Virtually Unrolling the Herculaneum Papyri by Diffeomorphic Spiral Fitting
Supplementary Material

S1. Transform parameterization

In this section we give additional details on the three components that make up the overall diffeomorphism $T_{S\to V}$ that maps the idealized spiral in canonical space to the highly deformed shape observed in the scan volume. The transform is a composition $T_{S\to V} = T_{\text{aff}} \circ T_{\text{flow}} \circ T_{\text{gap}}$, where: (i) T_{aff} is a per-slice affine (scaling and translation) transform; (ii) T_{flow} is the result of integrating a flow field; (iii) T_{gap} rescales inter-winding gaps.

S1.1. Per-slice affine transform

We non-isotropically scale and translate each slice in the xy-plane; the scale factors and translations vary wrt z. This stage is parameterized by N 2D log-scales and translations where N is a fixed number of keypoints along the z-axis; the translations and log-scales are linearly interpolated for other z values. Thus, for a given point $\mathbf{x} = (x_1, x_2, x_3)$, we have

$$T_{\text{aff}}(\mathbf{x}) = \begin{pmatrix} x_1 \exp(s_1) + t_1 \\ x_2 \exp(s_2) + t_2 \\ x_3 \end{pmatrix} \tag{S1}$$

where $(s_1, s_2) = (1-\alpha)\,\mathbf{s}^{(\lfloor\tau\rfloor)} + \alpha\,\mathbf{s}^{(\lfloor t\rfloor+1)}$, $\quad \alpha = \tau - \lfloor\tau\rfloor$, $\tau = N(x_3 - z_{\min})/(z_{\max} - z_{\min})$, and $\mathbf{S}^{(n)}, n = 1 \dots N$ are the N log-scales; translations are defined analogously.

S1.2. Integrated flow field

We introduce a flow velocity field $\mathbf{u}$, *i.e.* a 3D vector field of 3D velocities defined over the scan volume. This allows us to transform a given point by simulating its movement in flow of the given local velocities for a fixed length of time. This is done by solving an ODE defined by $\mathbf{u}$ and an initial condition (the starting point). Writing $\phi^{(t)}$ for the transform induced by simulating the flow $\mathbf{u}$ for time t, the final transformed position of a point $\mathbf{x}$ is given by $T_{\text{flow}}(\mathbf{x}) = \phi^{(1)}(\mathbf{x})$, where $\phi^{(t)}$ evolves according to

$$\frac{d}{dt}\phi^{(t)}(\mathbf{x}) = \mathbf{u}\left(\phi^{(t)}(\mathbf{x})\right) \tag{S2}$$

with the initial condition $\phi^{(0)}(\mathbf{x}) = \mathbf{x}$. The resulting transformation is guaranteed to be a diffeomorphism, provided the flow field $\mathbf{u}$ is itself smooth [1, 2]. Formally, the space of flow fields is the Lie algebra that generates the Lie group of diffeomorphisms; the former provides a straightforward way to parameterize the latter. Moreover, the inverse diffeomorphism is easily found by negating the flow field, then integrating backwards in time. We define the velocity field as the sum two underlying fields $\mathbf{u} = \mathbf{u}_{\text{coarse}} + \mathbf{u}_{\text{fine}}$, where $\mathbf{u}_{\text{coarse}}$ is more spatially smooth than $\mathbf{u}_{\text{fine}}$. Both fields are defined by trilinear interpolation of explicitly-represented discrete grids of 3D velocity vectors. The grid for $\mathbf{u}_{\text{fine}}$ is $48\times$ lower resolution than the original CT volume, and that for $\mathbf{u}_{\text{coarse}}$ is a further $6\times$ lower.

We solve the ODE using explicit Euler integration with 16 steps; more steps did not result in significantly different results, whereas fewer could cause $T_{\text{S}\to\text{V}} \circ T_{\text{V}\to\text{S}}$ to not be sufficiently close to the identity transform. Experiments with scale-and-square integration [5]—which requires fewer operations—showed that it is less stable in our setting than Euler integration. To update the flow field via gradient descent, we must differentiate back through the ODE solve, i.e. compute the Jacobian of $\phi^{(1)}(\mathbf{x})$ wrt $\mathbf{u}$. We achieve this by directly unrolling and back-propagating through the Euler updates; this is tractable due to the relatively small number of steps. With more steps or larger minibatches, it would become preferable to use the adjoint method, similar to continuous normalizing flows [3].

We also conducted initial experiments using a time-varying velocity field, defined explicitly by changing the 3D coarse and fine grids to be 4D. This has only a minor effect on compute cost but increases memory linearly with the number of timesteps. Despite such fields having intuitively simpler behavior, we did not notice any significant improvements to final results, hence all our main experiments use a time-constant field.

We also tested an implicit velocity field defined by a small neural network (a residual MLP) operating on 3D coordinates in the volume, and regressing these directly to the velocity vector, similar to neural radiance fields (NeRF) [4] and other neural field methods [6]. While this potentially allows allocating parameter capacity more efficiently where needed (instead of uniformly across the volume regardless of curvature), in practice we found the computational expense to be prohibitive.

S1.3. Inter-winding gap scaling

We define a scalar field over the 2D surface of the canonical scroll spiral, which specifies for each point, by how much the next winding should be pushed outward from its default distance (defined by being an Archimedean spiral), *i.e*. locally scaling the gaps between windings. Since this field is defined only on the 2D spiral surface (truncated at a certain number of windings), it has quadratic cost in memory. This allows using a finer resolution than the full-volume flow field, and is also more expressive since it can directly separate or push apart nearby sheets regardless of how their position relates to the fixed flow-field grids. The field itself is defined explicitly by scalar values on a 2D grid wrt the z and θ coordinates of eq. 1 in the main paper. This discrete representation is extended to the 2D plane by bilinear interpolation. The field values are interpreted as log scale factors multiplying the spacing between windings, originally defined by the global winding rate ω. Specifically, for a given point $\mathbf{x}$, we find the next spiral winding inwards from it, measure the radial distance to that winding, multiply this distance by the relevant scale factor, and move the point outwards correspondingly.

S2. Hyperparameters

Diffeomorphism parameterization	
number of Euler integration steps	16
voxel resolution of flow field	12
voxel resolution of gap expander	32
Optimization	
learning rate	5×10^{-4}
learning rate schedule	constant
number of optimization steps	20000
num. points sampled per path	100
num. points for relative winding number loss	2000
num. points for normals loss	2000
num. points for regularization	1500
path distance loss starting from step	10000
Loss weights	
surface normal	200
path radius	5
relative winding number	10
path distance (L1)	4
fiber direction	5
stretch regularization	200
centerline fixed	1

S3. Additional Visualizations

Figure S1a shows cross-sections of *PHerc. Paris 4*, with the original undeformed spiral (before optimization) overlaid (top), and the fitted version (bottom). We also display (Figure S1b) two slices that are transformed by the inverse diffeomorphism to map them back into canonical space, resulting in near-circular concentric windings. Similarly, Figure S2 shows cross-sections of *PHerc. 172* before and after fitting, while Figure S3 shows undeformed slices with circularized windings.

In Figure S4 we show the full unrolling of *PHerc. 172* produced by our method. In Figures S5–S11 we zoom in on parts of the unwrapped surface, with enhanced contrast. In each of these views, ink is visible forming Greek characters (light against the darker background) and sometimes complete words. Note also the individual fibers of papyrus visible forming continuous lines through the view, indicating the our method is correctly following the true sheet surfaces.

Alongside this PDF, we also include:

- a video interpolating between the idealized scroll with circular cross-section (placed on the centerline of the deformed scroll), and the deformed scroll shape defined by the diffeomorphism
- a video showing four cross sectional slices, interpolating between the spiral (red) in its undeformed and deformed states
- a high-resolution version of the virtually unrolled scroll, with the texture of the papyrus visible

References

[1] Vincent Arsigny, Olivier Commowick, Xavier Pennec, and Nicholas Ayache. A log-euclidean polyaffine framework for locally rigid or affine registration. In *Biomedical Image Registration*, pages 120–127, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[2] John Ashburner. A fast diffeomorphic image registration algorithm. *NeuroImage*, 38(1):95–113, 2007.

[3] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 32*, 2018.

[4] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.

[5] Cleve Moler and Charles Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review*, 45(1):3–49, 2003.

[6] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Computer Graphics Forum*, 41 (2):641–676, 2022.

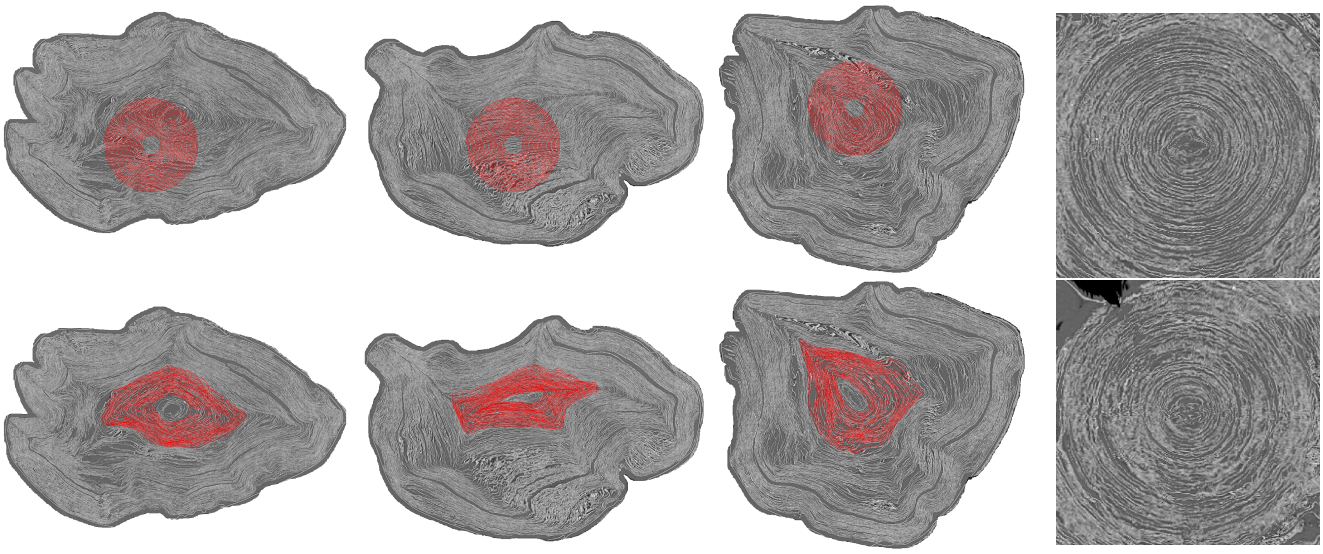

(a) (b)

Figure S1. **(a)** Cross-sections of *PHerc. Paris 4* with the spiral overlaid, before optimization (top) when the diffeomorphism is an identity transform; and after optimization (bottom). During fitting the spiral adapts to follow the distorted scroll surface. **(b)** Undeformed scroll cross-sections, given by transforming the scan with the inverse diffeomorphism. The distorted windings in (a) become near-circular in (b).

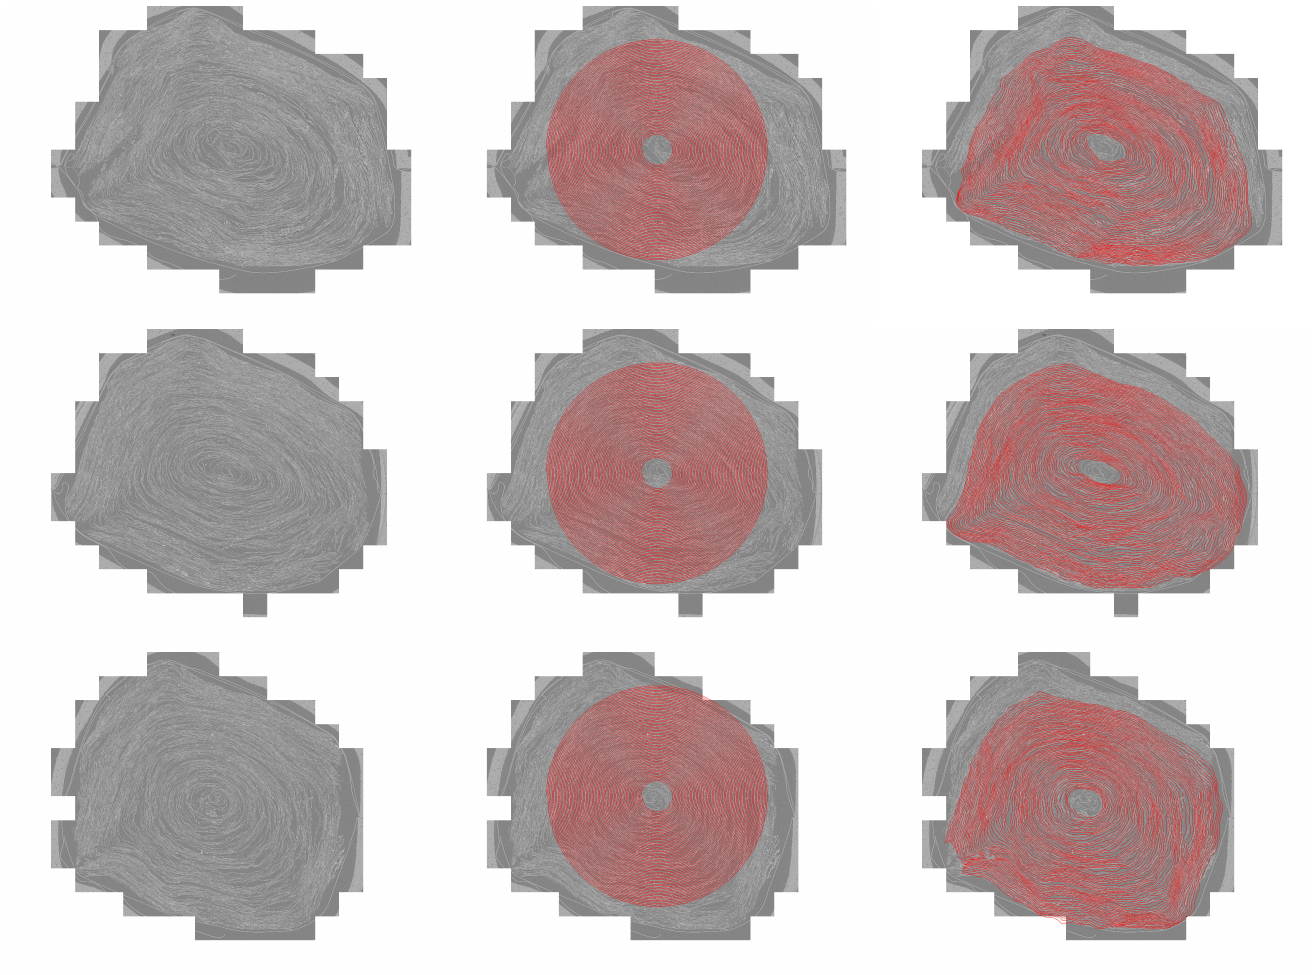

Figure S2. Cross-sections of *PHerc. 172*. **Left:** original slices, showing the highly deformed, tightly packed windings. **Middle:** spiral overlaid before optimization, when the diffeomorphism is an identity transform and the winding rate ω is arbitrarily chosen. **Right:** spiral overlaid after optimization. During fitting the spiral adapts to follow the distorted scroll surface.

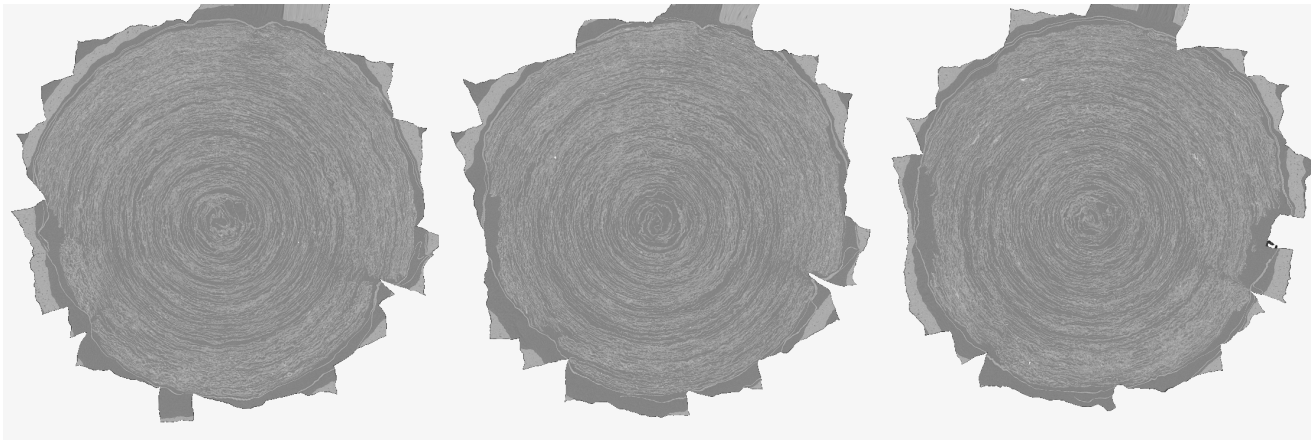

Figure S3. Undeformed cross-sections from *PHerc. 172*, given by transforming the scan with the inverse diffeomorphism. The distorted windings in Figure S2 become near-circular.

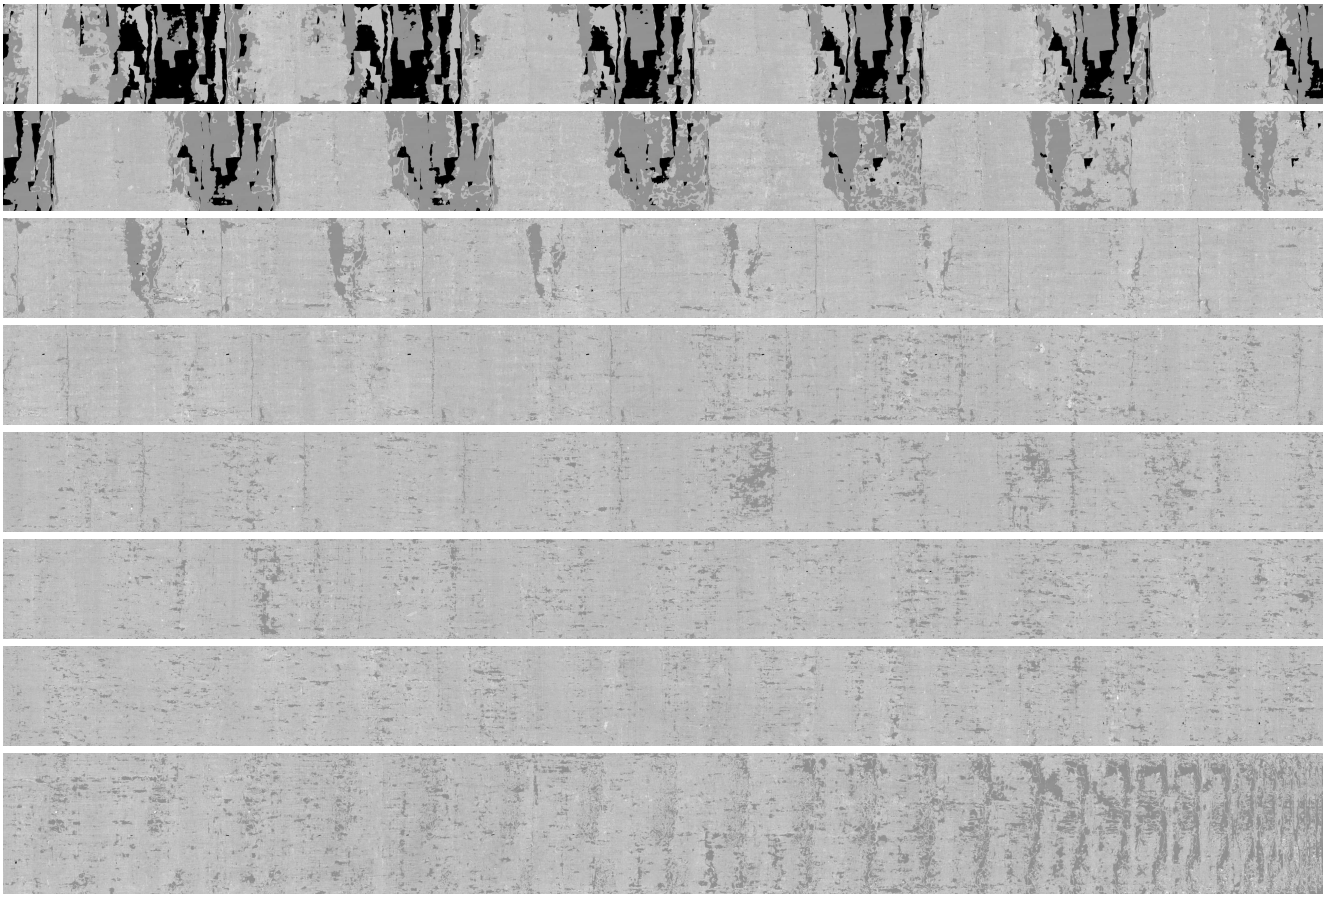

Figure S4. *PHerc. 172* unrolled using our method. Note the very long, continuous manifold output. Figures S5–S11 show higher-contrast zooms on part of this surface, with ink visible. Black and dark gray parts in the top two rows are where our method has tracked surfaces through regions where the scroll is missing (either burnt away, or damaged by parts flaking off).

Figure S5. Zoom on our virtual unrolling PHerc.172, with ink clearly visible as bright areas forming Greek characters

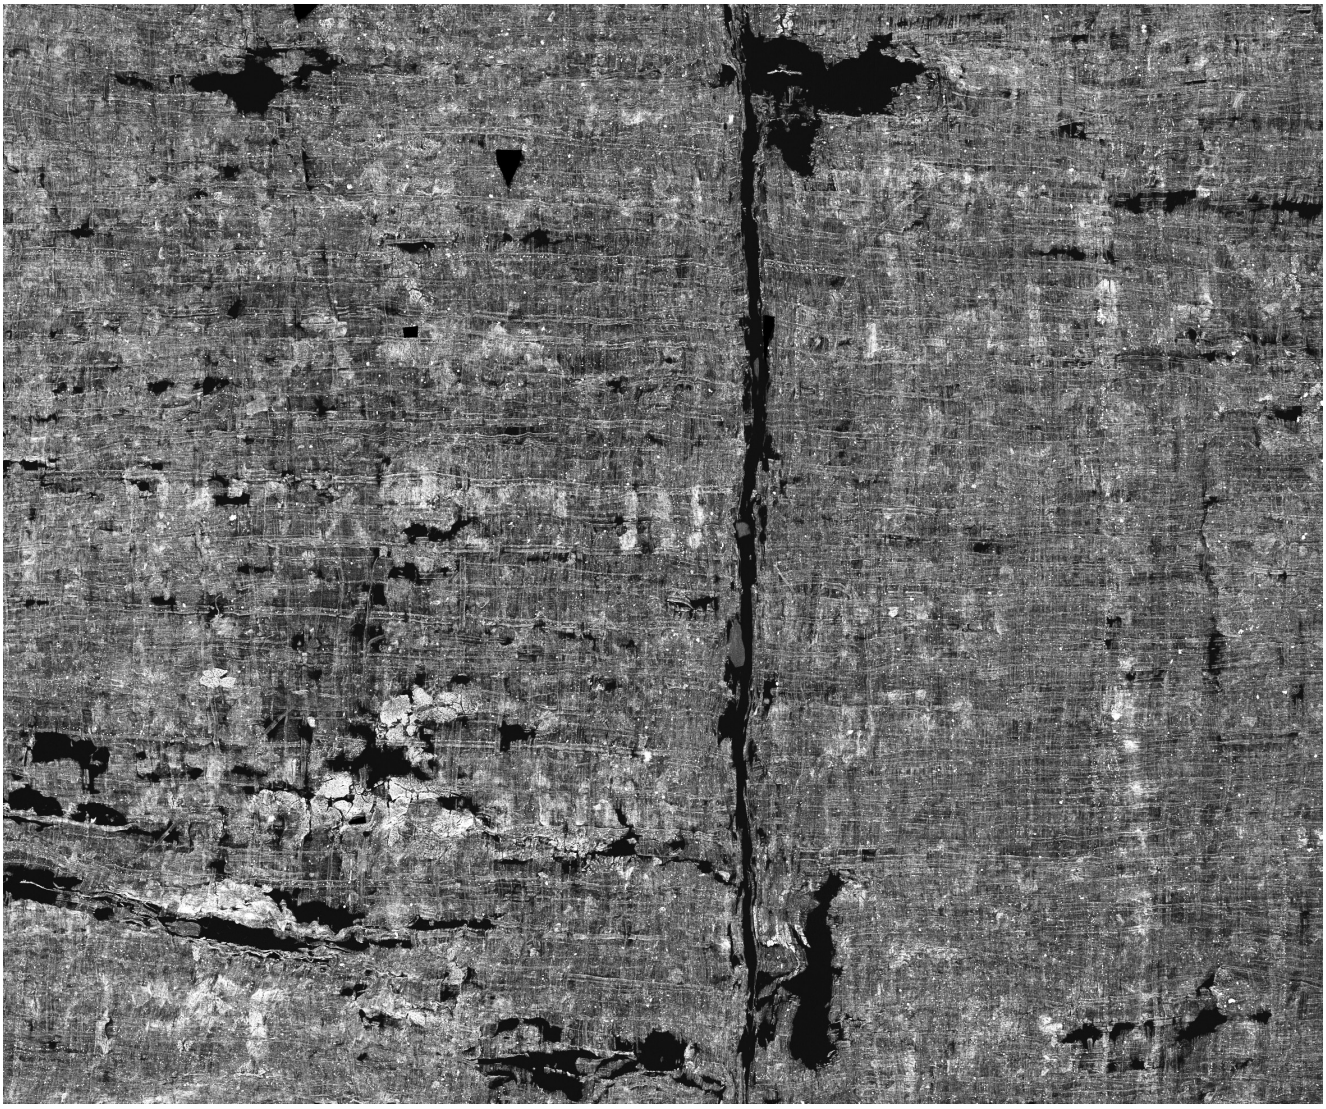

Figure S6. Zoom on our virtual unrolling PHerc.172, with ink clearly visible as bright areas forming Greek characters

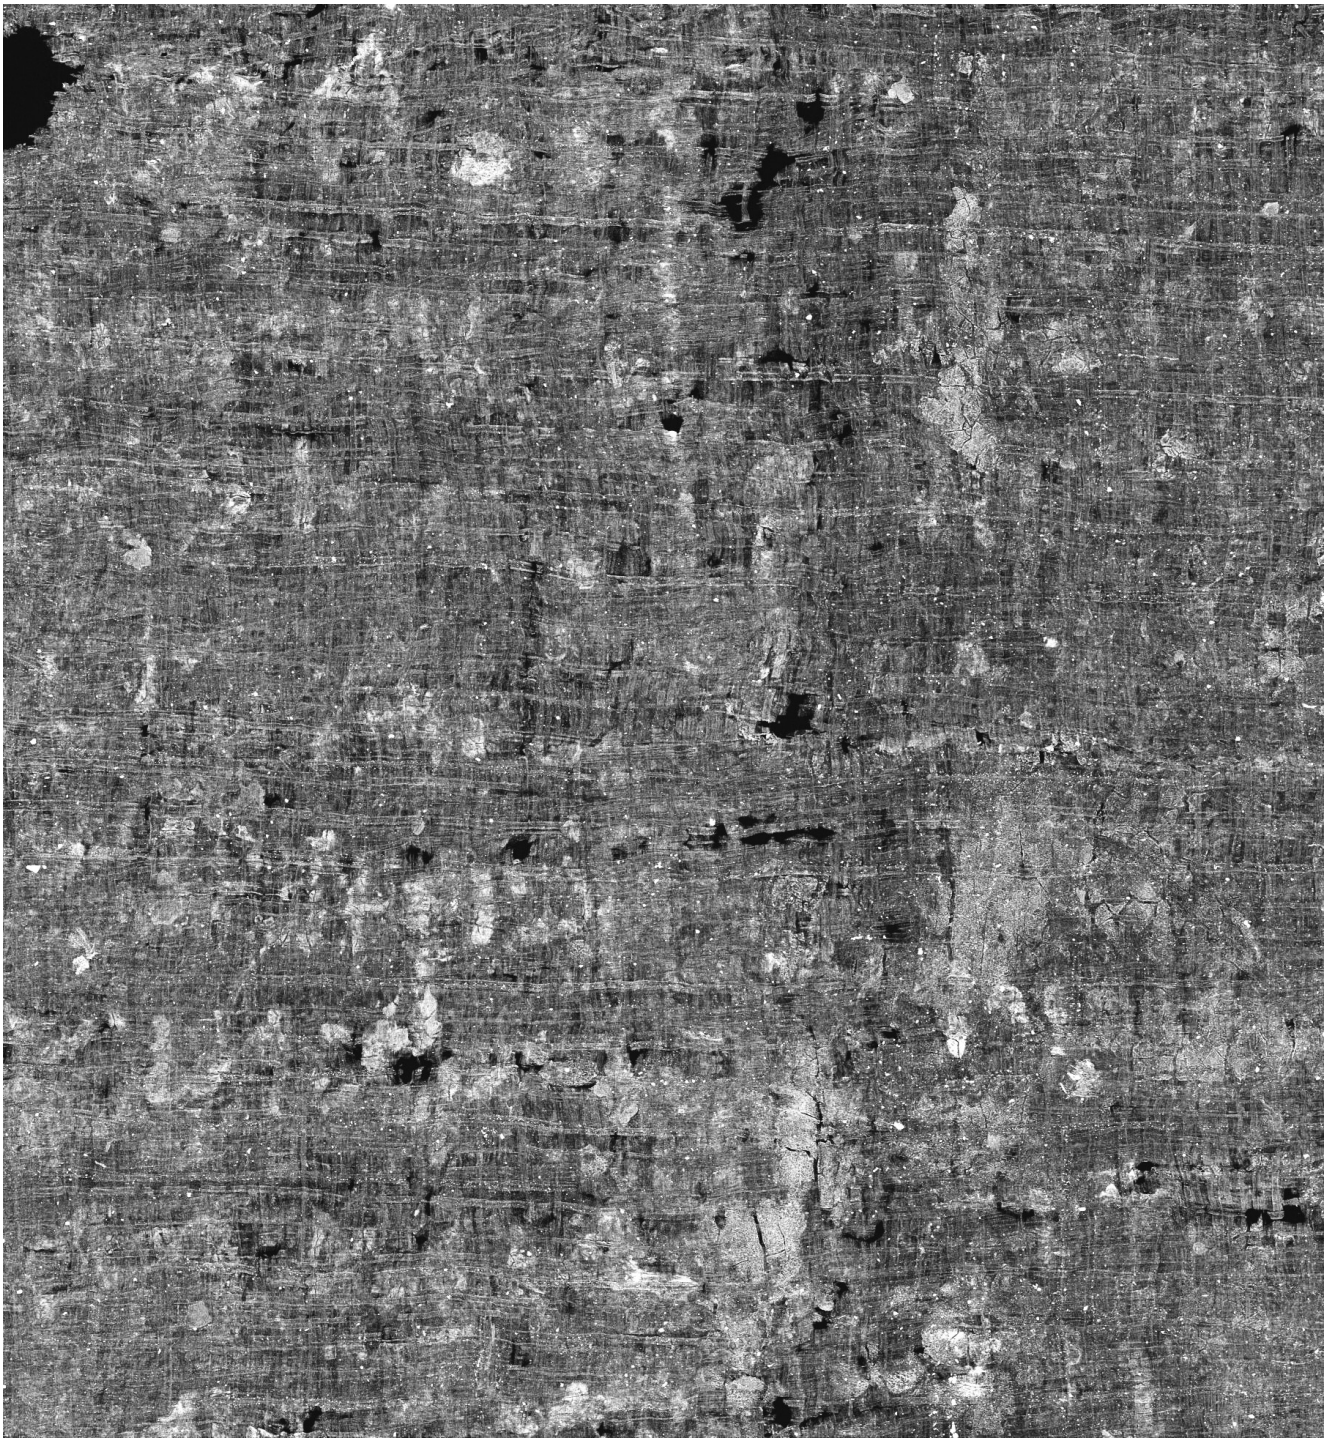

Figure S7. Zoom on our virtual unrolling PHerc.172, with ink clearly visible as bright areas forming Greek characters

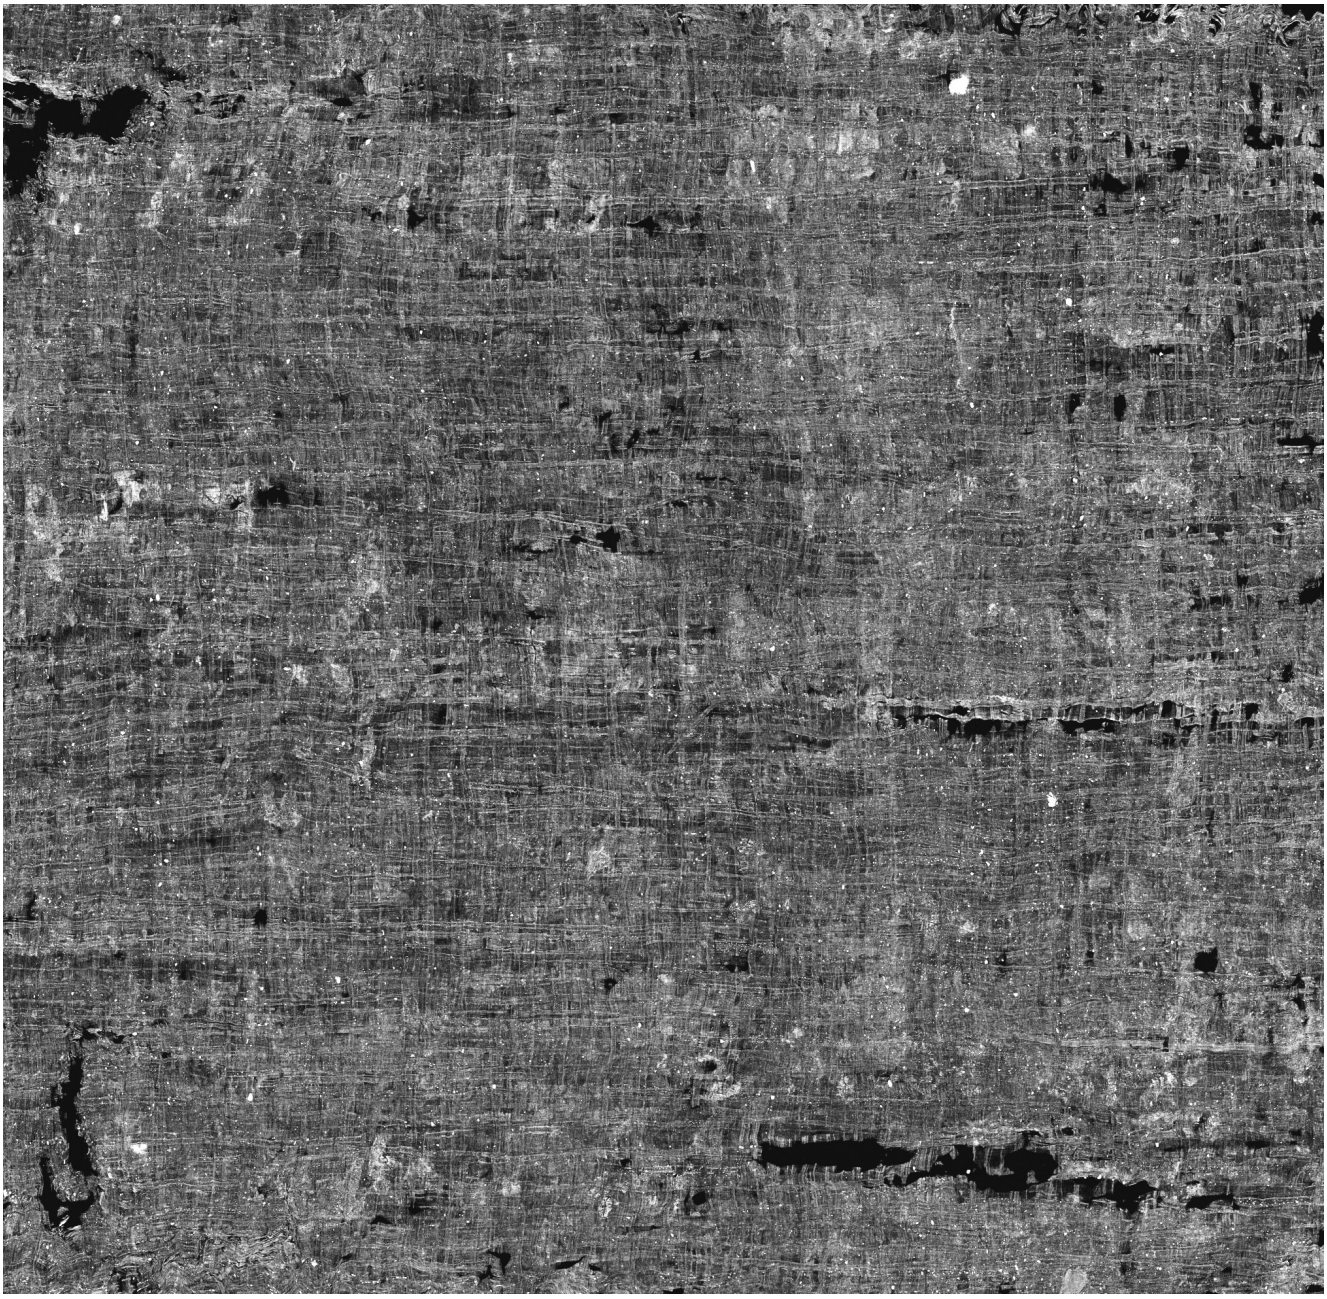

Figure S8. Zoom on our virtual unrolling PHerc.172, with ink clearly visible as bright areas forming Greek characters

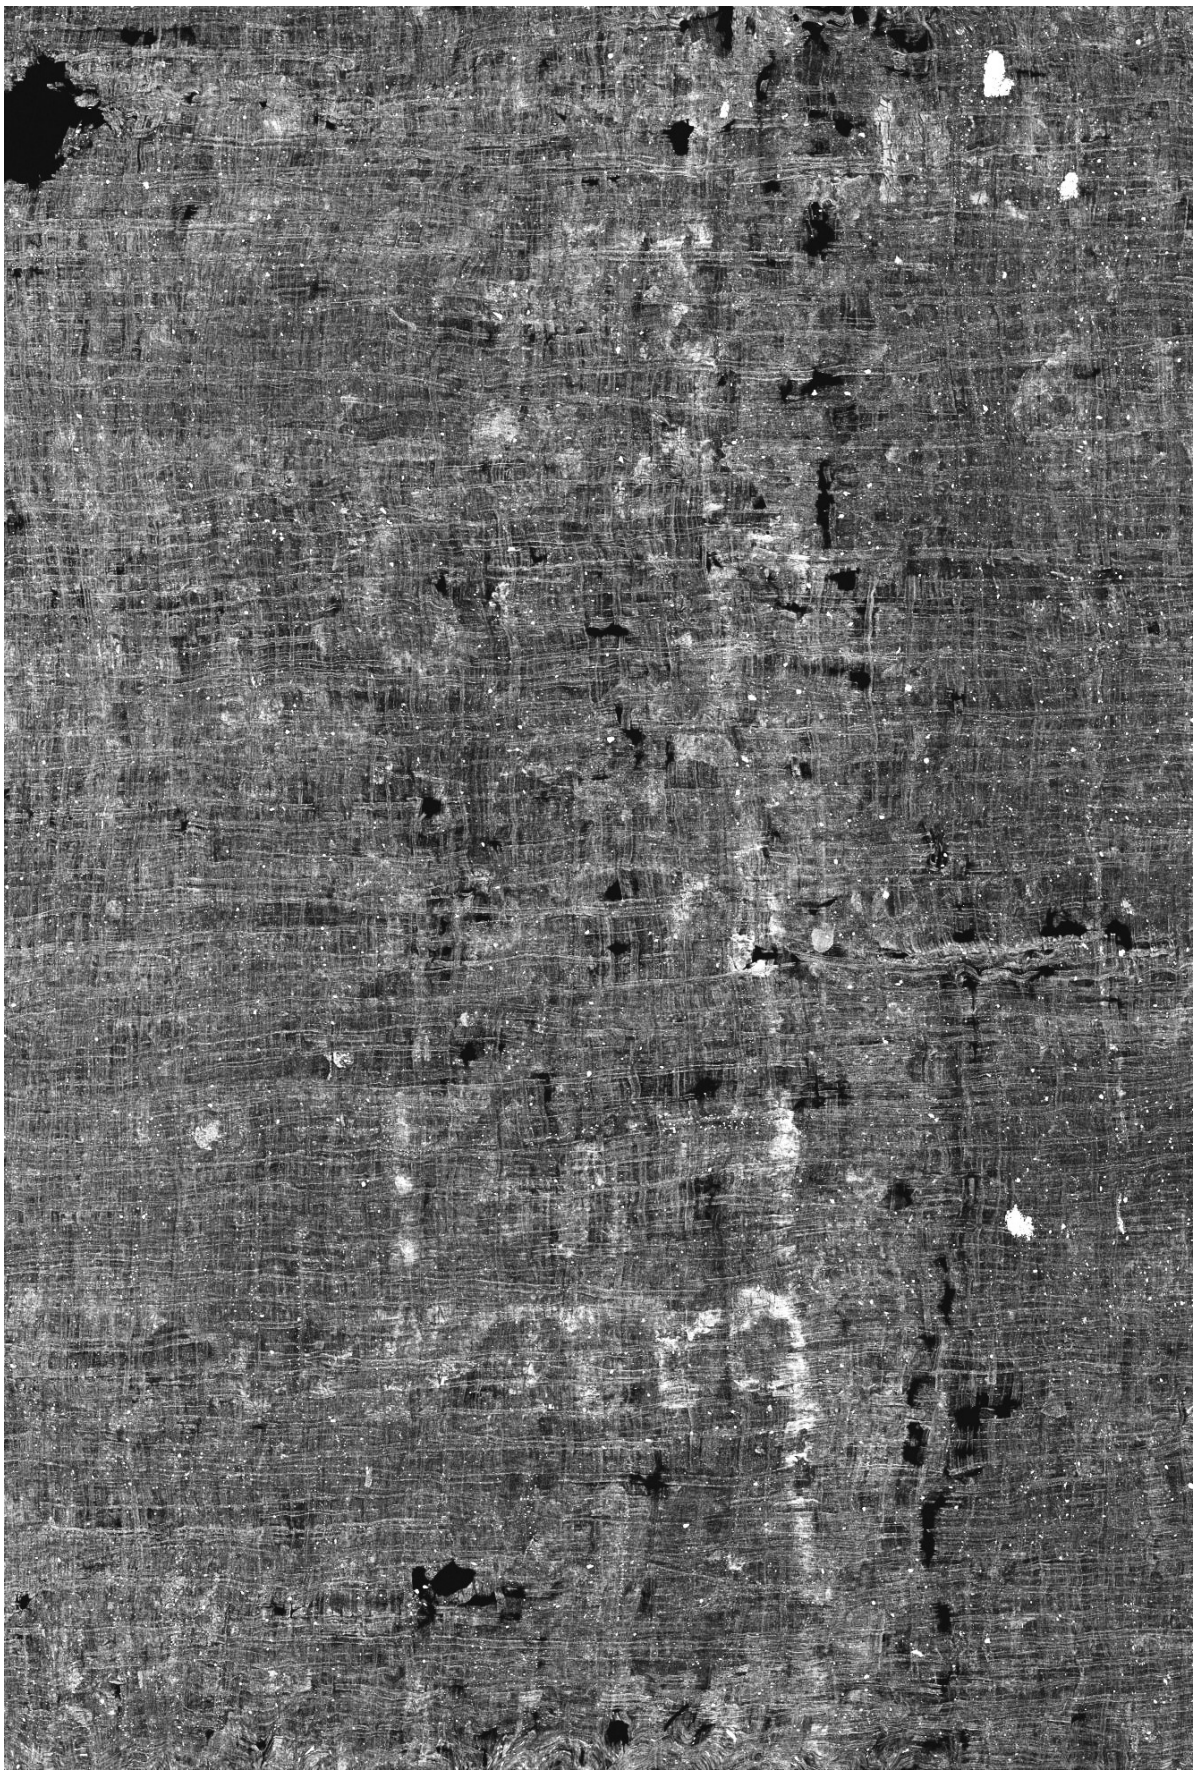

Figure S9. Zoom on our virtual unrolling PHerc.172, with ink clearly visible as bright areas forming Greek characters

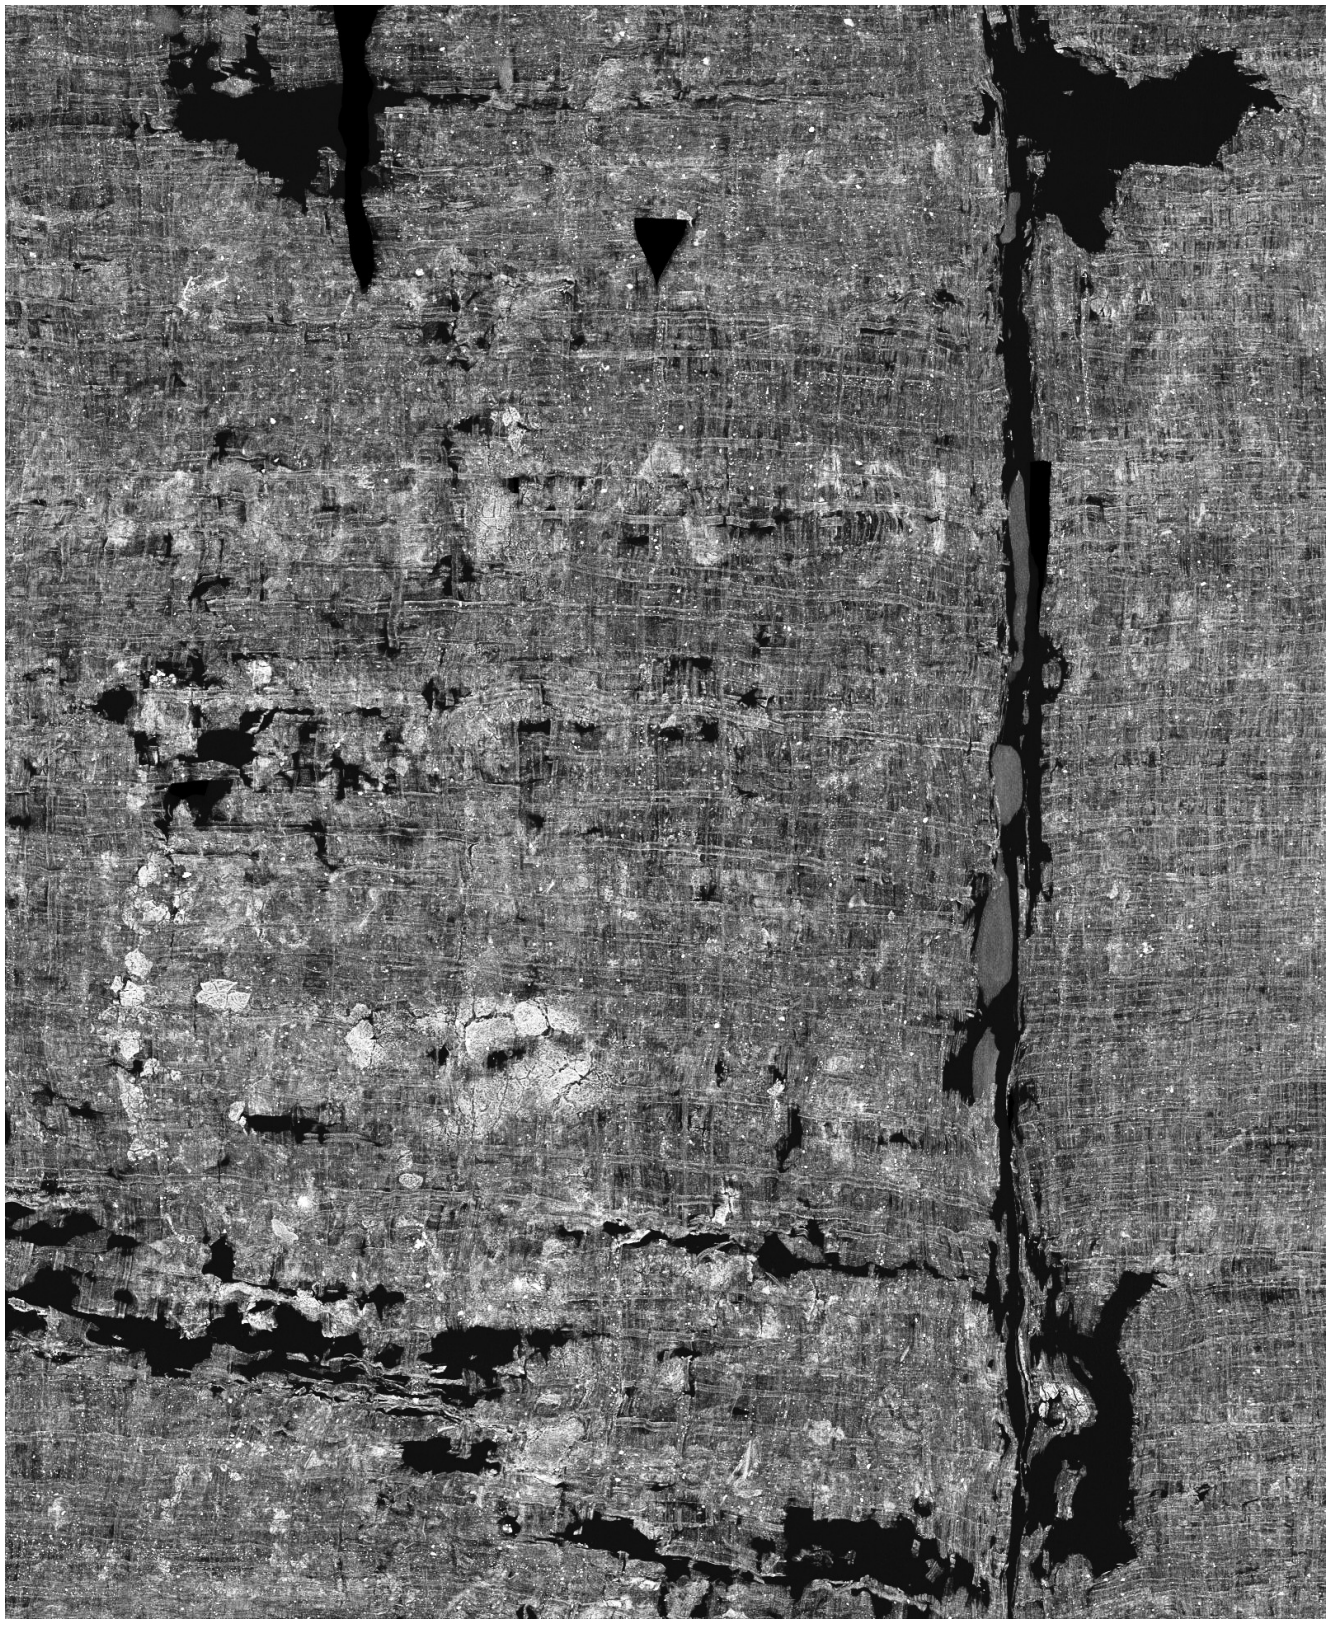

Figure S10. Zoom on our virtual unrolling PHerc.172, with ink clearly visible as bright areas forming Greek characters

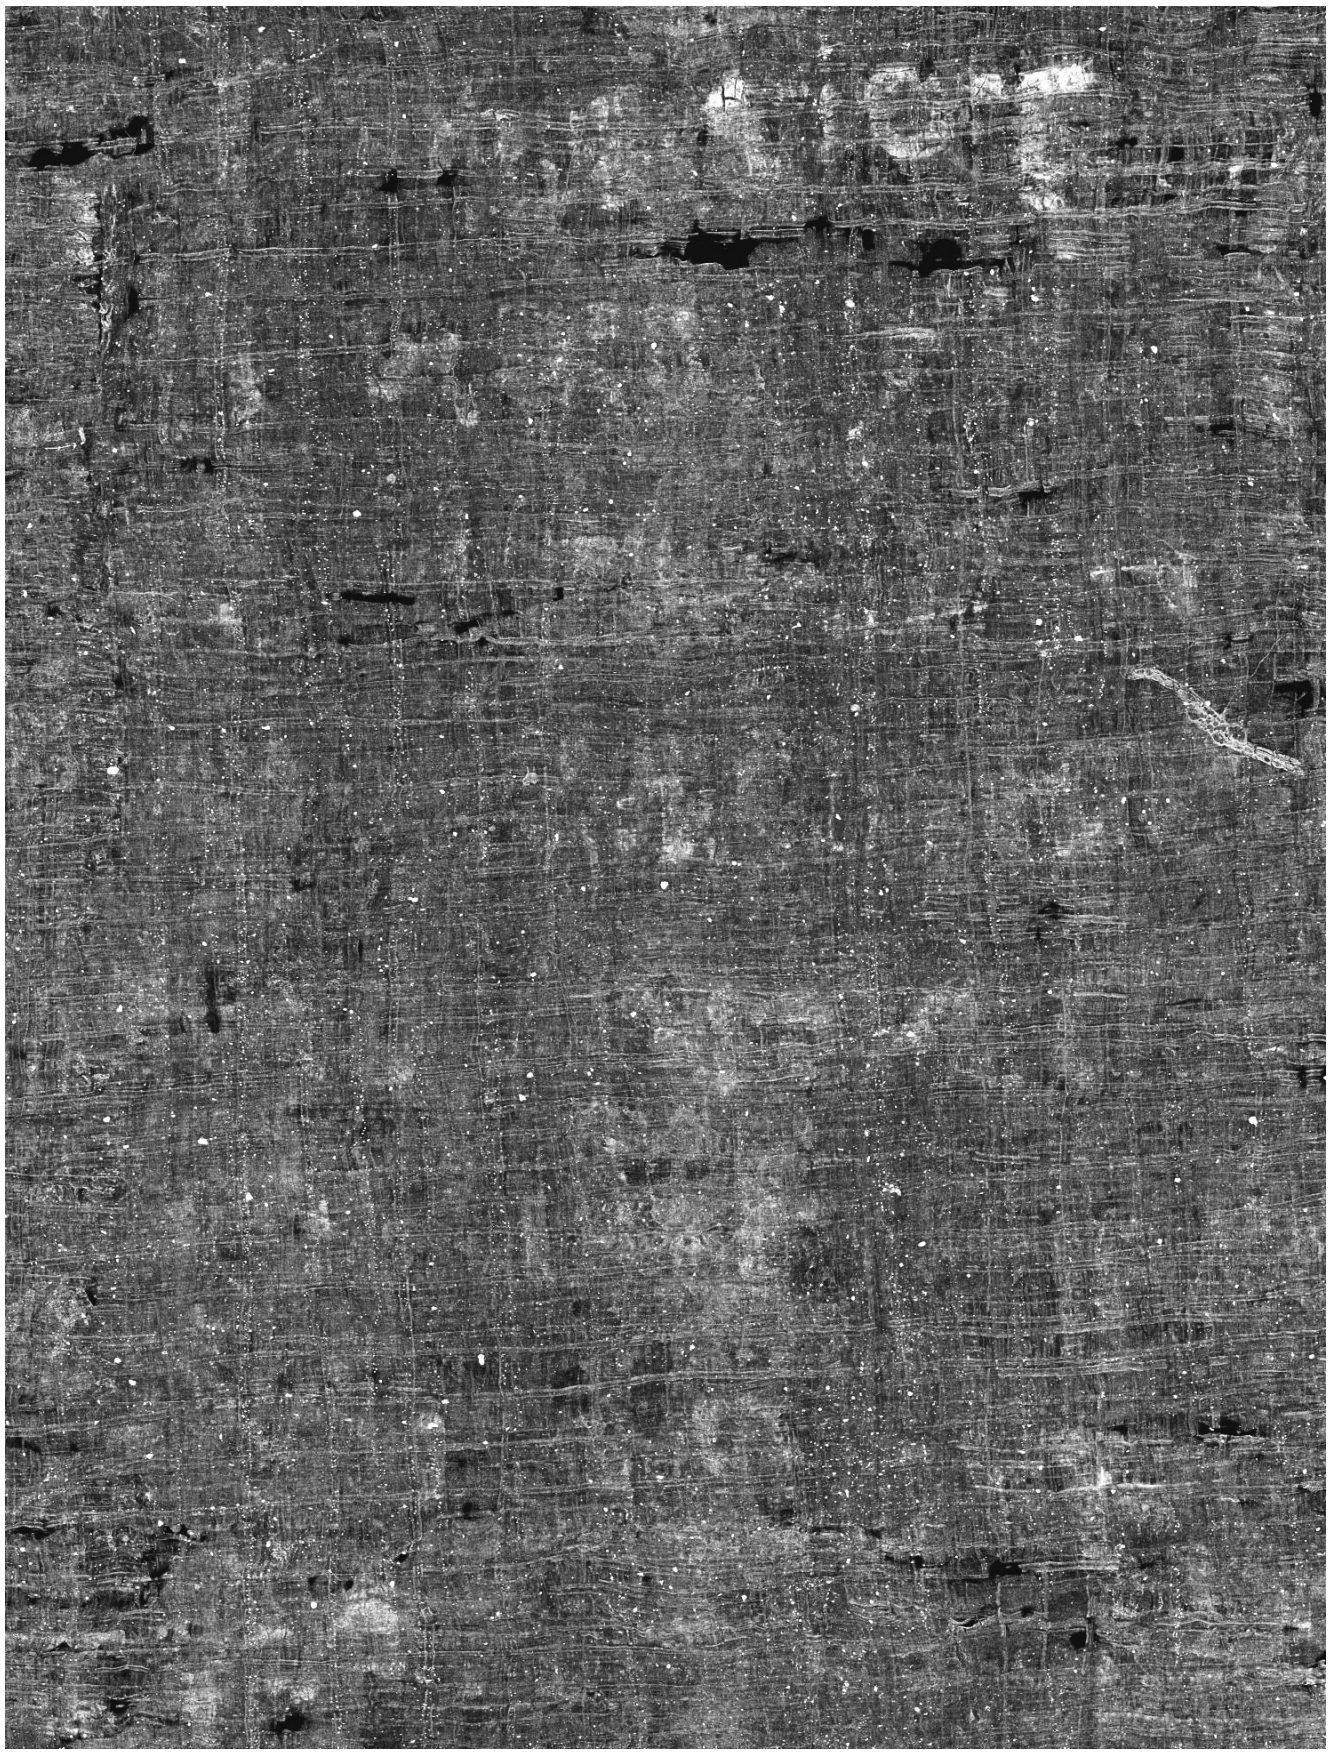

Figure S11. Zoom on our virtual unrolling PHerc.172, with ink clearly visible as bright areas forming Greek characters