

NeuralGarSim: Geometry-Agnostic Garment Simulation with Neural Fields

Arihant Gaur^{1,2} Navami Kairanda¹ Christian Theobalt¹ Vladislav Golyanik¹

¹MPI for Informatics, SIC ²Saarland University, SIC

Abstract. Most existing mesh-based methods for garment simulation suffer from sensitivities tied to mesh discretisation and resolution. Recent approaches such as NeuralClothSim employ continuous neural fields as a promising alternative, but remain constrained by a 2D curvilinear parameterisation that models a single cloth panel, limiting their ability to represent garments with multiple panels stitched together. This paper introduces **NeuralGarSim**, a quasistatic garment simulator formulated directly in 3D Euclidean space. Our method accepts diverse input representations—including distance fields, meshes, point clouds, and Gaussians—as undeformed garment states, making it compatible with off-the-shelf reconstruction frameworks without additional preprocessing. We then represent the garment deformation as a neural field and define a nonlinear Kirchhoff–Love shell model directly in $\mathbb{R}^3$ by applying tangential differential calculus on the garment surface. This ensures that the neural garment simulation is parameterisation-space-free and naturally extends to arbitrary topologies. By minimising a potential energy functional, our method learns a 3D neural deformation field that predicts physically consistent deformations across garments with multiple panels, seams, and holes. It supports multiple material models and trains $2\text{--}4\times$ faster than its predecessor, NeuralClothSim, while retaining the continuous, consistent, and memory-adaptive behaviour¹.

1 Introduction

Physics-based cloth simulation is central to modern computer graphics, underpinning applications ranging from visual effects and games to digital fashion design. Most existing systems rely on explicit and discrete representations, such as finite differences or finite elements on triangle meshes and particle systems [28, 46, 47, 51, 53, 84]. Despite their maturity, these discretisations impose fundamental limitations: the predicted cloth dynamics remain highly sensitive to mesh resolution and initial triangulation, often producing mesh-dependent wrinkle patterns [38, 91]. Because meshes are not inherently adaptive, simulators must either remesh—breaking the differentiability crucial for inverse problems [53]—or rerun the simulation on a denser mesh whenever a higher target resolution is desired, inflating both compute and memory budgets.

¹ Project page: <https://4dqv.mpi-inf.mpg.de/NeuralGarSim/>




Fig. 1: NeuralGarSim is the first neural garment simulator to operate directly on the 3D garment surface without a 2D parameterisation, enabling continuous simulation of complex garments. (a): The reference garment surface can be provided as any one of UDF, triangle mesh, point cloud, or Gaussians—each supplying only the normal field and curvature tensor to the simulator. **(b):** NeuralGarSim produces physically consistent drape and fine-scale wrinkles across garments of varying topology **(c):** Wrinkle patterns remain consistent across different mesh resolutions (top); independently, swapping material parameters yields perceptually distinct drape characteristics (bottom).

Neural fields have emerged as a compelling paradigm for compact 3D scene representation [80, 89] and surface modelling [61, 86], typically realised as coordinate-based MLPs. Bridging these continuous representations with cloth simulation is therefore a highly promising direction. NeuralClothSim [38] is a notable step along this path. It trains a neural deformation field using the quasi-static² physics of a parametric Kirchhoff–Love shell [57]. The method minimises potential energy via analytic derivatives of the neural field, sidestepping the resolution dependence of traditional finite elements. However, it relies on classical parametric surface differential geometry [3, 88], *i.e.*, on the existence of a well-defined 2D parametric domain³, which severely limits the complexity of garments it can handle.

The dependence on 2D parameterisation is particularly restrictive. Real garments comprise multiple panels, seams, and holes, making it difficult to construct a continuous parametric map whose derivatives remain analytic across boundaries. The classical finite element method (FEM) can cope with such discontinuities thanks to explicit meshes and local basis functions [65], but neural fields require global continuity and differentiability to evaluate energies via analytic derivatives [33, 41]. Hence, NeuralClothSim is limited to single-panel settings where suitable planar or periodic parameterisations are available, and extending neural cloth simulation to realistic garments demands a formulation that

² quasi-static denotes predicting the final equilibrium configuration under applied forces or boundary constraints.

³ by parametric domain, we mean the flat cloth panel that serves as the 2D input to the neural deformation field in Kairanda et al. [38], not the network parameters.

dispenses with 2D parameter domains. Furthermore, prior neural and learning-based simulators [38, 69] are effectively tied to a single input representation, requiring a high-quality mesh as input. This prevents direct use with point clouds, 3D Gaussians [42] or other 3D representations from off-the-shelf reconstruction pipelines [64, 87].

We address these limitations with NeuralGarSim, a new method that reformulates cloth simulation entirely in 3D. The key enabler of our approach is *Tangential Differential Calculus (TDC)* [17, 73], a mathematical framework for defining differential operators directly on embedded surfaces and shapes without any local 2D parameterisations. By casting the nonlinear Kirchhoff–Love shell model into TDC, all strain quantities are computed directly in $\mathbb{R}^3$ from differential quantities of a neural deformation field, which has two immediate consequences. First, NeuralGarSim handles garments of arbitrary topology—multiple panels, seams and holes—because there is no need for a parametric map introducing discontinuities at boundaries. Second, the simulation only requires a surface normal field and curvature tensor at each sample point. Therefore, any 3D representation that can supply these quantities—distance fields, meshes, point clouds, or Gaussian splats—integrates directly without UV unwrapping or topology-specific preprocessing. This enables NeuralGarSim to accept garment assets from off-the-shelf reconstruction pipelines [42, 87] that prior neural simulators cannot utilise. To summarise, this paper advances the state of the art with the following **technical contributions**:

- We introduce NeuralGarSim, a parameterisation-free and continuous neural quasi-static cloth simulator operating directly on embedded surfaces, avoiding the need for a 2D parametric domain (Sec. 4).
- We cast Kirchhoff–Love thin shells into the TDC formalism, enabling seamless support for multiple geometric representations under a unified physical model (Sec. 4.1).

In experiments, NeuralGarSim produces discretisation-consistent wrinkle patterns across different mesh resolutions, supports multiple material models, and trains 2–4× faster than NeuralClothSim (Sec. 5). Fig. 1 highlights representative results. Our project page provides further visualisations, video and the code.

2 Related Work

Classical Cloth Simulation has matured over several decades [1, 2, 9, 25, 34, 48, 79, 83], converging on a common recipe: discretise the cloth with FEM [20, 65], timestep implicitly for stability [1, 46], and resolve contact and collisions robustly [47, 60, 67, 78]. Constitutive models capture fabric-dependent stress response [13, 85], while recent differentiable solvers [51, 53] expose gradients that aid tasks such as reconstruction [40, 50]. Despite steady progress, FEM methods [53] remain sensitive to mesh discretisation, require extremely fine meshes for fidelity, and are expensive to run or integrate with modern learning systems, motivating alternative formulations relying on neural fields [38].

Neural Cloth Simulation. Another line of work replaces traditional solvers with neural models [6, 7, 11, 22, 23, 54, 71]. PBNS [6] tackles static draping on virtual characters through self-supervised losses that enforce physics without simulation ground truth, and later variants extend the idea to garment dynamics [7, 71]. The garments in these methods are skinned to the human body, and deformations are driven primarily by body shape and pose. While effective for body-driven apparel, these networks cannot serve as general cloth simulators and inherit the discretisation limitations of classical mesh-based simulators. Pfaff *et al.* [69] instead learn general-purpose dynamics with graph neural networks but rely on synthetic training data [65]; HOOD [23] and ContourCraft [22] retain the same discrete representation while incorporating physics-based supervision. NeuralClothSim [37, 38] departs from meshes by representing deformation as a continuous neural field and evaluating energy analytically, thereby mitigating discretisation sensitivity. Yet it requires a 2D parametric domain: single panels are supported, but seams, cuts, and multi-panel garments remain out of reach. This limitation motivates our parameterisation-free formulation. Thin-Shell-SfT [37] builds on the same neural thin-shell formulation for monocular non-rigid surface tracking, whereas ∂^∞ -Grid [39] accelerates neural differential-equation solving with differentiable feature grids. While these works address tracking and solver efficiency, we focus on parameterisation-free garment simulation across meshes, UDFs, point clouds, and 3D Gaussians.

3D Surface Representations for Garments. Neural implicit fields represent surfaces as zero-level sets of learned continuous functions [80, 89]. They have been adopted for geometry representation [12, 21, 56, 77], reconstruction [52, 86, 87], and geometric processing [61, 66, 90]. At the same time, explicit 3D representations such as 3D Gaussian splatting [42] and surface-aligned variants [26] are employed for dynamic garment modelling, for instance, MD-Splatting [19] and Thin-Shell-SfT [37]. GaussianGarments approach [70] further reconstructs simulation-ready clothing from multi-view video. In garments, neural and implicit fields have mainly supported representation, data-driven draping, and disentanglement: ULNeF [72] targets virtual try-on, Implicit Untangling [10] resolves garment-layer intersections, and DIG [49] and DrapeNet [15] model body-conditioned draping or garment generation. Compared to these methods, differentiable unsigned distance fields (UDFs) [21] improve open-surface regularity and make normals and curvature more reliable. Our contribution is complementary and multimodal: NeuralGarSim uses the differential quantities supplied by UDFs, meshes, point clouds, or 3D Gaussians as the simulation domain for a physics-based Kirchhoff–Love shell model.

Tangential Differential Calculus for Shells. TDC provides a coordinate-free framework for defining differential operators on embedded surfaces [16, 17], enabling shell and plate models [16, 17, 30–32, 74] to be formulated directly in $\mathbb{R}^3$ space. Prior work has applied TDC to Kirchhoff–Love [73] and related shell models, albeit with linear strain assumptions and non-neural classical FEM discretisations. *Our work is the first to couple TDC with neural fields, enabling*

fully nonlinear Kirchhoff-Love (KL) shell simulation directly in 3D space and avoiding the parametric domains required by earlier KL cloth simulators [13,38,81].

3 Preliminaries

We model garments as thin shells governed by the Kirchhoff-Love (KL) theory. We summarise the KL formulation (Sec. 3.1), review NeuralClothSim [38] as the closest prior work (Sec. 3.2), and outline Tangential Differential Calculus (TDC) [17,73] background for the differential geometry of surfaces (Sec. 3.3).

3.1 The Kirchhoff-Love Shell Theory

A *thin shell* is a volumetric object whose thickness h is much smaller than its characteristic in-plane dimension L (i.e., $h \ll L$). Under KL theory, 3D continuum mechanics is reduced to 2D kinematics posed on the *mid-surface* $\Gamma \subset \mathbb{R}^3$. Our formulation for shell deformation relies on the classical KL hypothesis: Straight lines that are initially normal to the mid-surface remain straight and normal to the mid-surface; they preserve their length during deformation, so the shell thickness stays constant.

Any point $\mathbf{x}_\Omega$ within the shell Ω is specified by its mid-surface projection $\mathbf{x}_\Gamma$ and a normal offset $\zeta \in [-h/2, h/2]$, as $\mathbf{x}_\Omega = \mathbf{x}_\Gamma + \zeta \mathbf{n}_\Gamma$ (see Fig. 2).

Under the KL hypothesis, the 3D displacement field $\mathbf{u}_\Omega$ decomposes linearly through the thickness (thus eliminating any transverse shear) during shell deformation:

$$\mathbf{u}_\Omega(\mathbf{x}_\Gamma, \zeta) = \mathbf{u}(\mathbf{x}_\Gamma) + \zeta \mathbf{w}(\mathbf{x}_\Gamma), \quad (1)$$

where $\mathbf{u} := \mathbf{x}_{\Gamma'} - \mathbf{x}_\Gamma$ is the mid-surface displacement with the deformed midsurface Γ' , and $\mathbf{w} := \mathbf{n}_{\Gamma'} - \mathbf{n}_\Gamma$ captures the change in the surface normal orientation.

3.2 NeuralClothSim [38]

NeuralClothSim parameterises the mid-surface via a 2D *curvilinear parametric domain* $\tau \subset \mathbb{R}^2$; see Fig. 2-(left). It then maps the coordinate samples $\xi =$




Fig. 2: Explicit vs implicit thin-shell modelling. (a): NeuralClothSim [38] parameterises the mid-surface via an explicit map $\mathbf{x}_\Gamma(\xi) : \tau \subset \mathbb{R}^2 \rightarrow \Gamma$. (b): NeuralGarSim represents instead Γ as a zero level-set of an implicit field ϕ . Under the Kirchhoff-Love thin-shell assumption, a shell point $\mathbf{x}_\Omega$ is specified by its mid-surface anchor $\mathbf{x}_\Gamma$, thickness coordinate ζ , and normal $\mathbf{n}_\Gamma$, providing the kinematic basis for formulating the deformation field $\mathbf{u}_\Omega$.

$(\xi^1, \xi^2) \in \tau$ to 3D deformation $\mathbf{u}(\boldsymbol{\xi}; \Theta)$ using a neural field implemented as a coordinate-based multilayer perceptron (MLP) with weights Θ . The deformed midsurface $\mathbf{x}_{\Gamma'}(\boldsymbol{\xi})$ reads as

$$\mathbf{x}_{\Gamma'}(\boldsymbol{\xi}) = \mathbf{x}_{\Gamma}(\boldsymbol{\xi}) + \mathbf{u}(\boldsymbol{\xi}; \Theta), \quad (2)$$

where $\mathbf{x}_{\Gamma}(\boldsymbol{\xi})$ is an analytical surface defined on the parametric domain.

Parametric Energy Minimisation. The internal hyperelastic energy $\Psi(\boldsymbol{\epsilon}, \boldsymbol{\kappa}; \boldsymbol{\Phi})$ depends on membrane strain $\boldsymbol{\epsilon}$, bending strain $\boldsymbol{\kappa}$, and material parameters $\boldsymbol{\Phi}$. NeuralClothSim evaluates $\boldsymbol{\epsilon}(\boldsymbol{\xi}; \Theta)$ and $\boldsymbol{\kappa}(\boldsymbol{\xi}; \Theta)$ via classical parametric differential geometry [3, 75, 76], *i.e.*, first and second fundamental forms of both reference and deformed surfaces are constructed from the partial derivatives of $\mathbf{x}_{\Gamma}(\boldsymbol{\xi})$ and $\mathbf{u}(\boldsymbol{\xi}; \Theta)$ with respect to ξ^α . The equilibrium configuration is found by minimising the total potential energy over τ :

$$\min_{\Theta} \int_{\tau} \Psi(\boldsymbol{\epsilon}(\boldsymbol{\xi}; \Theta), \boldsymbol{\kappa}(\boldsymbol{\xi}; \Theta); \boldsymbol{\Phi}) d\boldsymbol{\xi} - \int_{\tau} \mathbf{f} \cdot \mathbf{u}(\boldsymbol{\xi}; \Theta) d\boldsymbol{\xi}. \quad (3)$$

Here, $\boldsymbol{\Phi}$ denotes fixed material parameters, while optimisation is performed only over the deformation-field weights Θ . We refer to Appendix C for the explicit expressions. The parametric formulation requires a continuous 2D domain, which is difficult to obtain for multi-panel garments with seams and holes without introducing chart boundaries or stitching artifacts; Sec. 5.6 empirically studies this limitation.

3.3 Tangential Differential Calculus (TDC)

TDC [17, 29, 32, 73] is a comprehensive mathematical framework for defining differential operators directly on embedded manifolds $\Gamma \subset \mathbb{R}^3$, circumventing the need for curvilinear coordinates, metric tensors, and Christoffel symbols required by classical parametric differential geometry [43, 58]. Below, we outline three operators that underpin the strain computation of our method.

Projection Operator. Suppose $\mathbf{P}$ is the orthogonal *projector* onto the tangent space $T_{\mathbf{x}}\Gamma$ at point $\mathbf{x}$, it reads as

$$\mathbf{P} = \mathbb{I} - \mathbf{n}_{\Gamma} \otimes \mathbf{n}_{\Gamma}, \quad (4)$$

where $\mathbf{n}_{\Gamma}$ is the surface normal, $\mathbb{I}$ is the 3×3 identity tensor, and “ $\otimes$ ” denotes the outer vector product. Note that $\mathbf{P}$ projects any vector $\mathbf{v}$ onto $T_{\mathbf{x}}\Gamma$ so that $\mathbf{v}_t = \mathbf{P} \cdot \mathbf{v}$ is tangential (*i.e.*, $\mathbf{v}_t \cdot \mathbf{n}_{\Gamma} = 0$).

Tangential and Directional Gradients. For a scalar function $v : \Gamma \rightarrow \mathbb{R}$, its *tangential gradient* $\nabla_{\Gamma} v$ is given by the surface projection of its standard 3D gradient:

$$\nabla_{\Gamma} v(\mathbf{x}) = \mathbf{P}(\mathbf{x}) \cdot \nabla \tilde{v}(\mathbf{x}), \quad \mathbf{x} \in \Gamma, \quad (5)$$

where $\tilde{v}$ is any smooth extension of v into a neighbourhood of Γ . The tangential gradient, thus, lies in the tangent space: $\nabla_{\Gamma} v(\mathbf{x}) \in T_{\mathbf{x}}\Gamma$. For a vector field,




Fig. 3: Overview of NeuralGarSim. **Left (Sec. 4.1):** The input garment surface Γ is provided in one of four modalities—UDF, point cloud, mesh or 3D Gaussians. We initially fit UDF to an MLP $\phi(\mathbf{x}; \Theta_\phi)$ with weights Θ_ϕ . The derivatives of the UDF supply geometric priors, i.e., normals $\mathbf{n}_\Gamma$ and the Weingarten map $\mathbf{H}$. For all other modalities, these priors are extracted directly from the input geometry. **Right (Sec. 4.2):** The neural deformation field $\mathbf{u}(\mathbf{x}_\Gamma; \Theta_u)$ with weights Θ_u applies boundary conditions $\mathcal{B}(\mathbf{x}_\Gamma)$ and combines with the geometric priors to compute membrane ϵ_M and bending strains ϵ_B (Eq. (14)) directly in $\mathbb{R}^3$. Minimising the total potential energy Π (Eq. (15)) optimises Θ_u , balancing elastic terms against external forces $\mathbf{f}$ to recover the deformed surface $\mathbf{x}_{\Gamma'} = \mathbf{x}_\Gamma + \mathbf{u}$.

the *directional gradient* is formed by applying tangential gradients component-wise; tensor-valued extensions are given in App. A. In our strain formulation, the resulting directional membrane and bending strains are projected onto the tangent plane using $\mathbf{P}$.

Weingarten Map and Curvature. Surface curvature is captured by the *Weingarten map* $\mathbf{H}$, a symmetric in-plane tensor written as:

$$\mathbf{H} = \nabla_\Gamma^{\text{dir}} \mathbf{n}_\Gamma, \quad (6)$$

where $\nabla_\Gamma^{\text{dir}}$ denotes the directional surface gradient. $\mathbf{H}$ is required for bending strain computation in our method.

4 Our NeuralGarSim Approach

NeuralGarSim is a quasistatic neural cloth simulator that operates directly on reference geometries embedded in $\mathbb{R}^3$; see Fig. 3. Given an input garment $\mathcal{M}_\Gamma$ approximating the reference surface $\Gamma \subset \mathbb{R}^3$, we seek a neural deformation field (NDF) $\mathbf{u}(\cdot; \Theta_u)$ with weights Θ_u that produces the simulated surface Γ' under updated boundary conditions $\mathcal{B}$ and external forces $\mathbf{f}$. The method first extracts geometric operators ($\mathbf{n}_\Gamma, \mathbf{P}, \mathbf{H}$) from the input representation (Sec. 4.1) and learns $\mathbf{u}(\cdot; \Theta_u)$ by minimising the total potential energy under the prescribed boundaries and forces (Sec. 4.2). TDC provides the coordinate-free bridge to Kirchhoff–Love shell mechanics: it does not require 2D parameterisation, but when a smooth parameterisation exists, the same geometric quantities are recovered.

4.1 Garment Representation

We take the undeformed garment surface Γ as input. The surface can be represented in different modalities $\mathcal{M}_\Gamma$, such as discrete meshes, point clouds, 3D Gaussians, or implicit level sets. We prefer level-set representations, as the TDC-based formulation is well-explored [32, 73] for implicits. Our simulation stage (see Sec. 4.2) requires *only* a normal field $\mathbf{n}_\Gamma$ and a Weingarten map $\mathbf{H}$ at each surface point. These are extracted differently for each input modality.

Garment as a Neural Level Set. Garments are *open* surfaces with complex seams and holes; we represent the undeformed garment surface Γ implicitly as the zero level of a differentiable UDF (DUDF) $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}^+$ [21]:

$$\Gamma = \{\mathbf{x} \in \mathbb{R}^3 : \phi(\mathbf{x}) = 0\}, \quad \Omega = \left\{\mathbf{x} \in \mathbb{R}^3 : |\phi(\mathbf{x})| \leq \frac{h}{2}\right\}. \quad (7)$$

ϕ must be differentiable to compute higher-order differential geometric quantities. Therefore, we parametrise the distance field with an MLP $\phi(\mathbf{x}; \Theta_\phi)$ with weights Θ_ϕ and supervise it using the modified Eikonal problem of DUDF [21], which ensures that the distance field is differentiable with *hyperbolic scaling* (details in the Appendix; see Sec. E). With the learned $\phi(\mathbf{x}; \Theta_\phi)$, the unit surface normal $\mathbf{n}_\Gamma$ follows from the field gradient, from which the TDC projection operator $\mathbf{P}$ and Weingarten map $\mathbf{H}$ follow via Eqs. (4) and (6).

Garment as a Discrete Representation. If the garment is provided as a *triangle mesh*, we compute $\mathbf{n}_\Gamma$ as the area-weighted vertex normal, and $\mathbf{H}$ is derived via the cotangent-weighted Laplacian [62]. For a *point cloud*, $\mathbf{n}_\Gamma$ is estimated by principal component analysis (PCA) on a local neighbourhood [36]. $\mathbf{n}_\Gamma$ is set as the axis of smallest variance extracted from the covariance matrix of the 3D Gaussian. For sufficiently dense meshes, point clouds, and 3D Gaussians, $\mathbf{H}$ can be set to zero without affecting the thin-shell energy approximation.

4.2 Garment Deformation

We model garment deformation with an NDF $\mathbf{u}(\mathbf{x}_\Gamma; \Theta_\mathbf{u}) : \Gamma \rightarrow \mathbb{R}^3$ that maps each reference point $\mathbf{x}_\Gamma$ to its displaced location $\mathbf{x}_{\Gamma'}$ in 3D. The field is parametrised by a coordinate-based MLP $\mathcal{F}_{\Theta_\mathbf{u}}$ with Siren activations [77], whose smooth periodic activations support the higher-order derivatives required by the bending strain. Unlike NeuralClothSim, which operates on parametric coordinates (Sec. 3.2), our NDF $\mathbf{u}$ takes the 3D coordinates on Γ as input, eliminating any dependency on 2D parameterisation. Dirichlet (pinning) conditions can be imposed as hard constraints by modulating the NDF output, as in Refs. [55, 59]. In particular, for a constrained boundary subset $\partial\Gamma \subset \Gamma$, we introduce a scalar distance function $\mathcal{B}(\mathbf{x})$ that is zero on $\partial\Gamma$, and one elsewhere:

$$\mathbf{u}(\mathbf{x}_\Gamma; \Theta_\mathbf{u}) = \mathcal{F}_{\Theta_\mathbf{u}}(\mathbf{x}_\Gamma) \cdot \mathcal{B}(\mathbf{x}_\Gamma). \quad (8)$$

Boundary conditions are, therefore, specified directly in $\mathbb{R}^3$, avoiding the 2D points required by parametric formulations [38]. Moreover, periodic boundary

conditions as in NeuralClothSim are not required, as the surface is represented implicitly. Once $\mathbf{u}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}})$ is trained, inference can be performed at any spatial resolution supported by the network capacity of the NDF. Next, we describe how to train the NDF to generate quasistatic cloth deformation. Detailed algorithms for NDF training and inference are provided in Algs. 1 and 2 in the Appendix.

Strain Formulation. The core of our method is a nonlinear KL shell model formulated entirely in $\mathbb{R}^3$ using TDC, following classical shell mechanics and TDC-based shell formulations [32, 35, 73, 75]. Following the KL kinematic assumption (1), the 3D displacement vector $\mathbf{u}_\Omega$ at any point in the shell continuum is thus:

$$\mathbf{u}_\Omega(\mathbf{x}_\Gamma, \zeta; \Theta_{\mathbf{u}}) = \mathbf{u}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}}) + \zeta \mathbf{w}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}}), \quad (9)$$

where $\mathbf{u}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}})$ is the midsurface deformation and $\mathbf{w}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}})$ is the nonlinear director difference vector encoding orientation changes. For brevity, we drop the explicit dependence on $\mathbf{x}_\Gamma, \zeta$, and $\Theta_{\mathbf{u}}$ for further quantities in what follows. Because the KL model eliminates transverse shear, the director difference can be written as:

$$\mathbf{w} = \mathbf{H} \cdot \mathbf{u} - \nabla_\Gamma(\mathbf{u} \cdot \mathbf{n}_\Gamma). \quad (10)$$

To capture the large deformations and rotations happening in cloth, we employ the full Green-Lagrange strain tensor ϵ^{dir} , defined from the 3D directional deformation gradient $\nabla_\Gamma^{\text{dir}} \mathbf{u}_\Omega$ [35]. Neglecting the $\mathcal{O}(\zeta^2)$ term yields the standard split into membrane and bending strains:

$$\epsilon^{\text{dir}} \approx \epsilon_M^{\text{dir}} + \zeta \epsilon_B^{\text{dir}}, \quad (11)$$

where the *membrane strain tensor* capturing stretching reads

$$\epsilon_M^{\text{dir}} = \frac{1}{2} (\mathbf{F}_u + \mathbf{F}_u^\top + \mathbf{F}_u^\top \mathbf{F}_u), \quad (12)$$

and the *bending strain tensor* capturing curvature changes reads

$$\epsilon_B^{\text{dir}} = \frac{1}{2} (\mathbf{F}_w + \mathbf{F}_w^\top + \mathbf{F}_u^\top \mathbf{F}_w + \mathbf{F}_w^\top \mathbf{F}_u), \quad (13)$$

with $\mathbf{F}_u = \nabla_\Gamma^{\text{dir}} \mathbf{u}$ and $\mathbf{F}_w = \nabla_\Gamma^{\text{dir}} \mathbf{w}$ denoting the directional gradients of the midsurface displacement and orientation, respectively. The Green-Lagrange expansion leading to (11) is detailed in App. B. *Both tensors remain nonlinear due to the quadratic terms.* As they are defined in $\mathbb{R}^{3 \times 3}$, we project them onto the tangent plane using $\mathbf{P}$ to arrive at the in-plane strains:

$$\epsilon_M = \mathbf{P} \cdot \epsilon_M^{\text{dir}} \cdot \mathbf{P} \quad \text{and} \quad \epsilon_B = \mathbf{P} \cdot \epsilon_B^{\text{dir}} \cdot \mathbf{P}. \quad (14)$$

In the above formulation, the bending strain ϵ_B relies on the second-order derivatives of $\mathbf{u}$. This formulation aligns with our network’s continuity (C^∞), while remaining fully non-linear and parameterisation-free.

4.3 Cloth Material Models and Energy Minimisation

We find the equilibrium configuration of cloth encoded in $\Theta_{\mathbf{u}}^*$ by minimising the total potential energy $\Pi = \Pi_{\text{int}} + \Pi_{\text{ext}}$ directly on the 3D reference manifold Γ . Here, the internal elastic energy Π_{int} depends on the chosen material and the geometric strains (ϵ_M, ϵ_B) . The external potential Π_{ext} is given by loads such as gravity, *e.g.*, $\mathbf{f} = \rho \mathbf{g}$, where ρ is the surface mass density and $\mathbf{g}$ is the gravitational acceleration, leading to:

$$\Theta_{\mathbf{u}}^* = \arg \min_{\Theta_{\mathbf{u}}} \Pi[\mathbf{u}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}})], \quad \text{with} \quad \Pi_{\text{ext}} = - \int_{\Gamma} \mathbf{f} \cdot \mathbf{u} \, d\mathbf{x}_\Gamma. \quad (15)$$

Our formulation is orthogonal to the choice of constitutive material law as long as it can be written as a material energy density $\Psi(\epsilon_M, \epsilon_B; \Phi)$. We implement two material models for modelling Π_{int} : the isotropic model is used for controlled benchmark validation, while the orthotropic StVK model captures direction-dependent fabric stiffness for garment examples.

Isotropic Linear Elasticity. For simple simulation and validation (Belytschko benchmarks; see Sec. 5.3), we employ an isotropic model governed by Young’s modulus E and Poisson’s ratio ν . Mapping the projected strains to in-plane stress σ and pre-integrating through the thickness h analytically, yields the membrane resultant $\tilde{\mathbf{n}}(\epsilon_M)$ and bending resultant $\mathbf{m}(\epsilon_B)$. After the thickness integration, the internal energy on the mid-surface can be evaluated as:

$$\Pi_{\text{int}} = \int_{\Gamma} (\tilde{\mathbf{n}} : \epsilon_M + \mathbf{m} : \epsilon_B) \, d\mathbf{x}_\Gamma. \quad (16)$$

Here, “:” denotes double contraction between second-order tensors.

Orthotropic StVK Material Model. In realistic garment simulation, fabrics exhibit direction-dependent stiffness. We implement the orthotropic St. Venant–Kirchhoff (StVK) model defined by stiffness constants $\Phi = \{a_{11}, a_{22}, a_{12}, G_{12}\}$ [14] applied to strain components evaluated in a local material tangent frame. As this energy density Ψ_{StVK} is non-linear across the shell thickness, we evaluate Π_{int} by integrating explicitly through the thickness via Simpson quadrature:

$$\Pi_{\text{int}} = \int_{\Gamma} \int_{-h/2}^{h/2} \Psi_{\text{StVK}}(\epsilon_M + \zeta \epsilon_B; \Phi) \, d\zeta \, d\mathbf{x}_\Gamma. \quad (17)$$

Optimisation. We approximate the continuous surface integral in (15) via Monte Carlo quadrature. At each iteration, we uniformly sample points $\{\mathbf{x}_i\}_i$ on $\mathcal{M}_\Gamma$: from mesh faces in the case of mesh inputs, from the zero level set of ϕ for neural UDFs, and from mean positions for point clouds and 3D Gaussians. We then query the fixed geometric priors from the input representation (Sec. 4.1), and evaluate the displacements $\mathbf{u}(\mathbf{x}_i; \Theta_{\mathbf{u}})$. Depending on the chosen material, we compute the corresponding internal energy (Eq. (16) or Eq. (17)) alongside the external forces, and train the NDF end-to-end to minimise Π . For visualisation, we extract a mesh $\mathcal{M}_{\Gamma'}$ by evaluating $\mathbf{u}$ at the vertices of the reference $\mathcal{M}_\Gamma$ and displacing them accordingly (Fig. 3-(right)).

Table 1: Performance comparison with NeuralClothSim [38]. Training times correspond to results in Tab. 2 (top) and cloth experiments (bottom). Consistent speedups of up to $4\times$ are observed.

Dataset	NeuralClothSim	NeuralGarSim
Square Plate	13m 1s	6m 26s
Scordelis-Lo Roof	4h 49m	2h 21m
Pinched Cylinder	56m 4s	27m 17s
Napkin Moving Handles	6m 9s	1m 35s
Skirt Twist	21m 45s	7m 14s
Sleeve Buckle	11m 3s	5m 55s
Average	1h 6m 10s	31m 35s

Table 2: Maximum displacement at the target point for each Belytschko test. “*”: *Guo et al.* use different materials matching corresponding analytical values.

Method	Square Plate	Scordelis-Lo	Pinched Cyl.
Analytical [5, 82]	0.487	0.3024	1.825e-5
<i>Guo et al.</i> [27]	2.566*	(N/A)	(N/A)
<i>Bastek et al.</i> [4]	(N/A)	0.297	(N/A)
NeuralClothSim [38]	0.487	0.2965	1.825e-5
NeuralGarSim (Ours)	0.487	0.2987	1.820e-5

5 Experiments

5.1 Implementation Details

Garment Deformation. The deformation field $\mathbf{u}(\mathbf{x}_I; \Theta_{\mathbf{u}})$ maps each reference point to the deformed surface point $\mathbf{x}_{I'} = \mathbf{x}_I + \mathbf{u}(\mathbf{x}_I; \Theta_{\mathbf{u}})$. We use five hidden layers of 256 units. We use the Adam optimiser [44] for training. We start from a learning rate of 10^{-4} with an exponential decay schedule [38]. To evaluate the energy in Eq. (15), we uniformly sample 1000 points on the cloth mesh $\mathcal{M}_I$. For linear isotropic material, we set the parameters to representative textile values, $\Phi = (E, \nu, \rho, h) = (5000 \text{ Pa}, 0.25, 0.144 \text{ kg m}^{-2}, 1.2 \text{ mm})$. For the orthotropic StVK material model, we use cotton and silk stiffness constants following fabric tables from Refs. [13, 63]. Each simulation is trained for the number of optimisation steps in the range [2000 – 5000], depending on garment complexity.

NeuralGarSim is implemented in PyTorch [68]. For DUDF inputs, neural field fitting and NDF optimisation each complete in one to 15 minutes on a single NVIDIA H100, for a total of three to 17 minutes. For mesh, point cloud, and 3D Gaussian inputs, the end-to-end pipeline is shorter because the DUDF fitting stage is skipped. Consequently, only the shared NDF optimisation stage is run, which completes in 1 to 8 minutes depending on garment complexity.

5.2 Baselines and Datasets

We compare NeuralGarSim with NeuralClothSim [38], DiffCloth [51], DiffARC-Sim [53, 65], and MeshGraphNets [69]. MeshGraphNets is a learned autoregressive dynamics surrogate rather than a quasistatic energy minimiser, so we report it in App. G.1. Our experiments cover classical Belytschko benchmarks [5, 82] and multi-panel garments, including T-shirts, dresses, and jackets [8, 45].

5.3 Classical Benchmarks

Validation. We first validate the physical accuracy of our method for estimating quasistatic deformation. To this end, we reproduce the Belytschko benchmark

experiments [5, 24], which test different modes of shell deformation. We represent the benchmark reference geometries as triangle meshes and apply the prescribed forces and boundary conditions for the square plate (membrane-dominated), Scordelis-Lo roof (combined membrane/bending), and pinched cylinder (bending-dominated) tests, measuring the maximum displacement at the target point. As shown in Tab. 2, our method closely matches the analytical solutions, confirming that the TDC-based strain formulation captures both membrane and bending modes. Details are provided in App. F.

Speed. We compare training time against NeuralClothSim in Tab. 1. Thanks to the fully 3D strain formulation and the avoidance of parametrisation space, NeuralGarSim consistently trains faster, with speed-ups of roughly $2\text{--}4\times$. The gain comes from evaluating all differential quantities directly on embedded samples. We therefore avoid constructing and differentiating through Christoffel-symbol terms on a 2D parametric domain as in [38], while still optimising the same quasi-static energy objective with comparable Monte Carlo sampling.




Fig. 4: Given different input discretisations (top row) with an identical simulation scenario, our method produces consistent folds and wrinkles due to continuous NDF. On the other hand, the mesh-based simulator, DiffARCSim [53], is sensitive to initial discretisation, despite enabling adaptive remeshing.

5.4 Discretisation Independence

Setup. The key advantage of NeuralGarSim inherited from NeuralClothSim [38] is that the deformation is represented continuously (as a coordinate-based MLP; see Sec. 4.2) and, therefore, the simulation quality is independent of the input mesh discretisation $\mathcal{M}_I$. To demonstrate this, for each garment in the experiment, we generate three alternative discretisations via isotropic remeshing; see Fig. 4 (top row). Each mesh preserves the overall geometry and pinned boundary set. We then run the simulations under identical forces and the same pinned garment regions for ours, DiffARCSim [53] and DiffCloth [51]. Note that DiffARCSim is adaptive (*i.e.*, it supports adaptive remeshing), whereas DiffCloth is not and runs at the fixed initial resolution.

Quantitative Metric. We report full-surface Chamfer distance d_C and *wrinkle-region Chamfer distance* d_W , which quantifies the consistency of high-curvature regions across discretisations. Since the NDF provides correspondences, we extract




Fig. 5: Discretisation independence. Qualitative results at three mesh resolutions. NeuralGarSim produces consistent wrinkle patterns; DiffCloth shows sensitivity to mesh connectivity. **Fig. 6: Multi-modal consistency evaluation.** Qualitative results for four input representations. Small cross-modal distances confirm that simulation is not significantly affected by the choice of input.

Table 3: Wrinkle Chamfer distances d_W and full Chamfer distances d_C across three discretisations. NeuralGarSim achieves consistently lower values than DiffCloth.

Method	$d_W(\cdot, \cdot)^{\tau=75}$		$d_C(\cdot, \cdot)$	
	$(\mathcal{M}_{T'}^1, \mathcal{M}_{T'}^2)$	$(\mathcal{M}_{T'}^1, \mathcal{M}_{T'}^3)$	$(\mathcal{M}_{T'}^1, \mathcal{M}_{T'}^2)$	$(\mathcal{M}_{T'}^1, \mathcal{M}_{T'}^3)$
DiffCloth	0.035	0.034	0.020	0.021
Ours	0.011	0.013	0.009	0.009

Table 4: Cross-modal Chamfer distances relative to the mesh reference.

Modality	$d_W(\mathcal{M}_{T'}^{\text{ref}}, \mathcal{M}_{T'}^i)^{\tau=75}$	$d_C(\mathcal{M}_{T'}^{\text{ref}}, \mathcal{M}_{T'}^i)$
Mesh	ref.	ref.
UDF	0.017	0.013
Point cloud	0.	0.
Gaussians	0.017	0.014

$\mathcal{M}_{T'}$ using the initial mesh topology $\mathcal{M}_T$, whereas the baselines directly output meshes. The wrinkle skeleton of a mesh $\mathcal{M}$ is the vertex subset defined as

$$\mathcal{W}(\mathcal{M}) = \{v \in \mathcal{M} : |H(v)| \geq H^{(\tau)}(\mathcal{M})\}, \quad (18)$$

where $H(v)$ is vertex mean curvature and $H^{(\tau)}$ denotes the 75th percentile over vertices of $\mathcal{M}$. $d_W(\mathcal{M}_{T'}^i, \mathcal{M}_{T'}^j)$ is the bidirectional Chamfer distance between simulation wrinkles $\mathcal{W}(\mathcal{M}_{T'}^i)$ and $\mathcal{W}(\mathcal{M}_{T'}^j)$ from two initial discretisations $\mathcal{M}_{T'}^i$ and $\mathcal{M}_{T'}^j$. A lower value indicates more consistent wrinkles.

Results. Fig. 4 and Fig. 5 show results for T-shirts under three discretisations, while Tab. 3 reports the corresponding full-surface and wrinkle-region Chamfer distances. NeuralGarSim produces visually consistent fold patterns and achieves the lowest d_W across all methods. Both DiffARCSim (despite adaptive remeshing) and DiffCloth exhibit wrinkle patterns that shift substantially with input mesh connectivity, as expected from discretisation-dependent element formulations.

5.5 Multi-Modal Input Representations

To the best of our knowledge, NeuralGarSim is the only simulator in our comparisons that natively supports multiple input representations, while all baselines



Fig. 7: Multi-material simulation. We simulate the same garment with isotropic linear, cotton StVK, and silk StVK parameters. Anisotropic stiffness changes fold orientation and wrinkle density.




Fig. 8: Simulation on reconstructed 3D Gaussians. GaussianGarments [70] provides surface-registered 3D Gaussians from multi-view images. NeuralGarSim uses the extracted Gaussians directly as the simulation input. For clear visualisation, we show the same undeformed Gaussian asset as a reference garment and show the predicted quasistatic drape under pinned shoulders.

require triangle meshes. Fig. 6 shows NeuralGarSim applied to four input modalities: a differentiable UDF (primary representation), a triangle mesh, a point cloud, and 3D Gaussians encoding the same reference garment. The resulting deformations are perceptually consistent across modalities: large-scale fold structure and drape direction are preserved, and fine-scale wrinkles appear in the same surface regions. Cross-modal Chamfer distances reported in Tab. 4 are small, confirming that input representation does not significantly affect simulation fidelity.

Gaussians from Captured Data. We simulate a garment reconstructed from multi-view video using GaussianGarments [70]. We use its registered surface Gaussians as input and ignore appearance optimisation and learned dynamics, since NeuralGarSim requires only splat orientations for normals and is independent of colour or opacity. As shown in Fig. 8, NeuralGarSim predicts the quasistatic equilibrium under gravity and pinned shoulders without constructing a 2D parametric domain.


Multi-Material Simulation. Fig. 7 shows a garment simulated under isotropic linear elasticity and orthotropic StVK with cotton and silk parameters. The linear model produces uniform drape with equal resistance in all directions. The cotton StVK configuration, with higher warp stiffness ($a_{11} > a_{22}$), produces stiffer folds aligned with the warp direction, resulting in a more structured upright drape. The silk StVK configuration, with lower stiffness and higher shear compliance (G_{12}), drapes more loosely with softer folds. These distinct results confirm that

the orthotropic StVK model captures material anisotropy without changing the simulation architecture.

5.6 Ablation: Necessity of TDC

A natural question is whether a fully 3D TDC formulation is necessary, or whether the existing 2D parametrisation of NeuralClothSim can be extended to handle multi-panel garments by stitching UV charts at seams. To test this, we construct **NeuralClothSim+** by: (a) extending the single-panel UV domain of NeuralClothSim [38] to a multi-panel garment with a per-panel parameterisation and (b) replacing the 2D NDF input (*i.e.*, parametric coordinates ξ) with 3D positions $\mathbf{x}_F$. Since NeuralClothSim requires the strain expression to be defined in 2D parametric coordinates, we recover parametric derivatives from NDF through the chain rule, $\partial \mathbf{u} / \partial \xi = (\partial \mathbf{u} / \partial \mathbf{x}_F) J_{\text{rest}}$, where $J_{\text{rest}} = \partial \mathbf{x}_F / \partial \xi$ is the per-face Jacobian of the *rest* configuration. Although simulations seem plausible at first, NeuralClothSim+ suffers from the following fundamental limitations that TDC eliminates by construction.

Limitations of the Parametric Extension. With a conformal parameterisation, NeuralClothSim+ produces a plausible drape, but with a non-conformal parameterisation of the same garment, it fails to converge since distorted per-face Jacobians transfer inconsistent parametric gradients to the strain computation. Moreover, it cannot be applied to input modalities where a 2D parameterisation is unavailable or ill-defined, such as point clouds. In contrast, NeuralGarSim requires only surface normals and produces physically consistent simulations in all cases. Additional results are provided in App. C.1.



6 Discussion and Conclusion

NeuralGarSim enables garment simulation directly in 3D without an explicit 2D parameterisation, while preserving simulation fidelity and supporting different mesh extraction resolutions post-training. It matches NeuralClothSim on Belytschko benchmarks, trains 2–4× faster, and reduces wrinkle-region Chamfer distance compared to DiffCloth in discretisation experiments. It supports multi-panel garments across distance fields, meshes, point clouds, and 3D Gaussians, with consistent behaviour across geometric modalities and material-dependent drape. Current limitations are the quasistatic setting and the absence of inertia, damping, and collisions. Extending the method to dynamics, robust collision handling, faster optimisation, and inverse problems remains future work.

Acknowledgements

We thank Rhaleb Zayer and Briac Toussaint for insightful discussions on tangential differential calculus and shell theory.

References

1. Baraff, D., Witkin, A.: Large steps in cloth simulation. In: Conference on Computer Graphics and Interactive Techniques. SIGGRAPH (1998)
2. Barr, A.H.: Global and local deformations of solid primitives. In: Conference on Computer Graphics and Interactive Techniques. SIGGRAPH (1984)
3. Basar, Y., Krätzig, W.B.: *Mechanik der Flächentragwerke: theorie, berechnungsmethoden, anwendungsbeispiele*. Springer-Verlag (2013)
4. Bastek, J.H., Kochmann, D.M.: Physics-informed neural networks for shell structures. *European Journal of Mechanics-A/Solids* (2023)
5. Belytschko, T., Stolarski, H., Liu, W.K., Carpenter, N., Ong, J.S.: Stress projection for membrane and shear locking in shell finite elements. *Computer Methods in Applied Mechanics and Engineering* (1985)
6. Bertiche, H., Madadi, M., Escalera, S.: Pbns: physically based neural simulation for unsupervised garment pose space deformation. *ACM Transactions on Graphics* (2021)
7. Bertiche, H., Madadi, M., Escalera, S.: Neural cloth simulation. *ACM Transactions on Graphics* (2022)
8. Bhatnagar, B.L., Tiwari, G., Theobalt, C., Pons-Moll, G.: Multi-garment net: Learning to dress 3d people from images. In: International Conference on Computer Vision (ICCV) (2019)
9. Bridson, R., Fedkiw, R., Anderson, J.: Robust treatment of collisions, contact and friction for cloth animation. In: *ACM Transactions on Graphics* (2002)
10. Buffet, T., Rohmer, D., Barthe, L., Boissieux, L., Cani, M.P.: Implicit untangling: A robust solution for modeling layered clothing. *ACM Transactions on Graphics* (2019)
11. Chen, R., Chen, L., Parashar, S.: Gaps: Geometry-aware, physics-based, self-supervised neural garment draping. In: International Conference on 3D Vision (3DV) (2024)
12. Chibane, J., Mir, A., Pons-Moll, G.: Neural unsigned distance fields for implicit function learning. In: *Neural Information Processing Systems (NeurIPS)* (2020)
13. Clyde, D., Teran, J., Tamstorf, R.: Modeling and data-driven parameter estimation for woven fabrics. In: *ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA)* (2017)
14. Clyde, D.C.: *Numerical Subdivision Surfaces for Simulation and Data Driven Modeling of Woven Cloth*. University of California, Los Angeles (2017)
15. De Luigi, L., Li, R., Guillard, B., Salzmann, M., Fua, P.: DrapeNet: Garment Generation and Self-Supervised Draping. In: *Computer Vision and Pattern Recognition (CVPR)* (2023)
16. Delfour, M., Zolésio, J.: Tangential differential equations for dynamical thin/shallow shells. *Journal of Differential Equations* (1996)
17. Delfour, M.C., Zolesio, J.P.: *Shapes and Geometries: Metrics, Analysis, Differential Calculus, and Optimization*. SIAM (2011)
18. Desbrun, M., Meyer, M., Alliez, P.: Intrinsic parameterizations of surface meshes. In: *Computer Graphics Forum* (2002)
19. Duisterhof, B.P., Mandi, Z., Yao, Y., Liu, J.W., Song, S., Shou, M.Z., Ichnowski, J.: Md-splatting: Learning metric deformation from 4d gaussians in highly deformable scenes. *arXiv preprint arXiv:2312.00583* (2023)
20. Eitzmuß, O., Keckeisen, M., Straßer, W.: A fast finite element solution for cloth modelling. In: *Proc. of The Pacific Conference on Computer Graphics and Applications* (2003)

21. Fainstein, M., Siless, V., Iarussi, E.: Dufd: Differentiable unsigned distance fields with hyperbolic scaling. In: Computer Vision and Pattern Recognition (CVPR) (2024)
22. Grigorev, A., Becherini, G., Black, M., Hilliges, O., Thomaszewski, B.: ContourCraft: Learning to resolve intersections in neural multi-garment simulations. In: ACM SIGGRAPH Conference Papers (2024)
23. Grigorev, A., Black, M.J., Hilliges, O.: Hood: Hierarchical graphs for generalized modelling of clothing dynamics. In: Computer Vision and Pattern Recognition (CVPR) (2023)
24. Grinspun, E., Gingold, Y., Reisman, J., Zorin, D.: Computing discrete shape operators on general meshes. In: Computer Graphics Forum (2006)
25. Grinspun, E., Hirani, A.N., Desbrun, M., Schröder, P.: Discrete shells. In: ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA) (2003)
26. Guédon, A., Lepetit, V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In: Computer Vision and Pattern Recognition (CVPR) (2024)
27. Guo, H., Zhuang, X., Rabczuk, T.: A deep collocation method for the bending analysis of kirchhoff plate. *Computers, Materials & Continua* (2019)
28. Guo, Q., Han, X., Fu, C., Gast, T., Tamstorf, R., Teran, J.: A material point method for thin shells with frictional contact. *ACM Transactions on Graphics* (2018)
29. Gurtin, M.E., Ian Murdoch, A.: A continuum theory of elastic material surfaces. *Archive for rational mechanics and analysis* (1975)
30. Hansbo, P., Larson, M.G.: Finite element modeling of a linear membrane shell problem using tangential differential calculus. *Computer Methods in Applied Mechanics and Engineering* (2014)
31. Hansbo, P., Larson, M.G.: Continuous/discontinuous finite element modelling of kirchhoff plate structures in $\mathbb{R}^3$ using tangential differential calculus. *Computational Mechanics* (2017)
32. Hansbo, P., Larson, M.G., Larsson, F.: Tangential differential calculus and the finite element modeling of a large deformation elastic membrane problem. *Computational Mechanics* (2015)
33. Hao, Z., Liu, S., Zhang, Y., Ying, C., Feng, Y., Su, H., Zhu, J.: Physics-informed machine learning: A survey on problems, methods and applications. *arXiv preprint arXiv:2211.08064* (2022)
34. Harmon, D., Vouga, E., Tamstorf, R., Grinspun, E.: Robust treatment of simultaneous collisions. *ACM Transactions on Graphics* (2008)
35. Hjelmstad, K.D.: *Fundamentals of structural mechanics*. Springer (2005)
36. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Surface reconstruction from unorganized points. In: *Proceedings of the 19th annual conference on computer graphics and interactive techniques* (1992)
37. Kairanda, N., Habermann, M., Naik, S., Theobalt, C., Golyanik, V.: Thin-Shell-SfT: Fine-grained monocular non-rigid 3d surface tracking with neural deformation fields. In: *Computer Vision and Pattern Recognition (CVPR)* (2025)
38. Kairanda, N., Habermann, M., Theobalt, C., Golyanik, V.: NeuralClothSim: Neural deformation fields meet the thin shell theory. In: *Neural Information Processing Systems (NeurIPS)* (2024)
39. Kairanda, N., Naik, S., Habermann, M., Sharma, A., Theobalt, C., Golyanik, V.: ∂^∞ -Grid: A neural differential equation solver with differentiable feature grids. In: *International Conference on Learning Representations (ICLR)* (2026)

40. Kairanda, N., Tretschk, E., Elgharib, M., Theobalt, C., Golyanik, V.: ϕ -SfT: Shape-from-template with a physics-based deformation model. In: Computer Vision and Pattern Recognition (CVPR) (2022)
41. Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L.: Physics-informed machine learning. *Nature Reviews Physics* (2021)
42. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* (2023)
43. Kiendl, J., Hsu, M.C., Wu, M.C., Reali, A.: Isogeometric kirchhoff-love shell formulations for general hyperelastic materials. *Computer Methods in Applied Mechanics and Engineering* (2015)
44. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (ICLR) (2015)
45. Korosteleva, M., Lee, S.H.: Generating datasets of 3d garments with sewing patterns. In: Neural Information Processing Systems (NeurIPS) (2021)
46. Li, C., Tang, M., Tong, R., Cai, M., Zhao, J., Manocha, D.: P-cloth: interactive complex cloth simulation on multi-gpu systems using dynamic matrix assembly and pipelined implicit integrators. *ACM Transactions on Graphics* (2020)
47. Li, J., Daviet, G., Narain, R., Bertails-Descoubes, F., Overby, M., Brown, G.E., Boissieux, L.: An implicit frictional contact solver for adaptive cloth simulation. *ACM Transactions on Graphics* (2018)
48. Li, M., Kaufman, D.M., Jiang, C.: Codimensional incremental potential contact. *ACM Transactions on Graphics* (2021)
49. Li, R., Guillard, B., Remelli, E., Fua, P.: Dig: Draping implicit garment over the human body. In: Asian Conference on Computer Vision (ACCV) (2022)
50. Li, Y., Chen, H.y., Larionov, E., Sarafianos, N., Matusik, W., Stuyck, T.: Diffavatar: Simulation-ready garment optimization with differentiable simulation. In: Computer Vision and Pattern Recognition (CVPR) (2024)
51. Li, Y., Du, T., Wu, K., Xu, J., Matusik, W.: Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Transactions on Graphics* (2022)
52. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: Computer Vision and Pattern Recognition (CVPR) (2023)
53. Liang, J., Lin, M., Koltun, V.: Differentiable cloth simulation for inverse problems. *Neural Information Processing Systems (NeurIPS)* (2019)
54. Liao, Z., Wang, S., Komura, T.: Senc: Handling self-collision in neural cloth simulation. In: European Conference on Computer Vision (ECCV) (2024)
55. Liu, S., Zhongkai, H., Ying, C., Su, H., Zhu, J., Cheng, Z.: A unified hard-constraint framework for solving geometrically complex pdes. In: Neural Information Processing Systems (NeurIPS) (2022)
56. Long, X., Lin, C., Liu, L., Liu, Y., Wang, P., Theobalt, C., Komura, T., Wang, W.: Neuraludf: Learning unsigned distance fields for multi-view reconstruction of surfaces with arbitrary topologies. In: Computer Vision and Pattern Recognition (CVPR) (2023)
57. Love, A.E.H.: A treatise on the mathematical theory of elasticity. Cambridge university press (2013)
58. Lu, J., Zheng, C.: Dynamic cloth simulation by isogeometric analysis. *Computer Methods in Applied Mechanics and Engineering* (2014)
59. Lu, L., Pestourie, R., Yao, W., Wang, Z., Verdugo, F., Johnson, S.G.: Physics-informed neural networks with hard constraints for inverse design. *SIAM Journal on Scientific Computing* (2021)

60. Ly, M., Jouve, J., Boissieux, L., Bertails-Descoubes, F.: Projective dynamics with dry frictional contact. *ACM Transactions on Graphics* (2020)
61. Mehta, I., Chandraker, M., Ramamoorthi, R.: A level set theory for neural implicit evolution under explicit flows. In: *European Conference on Computer Vision (ECCV)* (2022)
62. Meyer, M., Desbrun, M., Schröder, P., Barr, A.H.: Discrete differential-geometry operators for triangulated 2-manifolds. In: *Visualization and mathematics III* (2003)
63. Miguel, E., Miraut, D., Otaduy, M.A.: Modeling and estimation of energy-based hyperelastic objects. In: *Computer Graphics Forum* (2016)
64. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: *European Conference on Computer Vision (ECCV)* (2020)
65. Narain, R., Samii, A., O’Brien, J.F.: Adaptive anisotropic remeshing for cloth simulation. *ACM Transactions on Graphics* (2012)
66. Novello, T., Da Silva, V., Schardong, G., Schirmer, L., Lopes, H., Velho, L.: Neural implicit surface evolution. In: *International Conference on Computer Vision (ICCV)* (2023)
67. Otaduy, M.A., Tamstorf, R., Steinemann, D., Gross, M.: Implicit contact handling for deformable objects. In: *Computer Graphics Forum* (2009)
68. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Neural Information Processing Systems (NeurIPS)* (2019)
69. Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.: Learning mesh-based simulation with graph networks. In: *International Conference on Learning Representations (ICLR)* (2021)
70. Rong, B., Grigorev, A., Wang, W., Black, M.J., Thomaszewski, B., Tsalicoglou, C., Hilliges, O.: Gaussian Garments: Reconstructing simulation-ready clothing with photorealistic appearance from multi-view video. In: *International Conference on 3D Vision (3DV)* (2025)
71. Santesteban, I., Otaduy, M.A., Casas, D.: SNUG: Self-Supervised Neural Dynamic Garments. In: *Computer Vision and Pattern Recognition (CVPR)* (2022)
72. Santesteban, I., Otaduy, M.A., Thuerey, N., Casas, D.: ULNeF: Untangled layered neural fields for mix-and-match virtual try-on. In: *Neural Information Processing Systems (NeurIPS)* (2022)
73. Schöllhammer, D., Fries, T.P.: Kirchhoff–love shell theory based on tangential differential calculus. *Computational mechanics* (2019)
74. Schöllhammer, D., Fries, T.P.: Reissner–mindlin shell theory based on tangential differential calculus. *Computer Methods in Applied Mechanics and Engineering* (2019)
75. Simo, J.C., Fox, D.D., Rifai, S.M.: On a stress resultant geometrically exact shell model. part ii: The linear theory; computational aspects. *Computer Methods in Applied Mechanics and Engineering* (1989)
76. Simo, J.C., Fox, D.D.: On a stress resultant geometrically exact shell model. part i: Formulation and optimal parametrization. *Computer Methods in Applied Mechanics and Engineering* (1989)
77. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: *Neural Information Processing Systems (NeurIPS)* (2020)
78. Tang, M., Liu, Z., Tong, R., Manocha, D.: Psc: Parallel self-collision culling with spatial hashing on gpus. *Proc. ACM Comput. Graph. Interact. Tech.* (2018)

79. Terzopoulos, D., Platt, J., Barr, A., Fleischer, K.: Elastically deformable models. SIGGRAPH Comput. Graph. (1987)
80. Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Yifan, W., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., et al.: Advances in neural rendering. Computer Graphics Forum (Eurographics State of the Art Reports) (2022)
81. Thomaszewski, B., Wacker, M., Straßer, W.: A consistent bending model for cloth simulation with corotational subdivision finite elements. In: ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA) (2006)
82. Timoshenko, S., Woinowsky-Krieger, S.: Theory of plates and shells. McGraw-hill New York (1959)
83. Volino, P., Thalmann, N.M.: Implementing fast cloth simulation with collision response. In: Proceedings Computer Graphics International (2000)
84. Wang, H.: Gpu-based simulation of cloth wrinkles at submillimeter levels. ACM Transactions on Graphics (2021)
85. Wang, H., O’Brien, J.F., Ramamoorthi, R.: Data-driven elastic models for cloth: modeling and measurement. ACM Transactions on Graphics (2011)
86. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: Neural Information Processing Systems (NeurIPS) (2021)
87. Wang, Y., Han, Q., Habermann, M., Daniilidis, K., Theobalt, C., Liu, L.: Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In: International Conference on Computer Vision (ICCV) (2023)
88. Wempner, G., Talaslidis, D.: Mechanics of solids and shells. CRC Press (2003)
89. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. In: Computer Graphics Forum (Eurographics State of the Art Reports) (2022)
90. Yang, G., Belongie, S., Hariharan, B., Koltun, V.: Geometry processing with neural fields. In: Neural Information Processing Systems (NeurIPS) (2021)
91. Zhang, J.E., Dumas, J., Fei, Y., Jacobson, A., James, D.L., Kaufman, D.M.: Progressive shell quasistatics for unstructured meshes. ACM Transactions on Graphics (2023)

NeuralGarSim: Geometry-Agnostic Garment Simulation with Neural Fields

Supplementary Material

Contents

A	Tangential Differential Calculus for Tensors	4
B	Strain Simplification	5
C	Details of NeuralClothSim	6
C.1	Necessity of TDC: Natural Garment	7
D	Kirchhoff-Love Shell Kinematics	8
E	Garment as Neural Level Set	8
E.1	Choice of Implicit Representation	9
F	Experimental Details	11
G	Additional Comparisons	11
G.1	Comparison with MeshGraphNets	11
G.2	Multiresolution Comparison with DiffARCSim	12
G.3	Convergence Speed: Ours vs NeuralClothSim	12
G.4	Runtime and Performance	13
H	NDF Robustness	13
H.1	Robustness to Degraded Dense Inputs	14
H.2	UDF Geometry Operators	14
H.3	Multiresolution NDF Sampling	14
H.4	NDF Editing	15
H.5	Network Capacity and Architecture	15
H.6	Sampling Strategy	17
I	Additional Ablations	18
I.1	Importance of Nonlinear Strain	18
I.2	Stochasticity and Reproducibility	19
J	Discussion and Future Work	19
K	Algorithm	19

Table I: Table of Notations.

Symbol	Description
<i>Cloth Representation</i>	
$\Gamma \subset \mathbb{R}^3$	Reference/undeformed cloth midsurface
$h \in \mathbb{R}$	Shell (cloth) thickness
$\Omega \subset \mathbb{R}^3$	Reference/undeformed shell continuum, $\Omega = \Gamma \times [-\frac{h}{2}, \frac{h}{2}]$
$\Gamma' \subset \mathbb{R}^3$	Deformed cloth midsurface

Continued on the next page

Table I – continued from previous page

Symbol	Description
$\Omega' \subset \mathbb{R}^3$	Deformed shell continuum, $\Omega' = \Gamma' \times [-\frac{h}{2}, \frac{h}{2}]$
$\mathbf{x}_\Gamma \in \Gamma$	A point on the reference midsurface
$\mathbf{x}_\Omega \in \Omega$	A point in the reference shell continuum
$\mathbf{x}_{\Gamma'} \in \Gamma'$	A point on the deformed midsurface, $\mathbf{x}_{\Gamma'} = \mathbf{x}_\Gamma + \mathbf{u}(\mathbf{x}_\Gamma)$
$\mathbf{x}_{\Omega'} \in \Omega'$	A point in the deformed shell continuum
$\zeta \in [-\frac{h}{2}, \frac{h}{2}]$	Coordinate along the shell's normal/thickness direction
$\partial\Gamma \subset \mathbb{R}^3$	Boundary of midsurface Γ
$T_{\mathbf{x}}\Gamma$	Tangent plane to the surface Γ at point $\mathbf{x}$
$\mathcal{M}_\Gamma$	Reference garment mesh discretising Γ
$\mathcal{M}_{\Gamma'}$	Deformed garment mesh, $\mathcal{M}_\Gamma$ displaced by $\mathbf{u}$
$N_\Gamma \in \mathbb{N}$	Number of Monte Carlo surface samples per training iteration
Network Parameters	
Θ_ϕ	Network weights for the UDF ϕ
$\Theta_{\mathbf{u}}$	Network weights for the NDF $\mathbf{u}$
Fields and Functions	
$\phi(\mathbf{x}; \Theta_\phi) : \mathbb{R}^3 \rightarrow \mathbb{R}^+$	An unsigned distance function (UDF) defining Γ
$\mathbf{n}_\Gamma(\mathbf{x}_\Gamma) : \Gamma \rightarrow \mathbb{R}^3$	Unit normal vector fields for midsurface
$\mathbf{u}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}}) : \Gamma \rightarrow \mathbb{R}^3$	Neural Deformation Field (NDF) on midsurface
$\mathbf{u}_\Omega(\mathbf{x}_\Omega, \zeta; \Theta_{\mathbf{u}}) : \Omega \rightarrow \mathbb{R}^3$	3D displacement field within the shell continuum
$\mathbf{n}_{\Gamma'}(\mathbf{x}_{\Gamma'}; \Theta_{\mathbf{u}}) : \Gamma \rightarrow \mathbb{R}^3$	Normal vector fields for deformed midsurface
$\mathbf{w}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}}) : \Gamma \rightarrow \mathbb{R}^3$	Non-linear director difference vector, $\mathbf{w}(\mathbf{x}_\Gamma) = \mathbf{n}_{\Gamma'}(\mathbf{x}_{\Gamma'}) - \mathbf{n}_\Gamma(\mathbf{x}_\Gamma)$
$\mathcal{B}(\mathbf{x}_\Gamma) : \Gamma \rightarrow \mathbb{R}$	Dirichlet boundary condition
$\mathbf{f}(\mathbf{x}_\Gamma) : \Gamma \rightarrow \mathbb{R}^3$	External force density
$\hat{d}_\Gamma(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}^+$	Differentiable UDF with hyperbolic scaling, $d_\Gamma(\mathbf{x}) \tanh(\alpha d_\Gamma(\mathbf{x}))$
Operators	
∇	Standard 3D gradient operator
$\mathbb{I}$	3×3 identity tensor
$\otimes$	Outer (dyadic) product
$\mathbf{P} : \Gamma \rightarrow \mathbb{R}^{3 \times 3}$	Projection operator onto tangent plane
∇_Γ	Tangential gradient operator for scalar function with respect to Γ
$\nabla_\Gamma^{\text{dir}}$	Directional gradient operator for a vector-valued function with respect to Γ

Continued on the next page

Table I – continued from previous page

Symbol	Description
∇_F^{cov}	Covariant gradient operator for vector- and tensor-valued functions with respect to F
$\mathbf{H} : F \rightarrow \mathbb{R}^{3 \times 3}$	Weingarten map (curvature operator) of F
Strain and Stress	
$\mathbf{F}_u : F \rightarrow \mathbb{R}^{3 \times 3}$	Midsurface displacement gradient, $\nabla_F^{\text{dir}} \mathbf{u}$
$\mathbf{F}_w : F \rightarrow \mathbb{R}^{3 \times 3}$	Director difference gradient, $\nabla_F^{\text{dir}} \mathbf{w}$
$\boldsymbol{\epsilon}^{\text{dir}} : F \rightarrow \mathbb{R}^{3 \times 3}$	Directional total strain tensor
$\boldsymbol{\epsilon}_M^{\text{dir}}, \boldsymbol{\epsilon}_M : F \rightarrow \mathbb{R}^{3 \times 3}$	Membrane strain before and after tangent-plane projection, respectively
$\boldsymbol{\epsilon}_B^{\text{dir}}, \boldsymbol{\epsilon}_B : F \rightarrow \mathbb{R}^{3 \times 3}$	Bending strain before and after tangent-plane projection, respectively
$\boldsymbol{\sigma}^{\text{dir}}, \boldsymbol{\sigma} : F \rightarrow \mathbb{R}^{3 \times 3}$	Directional total stress tensor (in $\mathbb{R}^{3 \times 3}$ and in-plane)
$\boldsymbol{\sigma}_M^{\text{dir}}, \boldsymbol{\sigma}_M : F \rightarrow \mathbb{R}^{3 \times 3}$	Directional membrane stress tensor (in $\mathbb{R}^{3 \times 3}$ and in-plane)
$\boldsymbol{\sigma}_B^{\text{dir}}, \boldsymbol{\sigma}_B : F \rightarrow \mathbb{R}^{3 \times 3}$	Directional bending stress tensor (in $\mathbb{R}^{3 \times 3}$ and in-plane)
$\tilde{\mathbf{n}}, \mathbf{m} : F \rightarrow \mathbb{R}^{3 \times 3}$	Membrane and bending moment resultant
$\Psi : F \rightarrow \mathbb{R}$	Strain energy density function
$\Pi, \Pi_{\text{int}}, \Pi_{\text{ext}} \in \mathbb{R}$	Total, internal, and external potential energy
$\boldsymbol{\epsilon}_T : F \rightarrow \mathbb{R}^{3 \times 3}$	Total in-plane strain at thickness coordinate ζ , $\boldsymbol{\epsilon}_M + \zeta \boldsymbol{\epsilon}_B$
Isotropic Material Parameters	
ρ	Mass area density of the cloth
E, ν	Young’s modulus and Poisson’s ratio
μ, λ	Lamé parameters for plane stress
Orthotropic StVK Parameters	
a_{11}, a_{22}	Warp and weft axial stiffness constants
a_{12}	Coupled axial stiffness constant
G_{12}	In-plane shear stiffness constant
$\mathbf{C} : \mathbb{R}^{3 \times 3} \rightarrow \mathbb{R}^{3 \times 3}$	Orthotropic material stiffness tensor in the local material frame
$(\mathbf{t}_1, \mathbf{t}_2) \in T_{\mathbf{x}} F$	Local material tangent frame defining warp and weft directions
Φ	Set of material parameters, <i>e.g.</i> , (E, ν) for isotropic or $(a_{11}, a_{22}, a_{12}, G_{12})$ for orthotropic StVK
Evaluation Metrics	

Continued on the next page

Table I – continued from previous page

Symbol	Description
$d_C \in \mathbb{R}^+$	Full-surface Chamfer distance between deformed meshes
$d_W \in \mathbb{R}^+$	Wrinkle Chamfer distance, restricted to high curvature surface regions

This appendix provides additional physical and geometrical details underlying our NeuralGarSim approach and the experiments. We begin by expanding the derivations and formalism of our method, including a comparison with a cloth simulator with an explicitly parametrised surface, NeuralClothSim [38]. We then present further qualitative results and ablations that substantiate the claims made in the main paper. App. A gives the tensor-valued TDC operators used by our strain formulation, App. B derives the Green–Lagrange membrane and bending components, and App. C summarises the parametric NeuralClothSim formulation used for comparison. App. C.1 further studies the necessity of TDC, App. D records complete KL shell kinematics, and App. E details DUDF fitting. The remaining sections provide experimental settings, additional comparisons, robustness studies, ablations, limitations, and implementation algorithms.

A Tangential Differential Calculus for Tensors

This section provides tensorial extensions of the Tangential Differential Calculus (TDC) operators used in the main paper. These operations appear in the derivation of membrane and bending strain measures (see App. B) and are required to express the directional and covariant gradients of surface tensor fields on the reference mid-surface.

For a vector field $\mathbf{v}(\mathbf{x}) : \Gamma \rightarrow \mathbb{R}^3$, this concept extends to *directional gradient* $\nabla_\Gamma^{\text{dir}} \mathbf{v}$,

$$\nabla_\Gamma^{\text{dir}} \mathbf{v}(\mathbf{x}) = \begin{bmatrix} (\nabla_\Gamma v_x)^\top \\ (\nabla_\Gamma v_y)^\top \\ (\nabla_\Gamma v_z)^\top \end{bmatrix}. \quad (19)$$

The rows of this 3×3 matrix are the tangential gradients of the components of $\mathbf{v}$. Because this tensor is generally not tangential, we project it to obtain the *covariant gradient*,

$$\nabla_\Gamma^{\text{cov}} \mathbf{v}(\mathbf{x}) = \mathbf{P}(\mathbf{x}) \cdot \nabla_\Gamma^{\text{dir}} \mathbf{v}(\mathbf{x}), \quad (20)$$

which is strictly in-plane (*i.e.*, $\nabla_\Gamma^{\text{cov}} \mathbf{v} \cdot \mathbf{n}_\Gamma = 0$).

For a second-order tensor $\mathbf{A} : \Gamma \rightarrow \mathbb{R}^{3 \times 3}$, its directional gradient can be obtained as the directional derivative of each column of the tensor. Let $\mathbf{A} = [\mathbf{A}_x, \mathbf{A}_y, \mathbf{A}_z]$, where $\mathbf{A}_i$ denotes the i th column. Then,

$$\nabla_\Gamma^{\text{dir}} \mathbf{A}(\mathbf{x}) = \begin{bmatrix} (\nabla_\Gamma^{\text{dir}} \mathbf{A}_x)^\top \\ (\nabla_\Gamma^{\text{dir}} \mathbf{A}_y)^\top \\ (\nabla_\Gamma^{\text{dir}} \mathbf{A}_z)^\top \end{bmatrix}. \quad (21)$$

Analogous to the covariant vector gradient in Eq. (20), the in-plane covariant gradient for a tensor field is defined as:

$$\nabla_F^{\text{cov}} \mathbf{A}(\mathbf{x}) = \mathbf{P}(\mathbf{x}) \cdot \mathbf{A}(\mathbf{x}) \cdot \mathbf{P}(\mathbf{x}). \quad (22)$$

Since the tensor is strictly in-plane it satisfies $\nabla_F^{\text{cov}} \mathbf{A} \cdot \mathbf{n}_F = \mathbf{0}_3$ where $\mathbf{0}_3 \in \mathbb{R}^{3 \times 1}$ is the zero vector.

For additional details on TDC and its applications to KL shells, we refer the reader to [32, 73].

B Strain Simplification

This section provides the full derivation of the simplified membrane and bending strains (see Eq. (12)-(main) and Eq. (13)-(main)).

To model large deformations, we employ the full Green-Lagrange (GL) strain tensor, defined from the 3D displacement gradient in the shell continuum [35]. Writing the directional displacement gradient on the shell as $\nabla_F^{\text{dir}} \mathbf{u}_\Omega$, the GL strain tensor is:

$$\boldsymbol{\epsilon}^{\text{dir}} = \frac{1}{2} (\nabla_F^{\text{dir}} \mathbf{u}_\Omega + (\nabla_F^{\text{dir}} \mathbf{u}_\Omega)^\top + (\nabla_F^{\text{dir}} \mathbf{u}_\Omega)^\top \nabla_F^{\text{dir}} \mathbf{u}_\Omega). \quad (23)$$

Once we substitute the KL displacement Eq. (9) into Eq. (23), the gradient becomes,

$$\nabla_F^{\text{dir}} \mathbf{u}_\Omega = \nabla_F^{\text{dir}} (\mathbf{u} + \zeta \mathbf{w}) = \nabla_F^{\text{dir}} \mathbf{u} + \zeta \nabla_F^{\text{dir}} \mathbf{w}. \quad (24)$$

Let $\mathbf{F}_u = \nabla_F^{\text{dir}} \mathbf{u}$ and $\mathbf{F}_w = \nabla_F^{\text{dir}} \mathbf{w}$. Substituting $\nabla_F^{\text{dir}} \mathbf{u}_\Omega = \mathbf{F}_u + \zeta \mathbf{F}_w$ into the GL tensor Eq. (23) and expanding gives:

$$\begin{aligned} \boldsymbol{\epsilon}^{\text{dir}} = & \frac{1}{2} (\mathbf{F}_u + \mathbf{F}_u^\top + \mathbf{F}_u^\top \mathbf{F}_u) \\ & + \zeta \frac{1}{2} (\mathbf{F}_w + \mathbf{F}_w^\top + \mathbf{F}_u^\top \mathbf{F}_w + \mathbf{F}_w^\top \mathbf{F}_u) \\ & + \zeta^2 \frac{1}{2} (\mathbf{F}_w^\top \mathbf{F}_w). \end{aligned}$$

We can neglect the higher-order term in ζ^2 , splitting the strain into membrane and bending components:

$$\boldsymbol{\epsilon}^{\text{dir}}(\mathbf{x}, \zeta) \approx \boldsymbol{\epsilon}_M^{\text{dir}}(\mathbf{x}) + \zeta \boldsymbol{\epsilon}_B^{\text{dir}}(\mathbf{x}), \quad (25)$$

where the *membrane strain tensor* is:

$$\boldsymbol{\epsilon}_M^{\text{dir}} = \frac{1}{2} (\mathbf{F}_u + \mathbf{F}_u^\top + \mathbf{F}_u^\top \mathbf{F}_u), \quad (26)$$

and the *bending strain tensor* is:

$$\boldsymbol{\epsilon}_B^{\text{dir}} = \frac{1}{2} (\mathbf{F}_w + \mathbf{F}_w^\top + \mathbf{F}_u^\top \mathbf{F}_w + \mathbf{F}_w^\top \mathbf{F}_u). \quad (27)$$

Note that the strain is non-linear due to the quadratic terms in both membrane and bending strains.

C Details of NeuralClothSim

For a comparison between our parameterisation-free approach and the parameterisation-based approach used in NeuralClothSim [38], we summarise the latter’s approach.

Let the reference surface be defined parametrically as $\mathbf{x}_r(\xi_1, \xi_2) : \tau \rightarrow \mathbb{R}^3$, and the deformed surface as $\mathbf{x}_{r'}(\xi^1, \xi^2) = \mathbf{x}_r + \mathbf{u}(\xi^1, \xi^2)$. The geometry of the shell is governed by the first fundamental form (metric tensor) and the second fundamental form (curvature tensor).

Parametric Energy Minimisation. Equilibrium is found by minimising:

$$\min_{\Theta} \int_{\tau} \Psi(\boldsymbol{\epsilon}(\boldsymbol{\xi}; \Theta), \boldsymbol{\kappa}(\boldsymbol{\xi}; \Theta)) d\boldsymbol{\xi} - \int_{\tau} \mathbf{f} \cdot \mathbf{u}(\boldsymbol{\xi}; \Theta) d\boldsymbol{\xi}, \quad (28)$$

where strains $\boldsymbol{\epsilon}, \boldsymbol{\kappa}$ are evaluated via classical parametric differential geometry [3, 75, 76].

For the reference surface, the tangent basis vectors are $\bar{\mathbf{a}}_{\alpha} = \frac{\partial \mathbf{x}_r}{\partial \xi^{\alpha}}$, $\alpha \in \{1, 2\}$ and the unit normal is $\mathbf{n}_r = \frac{\bar{\mathbf{a}}_1 \times \bar{\mathbf{a}}_2}{\|\bar{\mathbf{a}}_1 \times \bar{\mathbf{a}}_2\|}$. The coefficients of the first fundamental form, $\bar{a}_{\alpha\beta}$, and the second fundamental form, $\bar{b}_{\alpha\beta}$, are given by:

$$\bar{a}_{\alpha\beta} = \bar{\mathbf{a}}_{\alpha} \cdot \bar{\mathbf{a}}_{\beta}, \quad \bar{b}_{\alpha\beta} = -\bar{\mathbf{a}}_{\alpha} \cdot \frac{\partial \mathbf{n}_r}{\partial \xi^{\beta}} \quad \text{for } \alpha, \beta \in \{1, 2\}. \quad (29)$$

Similarly, for the deformed configuration $\mathbf{x}_{r'}$, the deformed tangents $\mathbf{a}_{\alpha} = \frac{\partial \mathbf{x}_{r'}}{\partial \xi^{\alpha}}$, the deformed metric $a_{\alpha\beta} = \mathbf{a}_{\alpha} \cdot \mathbf{a}_{\beta}$ and the deformed curvature $b_{\alpha\beta} = \mathbf{n}_{r'} \cdot \frac{\partial \mathbf{a}_{\alpha}}{\partial \xi^{\beta}}$ can be calculated.

NeuralClothSim computes the GL membrane strain components $\epsilon_{\alpha\beta}$ and bending strain components $\kappa_{\alpha\beta}$ from:

$$\epsilon_{\alpha\beta} = \frac{1}{2}(a_{\alpha\beta} - \bar{a}_{\alpha\beta}), \quad \kappa_{\alpha\beta} = \bar{b}_{\alpha\beta} - b_{\alpha\beta}. \quad (30)$$

A key limitation of the parametrisation-based formulation is the complexity of computing deformation gradients on a curved manifold. The deformation gradient $\mathbf{u}_{,\alpha}$ cannot be computed via simple element-wise differentiation of the displacement field. Instead, it must be decomposed into components relative to the reference basis vectors $\bar{\mathbf{a}}^{\lambda}$ and director $\bar{\mathbf{a}}^3$:

$$\mathbf{u}_{,\alpha} = \varphi_{\alpha\lambda} \bar{\mathbf{a}}^{\lambda} + \varphi_{\alpha 3} \bar{\mathbf{a}}^3, \quad (31)$$

where the scalar components $\varphi_{\alpha\lambda}$ and $\varphi_{\alpha 3}$ define the local deformation state. The in-plane components are defined as:

$$\varphi_{\alpha\lambda} = u_{\lambda|\alpha} - \bar{b}_{\alpha\lambda} u_3, \quad (32)$$

where $u_{\lambda|\alpha}$ is the covariant derivative of the displacement components. This operation requires precomputing Christoffel symbols of second kind, $\Gamma_{\alpha\beta}^{\lambda}$, to quantify the rate of change for basis vectors:

$$u_{\lambda|\alpha} = u_{\lambda,\alpha} - u_{\gamma} \Gamma_{\lambda\alpha}^{\gamma}, \quad \Gamma_{\alpha\beta}^{\lambda} := \bar{\mathbf{a}}^{\lambda} \cdot \bar{\mathbf{a}}_{\alpha,\beta}. \quad (33)$$

Using these derived quantities, the membrane strain $\epsilon_{\alpha\beta}$ in NeuralClothSim becomes:

$$\epsilon_{\alpha\beta} = \frac{1}{2}(\varphi_{\alpha\beta} + \varphi_{\beta\alpha} + \varphi_{\alpha\lambda}\varphi_{\beta}^{\lambda} + \varphi_{\alpha 3}\varphi_{\beta 3}). \quad (34)$$

The dependency on parametric coordinates becomes even more pronounced when calculating the change in curvature. NeuralClothSim approximates this with the nonlinear bending strain:

$$\kappa_{\alpha\beta} \approx -\varphi_{\alpha 3|\beta} - \bar{b}_{\beta}^{\lambda}\varphi_{\alpha\lambda} + \varphi_3^{\lambda} \left(\varphi_{\alpha\lambda|\beta} + \frac{1}{2}\bar{b}_{\alpha\beta}\varphi_{\lambda 3} - \bar{b}_{\beta\lambda}\varphi_{\alpha 3} \right). \quad (35)$$

Limitations. The parametric formulation has three principal limitations. (i) *Parameterisation dependence*: a continuous 2D domain τ is required, which is difficult to construct for multi-panel garments with seams and holes. (ii) *Seam handling*: NeuralClothSim proposes periodic boundary conditions for a single seam, but extending this to multiple seams or holes is non-trivial. (iii) *Metric distortion*: mapping a 3D garment to 2D invariably introduces distortion [18], which propagates error into the physical model.


In contrast, our proposed method utilises Tangential Differential Calculus to compute equivalent strain measures directly in Cartesian coordinates (Eq. (12)-(main) and Eq. (13)-(main)), bypassing the need for the domain coordinates $(\xi^1, \xi^2) \in \tau$ entirely. We evaluate the deformation directly in embedding space, eliminating the need for Christoffel systems and associated basis vectors.

C.1 Necessity of TDC: Natural Garment

On a box mesh, non-conformal parameterisation causes NeuralClothSim+ to improperly converge, while the conformal one succeeds. On a natural garment (a jacket with sewn panels), both conformal and non-conformal parameterisations fail, in that poor convergence is observed. While the external forces are clearly applied, the initial set of wrinkles from the reference state is preserved and exhibits minimal movement.

This confirms our central claim: UV parameterisation is garment-dependent and *unpredictable*. NeuralGarSim produces consistent, physically accurate results as the TDC operators depend only on the normal and curvature field, and carry no dependence on the parameterisation choice. This

establishes that the failure mode of parametric formulations is not tied to a specific UV algorithm but is an inherent consequence of the dependence of per-face Jacobian J_{rest} , whose quality cannot be guaranteed a priori.



D Kirchhoff–Love Shell Kinematics

For completeness, we define kinematic expressions for the entire KL shell rather than just the mid-surface. Under external loading, the shell deforms into a new mid-surface Γ' obtained by displacing every point on Γ :

$$\Gamma' = \{\mathbf{x}_{\Gamma'} \in \mathbb{R}^3 \mid \mathbf{x}_{\Gamma'} = \mathbf{x}_\Gamma + \mathbf{u}, \forall \mathbf{x}_\Gamma \in \Gamma\}. \quad (36)$$

The 3D shell displacement field decomposes linearly through the thickness $\zeta \in [-h/2, h/2]$ as:

$$\begin{aligned} \mathbf{x}_{\Omega'} &= \mathbf{x}_\Omega + \mathbf{u}_\Omega, \\ \mathbf{u}_\Omega &= \mathbf{u} + \zeta \mathbf{w}, \quad \mathbf{w} = \mathbf{n}_{\Gamma'} - \mathbf{n}_\Gamma. \end{aligned} \quad (37)$$

Here $\mathbf{u}$ captures mid-surface displacement and $\mathbf{w}$ captures changes in surface orientation. The Kirchhoff–Love (KL) assumption enforces a linear variation of $\mathbf{u}_\Omega$ across the thickness via $\mathbf{w}$, eliminating transverse shear. A point in the deformed shell volume is therefore $\mathbf{x}_{\Omega'} = \mathbf{x}_{\Gamma'} + \zeta \mathbf{n}_{\Gamma'}$, and the full 3D deformation is determined entirely by the deformation of the 2D mid-surface.

E Garment as Neural Level Set

Neural SDFs are ill-suited to garment modelling because garments are typically *open* surfaces with complex seams: the sign becomes ambiguous near boundaries, leading to gradient discontinuities that destabilise curvature estimates. We therefore adopt a UDF, which remains well-defined for open surfaces.

Standard UDFs are only C^0 continuous at the zero level set, producing two distinct surface responses in the normal and curvature fields and noisy geometric derivatives. We therefore employ the Differentiable UDF (DUDF) [21], which guarantees C^k continuity through *hyperbolic scaling*:

$$\hat{d}_\Gamma(\mathbf{x}) = d_\Gamma(\mathbf{x}) \tanh(\alpha d_\Gamma(\mathbf{x})). \quad (38)$$

This scaling contracts the distance field near the surface while preserving its sign-free character away from it. The resulting gradient and Hessian of $\hat{d}_\Gamma$ provide smooth, coherent estimates of $\mathbf{n}_\Gamma$ and the Weingarten map $\mathbf{H}_\Gamma$ at the zero level set, which are the only geometric quantities required by the NDF.

Modified Eikonal Training. The MLP $\phi(\mathbf{x}; \Theta_\phi)$ is trained to match $\hat{d}_\Gamma$ while satisfying a modified eikonal condition:

$$\begin{aligned} \Theta_\phi^* = \arg \min_{\Theta_\phi} & \int_\Gamma |\phi(\mathbf{x}; \Theta_\phi)| d\mathbf{x}_\Gamma + \int_\Gamma |\nabla \phi(\mathbf{x}; \Theta_\phi)| d\mathbf{x}_\Gamma \\ & + \int \left| \|\nabla \phi(\mathbf{x}; \Theta_\phi)\| - \|\nabla \hat{d}_\Gamma(\mathbf{x})\| \right| d\mathbf{x}. \end{aligned} \quad (39)$$

The first two terms enforce vanishing distance and unit-magnitude gradient on Γ ; the third term matches the gradient norm of the learned field to that of $\hat{d}_\Gamma$.

everywhere in the ambient volume, replacing the standard eikonal constraint $\|\nabla\phi\| = 1$ with one that respects the hyperbolic scaling. In practice we approximate these integrals with on-surface samples $\{(\mathbf{x}_i, d_\Gamma(\mathbf{x}_i))\}_i$ drawn on $\mathcal{M}_\Gamma$ and off-surface samples within a surrounding unit cube.

Geometric Priors. Once trained, the unit normal is recovered as:

$$\mathbf{n}_\Gamma(\mathbf{x}_\Gamma; \Theta_\phi) = \frac{\nabla\phi(\mathbf{x}_\Gamma; \Theta_\phi)}{\|\nabla\phi(\mathbf{x}_\Gamma; \Theta_\phi)\|}, \quad (40)$$

and the projection operator $\mathbf{P}$ and Weingarten map $\mathbf{H}_\Gamma$ follow from Eqs. (4) and (6). Our TDC-based strain formulation requires only $\mathbf{n}_\Gamma$ and $\mathbf{H}_\Gamma$ at each surface point; any input modality that supplies these two quantities integrates directly with our pipeline without further modification.




Fig. I: Deformed Garment with different distance fields. Given an initial garment configuration, we optimise the MLP using SDF, standard UDF and DUDF representations. The SDF representation produces overly smoothed surfaces. While UDF represents open surfaces, it creates two distinct surface responses at the zero crossing, causing uncertain geometric derivatives. In contrast, DUDF ensures C^k continuity, enabling sharp normal and curvature fields ideal for stable optimisation.

E.1 Choice of Implicit Representation

Our method requires geometric derivatives (surface normals and curvature) at each sample point. For neural implicit inputs, the quality of these derivatives




Fig. II: Belytschko obstacle experiment [5]. (a): Displacement magnitude fields for the three benchmark tests, evaluated at convergence. **(b):** Error over training. NeuralGarSim converges faster than NeuralClothSim while closely matching analytical displacements.

depends directly on the choice of implicit representation. We compare three candidates trained on the same garment geometry to justify our use of DUDF.

Figure I shows the reconstructed distance and normal fields. For the curvature field, we compute and visualise the mean curvature as the average of the principal curvatures, obtained from the eigenvalues κ_1 and κ_2 of the Weingarten map $\mathbf{H}_F$, (*i.e.*, $(\kappa_1 + \kappa_2)/2$).

- **SDF:** SDFs are ill-defined on open surfaces, causing physically implausible elongation. The learned field is overly smoothed in the neighbourhood of the surface. This produces blurred normal and curvature estimates, visible as thickened surface responses.
- **UDF:** UDFs correctly represent open surfaces but are only C^0 at the zero level set. While the overall drape is plausible, the inconsistent derivatives still introduce residual stretching artefacts.
- **DUDF:** The hyperbolic scaling of DUDF ensures C^k continuity around the zero level set. As seen in Fig. I, the normals are smooth and coherent, with sharp curvature fields. This produces unambiguous geometric derivatives at the boundaries, allowing NDF to minimise the energy density reliably. Therefore, DUDF provides stable optimisation and high-fidelity garment simulation.

These results justify our choice of DUDF as the implicit representation for neural inputs. Standard UDF and SDF inputs are not excluded by the method

formulation but produce degraded results in practice due to derivative quality, not architectural incompatibility.

F Experimental Details

Garment Representation. When the input is a UDF, we fit a scalar MLP $\phi(\mathbf{x}; \Theta_\phi)$ following the DUDF protocol of *Fainstein et al.* [21]: a Siren network with eight hidden layers of 256 units trained with Adam [44]. Each batch contains roughly 100k spatial samples drawn near and away from the zero level set. Training runs for 3000 iterations with a piecewise-constant learning rate: 1×10^{-4} for the first 1000 steps, 1×10^{-5} for the next 1000, and 1×10^{-7} for the final 1000. Once trained, $\phi(\mathbf{x}; \Theta_\phi)$ is frozen, and its gradient and Hessian are queried to obtain $(\mathbf{n}_T, \mathbf{P}, \mathbf{H})$ for NDF training. For triangle mesh, point cloud and Gaussian inputs, geometric priors are extracted directly from the input geometry as described in Sec. 4.1, and no neural field fitting is required.

Belytschko Benchmarks. Figure II supplements Tab. 2-(main) in the main paper with displacement field visualisations and convergence curves for all three tests.

Figure II-(a) shows displacement error versus training iterations for both NeuralGarSim (blue) and NeuralClothSim (red). NeuralGarSim converges faster on all three tests and reaches a lower final displacement error, confirming that the TDC-based strain formulation is both physically correct and numerically efficient relative to the parametric alternative.

Square plate (Fig. II-(b), left): a clamped square plate under uniform transverse pressure, dominated by membrane action. The displacement field shows the characteristic circular symmetry with maximum deflection at the centre, closely matching the reference solution.

Scordelis-Lo roof (Fig. II-(b), middle): a cylindrical shell roof with free longitudinal edges and rigid end-diaphragms under self-weight. This test combines membrane and bending contributions; the displacement field shows the expected saddle-shaped deformation with peak displacement at the mid-span free edge.

Pinched cylinder (Fig. II-(b), right): a thin cylindrical shell with point loads at opposing poles and symmetry boundary conditions, dominated by bending. The displacement field shows localised deformation at the load points, consistent with the high bending-to-membrane stiffness ratio of the test configuration.

G Additional Comparisons

G.1 Comparison with MeshGraphNets

Experimental Setup. MeshGraphNets [69] is a learned graph-based dynamics simulator. We train it using the official implementation released by the authors on the FlagSimple dataset, which consists of flag-like planar meshes simulated under wind forces. In their proposed Encode-Process-Decode architecture, we




Fig. III: Our approach gives consistent simulation (folds and wrinkles) irrespective of the initial discretisation of the garment mesh.

use 15 message-passing steps and a latent size of 128. Since the official open-source release did not contain remeshing support, we therefore evaluate the standard model without remeshing-aware training. At inference, we provide garment meshes with different discretisations as input, and autoregressively roll out 300 steps and take the final frame for comparison. To ensure reproducibility across discretisations, we fix a random seed for all experiments, controlling both the network initialisation and the autoregressive rollout.

Methodological Distinction. MeshGraphNets is a data-driven surrogate: it learns to approximate dynamics from training trajectories and does not strictly enforce physical law at inference. A direct comparison is therefore methodologically distinct; we include it for completeness.

Results. Similar to Fig. 4-(main), we provide discretisation comparisons against MeshGraphNets in Fig. III. Across all tested meshes, NeuralGarSim consistently converges to similar equilibrium configurations with comparable fold density and wrinkle patterns. MeshGraphNets produces plausible results within its training distribution but exhibits sensitivity to mesh resolution changes, as the graph connectivity changes with remeshing and falls outside the distribution seen during training.

G.2 Multiresolution Comparison with DiffARCSim

As shown in Fig. IV, while DiffARCSim [53] yields inconsistent converged states across different mesh resolutions, our method maintains stable and visually consistent fold and wrinkle characteristics. Similar to Fig. IV-(main), this highlights the resolution-independent nature of our deformed representation.

G.3 Convergence Speed: Ours vs NeuralClothSim

We compare the optimisation efficiency of our method against NeuralClothSim [38]. Figure V shows intermediate cloth states at various timestamps for both approaches.

While both methods eventually produce comparable fold and wrinkle patterns, our formulation exhibits noticeably faster convergence. In practice, our model recovers coarse drape within the first few seconds. By contrast, NeuralClothSim requires more optimisation steps to reach a similar shape.




Fig. IV: Comparison between DiffARCSim and our approach across different mesh resolutions. While DiffARCSim yields inconsistent states, our approach yields consistent folds and wrinkles regardless of the input mesh resolution.

G.4 Runtime and Performance

Direct wall-clock comparisons across all methods in our evaluation are complicated due to the following factors.

Problem Formulation. DiffCloth [51] and DiffARCSim [53] are dynamic simulators: they integrate equations of motion over time and report per-step costs across a trajectory. NeuralGarSim solves a quasistatic problem, minimising total potential energy to a single equilibrium configuration. These are fundamentally different computational tasks, and per-iteration comparison conflates the cost of solving distinct problems.

Hardware and solver differences. DiffCloth and DiffARCSim are CPU-based simulators using sparse linear solvers, whereas NeuralGarSim runs entirely on GPU via automatic differentiation through a neural field. Comparing wall-clock times across CPU and GPU implementations without a careful hardware-controlled study would not produce meaningful conclusions.

Since NeuralGarSim and NeuralClothSim [38] share the same paradigm, a direct wall-clock comparison between them is meaningful. On this basis, NeuralGarSim trains $2\text{--}4\times$ faster than NeuralClothSim, as reported in Tab. 1-(main). This speed-up is attributable to the simpler optimisation landscape of the TDC formulation, which avoids computing Christoffel symbols and parametric basis vectors at each iteration.

Establishing a principled computational comparison against physics-based dynamic simulators—accounting for problem formulation, hardware, and solver differences—remains an open direction for future work.

H NDF Robustness

We analyse the robustness of NeuralGarSim by testing out a finite set of choices: degrading input samples, approximating geometric operators, querying the NDF at different output resolutions during inference and warm-start NDF training.

H.1 Robustness to Degraded Dense Inputs

Table II degrades the dense input geometry (represented as point cloud or Gaussians) while keeping the material parameters, optimiser, NDF architecture, and loading fixed. We either randomly reduce the density of the input geometry, displace the points using Gaussian noise, or positionally degrade the normal orientations. Gaussian inputs degrade more gracefully than point clouds under severe sparsity and positional noise because each splat carries a coherent local orientation, whereas point-cloud normals must be re-estimated from a poorer neighbourhood. Normal perturbations up to 20° change the final drape only mildly for both modalities, indicating that the energy minimisation tolerates moderate orientation noise.




Table II: Robustness to degraded dense inputs.

Deg.	Level	PCD		Gauss.	
		d_W/d_C		d_W/d_C	
Sparsity	50%	.021/.014		.020/.016	
	10%	.082/.070		.034/.023	
Noise	0.5%	.031/.020		.018/.014	
	2%	.055/.038		.027/.017	
Normal	10°	.010/.008		.009/.009	
	20°	.015/.012		.012/.011	

H.2 UDF Geometry Operators

Table III ablates on the effect of TDC operators on UDF input geometry. Setting $\mathbf{H} = \mathbf{0}$ keeps the projection operator $\mathbf{P}$ intact and nearly matches full TDC, which explains why the same local-planar approximation is acceptable for dense point clouds and Gaussians. Replacing $\mathbf{P}$ with $\mathbb{I}$ removes the tangent-plane projection and lets ambient-space derivatives enter the shell strain, degrading the final output and confirming that tangent-space calculus is essential. The no-bending variant remains close on this garment because this test is not strongly bending-dominated. In contrast, the Scordelis-Lo roof benchmark requires bending stiffness: the same variant gives displacement 0.7882, compared with 0.2987 for full NeuralGarSim and 0.3024 for the analytical solution.




Table III: UDF geometry operators relative to full TDC.

Variant	d_W	d_C
Full TDC	ref.	ref.
$\mathbf{H} = \mathbf{0}$	$< 10^{-3}$	$< 10^{-3}$
No bending	.003	.003
$\mathbf{P} = \mathbb{I}, \mathbf{H} = \mathbf{0}$	.025	.021

H.3 Multiresolution NDF Sampling

The optimised NDF is a continuous displacement field over the garment surface rather than displacements stored only at training




Fig. VI: Multiresolution NDF sampling. A trained NDF can be queried at different mesh resolutions without retraining.

vertices. Figure VI therefore evaluates the same trained field on several output meshes without retraining or changing the equilibrium solve. The coarse mesh captures the global drape, while denser queries reveal finer folds where the same continuous deformation is sampled more densely.

H.4 NDF Editing

Following NeuralClothSim [38], we test whether a trained NDF can be reused when the same asset is locally edited. Starting from a garment trained from scratch in 7 min 34s, we warm-start the NDF weights for three changes: adding a lateral force, adding an additional boundary constraint, and switching the material to cotton StVK. Each edited configuration converges in approximately one minute because the previous solution already encodes the garment’s coarse drape and only the changed equilibrium must be corrected. This amortises repeated design edits, although NeuralGarSim remains a quasistatic solver and does not model transient motion between edits.




Fig. VII: NDF editing. Warm-start edits for force, boundary, and material changes.

H.5 Network Capacity and Architecture

To assess architectural sensitivity, we vary NDF depth, width, and the Siren frequency parameter ω_0 while keeping the garment, material, sampling budget, and optimiser fixed. These ablations test whether failures are caused by insufficient representation capacity, excessive capacity, or inappropriate spectral bias.

Number of Hidden Layers. We evaluate depths of 3, 5 (default) and 8 hidden layers; see Fig. VIII.

- **3 Layers:** The field still captures the large-scale drape, but the reduced capacity smooths some local folds and makes high-curvature regions less sharp.
- **5 Layers (Ours):** This depth provides enough capacity for sharp folds while preserving stable optimisation, which is why we use it by default.
- **8 Layers:** The larger network can represent high-frequency details, but the additional degrees of freedom slow convergence and do not produce a clear visual gain over five layers.




Fig. VIII: NDF depth ablation. Five layers balance expressivity and efficiency.




Fig. IX: NDF width ablation for $d \in \{128, 256, 512, 1024\}$.




Fig. X: Siren frequency ablation. $\omega_0 = 30$ gives the most stable optimisation.




Fig. XI: Sampling-density ablation. $N_\Gamma = 800$ balances convergence and cost.

Hidden Dimension. In addition to network depth, we analyse the impact of the network width (number of neurons per layer) on simulation quality. We test hidden dimensions $d \in \{128, 256, 512, 1024\}$; see Fig. IX.

- $d = 128$: The solution is physically plausible, but the limited width biases the field toward smoother and stiffer deformations.

- $d = 256$: The global drape is recovered, but fine wrinkles remain under-resolved.
- $d = 512$ (**Ours**): This setting captures both global drape and local buckling patterns without excessive memory use.
- $d = 1024$: The wider model adds training time and memory cost but gives no consistent visual improvement over $d = 512$.

Siren Frequency (ω_0). The sinusoidal activation function $\sin(\omega_0 \mathbf{x})$ used in SIREN [77] controls the spatial frequency of the learned function through the parameter ω_0 . In our case, this directly affects the expressiveness of the NDF. We ablate $\omega_0 \in \{1, 20, 30, 60\}$; see Fig. X.

- $\omega_0 = 1$: A very low frequency restricts the spectral content of the Siren field. The resulting deformation is overly smooth, external forces are absorbed mainly by coarse motion, and new folds do not form reliably.
- $\omega_0 = 20$: Increasing the frequency restores medium-scale deformation, but high-frequency wrinkles remain weaker than in the default setting.
- $\omega_0 = 30$ (**Ours**): This setting balances expressivity and optimisation stability. It captures sharp folds and realistic wrinkle deformation while maintaining stable convergence.
- $\omega_0 = 60$: A higher frequency can represent fine details, but optimisation becomes less stable. In our runs, the sleeves converge more slowly and do not descend fully within the same budget, so we avoid excessively high ω_0 values.

H.6 Sampling Strategy

Sampling Density. The accuracy of Monte Carlo integration used to compute the total potential energy depends on the number of surface samples N_F drawn per training iteration. We evaluate convergence for $N_F \in \{50, 200, 400, 800, 1000\}$ on a representative garment mesh.

As shown in Fig. XI, extremely low sampling rates ($N_F = 50$) provide insufficient coverage of the cloth surface, leading to noisy gradients and incomplete convergence. Increasing the number of samples progressively stabilises convergence. We find that $N_F = 800$ offers an effective balance between gradient quality and per-iteration cost for this garment. Using $N_F = 1000$ produces only marginal improvements on this garment.

In practice, the appropriate N_F depends on garment complexity. Denser or more geometrically intricate meshes benefit from high sample counts to maintain adequate surface coverage. A higher sample count increases per-iteration cost, but it can reduce the total number of iterations required for convergence.

Explicit vs Implicit Sampling.

A key advantage of our method is its independence from mesh connectivity. We compare mesh sampling, which draws samples uniformly from the input cloth mesh $\mathcal{M}_F$, with UDF sampling. The latter uniformly samples the bounding volume and retains points where the predicted distance satisfies $|\phi(\mathbf{x})| < \epsilon$, effectively treating the garment as a purely implicit surface.




Fig. XII: Mesh versus UDF sampling. Both sampling strategies produce visually similar equilibria.

Figure XII shows that both strategies lead to visually indistinguishable equilibrium states. This confirms that while mesh-based sampling is computationally efficient during training, our formulation does not rely on the mesh discretisation quality of the input mesh.

I Additional Ablations

I.1 Importance of Nonlinear Strain

Our method uses a non-linear Green-Lagrange strain tensor. To validate this design choice, we ablate our model by replacing the GL tensor with just the linear approximation. Specifically, we remove the quadratic terms $\mathbf{F}_u^T \mathbf{F}_u$ and $\mathbf{F}_u^T \mathbf{F}_w + \mathbf{F}_w^T \mathbf{F}_u$ from Eq. (12)-(main) and Eq. (13)-(main), respectively.

As seen in Fig. XIV, the linear formulation is valid only for small deformations. Under large deformations, the linear model generates spurious volume expansion and fails to exhibit transformation invariance. The non-linear model correctly constrains the cloth, leading to realistic fold patterns.




Fig. XIII: Simulation results with three different random seeds. While precise fold locations vary slightly due to stochastic nature of optimisation, the overall physical behaviour and energy states remain consistent.




Fig. XIV: Comparison of simulation using only linear strain versus a full nonlinear GL strain. The linear approximation leads to unrealistic volume expansion, while the full nonlinear strain ensures constrained optimisation and forms a realistic drape.

I.2 Stochasticity and Reproducibility

Since our method relies on Monte Carlo integration and neural network initialisation, we simulate and visualise the solution across three random seeds. As shown in Fig. XIII, different seeds lead to variations in the locations and shapes of folds, which is expected given the inherent stochasticity of the optimisation process and lack of single equilibria in garment simulation. Importantly, these differences do not reflect instability. Across all seeds, the cloth settles into a physically consistent equilibrium state with nearly identical total potential energy. This demonstrates that, while fine-scale wrinkle patterns may differ, the underlying quasistatic behaviour is robust and primarily governed by the physics-based loss rather than by random initialisation.

J Discussion and Future Work

NeuralGarSim is currently quasistatic: inertia, damping, and collisions are not modelled; extending to dynamic effects and collision handling are the most direct directions for future work, with level-set representations already proving useful for collision-aware garment untangling [15, 72]. Training cost is independent of mesh resolution since samples are drawn from the continuous surface rather than a fixed vertex set, unlike mesh-based simulators whose cost scales with vertex count; direct timing comparisons are omitted as they are confounded by hardware differences. Faster convergence via second-order optimisers or better initialisations, and inverse estimation of fabric constitutive parameters from captured deformations, remain open problems for future work.

K Algorithm

For reproducibility, we provide the detailed training procedure for NDF in Algorithm 1. After training, the deformation can be evaluated efficiently by applying the NDF for each vertex of the reference mesh, as shown in Algorithm 2.

Inference is performed in a single forward pass over all mesh vertices and does not require any geometric quantities or energy terms.

Sampling During Training. Following NeuralClothSim [38], we sample points $\mathbf{x}_r$ on the reference surface $\mathcal{M}_r$ at each iteration. This ensures unbiased gradient estimates and avoids overfitting to a fixed set of points. These samples are used to approximate the surface integrals in Eq. (15)-(main).

Sampling During Inference. Inference does not require Monte Carlo sampling. We simply evaluate the trained NDF at the vertices of the reference mesh. This provides the deformed garment directly, without the need to resample from $\mathcal{M}_r$ or compute geometric quantities.




Fig. V: Convergence Comparison. Comparison of the convergence speed between NeuralClothSim and ours. While both approaches converge to a physically plausible cloth, our method converges substantially faster.



Fig. XV: Qualitative Results. Visualisations of our method across different training iterations. Early iterations show coarse deformations, which are progressively refined over time, eventually producing smooth and physically consistent garment folds.

Algorithm 1: NDF Training in NeuralGarSim

Input : Reference surface Γ with geometric priors $\{\mathbf{n}_\Gamma^{(i)}, \mathbf{H}^{(i)}, \mathbf{P}^{(i)}\}$ (from UDF ϕ , mesh, point cloud, or Gaussians)
 NDF $\mathbf{u}(\mathbf{x}_\Gamma; \Theta_{\mathbf{u}})$ with randomly initialised $\Theta_{\mathbf{u}}$
 Material parameters Φ
 External force $\mathbf{f}(\mathbf{x}_\Gamma)$
 Number of Monte Carlo samples N_Γ
 Number of iterations T
 Learning rate η

Output: Optimised NDF parameters $\Theta_{\mathbf{u}}^*$

for $t = 1$ **to** T **do**

Sample $B \subset \mathcal{M}_\Gamma$ uniformly, $|B| = N_\Gamma$

Query geometric priors at $\{\mathbf{x}_\Gamma^{(i)}\}_{i \in B}$ (if UDF input):

$\mathbf{n}_\Gamma^{(i)} \leftarrow \nabla \phi(\mathbf{x}_\Gamma^{(i)}) / \|\nabla \phi(\mathbf{x}_\Gamma^{(i)})\|$

$\mathbf{H}^{(i)} \leftarrow -(\mathbf{P}^{(i)} \cdot \nabla^2 \phi(\mathbf{x}_\Gamma^{(i)}) \cdot \mathbf{P}^{(i)}) / \|\nabla \phi(\mathbf{x}_\Gamma^{(i)})\|$

$\mathbf{P}^{(i)} \leftarrow \mathbb{I} - \mathbf{n}_\Gamma^{(i)} \otimes \mathbf{n}_\Gamma^{(i)}$

Otherwise extract directly from input geometry.

foreach $i \in B$ **do**

Deformation:

$\mathbf{u}^{(i)} \leftarrow \mathbf{u}(\mathbf{x}_\Gamma^{(i)}; \Theta_{\mathbf{u}})$

Apply boundary mask: $\mathbf{u}^{(i)} \leftarrow \mathcal{B}(\mathbf{x}_\Gamma^{(i)}) \cdot \mathbf{u}^{(i)}$ $\mathbf{F}_u^{(i)} \leftarrow \nabla_\Gamma^{\text{dir}} \mathbf{u}^{(i)}$

$\mathbf{w}^{(i)} \leftarrow \mathbf{H}^{(i)} \mathbf{u}^{(i)} - \nabla_\Gamma(\mathbf{u}^{(i)} \cdot \mathbf{n}_\Gamma^{(i)})$

$\mathbf{F}_w^{(i)} \leftarrow \nabla_\Gamma^{\text{dir}} \mathbf{w}^{(i)}$

Strains:

$\epsilon_M^{(i)} \leftarrow \mathbf{P}^{(i)} \frac{1}{2} (\mathbf{F}_u^{(i)} + \mathbf{F}_u^{(i)\top} + \mathbf{F}_u^{(i)\top} \mathbf{F}_u^{(i)}) \mathbf{P}^{(i)}$

$\epsilon_B^{(i)} \leftarrow \mathbf{P}^{(i)} \frac{1}{2} (\mathbf{F}_w^{(i)} + \mathbf{F}_w^{(i)\top} + \mathbf{F}_u^{(i)\top} \mathbf{F}_w^{(i)} + \mathbf{F}_w^{(i)\top} \mathbf{F}_u^{(i)}) \mathbf{P}^{(i)}$

Stresses (isotropic linear, $\Phi = \{E, \nu\}$):

$\sigma_M^{(i)} = 2\mu \epsilon_M^{(i)} + \lambda \text{tr}(\epsilon_M^{(i)}) \mathbf{P}^{(i)}$

$\sigma_B^{(i)} = 2\mu \epsilon_B^{(i)} + \lambda \text{tr}(\epsilon_B^{(i)}) \mathbf{P}^{(i)}$

Stresses (orthotropic StVK, $\Phi = \{a_{11}, a_{22}, a_{12}, G_{12}\}$):

Project strains to local material frame $(\mathbf{t}_1^{(i)}, \mathbf{t}_2^{(i)})$:

$\epsilon_{M,\text{loc}}^{(i)} \leftarrow \mathbf{T}^{(i)\top} \epsilon_M^{(i)} \mathbf{T}^{(i)}, \quad \epsilon_{B,\text{loc}}^{(i)} \leftarrow \mathbf{T}^{(i)\top} \epsilon_B^{(i)} \mathbf{T}^{(i)}$

$\sigma_{M,\text{loc}}^{(i)} \leftarrow \mathbf{C} \epsilon_{M,\text{loc}}^{(i)}, \quad \sigma_{B,\text{loc}}^{(i)} \leftarrow \mathbf{C} \epsilon_{B,\text{loc}}^{(i)}$

Project back: $\sigma_M^{(i)} \leftarrow \mathbf{T}^{(i)} \sigma_{M,\text{loc}}^{(i)} \mathbf{T}^{(i)\top}, \quad \sigma_B^{(i)} \leftarrow \mathbf{T}^{(i)} \sigma_{B,\text{loc}}^{(i)} \mathbf{T}^{(i)\top}$

Resultants & energy (both material models):

$\tilde{\mathbf{n}}^{(i)} = h \sigma_M^{(i)}$

$\mathbf{m}^{(i)} = \frac{h^3}{12} \sigma_B^{(i)}$

$\Pi_{\text{int}}^{(i)} = \tilde{\mathbf{n}}^{(i)} : \epsilon_M^{(i)} + \mathbf{m}^{(i)} : \epsilon_B^{(i)}$

$\Pi_{\text{ext}}^{(i)} = -\mathbf{f}(\mathbf{x}_\Gamma^{(i)}) \cdot \mathbf{u}^{(i)}$

Total potential energy:

$\Pi = \frac{|\Gamma|}{|B|} \sum_{i \in B} (\Pi_{\text{int}}^{(i)} + \Pi_{\text{ext}}^{(i)})$

Gradient descent:

$\Theta_{\mathbf{u}} \leftarrow \Theta_{\mathbf{u}} - \eta \nabla_{\Theta_{\mathbf{u}}} \Pi$

return $\Theta_{\mathbf{u}}^*$

Algorithm 2: NDF Inference in NeuralGarSim

Input : Reference mesh vertices $\{\mathbf{x}_r^{(i)}\}_{i=1}^{|\mathcal{M}_r|}$
Trained NDF network $\mathbf{u}(\mathbf{x}_r; \Theta_{\mathbf{u}}^*)$
Output : Deformed mesh vertices $\{\mathbf{x}_{r'}^{(i)}\}_{i=1}^{|\mathcal{M}_r|}$
for $i = 1$ **to** $|\mathcal{M}_r|$ **do**
 Deformation: $\mathbf{u}^{(i)} \leftarrow \mathbf{u}(\mathbf{x}_r^{(i)}; \Theta_{\mathbf{u}}^*)$
 Deformed position: $\mathbf{x}_{r'}^{(i)} = \mathbf{x}_r^{(i)} + \mathbf{u}^{(i)}$
return $\{\mathbf{x}_{r'}^{(i)}\}_{i=1}^{|\mathcal{M}_r|}$
