

Enhancing Cloth 3D Simulation and Object Interaction via Machine Learning-Based Deformation Models

Roman Chekhmestruk^a, Olena Voitsekhovska^{*a}, Zbigniew Omiotek^b Tetiana Piliavoz^a,
Yuliia Kovalova^c

^aVinnitsia National Technical University, Vinnitsia, Ukraine; ^bLublin University of Technology, Lublin, Poland; ^cDesign agency Smart Alpaca LLC, USA

ABSTRACT

Realistic cloth simulation remains a challenging task in computer graphics, especially when simulating intricate interactions between deformable fabrics and complex three-dimensional geometries. Traditional physically based models, such as mass-spring or finite element methods, often suffer from high computational cost or numerical instability, limiting their applicability in real-time environments or high-fidelity rendering. In this work, we propose a novel hybrid framework that enhances cloth simulation accuracy and computational efficiency through the integration of machine learning-based deformation prediction. Our method employs a supervised deep learning architecture to approximate the nonlinear dynamics of cloth behavior under varying boundary conditions and interactions with rigid or soft 3D bodies. By training on physically simulated datasets that include diverse cloth types, contact scenarios, and environmental forces, our approach generalizes to unseen geometries and interaction modes. The learned model is embedded within a coarse physical simulation loop to preserve global consistency while accelerating local deformation computations. We evaluate the proposed method using both synthetic benchmarks and real-world datasets. Quantitative results demonstrate a significant reduction in simulation time – up to 60% – while maintaining a high degree of physical plausibility compared to traditional solvers. Qualitative experiments show improved handling of high-frequency wrinkles, collision resolution, and dynamic contact response. This framework paves the way for practical applications in virtual try-on systems, real-time animation, and haptic feedback environments. The proposed method contributes to the growing field of differentiable simulation and demonstrates the potential of data-driven models in addressing the limitations of conventional cloth physics engines. Future work includes extending the model to support multi-layered garments, anisotropic materials, and adaptive resolution strategies.

Keywords: cloth simulation; machine learning; deformable object interaction; deep learning; physics-based modeling; real-time graphics; collision detection

1. INTRODUCTION

The realistic simulation of cloth and its interaction with arbitrarily shaped three-dimensional (3D) objects has long been recognised as a pivotal yet persistently difficult research problem within computer graphics, virtual prototyping, and digital entertainment. Canonical physically based approaches – most notably mass-spring lattices¹, position-based dynamics², and finite element methods (FEM)³ – provide a principled way to model elastic and inelastic behaviour at multiple scales, but they also impose severe computational burdens when high spatial resolutions, complex collision topologies, or stiff material parameters are required. In practice, these burdens manifest as prohibitively small time-steps necessary for explicit integrators to remain stable or, conversely, the substantial non-linear solves demanded by implicit schemes. As a result, interactive rates remain elusive for many real-time applications, while film-quality accuracy necessitates massive offline computation⁴.

*vojcehovska.o.v@vntu.edu.ua

A further source of complexity arises from collision handling and contact response. Robust cloth–cloth and cloth–object interaction requires continuous collision detection and reliable response strategies to avoid inter-penetrations, self-intersections, or ill-conditioned constraint systems⁵. Despite decades of progress, spanning hierarchical bounding structures, adaptive remeshing, and frictional contact models, state-of-the-art solvers often exhibit either numerical instability under rapid deformation or excessive dissipation that damps visually important high-frequency wrinkles⁶. These effects are exacerbated in heterogeneous garments composed of anisotropic fabrics, layered composites, or non-woven technical textiles, whose constitutive laws are difficult to parameterise a priori.

Recent advances in machine learning (ML) offer an appealing alternative paradigm: rather than solving the full set of partial differential equations each frame, one can learn a low-dimensional surrogate that maps coarse physical states and boundary conditions to fine-scale deformations^{7–9}. Pioneering works on neural subspace dynamics¹¹ and graph-based message-passing networks¹¹ have demonstrated that data-driven predictors can capture complex non-linear strain–stress relationships, generalise to unseen poses, and significantly reduce per-frame computation. Nonetheless, existing ML-based simulators typically focus on either free-fall cloth or limited draping scenarios and thus fall short when rich, persistent contacts and deeply coupled cloth–object interactions are present. Moreover, pure data-driven approaches may violate conservation laws or fail to respect global constraints such as momentum preservation and volume exclusion.

To address these limitations, the present study introduces a hybrid simulation framework that tightly couples a lightweight, physically motivated coarse solver with a deep neural deformation module trained on a curated database of high-resolution cloth–object interactions. In contrast to prior surrogate schemes, our method exploits bidirectional feedback: the coarse solver enforces global consistency and boundary constraints, while the neural component injects high-frequency detail and corrects local collision artefacts. Crucially, the learning stage is formulated in a differentiable setting, facilitating end-to-end optimisation of network parameters with respect to physically meaningful loss functions, thereby ensuring adherence to energy conservation and contact invariants.

This paper contributes three key elements to the field:

1. **A differentiable data generation pipeline** that synthesises physically plausible, richly annotated cloth–object trajectories spanning diverse fabric types, material anisotropies, and interaction geometries, enabling scalable supervised training without manual intervention.
2. **A deformation-aware neural architecture** that blends graph-based relational reasoning with spectral attention to capture both global cloth dynamics and local wrinkle formation, while remaining agnostic to mesh resolution.
3. **A comprehensive experimental evaluation** demonstrating that the proposed hybrid approach yields a two-to-fourfold speed-up over state-of-the-art FEM solvers, with mean positional error below 2 mm across a suite of canonical benchmarks and real-world capture datasets.

By bridging rigorous physical modelling with data-driven prediction, the proposed framework advances the fidelity, stability, and efficiency of cloth simulation, opening new opportunities for real-time virtual try-on systems, physically informed avatar animation, and interactive digital twins in robotics and manufacturing.

2. Physics-Informed Hybrid Deformation Framework for Cloth–Object Interaction

Hybrid Simulation Paradigm. Figure 1 schematically summarises the proposed pipeline. A lightweight physically based coarse solver (left) advances the cloth state with a large time-step under gravity, wind drag, and user-specified constraints. Its low spatial resolution guarantees unconditional stability but sacrifices fine wrinkles and accurate contact. The resulting coarse state x_t^c is passed, together with collision candidates, to a neural deformation module (centre) that predicts a dense displacement field $\Delta x_t \in \mathbb{R}^{3N}$ over the high-resolution cloth mesh (NNN vertices). The corrected configuration $x_t = x_t^c + \Delta x_t$ is fed back to the next physics step, thereby producing bidirectional coupling that preserves large-scale momentum while injecting high-frequency detail. **Differentiable Data Generation.** To supervise learning we synthesise a

corpus $D = \{(s_i, x_i)\}_{i=1}^M$ of high-resolution cloth-object trajectories using an implicit finite-element solver with adaptive remeshing¹². Each scene parameter tuple s_i encodes material coefficients (Young's modulus E , Poisson ratio ν), anisotropic warp-weft orientation, environmental loads, and the signed-distance field of interacting rigid or deformable obstacles. Automatic differentiation through the simulator yields per-vertex stress tensors σ and penetration depths d that later act as auxiliary supervision signals. Coarse Physics Solver. We adopt projective dynamics¹³ in combination with edge-length, bending, and area constraints evaluated on 16 * decimated cloth mesh (aprox. 5% of original vertices). A single global projection-Gauss-Seidel iteration per frame suffices for visually plausible macroscopic motion, running at 1.3 ms on NVIDIA RTX 4090 for garments of $N > 50,000$ vertices.




Figure 1. Hybrid cloth simulation pipeline unifying coarse physics with neural refinement (generated diagram).

Figure 2 details the network topology. The input mesh is first embedded into a graph $G = (v, \varepsilon)$ whose nodes carry vertex positions, velocities, coarse residuals, and local contact features (d, n) . A message-passing graph neural network (GNN) of four layers encodes global stress propagation, while two stacked spectral attention blocks learn frequency-aware filters that accentuate fine wrinkles. Positional encodings derived from the eigenvectors of the cotangent Laplacian ensure resolution-independence¹⁹⁻²². The decoder predicts per-vertex offsets Δx and corrective damped velocities Δv .




Figure 2. Graph-based neural network with spectral attention for high-frequency deformation prediction.

Loss function:

$$\mathcal{L} = w_1 \sum_t \|\Delta \mathbf{x}_t - \widehat{\Delta \mathbf{x}}_t\|_2^2 + w_2 \sum_t \|\nabla_{\mathbf{x}_t} \mathcal{E} - \hat{\mathbf{f}}_t\|_2^2 + w_3 \sum_t \sum_j \max(0, d_t^{(j)})^2,$$

where $\widehat{\Delta \mathbf{x}}_t$ and $\hat{\mathbf{f}}_t$ denote ground-truth displacements and nodal forces; $\mathcal{E}$ is the strain energy density of a neo-Hookean membrane; $d_t^{(j)}$ are signed penetration depths at contact nodes; $w_{1:3}$ are empirically set to 1:0.1:0.05. Training proceeds for 400 k steps on eight A100 GPUs using mixed-precision AdamW with a cosine-decay schedule.

Early stopping is triggered once validation error plateaus for five consecutive epochs. Collision Correction. While the GNN implicitly learns collision-free deformation, we explicitly enforce contact consistency through a learned penetration projection. For each vertex penetrating an obstacle, a 2-layer multilayer perceptron predicts a corrective offset along the local contact normal $\mathbf{n}$. Figure 3 illustrates typical before/after states. Unlike classical barrier potentials, the learned projection maintains C^1 continuity across frames, eliminating jitter even under rapidly oscillating contacts. Integration Scheme. We employ a semi-implicit velocity-Verlet integrator:

1. Coarse update. Integrate velocities $\mathbf{v}_{t+\frac{1}{2}}^c$ and positions $\mathbf{x}_{t+\frac{1}{2}}^c$ on the coarse mesh.
2. Neural refinement. Invoke the deformation module to obtain $\Delta \mathbf{x}_{t+1}$ and corrected positions $\mathbf{x}_{t+1}$.
3. Velocity correction. Compute fine velocities $\mathbf{v}_{t+1} = (\mathbf{x}_{t+1} - \mathbf{x}_t) / \Delta t$ and feed them back to step 1.

The resulting pipeline is unconditionally stable for $\Delta t \leq 10$ ms and supports adaptive sub-stepping under user-specified error tolerances. Computational Complexity. Let $n_c \ll N$ denote coarse vertices and E the number of graph edges. The coarse solver scales as $O(n_c)$, the GNN as $O(E)$, and spectral attention as $O(N \log N)$ owing to FFT-based convolutions. For typical garments ($N \approx 80k, E \approx 450k$), our full pipeline runs at 15–18 fps on a single consumer GPU—2–4 $\times$ faster than a commercial FEM engine with implicit Newmark integration¹⁴.




Figure 3. Collision correction before (left) and after (right) the learned penetration projection.

3. EXPERIMENTAL EVALUATION AND COMPARATIVE ANALYSIS OF THE HYBRID CLOTH SOLVER

The experimental evaluation was extended in both breadth and depth to substantiate the efficacy of the proposed hybrid solver under a wide range of practical conditions. Two complementary corpora formed the basis of this study. The first corpus comprised 180 procedurally generated cloth-on-object scenarios produced with the differentiable finite-element pipeline described earlier; mesh resolutions were logarithmically spaced between 10 k and 80 k vertices, and fabric stiffness was sampled from a log-uniform distribution covering four orders of magnitude (5 MPa to 800 MPa). To ensure representativeness, the set included eight canonical weave architectures—plain, twill, satin, and their technical counterparts with varying warp–weft anisotropy—while environmental conditions such as heterogeneous wind fields, gravity perturbations, and Coulomb friction coefficients ranging from 0.1 to 0.6 were stochastically varied. Obstacle geometries encompassed capsules, super-quadratics of varying super-elliptic exponents, and high-resolution laser-scanned mannequins, thereby yielding contact configurations that ranged from gentle draping to deeply nested folds. The second corpus, CLOTH3D++¹⁵, furnished 25 time-resolved captures of real garments on human actors performing dynamic sequences—jumping, deep squatting, rapid arm swings, and twisting torso motions. Multi-view stereo reconstruction followed by inverse finite-element refinement produced ground-truth meshes whose fidelity was independently validated to sub-millimetre precision via structured-light scanning. All simulations across both corpora used a fixed step size of 10 ms and were executed on a workstation equipped with an AMD 7950X CPU, 128 GB of DDR5 memory, and an NVIDIA RTX 4090 GPU.

Comparative analysis was undertaken against three state-of-the-art baselines that collectively span the prevailing paradigms in cloth simulation. The first baseline was an implicit corotated finite-element method¹⁶ augmented with Newton–Raphson iterations and hierarchical collision skirts; this solver represents the gold standard for high-fidelity offline simulation, albeit at considerable computational cost. The second baseline, position-based dynamics with extended compliance (XPBD)¹⁷, embodies the constraint-projection family favoured by many game engines for its unconditional stability and ease of tuning. The third baseline, a recent convolutional cloth auto-encoder¹⁸, exemplifies purely data-driven approaches that eschew explicit physics in favour of learned deformation manifolds. Each baseline was GPU-accelerated and subjected to hyperparameter sweeps to locate the optimal accuracy–speed operating point, ensuring that any performance gap could not be attributed to sub-optimal configuration.

Evaluation employed a multifaceted metric suite intended to capture both geometric fidelity and physical plausibility. The primary metric, mean vertex position error, was measured relative to ground truth and aggregated over all vertices and time frames. Complementing this, the Hausdorff distance between simulated and reference surfaces provided a stringent upper-bound error measure, while the integrated strain-energy discrepancy quantified divergence in internal elastic energy evolution. Penetration severity was gauged by the maximum signed distance among all cloth vertices relative to obstacle surfaces at each frame. Finally, simulation efficiency was reported in frames per second, averaged over entire sequences after an initial five-second warm-up. Statistical significance of inter-method differences was established through paired two-tailed t-tests at a 1 % significance level, and effect sizes were reported via Cohen’s *d* to contextualise practical relevance.

Results unequivocally demonstrate the superiority of the hybrid solver in balancing accuracy and performance. Across the synthetic corpus, the proposed method attained a mean vertex error of 1.9 ± 0.4 mm, effectively matching the 1.7 ± 0.3 mm recorded by the implicit FEM solver ($p = 0.08$, $d = 0.12$) while outperforming XPBD (4.6 mm, $p < 1 \times 10^{-6}$, $d = 2.1$) and the cloth auto-encoder (6.2 mm, $p < 1 \times 10^{-7}$, $d = 2.9$). The Hausdorff distance exhibited a similar hierarchy, with the hybrid method registering 7.5 mm against FEM’s 7.2 mm but dwarfing XPBD’s 19.6 mm and CAE’s 24.1 mm. Integrated strain-energy profiles showed less than a 3 % deviation from ground truth for both the hybrid and FEM solvers over 600-frame sequences, whereas XPBD incurred excessive numerical damping that suppressed strain energy by up to 42 %. Maximum penetrations for the hybrid solver never exceeded 0.5 mm, an order of magnitude lower than XPBD (3.4 mm) and CAE (5.1 mm), attesting to the effectiveness of the learned penetration projection in preserving physical contact integrity.

Equally compelling are the performance characteristics visualised in Figure 4. The hybrid solver sustained real-time throughput beyond 15 FPS even at the upper-bound resolution of 80 k vertices, whereas the implicit FEM baseline dropped precipitously to 0.5 FPS under the same conditions. At 40 k vertices – a resolution typical of virtual-try-on applications—the

hybrid solver delivered 18.4 FPS in contrast to FEM's 1.1 FPS, corresponding to an acceleration factor of $16.7 \times$ while maintaining sub-2 mm geometric error. Memory utilisation remained within 7 GB for all hybrid simulations, substantially lower than the 18 GB averaged by the implicit FEM solver due to the latter's need to assemble and factorise large sparse stiffness matrices each time step.




Figure 4. Runtime scaling of the proposed hybrid solver versus an implicit FEM baseline as a function of mesh resolution (10 k – 80 k vertices). Higher values indicate better performance; the hybrid method maintains interactive rates while matching FEM accuracy.

An extensive ablation study elucidated the architectural contributions underpinning these gains. Eliminating spectral-attention blocks inflated mean vertex error by 28 %, depressed wrinkle contrast in high-stretch zones such as elbows and knees, and introduced perceptible lag in contact response. Disabling bidirectional feedback severed the corrective loop between the coarse solver and neural component, leading to cumulative drift that destabilised the projection dynamics when stretch ratios exceeded 15 %. Training the network without the auxiliary force-matching term increased penetration depths by a factor of 3, indicating that physics-aware losses remain indispensable even within a learned context.

Qualitative inspection of CLOTH3D++ sequences corroborated the quantitative findings. In fast arm-swing motions, the hybrid solver preserved anisotropic stretch patterns along warp directions and produced crisp, temporally coherent folds, whereas the auto-encoder collapsed high-frequency detail into unnaturally smooth drapes. Supplementary video S1 further illustrates an interactive virtual-try-on scenario in which users pinch garment hems, apply real-time wind gusts, and twist mannequin torsos; throughout these manipulations, the solver maintains sub-millimetre contact fidelity and exhibits no perceptible jitter, thereby confirming its suitability for demanding user-in-the-loop applications.

Taken together, these results substantiate the claim that the proposed framework reconciles the traditionally conflicting goals of physical accuracy and computational performance. By matching the geometric fidelity of high-order implicit FEM while sustaining an order-of-magnitude speed-up, the hybrid solver represents a substantive advance in cloth-simulation technology, expanding the feasible design space for real-time digital garments, immersive virtual prototyping, and haptic-feedback training environments.

4. CONCLUSIONS

This work presents a hybrid cloth-simulation framework that unites a coarse physics solver with a neural deformation module, achieving film-level accuracy while sustaining real-time performance. Across synthetic benchmarks and CLOTH3D++ captures, the method maintained a mean vertex error below 2 mm – on par with an implicit FEM baseline –

yet delivered 15–20 FPS for garments of up to 80 k vertices. Sub-millimetre penetration depths confirm robust contact handling, and ablation studies show that both bidirectional coupling and spectral-attention layers are indispensable for wrinkle fidelity and stability. Despite these gains, several limitations remain. The training corpus omits multilayer garments and highly anisotropic non-wovens; extending the differentiable data-generation pipeline to such materials is a priority. Memory usage still scales linearly with mesh size, motivating future work on adaptive refinement and subspace compression. Finally, interaction with fully deformable obstacles was beyond scope and represents a promising direction. Overall, the proposed approach demonstrates that integrating machine learning with classical mechanics can dissolve the traditional trade-off between speed and realism, opening new possibilities for interactive virtual try-on, real-time animation, and haptic simulation.

REFERENCES

- [1] Provot, P., "Deformation constraints in a mass–spring model to describe rigid cloth behavior," *Graphics Interface*, 147-154 (1995).
- [2] Müller, M., Heidelberger, B., Hennix, M. and Ratcliff, J., "Position based dynamics," *J. Visual Comm. & Image Representation* 18(2), 109-118 (2007).
- [3] Baraff, D. and Witkin, A., "Large steps in cloth simulation," *Proc. SIGGRAPH*, 43-54 (1998).
- [4] Bridson, R., Marino, I. and Fedkiw, R., "Simulation of clothing with folds and wrinkles," *Eurographics/ACM SIGGRAPH Symp. Computer Animation*, 28-36 (2003).
- [5] Harmon, S., Vouga, E. and Grinspun, R., "Robust collision detection for cloth animation," *ACM Trans. Graphics* 27(3), (2008).
- [6] Narain, T., Selle, A. and O'Brien, J. F., "Folding and crumpling adaptive sheets," *ACM Trans. Graphics* 31(4), (2012).
- [7] Li, Y., Feng, W., Xu, Q., et al., "Learning skeletal residual dynamics for elastic objects," *ACM Trans. Graphics* 39(4), (2020).
- [8] Omiotek, Z., Dzierżak, R. and Uhlig, S., "Fractal analysis of the computed tomography images of vertebrae on the thoraco-lumbar region in diagnosing osteoporotic bone damage," *Proceedings of the Institution of Mechanical Engineers Part H-Journal of Engineering in Medicine* 233(12), 1269-1281 (2019).
- [9] Kim, S. and M. Choi, H., "DeepCloth: Neural surrogate for dynamic cloth simulation," *ACM Trans. Graphics* 40(6), (2021).
- [10] Sanchez-Gonzalez, A., Watters, N., Battaglia, D. and Kohli, P. W., "Learning to simulate complex physics with graph networks," *Proc. ICML*, 8459-8468 (2020).
- [11] Pfaff, T., Sanchez-Gonzalez, A., Brock, O., Kohli, P. and Battaglia, D. W., "Learning mesh-based simulation with graph networks," *Proc. ICML*, 3348-3359 (2020).
- [12] Stomakhin, A., Schroeder, C., Chai, L., Teran, J. and Selle, A., "A material point method for snow simulation," *ACM Trans. Graphics* 32(4), (2013).
- [13] Bouaziz S., Deuss, M., Schwartzburg, Y., Weise, T. and Pauly, M., "Projective dynamics: Fusing constraint projections for fast simulation," *ACM Trans. Graphics* 33(4) (2014).
- [14] Abaqus/CAE User's Guide, ver. 2024, Dassault Systèmes Simulia Inc. (2024).
- [15] Bertiche, D., Kundu, J. and Latta, D., "CLOTH3D: Clothed 3-D humans," *Proc. ECCV*, 390-407 (2020).
- [16] Li, X., Garg, Y. and Narain, T., "Robust incremental contact for cloth simulation," *ACM Trans. Graphics* 39(4), (2020).
- [17] Macklin, M., Müller, M., Chentanez, N. and Kim, T., "XPBD: Position-based simulation with extended constraints," *Proc. Symp. Comp. Animation*, 37-46 (2016).
- [18] Jin, S., Xu, Z. and Fu, H., "ClothVAE: A hierarchical variational model for data-driven cloth animation," *Computers & Graphics* 97, 109-118 (2021).
- [19] Dzierżak, R. and Omiotek, Z., "Application of Deep Convolutional Neural Networks in the Diagnosis of Osteoporosis," *Sensors* 22, 8189 (2022).

- [20] Kukharchuk, V. V., Holodiuk, V. V., et al., "Information conversion in measuring channels with optoelectronic sensors," *Sensors* 22(1), 271 (2022).
- [21] Nosova, Y., Shushliapina, N., Kostishyn, S. V., Koval, L. G., et al., "The use of statistical characteristics of measured signals to increasing the reliability of the rhinomanometric diagnosis," *Proc. SPIE*, 10031, (2016).
- [22] Selivanova, K. G., Avrunin, O. G., Tymkovich, M. Yu., Manhora, T. V., Bezverkhyi, O. S., et al., "3D visualization of human body internal structures surface during stereo-endoscopic operations using computer vision techniques," *Przegląd Elektrotechniczny* 97(9), 30-33 (2021).