Supplementary Material:Physics-Guided Multistep Deformation Reversal for Ancient Bamboo Slip Restoration

1. Dataset Background

Archaeological Context. The bamboo slips in this study originate from Tomb 77 at Shuihudi, Yunmeng County, Hubei Province, excavated in 2006. The tomb owner, 'Yue Ren', served as a local official in Anlu County. The slips document a rich corpus of legal codes, administrative records, work logs, mathematical texts, and almanacs dating to the Western Han dynasty during Emperor Wen's reign (179-157 BCE).

The Restoration Challenge. These artifacts suffer from fractures due to external forces, and the waterlogged soil environment has caused mild corrosion through microbial activity and moisture-induced degradation of bamboo fibers. This preservation status presents significant restoration challenges. First, the artifacts are extremely fragile, making traditional manual restoration (e.g., mechanical flattening) a time-consuming and high-risk process that can cause irreversible damage, such as structural breakage. Second, the deformation itself is highly complex, involving multi-directional, non-uniform warping governed by the material's lignocellulosic fiber structure in response to moisture and stress.

Dataset Construction. To develop a computational restoration method, we constructed a specialized dataset. We first curated a dataset of 2,000 'pristine' (straight and well-preserved) slips from the high-resolution digital archives of the **Shuihudi Tomb 77 artifacts**. To create the necessary paired training data (which does not exist naturally), we utilized the physics-based forward process (detailed in the main paper, Sec. 4.2) to apply physically-plausible deformations to these pristine slips, dynamically generating a large-scale training set. Furthermore, to ensure an objective evaluation, we collected an additional 500 slips with authentic, natural deformations from this excavation, which were held out as an independent test set.

2. Detailed Metric Definitions

As stated in the main paper, a successful restoration must satisfy three core criteria: (1) **Geometric Fidelity** (for rejoining), (2) **Textual Preservation** (for legibility), and (3) **Physical Plausibility** (for credible results). The following metrics are designed to evaluate these criteria.

2.1. Straightness Metric

To evaluate Geometric Fidelity, this metric quantifies the straightness of the restored slip's boundaries [1].

$$S = 1 - \frac{\sigma_{\text{curvature}}}{\mu_{\text{curvature}} + \epsilon} \tag{1}$$

Higher scores indicate better geometric restoration.

2.2. Perceptual Quality (LPIPS)

This metric assesses general visual fidelity using Learned Perceptual Image Patch Similarity [14]. Lower scores indicate better perceptual quality.

2.3. Text Readability Quality (TRQ)

To evaluate Textual Preservation, we propose TRQ (Text Region Quality) [13]. We use this heuristic metric as reliable open-source OCR models for Qin dynasty seal script are lacking. The three components of TRQ are grounded in validated document quality indicators: Kumar and Ramakrishnan [7] systematically evaluated document quality metrics against human annotations across 132 multi-script document images annotated by 6 experts, establishing that metrics with correlation coefficients exceeding ± 0.3 are valid predictors of perceived quality. Their results demonstrated that **sharpness** ($\rho \approx -0.55$), **continuity** ($\rho \in [-0.30, -0.47]$), and **contrast** ($\rho = -0.49$) all significantly correlate with mean subjective scores, providing strong statistical support that our EdgeSharpness, StrokeContinuity, and TextClarity components are recognized indicators of text quality aligned with expert perception.

TRQ is computed as an improvement ratio:

$$\text{TRQ} = Q_{restored}/Q_{deformed} \tag{2}$$

A score > 1 indicates improvement. The quality score Q is a weighted sum of three sub-metrics:

$$Q = 0.4 \cdot \text{ES} + 0.3 \cdot \text{SC} + 0.3 \cdot \text{TC} \tag{3}$$

The sub-metrics are defined as: (1) **EdgeSharpness (ES)**: The mean Sobel gradient magnitude within the text region

R.

$$\text{ES} = \frac{1}{255|R|} \sum_{(x,y)\in R} \sqrt{G_x^2 + G_y^2} \tag{4}$$

(2) **StrokeContinuity (SC)**: The ratio of valid text connected components c (area $\in [10, 500]$ pixels) to filter noise.

$$\text{SC} = \frac{\#\{c : 10 \leq \text{Area}(c) \leq 500\}}{\#\{\text{all connected components}\}} \tag{5}$$

(3) **TextClarity (TC)**: The Weber contrast between text (μ_{text}) and background (μ_{bg}) pixels.

$$\text{TC} = \frac{|\mu_{text} - \mu_{bg}|}{255} \tag{6}$$

2.4. Deformation Consistency Index (DCI)

To evaluate **Physical Plausibility**, we propose the DCI. The final $DCI \in [1, 10]$ is a linear normalization of a base score DCI_{base}:

$$DCI_{\text{base}} = \exp(-\alpha E_{\text{cont}} - \beta E_{\text{energy}}) \tag{7}$$

where $\alpha = 1.0$ and $\beta = 1.0$. The energy terms are:
(1) **Continuity Energy (E_{cont})**: Penalizes discontinuity (high strain) in the deformation field $F = (F_x, F_y)$.

$$E_{\text{cont}} = \frac{1}{HW} \sum_{x,y} \left[\left(\frac{\partial F_x}{\partial x}\right)^2 + \left(\frac{\partial F_x}{\partial y}\right)^2 + \left(\frac{\partial F_y}{\partial x}\right)^2 + \left(\frac{\partial F_y}{\partial y}\right)^2 \right] \tag{8}$$

(2) **Deformation Energy (E_{energy})**: Penalizes high displacement magnitudes (high energy state).

$$E_{\text{energy}} = \frac{1}{HW} \sum_{x,y} ||F(x, y)||^2 \tag{9}$$

The gradient terms are computed using finite differences. A higher DCI score indicates a more physically consistent restoration.

3. Complete Physical Deformation Model

This section provides the complete mathematical derivations of the physical deformation model presented in Section 3 of the main paper. We establish the continuous formulations that describe bamboo slip deformation under moisture-induced stress, which serve as the theoretical foundation for our discrete implementation.

3.1. Continuous Physical Theory

We begin by formulating the complete continuous physical model that governs bamboo slip deformation. This theoretical framework integrates fiber bundle mechanics, viscoelastic behavior, and moisture-stress coupling effects.

3.1.1. Fundamental Deformation Model

The deformation of bamboo slips is primarily driven by changes in moisture content. For each fiber bundle, the length change due to moisture variation is given by:

$$x_i' = x_i(1 + \alpha_i \cdot \Delta W), \tag{10}$$

where x_i' is the new length of the i-th fiber bundle, x_i is the original length, α_i is the moisture expansion coefficient, and ΔW denotes the change in moisture content. This moisture-induced deformation leads to geometric incompatibility, which results in a curvature described by:

$$\kappa \approx \frac{\Delta L}{L \cdot h}, \tag{11}$$

where ΔL is the length difference between inner and outer fiber bundles, L is the bamboo slip length, and h is its thickness. The total energy of the system is the sum of bending and stretching energies:

$$E_{total} = E_{bend} + E_{stretch}, \tag{12}$$

where $E_{\text{bend}} \propto \kappa^2$ represents bending energy and $E_{\text{stretch}} \propto \sum_i (x_i' - x_i^{\text{actual}})^2$ represents the stretching/compression energy. These energy components are critical in guiding the deformation and restoration processes.

3.1.2. Fiber Bundle Interaction and Viscoelastic Effects

The deformation process is further complicated by interactions between adjacent fiber bundles [5, 8] and viscoelastic behavior of the bamboo material [9]. When one fiber bundle deforms, it influences neighboring bundles [5] through the matrix material that connects them. This mechanical coupling can be modeled as a viscoelastic interaction [9].

For each fiber bundle, the stress transfer to adjacent bundles can be described by:

$$\sigma_{i\to j}(r) = \sigma_i \cdot e^{-r/\lambda}, \tag{13}$$

where $\sigma_{i\to j}$ represents the stress transferred from bundle i to bundle j, σ_i is the stress in bundle i, r is the distance between bundles, and λ is a characteristic length parameter that depends on the matrix properties.

The viscoelastic response of each fiber bundle can be modeled using a modified Kelvin-Voigt formulation [4]:

$$\varepsilon_i(t) = \frac{\sigma_i}{E_i} + \frac{\sigma_i}{\eta_i}(1 - e^{-t/\tau_i}), \tag{14}$$

where $\varepsilon_i(t)$ is the time-dependent strain, σ_i is the applied stress, E_i is the elastic modulus, η_i is the viscosity coefficient, and $\tau_i = \eta_i/E_i$ is the relaxation time for the i-th fiber bundle.

The orientation of fiber bundles changes gradually under these interactions, following:

$$\frac{d\theta_i}{dt} = k \cdot \sum_{j\neq i} \sigma_{j\to i} \cdot \sin(\theta_j - \theta_i), \tag{15}$$

where θ_i is the orientation angle of the i-th fiber bundle, and k is a rotational response coefficient.

3.1.3. Simulation of Temporal Degradation Factors

The deformation of bamboo slips is not instantaneous but evolves over time [9]. To model this temporal evolution, we account for spatial variability in the moisture expansion coefficient. As introduced in the main paper, $\alpha_i = \bar{\alpha} + \Delta\alpha_i$, where $\bar{\alpha}$ is the mean expansion coefficient and $\Delta\alpha_i$ captures local deviations across fiber bundles due to material heterogeneity and burial-induced degradation.

Similarly, degradation in fiber bundle material properties can be modeled as:

$$E_i(t) = E_i(0) \cdot f_E(t), \tag{16}$$

$$\eta_i(t) = \eta_i(0) \cdot f_\eta(t), \tag{17}$$

where $E_i(t)$ represents elastic modulus over time, $\eta_i(t)$ denotes viscosity variation, and $f_E(t)$ and $f_\eta(t)$ are degradation functions that depend on burial conditions and time.

Water absorption spatial randomness across fiber bundles further contributes to the complexity of deformation:

$$\frac{dW_i}{dt} = \overline{\frac{dW}{dt}} + \Delta\left(\frac{dW_i}{dt}\right). \tag{18}$$

Combining these factors, the time-dependent strain in each fiber bundle can be expressed as:

$$\varepsilon_i(t) = \varepsilon_i + (\bar{\alpha} + \Delta\alpha_i) \times \int_0^t (1 + \gamma_i\sigma_i)\left(\overline{\frac{dW}{dt}} + \Delta\left(\frac{dW_i}{dt}\right)\right) dt. \tag{19}$$

This comprehensive physical model based on fiber bundle interactions captures how water content changes drive bamboo slip deformation, with deterministic patterns [11] and stochastic variations [10] due to material heterogeneity. By simulating the forward deformation process through incremental moisture parameter adjustments, we establish a pathway for the inverse restoration algorithm described in the following section, which effectively reverses the moisture-induced warping that has occurred over centuries.

3.2. From Continuous Physics to Discrete Implementation

Having established the continuous physical theory, we now describe how these principles are implemented in our CreepDeformationEngine through a discrete control point representation. This subsection establishes the principled connection between the theoretical formulation and our numerical implementation.

3.2.1. Control Point Grid Design

The control point grid is defined automatically from the slip's contour, requiring no manual annotation. Following ablation studies on grid density (see Appendix sec5 for details), we selected an optimal 32×4 grid. This configuration divides the image into 4 vertical fiber bundles (reflecting the physical structure) and samples 32 points along each, totaling 128 control points. These sparse displacements are mapped to a dense field via bilinear interpolation, which ensures computational efficiency (>98% parameter reduction) and spatial smoothness consistent with continuous media mechanics.

3.2.2. Fiber Elongation: From Curvature to Displacement

To implement the bending physics (E_{bend}) on our discrete control points, we derive the displacement field from classical beam theory. For a beam undergoing bending, the axial strain ε at distance y from the neutral axis is linearly proportional to the curvature κ:

$$\varepsilon(y) = \kappa \cdot y. \tag{20}$$

Under the small deformation assumption, strain relates to displacement (where u **is the axial displacement**) as:

$$\varepsilon = \frac{\partial u}{\partial x} \approx \frac{\Delta u}{L}. \tag{21}$$

Combining these relations, the axial displacement at control point i is:

$$\Delta u_i = \varepsilon(y_i) \cdot L = \kappa \cdot (y_i - y_{\text{center}}) \cdot L. \tag{22}$$

In our coordinate system, axial displacement corresponds to the vertical component. To model the progressive accumulation of creep deformation (analogous to the time-dependent term in the Kelvin-Voigt model, Eq. 5 of the main paper), we introduce a temporal factor:

$$\Delta d_{i,y}^{\text{elong}} = \frac{t}{T} \cdot A_\kappa \cdot (y_i - y_{\text{center}}), \tag{23}$$

where $A_\kappa = \kappa \cdot L$ is the curvature amplitude parameter integrating the curvature magnitude and slip length, and t/T models the gradual nature of creep. The small deformation assumption is valid for archaeological bamboo slips, where typical curvature radii exceed 50 cm, ensuring $\kappa \cdot h \ll 1$.

3.2.3. Stress Transfer: Approximating Exponential Decay

The physical model (Eq. 4 of the main paper) describes stress transfer as an exponential decay $\sigma_{i\to j}(r) = \sigma_i \cdot e^{-r/\lambda}$. For short-range interactions within a local neighborhood, this decay can be efficiently approximated using a Gaussian

kernel. The Taylor expansion of the exponential function in the regime $r \ll \lambda$ yields:

$$e^{-r/\lambda} \approx 1 - \frac{r}{\lambda} + \frac{r^2}{2\lambda^2} + O(r^3), \tag{24}$$

which closely matches the profile of a Gaussian kernel for small distances. We implement this using a 3×3 Gaussian kernel:

$$K(r) = \frac{1}{2\pi\sigma_{kernel}^2} \exp\left(-\frac{r^2}{2\sigma_{kernel}^2}\right), \tag{25}$$

where σ_{kernel} is chosen such that the kernel effectively captures the stress distribution within the nearest-neighbor control points. For our grid spacing ($\Delta x \approx 6$ mm) and material parameter ($\lambda = 0.3$ corresponding to approximately 10 pixels), the maximum distance within the 3×3 neighborhood ($r_{\max} = \sqrt{2}\Delta x$) satisfies $r_{\max}/\lambda \approx 0.28 < 1$, ensuring the Gaussian approximation remains accurate. This localized convolution provides a computationally efficient implementation while preserving the essential physics of short-range stress coupling.

3.2.4. Moisture-Stress Coupling: Displacement Gradient as Stress Proxy

To implement the moisture-stress coupling $(1 + \gamma_i\sigma_i)$ from Eq. (8) of the main paper, we require a computationally tractable proxy for local stress σ_i. In continuum mechanics, stress is fundamentally related to strain through the constitutive law:

$$\sigma \propto \varepsilon, \tag{26}$$

and strain is defined as the spatial gradient of the displacement field:

$$\varepsilon = \nabla\mathbf{u}. \tag{27}$$

Therefore, stress is directly proportional to the displacement gradient: $\sigma \propto \nabla\mathbf{d}$. We use the magnitude of the numerically computed displacement gradient as a physically grounded stress proxy:

$$\sigma_i \approx \|\nabla\mathbf{d}_i\| = \sqrt{\left(\frac{\partial d_x}{\partial x}\right)^2 + \left(\frac{\partial d_y}{\partial y}\right)^2}, \tag{28}$$

where gradients are computed using finite differences on the control point grid. This formulation captures the dominant stress patterns arising from bending deformation. The moisture-stress amplification is then implemented as:

$$(1 + \gamma \cdot \|\nabla\mathbf{d}\|). \tag{29}$$

Material heterogeneity, arising from variations in fiber density and moisture distribution, is incorporated by adding Gaussian noise $\mathcal{N}(0, 0.05 \cdot t/T)$, whose variance increases with time to reflect accumulated environmental variability.

3.2.5. Progressive Cumulative Deformation

The forward process executes over T timesteps, where the CreepDeformationEngine numerically implements the simplified physics from Section 3. It computes incremental micro-deformations $\mathbf{d}_t^{\text{fw}}$ via three coupled sub-processes at each control point i (derivations provided in the preceding subsections):

(1) *Fiber Elongation*: From classical beam theory, bending strain is linearly proportional to the distance from the neutral axis. Under small deformation assumptions, this yields: $\Delta d_{i,y}^t = \frac{t}{T} \cdot A_\kappa \cdot (y_i - y_{\text{center}})$, where $A_\kappa = \kappa \cdot L$ integrates curvature and slip length, and t/T models progressive creep.

(2) *Force Balance*: Stress transfer follows the exponential decay $e^{-r/\lambda}$ from Eq. (4) of the main paper. For computational efficiency, we approximate this using a 3×3 Gaussian kernel $K(r)$ on the discrete grid, which accurately captures short-range interactions when $r/\lambda < 1$. Orientation evolution uses a tanh function to ensure coherent motion of adjacent control points.

(3) *Moisture-Stress&Heterogeneity*: From continuum mechanics, stress is proportional to strain, which is defined as the displacement gradient: $\sigma \propto \varepsilon = \nabla\mathbf{d}$. We implement the stress amplification $(1 + \gamma_i\sigma_i)$ from Eq. (8) of the main paper using $\|\nabla\mathbf{d}\|$ as a physically grounded stress proxy: $(1 + \gamma \cdot \|\nabla\mathbf{d}\|)$. Material heterogeneity is simulated by adding Gaussian noise $\mathcal{N}(0, 0.05 \cdot t/T)$.

This simplified, differentiable implementation retains the core physical mechanisms while avoiding coupled PDEs.

3.2.6. State Evolution Equations

At timestep t, the bamboo slip deformation state is completely described by the state tuple (x_t, D_t), where x_t represents the current deformed image and D_t denotes the cumulative control point displacement field. We define $\mathcal{W}(x, D)$ as the warping function that applies the dense displacement field D to image x via bilinear interpolation. State updates follow the iterative equations:

$$\begin{aligned} \mathbf{d}_t^{\text{fw}} &= \text{CreepEngine}(x_{t-1}, D_{t-1}), \\ D_t &= D_{t-1} + \mathbf{d}_t^{\text{fw}}, \\ x_t &= \mathcal{W}(x_0, D_t). \end{aligned} \tag{30}$$

This process starts from the initial state $(x_0, D_0 = \mathbf{0})$, generating a unique physical deformation trajectory that culminates in the fully deformed state (x_T, D_T).

3.3. Conclusion

The integration of physical deformation modeling with deep learning provides a powerful framework for simulating and reversing the deformation of bamboo slips. By encoding physical principles such as moisture-induced expan-

sion and stress transfer mechanisms into the data generation process through the CreepDeformationEngine, the network learns a physically plausible restoration process. The MSE loss, while simple in form, implicitly enforces physical constraints through the structure of the physics-derived ground truth $\mathbf{d}_t^{\text{true}}$. This approach transforms the restoration task from a black-box restoration problem into a physically grounded parameter prediction task, enabling more accurate and interpretable results for archaeological restoration.

4. Baseline Implementation Details

To ensure a rigorous and fair comparison, all models were evaluated on the same hardware environment (NVIDIA RTX 5090 GPU) and the same 200-image test set. For methods requiring training, we further ensured fairness by using identical data sources and preprocessing. All input images were standardized to 320×32 resolution, and all models were trained using the same 1,800 pristine slips (and their corresponding physical deformations, where applicable). We followed the official implementations for all baselines. For models we trained (DDRM, CycleGAN), hyperparameters were either matched to our own (for direct comparison) or set as recommended by their authors, as detailed below.

Models Trained on Our Paired Dataset

To create a direct, apples-to-apples comparison of supervised restoration algorithms, we trained the following model from scratch on our own paired dataset (the 1,800 pristine and synthetic-deformed pairs).

- **DDRM** [6]: To directly compare our physics-guided approach against a representative learning-based iterative restoration model, we trained the official DDRM implementation using the exact same hyperparameters as our own model. This includes 100 epochs, an AdamW optimizer with a learning rate of 1×10^{-4}, and a batch size of 8.

Models Trained with Unpaired Data

This model represents the alternative paradigm of training without paired data.

- **CycleGAN** [15]: We used the official implementation and trained it on our data, defined as two unpaired domains (the 1,800 pristine slips and the 1,800 deformed slips). To ensure the model reached full convergence as recommended by its authors , we trained it for 200 epochs using an Adam optimizer (lr=2×10^{-4}, $\beta_1 = 0.5, \beta_2 = 0.999$) and a batch size of 8 .

Pre-trained Foundational & Specialized Models

This category evaluates the generalization capability of SOTA models pre-trained on other large-scale datasets.

- **DewarpNet** [2]: As a SOTA method for document unwarping, we used the official pre-trained model (trained on the DocUNet dataset) and applied it directly to our test set.
- **Stable Diffusion** [12]: To assess the capability of large-scale foundational models, we used the pre-trained v1.5 model in its image-to-image (img2img) pipeline. The model was guided by the following text prompts:
  - Prompt: "a corrected, straight, flat bamboo slip, artifact photograph"
  - Negative Prompt: "warped, distorted, bent"

  The denoising strength was set to 0.75.

5. Control Point Grid Ablation

We conduct ablation experiments to validate our choice of the 32×4 control point grid configuration. The 4-column structure is fixed across all variants, as it directly corresponds to the anisotropic fiber bundle organization of bamboo material [3]. We vary the longitudinal resolution by testing 16×4, 32×4, and 64×4 grids to examine the trade-off between deformation modeling capacity and computational efficiency.

5.1. Experimental Setup

All variants are trained with identical physics parameters (curvature amplitude $A_\kappa = 0.15$, force coupling strength $\lambda = 0.3$, moisture diffusion coefficient $D_{moist} = 0.15$) for 100 epochs on the same dataset.

5.2. Results and Analysis

Table 1 presents the quantitative results. The 16×4 configuration is clearly suboptimal. Moving from 16 to 32 rows yields significant gains across all metrics, including a **+4.6%** increase in Straightness (0.283 $\rightarrow$ 0.296) and a **+6.6%** increase in physical plausibility (DCI: 7.45 $\rightarrow$ 7.94). This demonstrates that 16 longitudinal points are insufficient to capture the slip's complex curvature.

Conversely, increasing the resolution from 32×4 to 64×4 demonstrates **diminishing returns**. The gains are marginal: Straightness improves by only 0.7% (0.296 $\rightarrow$ 0.298) and DCI by 1.0% (7.94 $\rightarrow$ 8.02), while TRQ (text readability) slightly decreases. This performance plateau indicates that 32 longitudinal points are sufficient to model the deformation manifold. Given that the 64×4 grid incurs $\sim$35% longer training time for no commensurate benefit, our 32×4 configuration is validated as the optimal balance between modeling capacity and efficiency.

6. Physical Parameter Validation

To validate the rationality of our chosen physical parameters, we conduct a systematic sensitivity analysis following the methodology of distribution matching. We evaluate seven parameter configurations by computing the Fréchet Curve Distance (FCD) between synthetic deformed samples

Grid Resolution	Straightness↑	LPIPS↓	TRQ↑	DCI↑
16×4 (64 pts)	0.283 ± 0.0765	0.237 ± 0.0668	1.005 ± 0.1087	7.45 ± 2.2341
32×4 **(128 pts)**	0.296 ± 0.0737	0.232 ± 0.0641	$\mathbf{1.018} \pm 0.1061$	7.94 ± 2.1738
64×4 (256 pts)	$\mathbf{0.298} \pm 0.0729$	$\mathbf{0.231} \pm 0.0652$	1.016 ± 0.1073	$\mathbf{8.02} \pm 2.1895$

Table 1. Ablation study on control point grid resolution. The 32×4 configuration achieves the best balance between restoration quality (highest TRQ) and efficiency. Doubling the resolution to 64×4 yields only marginal gains ($<$1% in Straightness/DCI) at a $\sim$35% increase in training time. Optimal **trade-off** is 32×4.

and real archaeological deformed bamboo slips. The FCD measures the similarity between the centerline curves extracted from bamboo slip images, where lower values indicate closer alignment with real-world deformation patterns.

For each configuration, we extract centerline curves from 50 synthetic and 50 real deformed samples using B-spline interpolation with 64 sampling points, yielding approximately 2,000 pairwise FCD measurements. The distribution statistics provide a robust evaluation of how well each parameter set replicates authentic physical deformation mechanisms.

As shown in Table 2, our chosen parameters consistently outperform alternative configurations. The **fiber elongation factor** ($A_\kappa = 0.15$), defined in Eq. 9 of the main paper, governs the bending deformation magnitude. Values too low (0.05) produce insufficient bending (FCD=0.723), and values too high (0.30) create unrealistic over-deformation (FCD=0.714). The **force coupling strength** ($\lambda = 0.3$) balances the stress transfer mechanism in Eq. 4 of the main paper—weak coupling (0.1) causes unnatural local discontinuities (FCD=0.715), while strong coupling (0.6) oversmooths spatial heterogeneity (FCD=0.720). Finally, the **moisture diffusion coefficient** ($D_{moist} = 0.15$) replicates the gradual water content changes during burial; lower values (0.05, FCD=0.722) create abrupt moisture gradients, whereas higher values (0.30, FCD=0.723) fail to preserve the localized drying patterns observed in archaeological contexts.

The optimal configuration achieves a mean FCD of 0.707, demonstrating systematic optimization to reflect authentic bamboo slip deformation mechanisms rather than arbitrary parameter selection.

7. Physical Realism Validation

Before utilizing the synthetic data generated by the CreepDeformationEngine for training, it is crucial to validate its physical plausibility and consistency with real-world deformed bamboo slips. This validation is conducted from two perspectives: the quantitative consistency of physical parameters and the statistical consistency of feature space distributions.

Consistency of Physical Parameters. We extract key physical parameters—including boundary curvature κ, fiber bundle length difference Δx, and displacement gradient $\|\nabla \mathbf{d}\|$—from 500 real deformed bamboo slips. These are compared against synthetic data generated by the CreepDeformationEngine under the same optimal parameter configuration. Table 3 presents the statistical comparison of these parameter distributions. The curvature distribution of the synthetic data (mean 0.142, std 0.051) aligns closely with the real data (mean 0.138, std 0.048), and the ranges of fiber length variations substantially overlap. This indicates that our physical model, grounded in bamboo rheology, accurately captures the authentic deformation characteristics of real bamboo slips.

Consistency of Data Distribution. To further verify the distributional consistency between synthetic and real data in the feature space, we employ the Silhouette Score for quantitative evaluation. Specifically, we extract boundary curve features from both real and synthetic deformed slips, apply UMAP (Uniform Manifold Approximation and Projection) to reduce the high-dimensional features to a 2D space for visualization, and calculate the Silhouette Score s:

$$s(j) = \frac{b(j) - a(j)}{\max\{a(j), b(j)\}}, \tag{31}$$

where $a(j)$ is the mean distance between data point j and other points within the same cluster, and $b(j)$ is the mean distance between j and points in the nearest different cluster. The score ranges from $[-1, 1]$, where a value closer to 0 indicates that the two data distributions are highly overlapping and statistically indistinguishable.

Figure 1 illustrates the UMAP visualizations of data generated using optimal versus random physical parameters. When employing the optimal parameters ($A_\kappa = 0.15$, $\lambda = 0.3$, $D_{\text{moist}} = 0.15$), the synthetic data heavily overlaps with the real data in the feature space, yielding a Silhouette Score of $s = 0.003$. This near-zero score demonstrates that the two distributions are statistically indistinguishable. In contrast, data generated with random parameters clearly separates from the real data ($s = 0.430$), validating the necessity of our physical parameter optimization. Furthermore, compared to purely data-driven generation methods (e.g., GAN: $s = 0.410$, DDPM: $s = 0.429$), our physics-guided approach exhibits a significant advantage in data distribution fidelity.

Configuration	A_κ	λ	D_{moist}	Mean FCD↓	Median↓	Std Dev	Δ vs. Opt.
Ours (Optimal)	**0.15**	**0.3**	**0.15**	**0.707** $\pm$ 0.278	**0.748**	0.278	–
Low Fiber	0.05	0.3	0.15	0.723 $\pm$ 0.280	0.800	0.280	+2.2%
High Fiber	0.30	0.3	0.15	0.714 $\pm$ 0.279	0.800	0.279	+1.0%
Weak Coupling	0.15	0.1	0.15	0.715 $\pm$ 0.275	0.755	0.275	+1.1%
Strong Coupling	0.15	0.6	0.15	0.720 $\pm$ 0.280	0.791	0.280	+1.8%
Low Diffusion	0.15	0.3	0.05	0.722 $\pm$ 0.276	0.800	0.276	+2.1%
High Diffusion	0.15	0.3	0.30	0.723 $\pm$ 0.270	0.769	0.270	+2.2%

Table 2. Parameter sensitivity analysis using Fréchet Curve Distance (FCD). A_κ: curvature amplitude parameter, λ: force coupling strength , D_{moist}: moisture diffusion coefficient. Lower FCD indicates better alignment with real deformation patterns. Our optimal parameters achieve the best performance across both mean and median metrics, with up to 2.2% improvement over suboptimal configurations. Best results are highlighted in **bold**.

Parameter	Real Data	Synthetic Data	Relative Error
Mean Curvature μ_κ	0.138	0.142	2.9%
Curvature Std. Dev. σ_κ	0.048	0.051	6.3%
Max Displacement (pixels)	24.3	23.7	2.5%
Displacement Gradient $\|\nabla \mathbf{d}\|$	0.087	0.091	4.6%

Table 3. Comparison of physical parameters between real deformed bamboo slips and synthetic data.

These validation results confirm that the deformation model established by the CreepDeformationEngine, rooted in the physical properties of bamboo, successfully generates synthetic training data that is highly consistent with real-world deformations in both physical parameters and feature distributions. This provides a reliable guarantee of physical plausibility for the subsequent network training, experimentally verifying the superiority of "physics-guided data generation" over purely data-driven generative approaches.

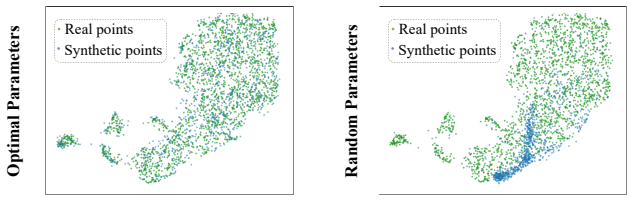

Figure 1. UMAP between real and synthetic data distribution.

8. Supplementary Qualitative Results

This section provides additional qualitative results to supplement the analysis in the main paper. Figure 2 presents four more restoration examples to further visually validate the performance of our proposed method against the baselines.

From these comparisons, it can be more clearly observed that baseline methods lacking physical priors produce various specific restoration deficiencies when faced with real, complex bamboo slip deformations. For example, the results from **CycleGAN (e)** exhibit the most severe distortion; not only does it fail to accurately correct the curvature, but it also generates numerous non-existent artifacts and spots on the surface, severely damaging the original text strokes. The surface texture generated by **Stable Diffusion (d)** is entirely unrealistic, looking more like a digital 'painting' than a restoration of a real artifact. **DewarpNet (c)**, specifically designed for document image unwarping, is unable to handle the non-uniform, complex warping with minor undulations characteristic of bamboo slips, resulting in unnatural edge lines after restoration. **DDRM (b)**, the method most similar to our own in terms of iterative processing, successfully straightens the slip through progressive refinement, but its purely data-driven approach treats fine textural details as reconstruction artifacts. While correcting the geometry, it simultaneously degrades this critical high-frequency information, causing the text to become blurry and illegible.

In contrast, **our method (f)** demonstrates a clear advantage across all comparison groups, which stems from its core physics-guided idea. Because our forward process simulates real physical deformation and the inverse process learns to 'undo' these physical steps, the model effectively models the material properties and deformation mechanisms. As can be seen from the results, the bamboo slips processed by our model achieve correction in

their macroscopic geometry, and the deformation is partially recovered. Although a complete restoration is not yet achieved at this stage, this work represents a key step from models that completely ignore physical structure, to one that can learn and adhere to certain physical properties. We will further optimize the model to improve restoration performance in future work.

References

[1] Partha Bhowmick and Bhargab B Bhattacharya. Fast polygonal approximation of digital curves using relaxed straightness properties. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29 (9):1590–1602, 2007. 1

[2] Sagnik Das, Ke Ma, Zhixin Shu, Dimitris Samaras, and Roy Shilkrot. Dewarpnet: Single-image document unwarping with stacked 3d and 2d regression networks. *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 131–140, 2019. 5

[3] Patrick G Dixon and Lorna J Gibson. The structure and mechanics of moso bamboo material. *Journal of the Royal Society Interface*, 11(99), 2014. 5

[4] William N Findley and Francis A Davis. *Creep and relaxation of nonlinear viscoelastic materials*. Courier corporation, 2013. 2

[5] Robert M Jones. *Mechanics of composite materials*. CRC press, 2018. 2

[6] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Proc. Adv. Neural Inform. Process. Syst.*, 35:23593–23606, 2022. 5

[7] Deepak Kumar and AG Ramakrishnan. Quad: Quality assessment of documents. *CBDAR*, pages 79–84, 2011. 1

[8] Leon Mishnaevsky Jr and Hai Qing. Micromechanical modelling of mechanical behaviour and strength of wood: State-of-the-art review. *Computational Materials Science*, 44(2): 363–370, 2008. 2

[9] Parviz Navi and Stefanie Stanzl-Tschegg. Micromechanics of creep and relaxation of wood. a review. *Holzforschung*, 63:186–195, 2009. 2, 3

[10] Raviduth Ramful and Atsushi Sakuma. Investigation of the effect of inhomogeneous material on the fracture mechanisms of bamboo by finite element method. *Materials*, 13 (21):5039, 2020. 3

[11] Junuthula Narasimha Reddy. *Energy principles and variational methods in applied mechanics*. John Wiley & Sons, 2017. 3

[12] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 10684–10695, 2022. 5

[13] Peng Ye and David Doermann. Document image quality assessment: A brief survey. *Proceedings of International Conference on Document Analysis and Recognition*, pages 723–727, 2013. 1

[14] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 586–595, 2018. 1

[15] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 2223–2232, 2017. 5

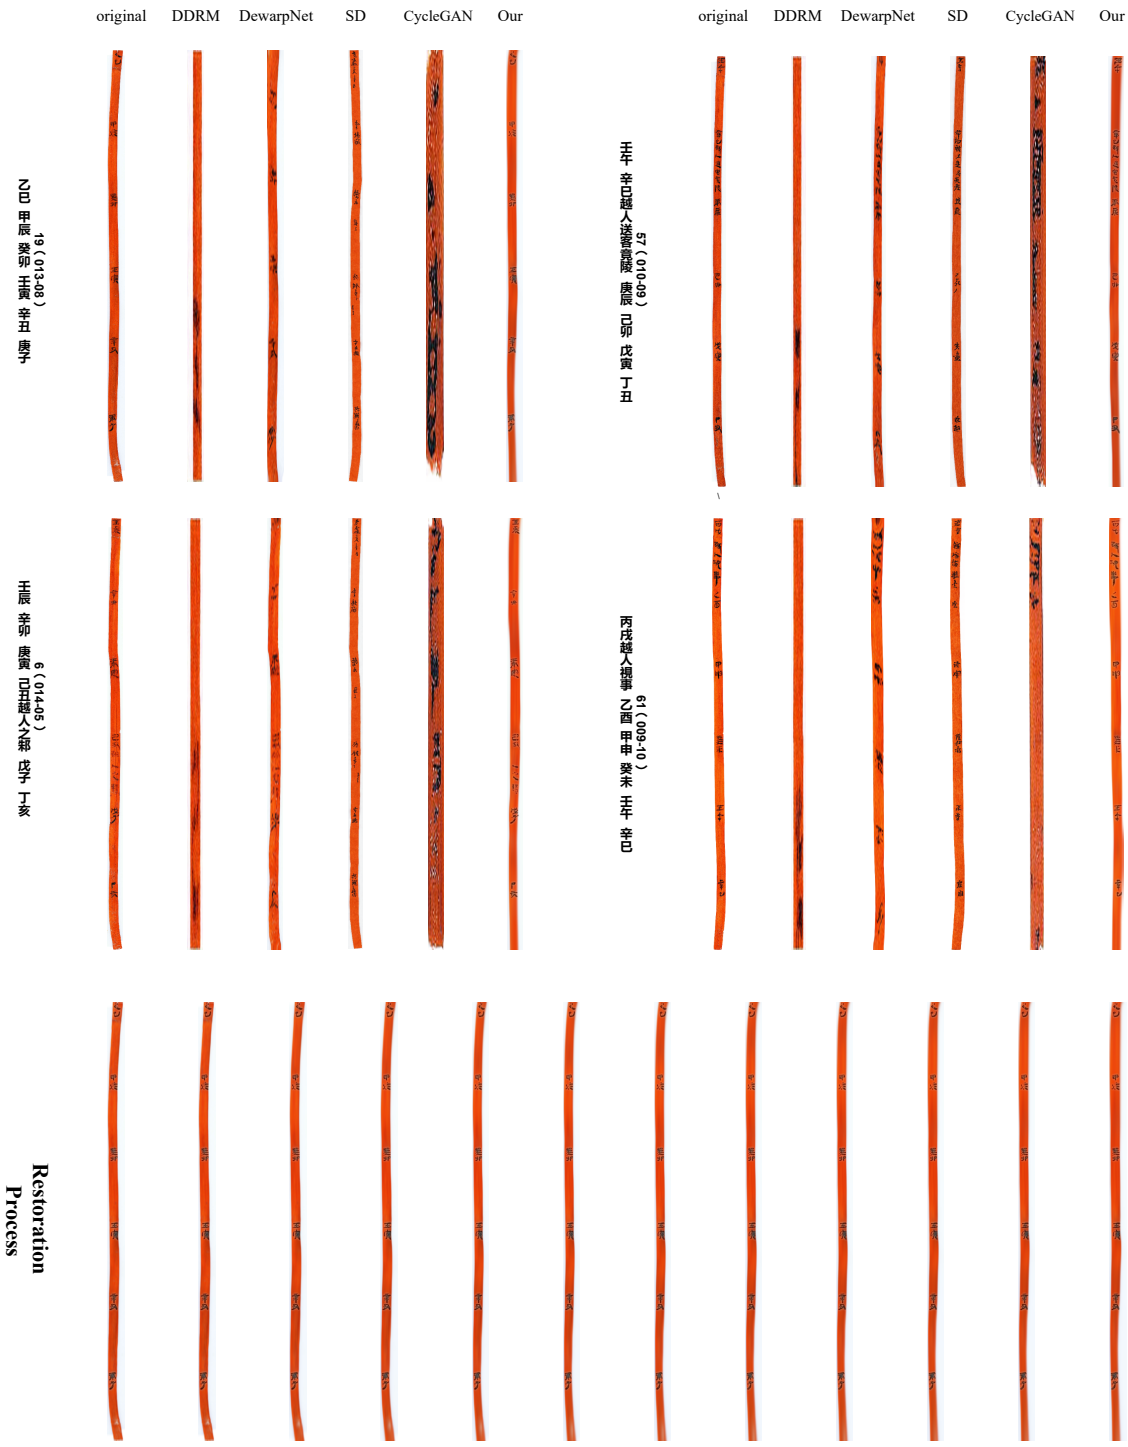

Figure 2. Supplementary qualitative comparison results. This figure presents four additional groups of restoration comparisons on different bamboo slips. **For each comparison group (upper rows):** we display the original deformed slip (a), alongside the restoration results from baseline methods including DDRM (b), DewarpNet (c), Stable Diffusion (SD) (d), and CycleGAN (e), compared with our final result (f). **Iterative Process Visualization (bottom row):** To further illustrate our method's mechanism, the bottom row showcases the step-by-step restoration process for one of the examples, visualizing key intermediate states as our model progressively reverses the physical deformation. These examples further highlight our method's improved ability in both geometric correction and preservation of fine textual details.