 μm (Extended Data Fig. 6). This input size is smaller than all letters detected to push the model to learn the underlying signal and restricts the model’s context about the underlying text. Since the training and inference predictions have no shared context, each forward pass through the model can only rely on local information, making full-letterform hallucinations from linguistic or word-level priors practically impossible.

This approach proves successful in detecting initial ink traces on both PHerc. 139 and PHerc. 1667. Without being trained on these scrolls, the model can already detect ink traces, demonstrating generalization capabilities. The fragments model becomes the pretrained backbone for further pseudo-labels training rounds.

No OCR or language model was used at inference time or in the generation of pseudo-labels. Instead, we rely solely on the model’s own signal, fine-tuning the pretrained backbone to extend its representation to the new scroll we want to uncover. The pseudo-labeling rounds aim to generalize the model’s representation of the underlying ink signal to a specific scroll through a simple bootstrapping loop. Starting from an initial set of labeled tiles, we fine-tune the model and run inference across the scroll. Regions in which the fine-tuned model uncovers new text are cropped into additional training tiles, with pseudo-labels derived from the model’s own predictions. Each round thus grows the dataset, and fine-tuning on this enlarged dataset yields a model that increases the consistency and visibility of candidate ink signal later reviewed by papyrologists. When successful, each round reveals more text than was previously detected, and these newly legible regions seed the next iteration. Extended Data Fig. 7 shows an ablation of this process: a single column of text from PHerc. 1667 is used for iterative data labeling, while a separate unrolled column of possible text is held out as validation and never seen by the model. As the iterations progress, text near the top of the validation column is gradually revealed, with performance saturating around iteration 5. To fully read PHerc. 1667, we applied the same process at scale, using three columns for iterative labeling and performing five iterations. We trained with a

mixture of Dice loss and binary cross-entropy loss with strong label smoothing to prevent overfitting to the signal and to reduce the impact of label noise. Further details on training, together with the model architecture, are provided in the Supplementary Information.

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

VOLUMETRIC INK SEGMENTATION IN PHERC. PARIS 4

In PHerc. Paris 4, the 2.4 μm -voxel reconstruction preserved enough ink–substrate contrast to support direct visual identification of ink. We trained a 3D residual U-Net to predict per-voxel ink probability from 256³-voxel patches. A custom 3D DINOv2[24] representation was used to derive ink labels. Architectures and training procedures are described in the Supplementary Information.

Dense supervision was obtained primarily from the 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 took the intersection of two independent signals: the DINO ink-likeness map and the output of a separately trained volumetric ink detector. This detector was trained from surface-conditioned 2D annotations projected into 3D and was therefore useful mainly as a coarse support constraint: it reduced spurious DINO responses but did not drive localization. In practice, the two maps were multiplied and thresholded, with low-intensity air and void voxels masked out so that labels could only occur nearby papyrus material.

A student detector was trained on these intersection labels, then the procedure was repeated using the guided student as the coarse support model. This iteratively cleaned 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 sharpened boundaries and suppressed some residual artifacts.

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 19-24cm of a full roll) and diameter (2 cm out of 4-6 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 made in November 1969 by the conservator A. Fackelmann, but it did not reveal any visible traces of writing and was discontinued [34], producing only a few small and largely unreadable portions of text while substantially reducing the dimensions of the surviving portion of the roll. 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 [35]. 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 artefact, 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 palaeographical 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 2nd century BC, while not excluding the very late 3rd century BC.

While a comprehensive commentary lies beyond the scope of the present paper, we present here a transcription of the Greek text together with an English translation.

Coll. 1-4 traces

Col. 5 lost margin

ll. 1-2 *traces*

]. τομ[
]τῆθ̄εσθαι ψ̄εϋ. [
 5]ς τὸ ὅμοιον. . . [

3 sq. τὸ μὴ [συγχατα]|τῆθ̄εσθαι ψ̄εϋθ̄[ῶς?

« ... the similar ... »

Col. 6 lost margin

ll. 1-5 *traces*

ὁρμὰς περὶ[.]
 ας τομ̄η̄σϋ. [.]
 θεσθαι τῶι [.]

7 sq. τὸ μὴ συγ[χατατι]|θεσθαι

« ... impulses ... »

Col. 7 *lost margin*

ll. 1-6 *traces*

..ξ...εψ[.....]
 ρει. προσγ.[.....]
το φουρ[.....]
 10 την άπροπ.[.....]

10 fort. άπροπτ[ωσί]αγ

Col. 8 *lost margin*

ll. 1-7 *traces*

χα. η...[.....]
 τιν δε...[.....]
 10 γος παρα...[.....].
 άλλα. π. υ[.....]

Col. 9 *lost margin*

[.....]...[...]
 [.....], ει εν[...]
[.....] καθ' ὅσον
 5 ῥεξα. [.....] τοσ ἔς-
 τιν...[.....]
 ὑποκατα. [.....]
 τοσ τον[.....]
 τον άνθρ[ωπ.]ειν
 αὐτην ἢ εχέιν. [..][...]
 10 ναι ὅτι δὲ χα.....οις

« ... so far as ... this or to have ... that ... »

Col. 10 *lost margin*

[.....]...[...]
 [.....]...[...]
 τῆ...[...], τω.....
 5 τῶ ετ...[.....]
 ευνκριε.[.....]
 νησπο...[.....]τά-
 ναι προσῆκον ἐπὶ τὸ
 πᾶν ἔτι. ο. [.....]ηρου ἔς-
 ται ῥέρος καὶ ἐν...ωι
 10 τὸ μέγα καὶ μακρὸν

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

4 fort διπλῆ conspic potest 8]ηρου or]πρου fort. τῶ [πον]ηροῦ? 9 ῥέρος dubitanter ἐν πρόγῳι?

Col. 11 *lost margin*

[.] . αἰ . [.] . .
 [.] . εἰ θυ . . . α . α .
 τῶν [. . .] . ἐκτάτων τ . .
 . . . [. . .] . τήν και τήν
 5 ὀρμ[ήν] . ατα
 τούτ[.] . ἀρίστη .
 πρὸς [δ] ἐ τὰ οὕτως ἕκασ-
 τα πο . . . φε πεφύκα-
 10 μεν και πρὸς τὸ ταῦ-
 τα ποιεῖν ἃ φαίνεῖν

6 μαχαρίστη? 7 fort. τὸ 8 πονηρῶς? ποιεῖν ὅς? 9 sq. ταῦ|τα vel ταῦ|τά

« ... and the impulse ... For/towards each of these things in this way ... we are by nature ... and for doing these things that ... seem ... »

Col. 12 *lost margin*

[.] ε . τῶ [. . .]
 [. . .] τῶ . . . ε . . . τῶ . . . εἰ
 [. . .] γότατα δ' ἀνθρώ-
 5 [πο]ῖς και θηρεῖν, τάδε
 [.] . κνα
 [.] . (.) αιτασ
 [. (.)] πολὺ δὲ μάλιτα τὰ
 κοινότατα ἕκαστα
 ταῦτα συνίστην
 10 . ηφ αι γὰρ ἀναγ-

2 τοῦ ποιεῖν? 3 fort. [και]γότατα 10 fort. ἴμ' ἠφ

« ... to men and beasts, these things ... And above all, each of the most common things constitutes these ... For, [necessity? necessary?] ... »

Col. 13 *lost margin*

φυσικὴν ὑπ[.]
 διὸ και οὐ [. . .] . . .
 κατα τήν [.] . . .
 ταύτην τα . . . [. . .]
 5 [. . .] . α εὐρεθήσετ[αι]
 και οἱ βίοι προάξουσ[ιν]
 οὐθέν, ἡμῶν οὔτε τῆς
 ἡδονῆς οὔτε τοῦ πό-
 νου προσδεομένων.
 10 ὡσαύτως δὲ και τ . (.)

« ... 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 ... »

Col. 14 *lost margin*

[.] . τὰτα
 ταυτ[.]ε . . ς και ἐλ-
 λείπον οὕτως βούλο-
 μαί λέγειν α
 5 π . . το . πα . . τέλεον
 ἐλλείπειν [.]αριςτ . . .
 και ὡσαυτ[. .]ελξῖν
 ἐπί τῶν δεξιῶν μέ-
 ρων πρὸς τὰ ἀριστερά.
 10 γίνεται δ' ὁ πλεονασ-
 μὸς κατὰ τὴν ὀρμὴν

6 τ[ἀ] ἀριστερά? 7 ὡσαύτ[ως τ]ελξῖν?

« ... and thus lacking, I want to say ... perfect ... to lack ... and ... on the right towards the left. There is an excess in the impulse ... »

Col. 15 *lost margin*

[. .]λογ κα[. .]
 ὁμοῖος και τῶν
 παραπληρείων πάν-
 των. καθ' οἶον γὰρ εἶδος
 5 αἱ ὀρμαὶ πεφύκασι
 γίνεσθαι, κατὰ τοῦ-
 τό ἐστιν τὸ ἀνελλι-
 πές ὥστε μηθὲν ἐπι-
 ζητεῖν ἔτι, ἀλλὰ ἀ-
 10 ναπεπληρῶσθαι κα-
 τὰ πᾶν ὡς ἂν κατα

1 sq. fort. καθά[περ] | μεθόμοιος

« ... and of all similar things. For, according to this kind, according to which impulses exist by nature, there will be that which lacks nothing, so that one seeks nothing more, but completes in every respect as ... »

Col. 16 *lost margin*

γονε), δι[ἀ τὴν φρό]νη-
 ριν ἐγγίζουσι τῆς
 συντελείας. ἀ[π]ὸ τοῦ-
 των δὲ μεταβαί-
 5 γουσα ἐπὶ τογον
 τὰ μὲν καθ' αὐτήν
 ἐν ἡμῖν συντετέλε-
 κεν πάντα οὐ δυνα-
 μένη συμπληρώ{ε}-
 10 ραι τὴν φύσιν, πα-
 ρέδωκε δὲ ὑπὸ πρα-

5 τὸν λόγον? τὸ μεμικρόν? τὸ νεῦρον? τὸν πρῶτον? 11 ὑπὸ
 πρα||[γμάτων]

« ... by means of practical wisdom they approach perfection. Moving from these things to the ... , it accomplishes within us all that pertains to it, even though it cannot fully complete nature, and it provided ... »

Col. 17 *lost margin*

[. . .] οἱ ζητήσο-
 μέν τι, ἀλλ' οὐχ ἔξο-
 μεν ἕαν τρόπον τι-
 νὰ ἑαυτῶν ἀποστῶ-
 5 μέν καὶ τῆς φύσεως
 ἡμῶν, ἔτι τε ὃν τρό-
 πον ἂν ῥηθείησαν αἰ-
 λοιπαὶ τέχναι κατὰ
 τι μὲν τέλειαι εἶναι,
 10 κατὰ τι δὲ ἐνλείπειν,
 τῆς φρονήσεως 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 besides, in the same way as the remaining arts may be said to be perfect in one respect, but to be deficient in another respect, [being] practical wisdom ... »

Col. 18 *lost margin*

περὶ ἐ-]
 κείνην τῆ[ν φ]ρόνη-
 σιν οὔσαν, καὶ περὶ αὐ-
 τὴν εἶναι. τοῦτον δὲ
 5 περὶ τὰς βαναύσους
 ὄντα τέχνας ἀπολε-
 λεῖσθαι πολὺ τοῦ τοι-
 ούτου καὶ ὡς περὶ χω-
 λὸν τὸ τεχνικὸν ἔ-
 10 χεῖν καὶ ἐνλείπει-
 μένον τοιοῦτό τι μοι
 δοκεῖ, καὶ ἐπὶ τῶν

« ... being practical wisdom about that, is also about this. This [person? aspect?] that concerns the mechanical arts seems to me to be very distant from such a [conception?], and to have the technical fulfilment that is, so to speak, lame and something of such type lacking, and in regard to the ... »

Col. 19 *lost margin*

[. . .] . . . [.]
 . [.] . δεῖσθαι μη[δ]ε-
 μίας. ἐκπονηθέν-
 των μέγτοι ἡμῶν
 5 διὰ τῆς ζητήσεως
 ἢ τῆς μαθήσεως, κα-
 τ' οὐθὲν ἔτι λειφθηρό-
 μεθα αὐτῶν ὁμοίως
 τέλεα τὰ καθήκον-
 10 τ' αὐτοῖς ποιῶντες
 καὶ τὴν αὐτὴν ἔχον-
 τες αὐτοῖς φρόνησιν

« ... need none. Having certainly been instructed 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 practical wisdom as they do ... »

7 sq. ἔτι λειφθηρό|μεθα vel ἐπιλειφθηρό|μεθα

Col. 20 -----

. [. . .]
 ευ ν τὸ συντυγ-
 χάνειν. τοιοῦτων
 5 δὲ τῶν ἀγαθῶν ὄν-
 των ἡμῖν, καὶ ἐκ
 τῶν ἐναντίων
 τῶν κακῶν οὔτε
 ἀγαθόν τι ἔσται – οὐ
 10 καλὸν δέ – οὔτε κα-
 κὸν – οὐκ αἰσχρὸν – οὔ-
 τε τὸ εὐδαιμονεῖν

2 fort. εὐλογον 11 ex. gr. τὸ εὐδαιμονεῖν || [ἐπιταθῆσεται ἢ ἀνεθήσεται]

« ... to happen. And such being the goods for us, even from the things that are opposed to evils there will be neither anything good—let alone beautiful—nor anything bad—let alone ugly—, nor happiness [will be increased or diminished] ... »

Col. 21 *lost margin*

. [. . .] . . . [. φ] ρο-
 νοῦντες μέγα καὶ
 11. 4-5 *traces*
 γοι . . . αἰτε ἀπὸ
 τοῦ οὔτε . . . ἐπαι-
 νεῖν τ . ν
 l. 9 tantum vestigia
 10 μένου οἶον κατὰ τὰς
 ὑμνήσεις ευτῶνι . .

« ... showing off and being revered and ... to praise ... as according to the eulogies ... »

Col. 22 *lost margin*

γων εἰς ἐτι συνα . . . [. .]
 . (.) Ἀριστοκρέων τὰς [.]
 . . . (.) τοῖς ἐχομένοις . .

1 fort. συναγ vel συναπ

« ... still ... Aristocreon ... to possessed/following things ... »

VALIDATION AND PAPHYROLOGICAL REVIEW

Machine output was not treated as a substitute for reading. The PHerc. 1667 transcription was produced and interpreted at column level by eight 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.

STATISTICS AND REPRODUCIBILITY

The complete virtual unwrapping of PHerc. 1667 amounted to 31 wraps and 1231 cm² of papyrus surface. The unwrapping was completed using the most efficient semi-automated segmentation tooling available at the time, a wrap by wrap copy tool combined with ~25 hours per wrap of manual annotation.

USE OF LARGE LANGUAGE MODELS

Large language model assistance was used during drafting to support language editing, structural revision and consolidation of cited background literature. All scientific claims, interpretations, methodological descriptions and final wording were reviewed and approved by the human authors. No large language model is listed as an author.

ACKNOWLEDGEMENTS

We thank the curators and conservators responsible for the PHerc. material at the National Library of Naples “Vittorio Emanuele III” for access and guidance, in particular Silvia Scipioni, Giovanni Bova and Valentina De Martino.

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 a first version 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 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, and led the annotation efforts since February 2025.

- 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 and contributed to the 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 and contributed to the photogrammetric acquisition and multispectral and infrared imaging.
- C.V. contributed to papyrological review with particular focus on PHerc. 1667 coll. 1-12, manuscript annotation and contributed to the 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 and contributed to the photogrammetric acquisition and multispectral and infrared imaging.
- M.D.A. contributed to papyrological review with particular focus on PHerc. 139, manuscript annotation and contributed to the 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. 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. 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 or W. Brent Seales.

DATA AVAILABILITY

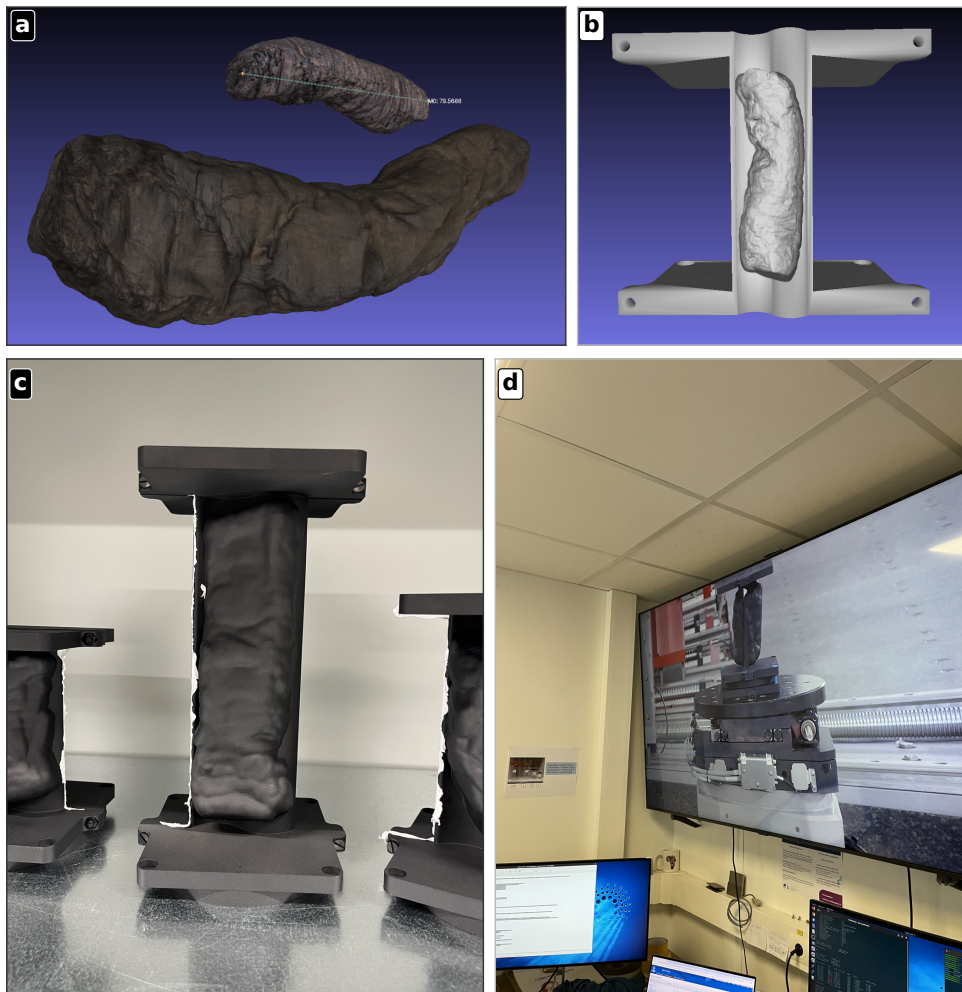
Tomographic volumes, reconstructed surfaces, flattened renderings, figure source data and transcription artifacts are available at https://scrollprize.org/data_browser. Data is shared under the Creative Commons 4.0 - BY NC License unless otherwise specified. Jp2 stacks versions of the tomographic volumes are also available on the ESRF open access database for art, history and archaeology <https://cultural-heritage.esrf.fr/tomo> 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 1–4; model weights will be made 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

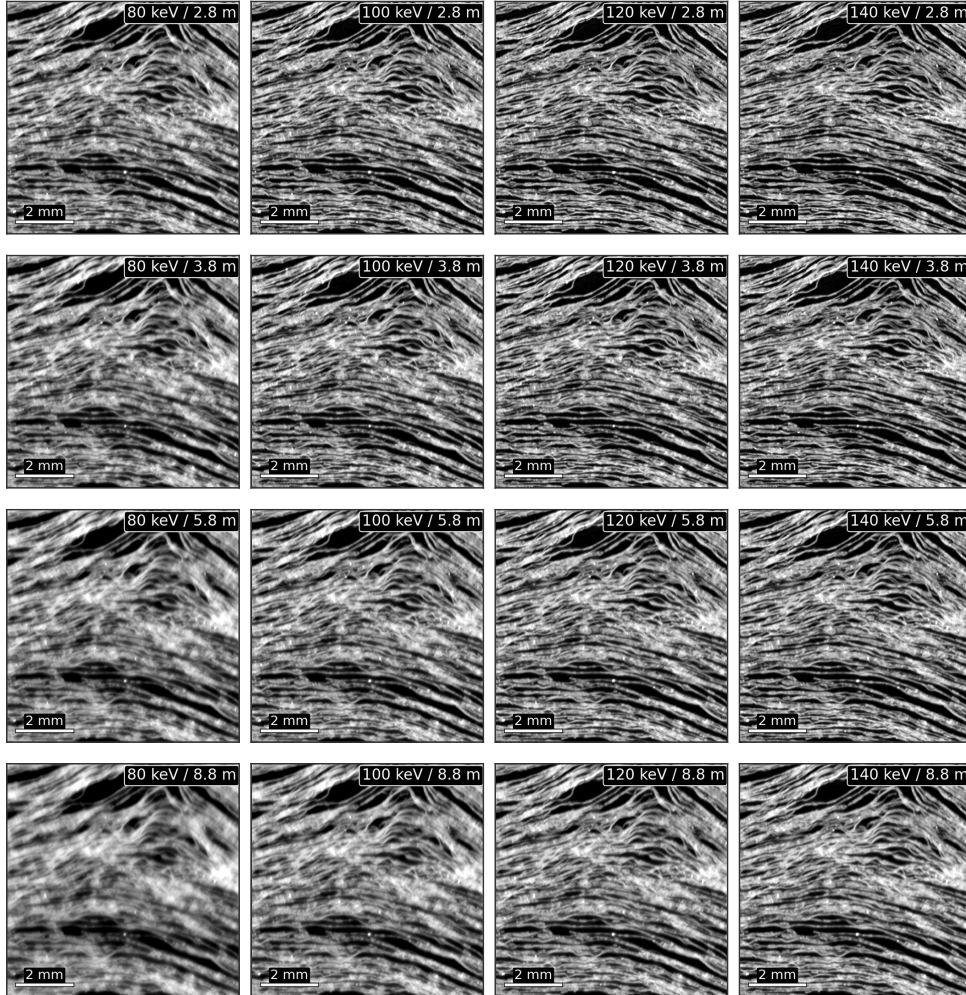


Extended Data Fig. 1 | Sample preparation and scan cases. a, Photogrammetric meshes of PHerc. 1667 (top) and PHerc. Paris 4 (bottom), reconstructed from multi-view photographs; the M0 label gives the maximum chord length of PHerc. 1667 in millimetres. b, Automatically generated half-case model for PHerc. 1667 (non-textured render), with the photogrammetric scroll silhouette visible inside the cavity; the case is fitted to the outer morphology to minimise 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.

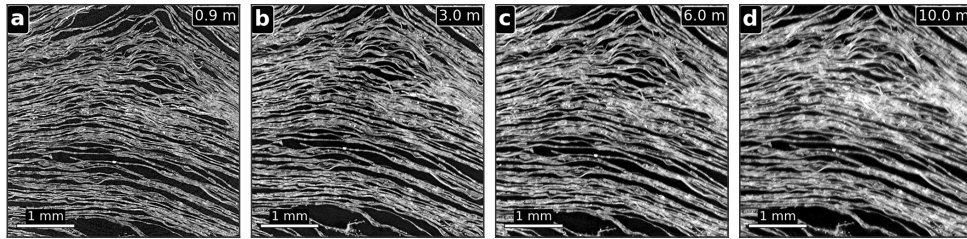
| Sample ID | Role in manuscript | Beamline | Avg. energy (keV) | Voxel / pixel size (μm) | Sample-to-detector distance (m) | Phase retrieval δ/β | Unsharp coeff. | Unsharp σ |
|-----------------------|---|----------|-------------------|--------------------------------------|---------------------------------|--------------------------------|----------------|------------------|
| 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-optimization reference | BM18 | 77 | 2.403 | 0.22 | 1000 | 4 | 1.2 |
| PHerc. 500P2 fragment | Fragment-derived ink supervision | BM18 | 110 | 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 | 1000 | 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 | 78 | 2.399 | 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 *Herculeum papyrus*. Representative reconstructions showing the effects of pixel size, propagation distance, incident energy and phase retrieval on layer separability and haze-like fiber effects.

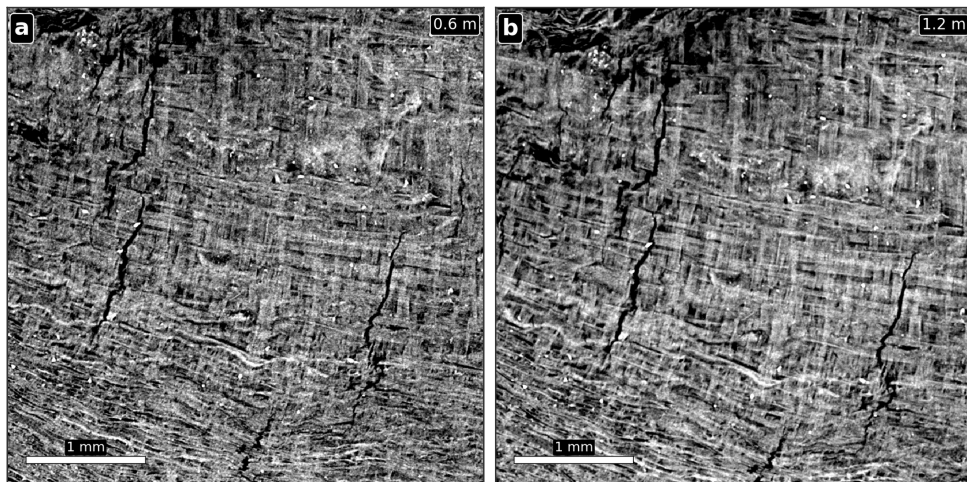


a. Joint propagation-distance \times incident-energy sweep at fixed pixel size. Sixteen axial reconstructions of the same PHerc. 175A slice in a 4×4 matrix: rows (top \rightarrow bottom) are sample-to-detector propagation distance 2.8 / 3.8 / 5.8 / 8.8 m; columns (left \rightarrow right) are mean incident 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 \times 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 $8\ \mu\text{m}$ the empiric sweet spot for the energy is between 100 and 120 keV with propagation distances of about 3m.

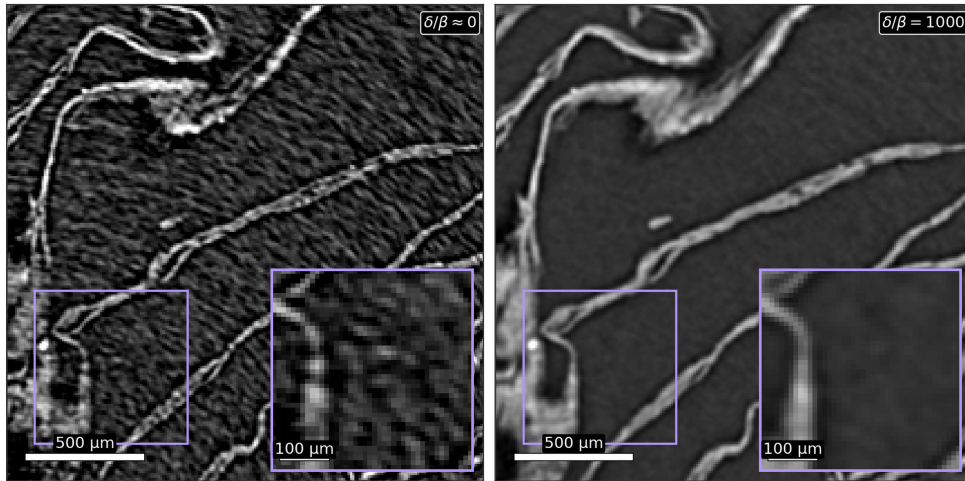


b, Sample-to-detector propagation distance effect on retrieved phase-contrast at fixed pixel size and energy. Four reconstructions of the same PHerC. 175A slice at fixed mean energy (133 keV), fixed Paganin filter ($\delta/\beta = 1000$) and effective pixel size $4.327\ \mu\text{m} \rightarrow 4.116\ \mu\text{m}$ as detector magnification scales with distance: 0.9 m; 3.0 m; 6.0 m; 10.0 m. Physical field of view $4\ \text{mm} \times 4\ \text{mm}$. Image quality increases with decreasing propagation distance due to the progressively increasing decoherence effect with increasing distance.

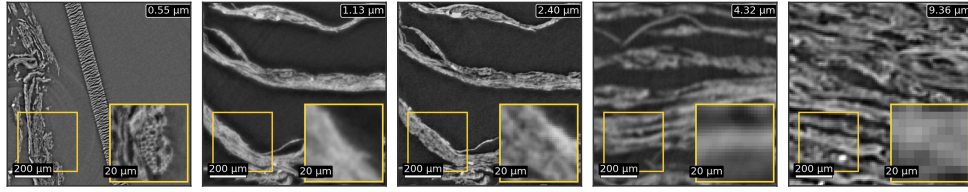


c, Sample-to-detector propagation distance at fixed pixel size and energy. Two reconstructions of slices from PHerC. 268 slices at fixed mean energy (133 keV), fixed Paganin filter ($\delta/\beta = 1000$), and effectively matched pixel size ($4.334\ \mu\text{m}$ at 0.6 m, $4.320\ \mu\text{m}$ at 1.2 m): 0.6 m; 1.2 m. Physical field of view $4\ \text{mm} \times 4\ \text{mm}$; 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 looks 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

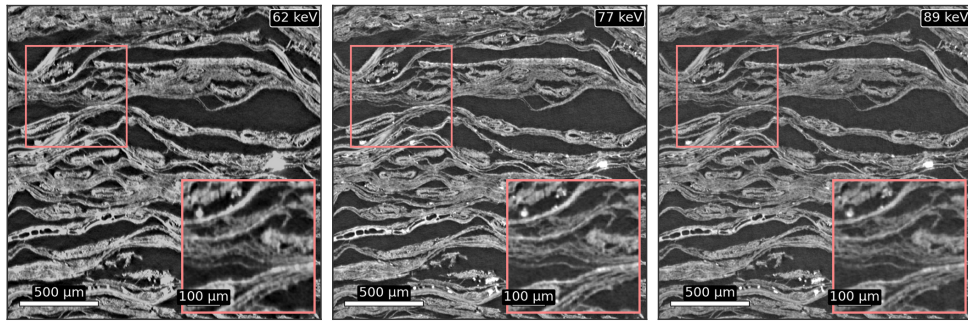
low as the samples are mostly composed of carbon that absorbs faintly in the hard X-ray regime. The picture on the right gives an overall better visibility of the cracks and of some other structures, but at the expense of the sharpness in many places. Once again, this is due to the decoherence effect that starts to affect the resolution at 1.2m. 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. As a result, the optimal propagation distance has been selected as the one that does not produce decoherence effect on the largest scrolls scanned so far, but that still gives good contrast, i.e. 220mm for the 2.4 μm pixel size setup. This distance can be longer when larger pixel sizes are used (about 0.5m at 4.6 μm and 2m 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.



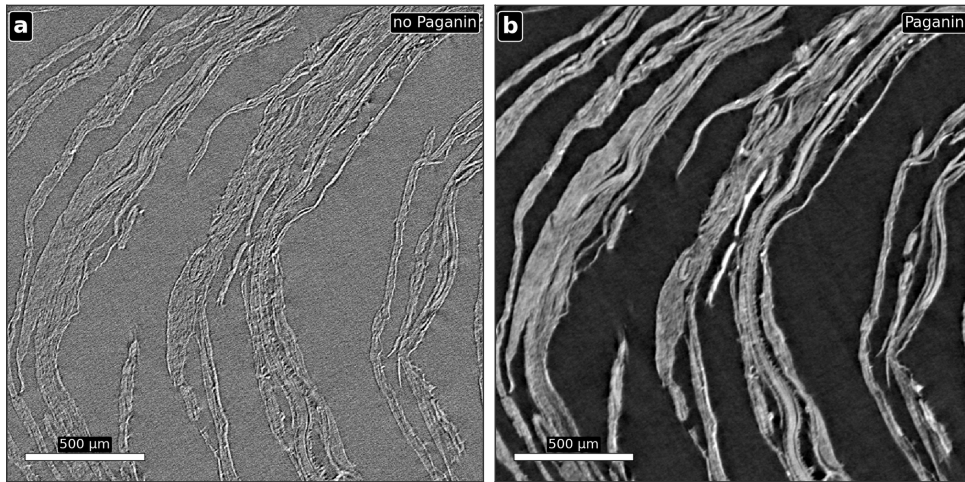
d, Companion view of the Paganin phase-retrieval trade-off on PHerc. 139 (9.362 μm pixel size) PHerc. 139, 113 keV mean, 1.2 m sample-to-detector, 9.362 μm effective pixel size, binmean2 of a 4.681 μm acquisition processed with two Paganin δ/β values illustrate the phase-retrieval trade-off. From left: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



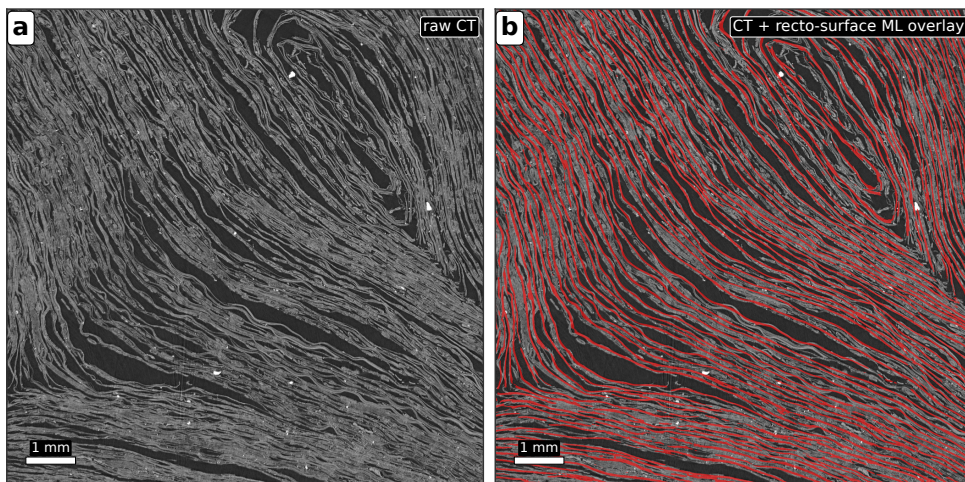
e, Detector pixel size determines what fraction of the layered papyrus structure can be resolved. Reconstructions of densely packed Herculanum papyrus at five detector pixel sizes show the progressive loss of layer separability as the pitch coarsens. From left: 0.55 μm (PHerc. 500P2 ID11 monochromatic small-ROI scan, 65 keV, 0.07 m propagation, $\delta/\beta = 20$), 1.129 μm (PHerc. 1667, BM18 59 keV polychromatic, 0.22m 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 \times 1.0 mm axial field of view at its native pyramid level 0; the yellow rectangle marks a 100 \times 100 μm sub-ROI and the bottom-right inset is that ROI at 10 \times 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.; the 2.4 μm appears as the best possible resolution on a setup compatible with large scrolls complete imaging; the 4.317 μm panel is haze-limited and the 9.362 μm panel is pixel-limited, the inset shows the bare voxel grid.



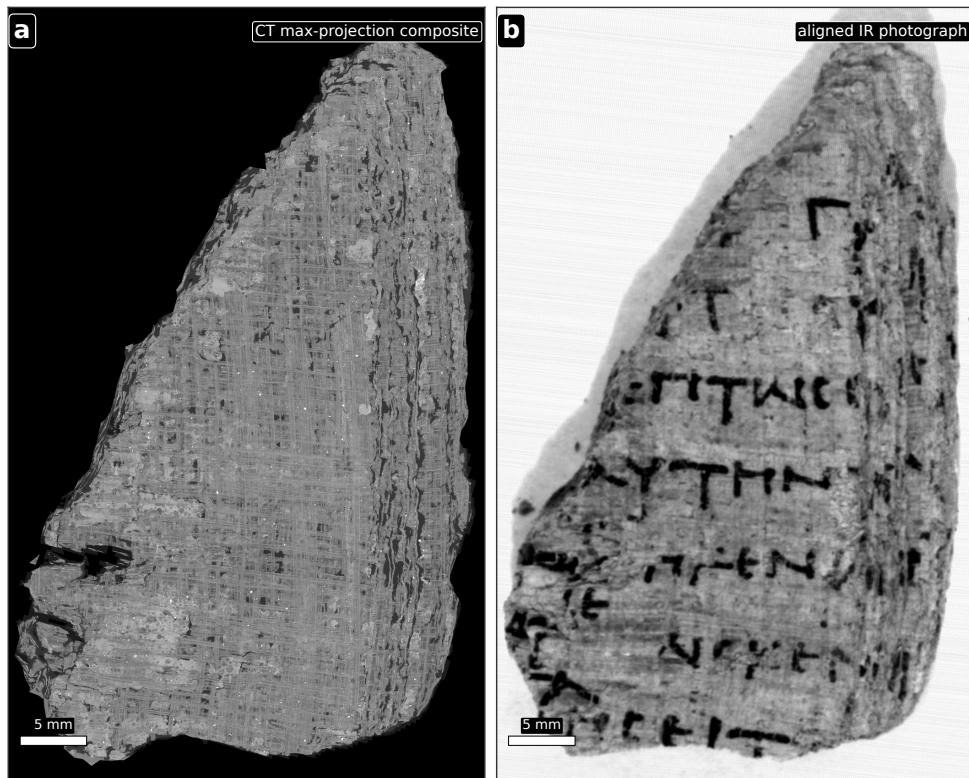
f, Incident X-ray energy trades contrast magnitude and beam decoherence. Three slice reconstructions of PHerc. 343P 2.4 μm / 0.22 m / $\delta/\beta = 1000$, 62 keV; 77 keV; 89 keV. Physical size 2.0 mm \times 2.0 mm FOV. Across 62 \rightarrow 89 keV the general contrast (faint absorption with superimposed propagation phase-contrast) drop is visible but gentle, and layer boundaries read crisply at every column, however the layer separability increases. We found the empirical sweet spot to be at 77 keV.



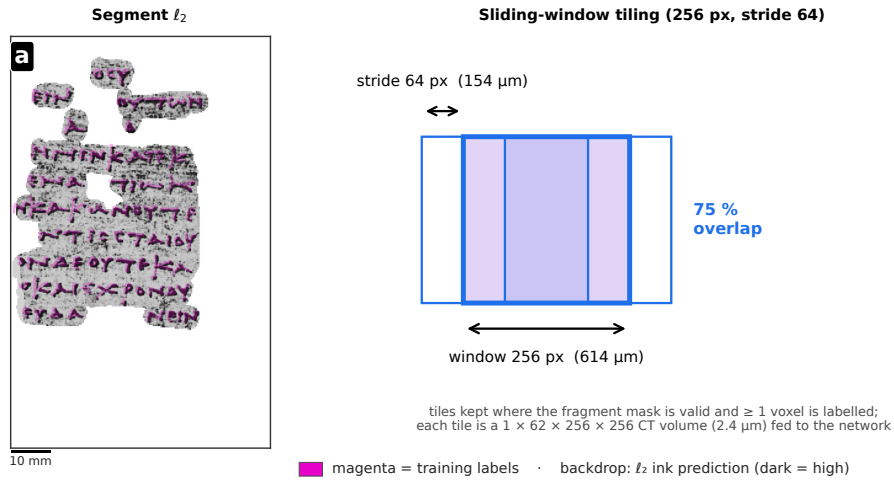
Extended Data Fig. 3 | Ink directly visible in PHERC. Paris 4. Comparison of the same region of interest in the PHERC. Paris 4, 78 keV, 2.4 μm pixel size, propagation distance 0.22 m **a**, No Paganin (or more accurately $\delta/\beta = 0.0001$). The slice is dominated by phase contrast fringes. Contrast 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.



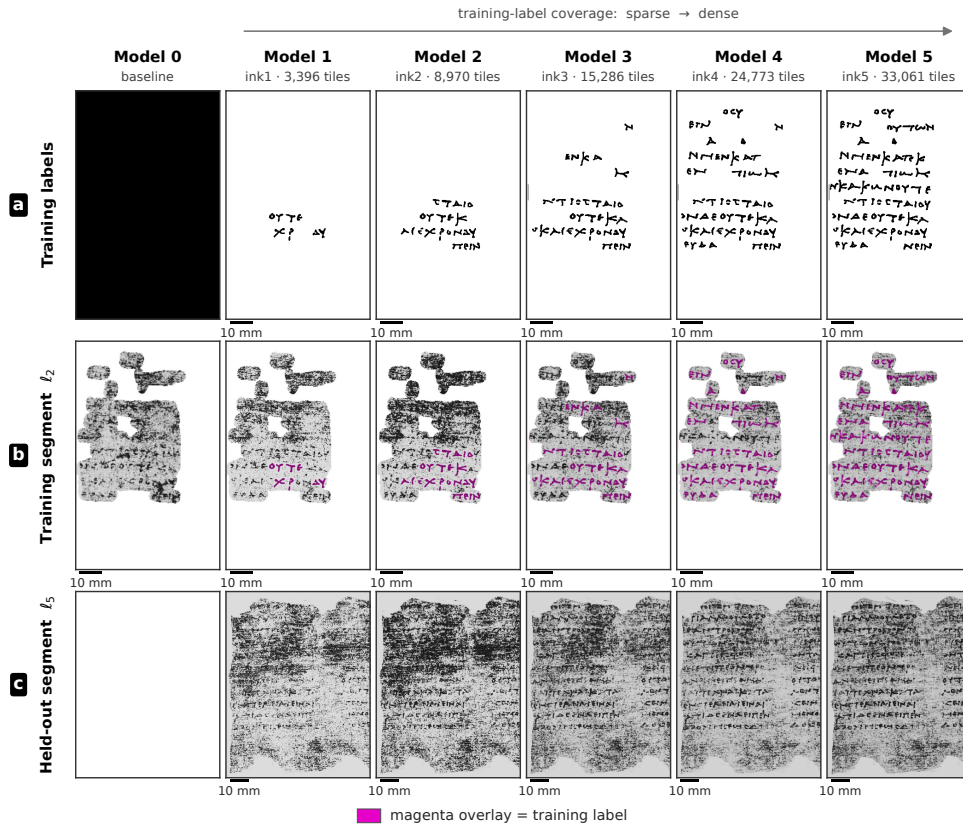
Extended Data Fig. 4 | Surface predictions. Examples of recto-surface prediction in PHERC. Paris 4.



Extended Data Fig. 5 | 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 \times 58.1 mm physical FOV. Scale bars 5 mm.



Extended Data Fig. 6 | Sliding-window sampling and model input geometry **a**, A region of scroll PHerC. 1667 (4.30×4.53 mm; $2.4 \mu\text{m}/\text{voxel}$) centred on a representative ink-labelled connected component (teal). The faint white grid traces every 256-voxel training window positioned at stride 64 ($8\times$ spatial oversample; 624 unique positions in this region). Five highlighted windows on one row, separated by the 64-voxel stride ($154 \mu\text{m}$), illustrate how adjacent samples overlap by 75% along each axis; the bold pink box marks the single window enlarged in **(b)**. **b**, The model input is a $256 \times 256 \times 62$ -voxel cuboid. Greyscale shading is z-slice 31 of the volume; the binary training label is overlaid in magenta. Scale bar: $200 \mu\text{m}$.



Extended Data Fig. 7 | Pseudo-labeling process. Each column is an iteration of pseudo-labeling. Model 0 is the starting point not trained on PHerc. 1667. Row a shows the training labels created from the previous iteration’s inference. Row b shows prediction from this iteration’s model, labels used are overlaid in magenta. Row c is the validation region, never used for creating labels.

SUPPLEMENTARY INFORMATION

Supplementary Table 1 | 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 |
| | Normalisation | 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 normalisation | Per-volume 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 |
| Optimisation | 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) |
| | Batch size | 2 |
| | Mixed precision | AMP fp16 |
| | Gradient clipping (L2 norm) | 12 |
| | Validation cadence | every 5 epochs (\leq 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 |

Values pulled from the wandb run vesuvius-challenge/paper/l2_ink5_l5infer (rd0qz8ps), created 2026-06-15, fine-tuned from the Kinetics-700 ResNet3D-50 release r3d50_KM_200ep.pth[36].

Supplementary Table 2 | Training configuration of the ink detection model ablation

| Category | Parameter | Value |
|--|--|---|
| Architecture | Backbone | ResNet3D-50[36] – 3D conv, BN, ReLU residual blocks |
| | Decoder | 2-D U-Net: 3 upsampling blocks of (bilinear $\times 2 \rightarrow$ concat skip $\rightarrow 3 \times 3$ conv \rightarrow BN \rightarrow ReLU) |
| | Encoder–decoder bridge | per-stage max over the z (depth) axis \rightarrow 4-D feature pyramid |
| | Convolution | nn.Conv3d (encoder), nn.Conv2d (decoder) |
| | Normalisation | nn.BatchNorm3d (encoder), nn.BatchNorm2d (decoder) |
| | Encoder init | r3d50_KM_200ep.pth (Kinetics-700, Hara et al.); conv1 weights summed across RGB \rightarrow 1 grayscale channel |
| Input | Output channels | 1 (sigmoid binary ink probability) |
| | Patch size (D \times H \times W) | 62 \times 256 \times 256 voxels |
| | 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) |
| Training data (<i>ink5 row</i>) | Segment | l_2 (PHerc. fragment) |
| | Label file | labels/l_2_inklabels5.png |
| | Train tiles | 33,061 (256 \times 256 sub-tiles at stride 64) |
| | Min labeled-ratio gate | label must have ≥ 1 non-zero pixel after binarisation |
| | 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 \times oversample) |
| | Inference tiles on l_2 | 22,204 (stride 128, after fragment-mask filter) |
| Augmentations (<i>training only</i>) | Flip | HorizontalFlip(p=0.5), VerticalFlip(p=0.5) |
| | Affine | ShiftScaleRotate(rot $\pm 360^\circ$, 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) |
| Optimisation | Loss | 0.5 \times Dice + 0.5 \times SoftBCE (smooth = 0.25), against label down-interpolated to 64 \times 64 |
| | Optimizer | AdamW |
| | Initial LR | 2 $\times 10^{-5}$ |
| | Max LR (OneCycle) | 3 $\times 10^{-4}$ |
| | 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 | cuda.deterministic = True, cudnn.benchmark = False | |
| Compute | GPU: 1 \times NVIDIA H100 80 GB | |
| Compute | Per-run memory | \approx 12 GB GPU, \approx 160 GB host RAM (l_2 + l_5 layer stacks held in process memory) |
| | Per-run wall time | \approx 2 h end-to-end (load + train + l_5 inference) |
| | Parallelisation | 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 (<i>ink5, step 12,396</i>) | Train total loss | 0.5299 |

Supplementary Table 3 | 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`.

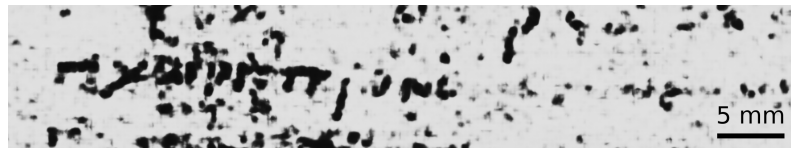
| Category | Parameter | Value |
|---------------------------|---|--|
| Architecture | Backbone | 3D Vision Transformer (DINOv2 / EVA-style, "v2" config), single-channel volumetric input |
| | Patch embedding | nn.Conv3d, $8 \times 8 \times 8$ -voxel patches (non-overlapping) |
| | Embedding dimension | 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) |
| | Normalisation | LayerNorm (pre-norm) |
| RoPE settings | Register tokens | 4 |
| | base 100; per-axis ("separate") normalisation | train-time coord. aug. shift 0.05 / jitter 1.05 / rescale 2 |
| Input | Stochastic depth (drop-path) | 0.2 |
| | Input channels | 1 (grayscale CT) |
| | Global crops ($Z \times Y \times X$) | $128 \times 128 \times 128$ voxels, $\times 2$ |
| | Local crops | $64 \times 64 \times 64$ voxels, $\times 8$ (scale range 0.02–0.25) |
| | Source sampling region | $256 \times 256 \times 256$ voxels |
| | Intensity normalisation | Per-volume robust scaling (median / IQR) |
| Self-supervised objective | Voxel size | $\approx 2.4 \mu\text{m}$ isotropic |
| | Losses | DINO (CLS-token) + iBOT (masked-patch) + KoLeo regulariser (Gram term present, weight 2, inactive) |
| | Centering | Sinkhorn–Knopp |
| | Teacher temperature | 0.04 \rightarrow 0.07 (warmup 25,000 steps) |
| Training data | Teacher EMA momentum | 0.994 \rightarrow 1.0 |
| | Pretraining volumes (10 scrolls, ESRF $\approx 2.4 \mu\text{m}$) | PHerc. 9B, 500P2, 814, 1299, 343P, 332, 139, MAN5, 1451, MANBp |
| | Validation volumes | PHerc. 332, PHerc. 139 |
| Optimisation | PHerc. Paris 4 finetuning volume | PHerc. Paris 4 (20260411..., $2.400 \mu\text{m}$) — added at finetuning |
| | Optimizer | AdamW |
| | Learning rate | 2.5×10^{-5} (cosine; warmup 35,000 steps \rightarrow min 1×10^{-6}) |
| | Patch-embedding LR multiplier | 0.2 |
| | Weight decay | 0.04 \rightarrow 0.2 (cosine) |
| | Gradient clipping (L2 norm) | 5 |
| | Batch size | 3 (per device) |
| | Mixed precision | AMP |
| Pretraining end-state | Max iterations (scheduled) | 1,000,000 |
| | Checkpoint / validation cadence | every 2,500 / 500 steps |
| | Final logged step | 342,558 (end-state: stopped; checkpoint_step_342500 used downstream) |
| PHerc. Paris 4 finetuning | 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 \rightarrow final 353,686 ($\approx 11,200$ steps) |
| | Validation loss (total) | 15.076 |

Supplementary Table 4 | Training configuration of the PHerc. Paris 4 volumetric 3D ink-detection models. The four models share the architecture and optimisation 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).

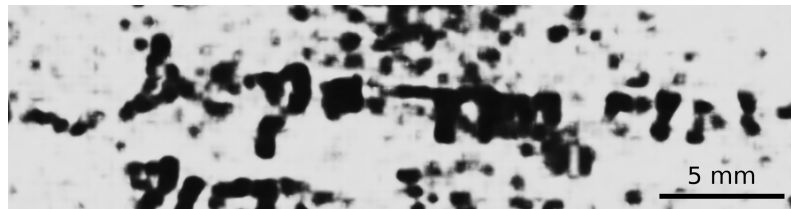
| Category | Parameter | Value |
|---------------------------------|---|---|
| Architecture | Backbone | nnU-Net-style 3D residual-encoder U-Net (vesuvius NetworkFromConfig, auto-configured) |
| | Convolution | nn.Conv3d |
| | Normalisation | nn.InstanceNorm3d |
| | Encoder stages / feature widths | 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 ($1 \times 1 \times 1$ conv) |
| | Parameters | 1.42×10^8 (≈ 142 M) |
| Input | Output channels | 1 (ink; sigmoid + BCE) |
| | Patch size ($Z \times Y \times X$) | $256 \times 256 \times 256$ voxels |
| | Input channels | 1 |
| | Intensity normalisation | Per-volume percentile min–max |
| | Mode | full-3D (surface annotations placed ± 3 voxels about the reference surface) |
| Training data | Min labeled-coverage gate | 0.02 |
| | 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 |
| Optimisation | Seed | 27 |
| | 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 |
| DINO-guided rounds | Weight EMA | decay 0.9995 (from step 1,000); used for validation and as the frozen guidance teacher |
| | Max iterations (scheduled) | 250,000 |
| | DINO model | Paris 4-finetuned dinovol (checkpoint_step_352500); 128^3 sliding windows (stride 128) \rightarrow 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-normalised |
| | Threshold τ | 0.5 (guided rounds) |
| Pipeline stages & run end-state | Background mask | input intensity ≥ 50 (added from round 2 onward) |
| | Threshold τ | 0.5 |
| | 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 = $\sigma(\text{teacher})\text{-minmax}(\text{cosine} \rightarrow \text{ink}) > 0.5$; final step 62,950; val balanced-acc. 0.715; val loss 0.317 (EMA 0.131) |
| 44 | 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) |

Supplementary Table 5 | 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 \times 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 | 4.327 \rightarrow 4.116 | 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 | 1000 | Used as the 1.129 μm panel in the pixel-size comparison. Confirm δ/β and unsharp settings. |
| 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 $\delta/\beta = 1000$. |
| 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 $\delta/\beta = 1000$. |
| 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. |



a w49: χωρὶς προνο[ί]ας ("without providence")



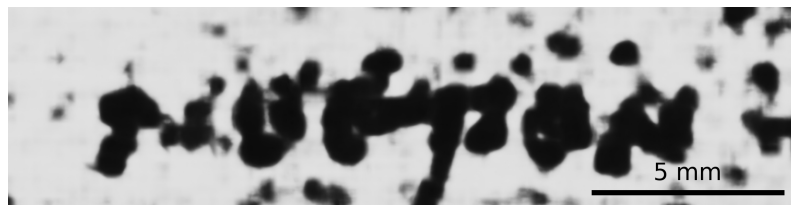
b w25: ἀόρατα ("invisible entities")



c w34: καὶ τὸ κατὰ φύσιν ("according to nature")



d w47: καὶ πάντα χωρὶς πόνων ("and everything free from troubles")



e w49: νοερόν ("intellectual")

Supplementary Fig. 1 | Readings from PHerc. 139. Scale bars, 5 mm.