Journal of Fiber Bioengineering and Informatics 19:1 (2026) 93–112

Cite as: Q. Yang et al.
doi:10.3993/jfbim26007

Robust Real-Time Garment Fitting from 3D Point Clouds with Physics-Guided Uncertainty and Reliability Monitoring

Qian Yang [a,b,*], Ming Li [c]

[a] *Universiti Kuala Lumpur, 1016, Jalan Sultan Ismail, 50250, Kuala Lumpur, MALAYSIA*

[b] *Minjiang University, Faculty of Clothing and Design, No.200, Xiyuangong Road, Shangjie Town, Minhou County, Fuzhou City, Fujian Province, China*

[c] *Lam Family College of Business, San Francisco State University, San Francisco, USA*

Abstract

Real-time garment fitting from 3D point clouds is a fundamental capability for telepresence, virtual try-on, and augmented reality. Yet, it remains difficult due to high-dimensional non-rigid deformation, frequent self-occlusion, and the need to maintain physically plausible cloth behavior under noisy measurements. Existing studies have improved geometric registration, learning-based reconstruction, or cloth simulation in isolation, but still lack a unified real-time framework that combines dense vertex-level tracking, uncertainty-aware sequential estimation, and physics-guided regularization for robust garment fitting under noisy and incomplete 3D observations. To address this gap, we propose Dense-Decoupled Adaptive Kalman Filtering (DD-AKF), an integrative sequential estimation framework that tracks dense garment vertices in real time while explicitly propagating uncertainty. To make high-dimensional filtering tractable, DD-AKF employs a block-diagonal (decoupled) covariance approximation that reduces per-frame complexity to linear time in the number of vertices while retaining per-vertex uncertainty estimates. Physical plausibility is incorporated by introducing differentiable stretch, bending, and collision energies as soft constraints in a Gaussian-approximate (sequential MAP) update, enabling an adaptive trade-off between sensor fidelity and physics under degraded observations. We also compute an online reliability score based on innovation (residual) statistics to monitor tracking quality and trigger robustification in the presence of occlusion or sensor corruption. Experiments on FAUST, CLOTH3D, and real capture sequences show reduced temporal flicker and penetration compared with representative optimization-based, stochastic filtering, and learning-based baselines, while maintaining competitive geometric accuracy at interactive rates (30+FPS). These results support reliability-aware deployment in practical virtual try-on and AR garment-fitting pipelines using commodity 3D sensors.

Keywords: real-time garment fitting; 3D point cloud processing; cloth deformation modeling; physics-guided Kalman filtering; uncertainty-aware estimation; reliability monitoring

*Corresponding author.
Email address: yang.qian@s.unikl.edu.my (Qian Yang).

1 Introduction

Real-time digitization of dynamic, clothed humans enables telepresence, virtual try-on, and augmented reality (AR) retail [1–3]. While parametric body models such as SMPL and GHUM are widely adopted [4–6], stable garment fitting from raw 3D point clouds remains difficult. Clothing exhibits high-dimensional non-rigid deformation, complex topology, and self-contact, and visually plausible motion must respect physical regularities such as limited stretch, bending resistance, and non-penetration [7–9].

In practical capture pipelines that use consumer depth cameras or LiDAR, garment tracking is further degraded by measurement noise, partial observations due to self-occlusion (e.g., crossed arms), and abrupt changes in the visible surface support [10, 11]. For interactive use, a fitting method should therefore (i) scale to dense garment states, (ii) remain robust when observations are missing or corrupted, and (iii) encourage physical plausibility without exceeding real-time budgets.

Existing approaches broadly follow two paradigms: optimization-based registration and learning-based inference. Classical non-rigid registration methods—such as Non-Rigid ICP [12] and Embedded Deformation [13]—can achieve good per-frame geometric alignment when correspondences are reliable [14]. However, solving a new optimization at each frame can lead to drift under occlusion and high-frequency jitter ("temporal flicker") under noisy measurements [15, 16]. Incorporating strong physical constraints (especially collision handling and self-contact) into per-frame optimization further increases computational cost, often beyond real-time budgets [17, 18]. Learning-based methods infer garment geometry from pose cues or sparse observations [19, 20]. They can produce visually plausible reconstructions on in-distribution data (e.g., DeepFashion3D [21], SNUG [22]), but they may degrade under domain shift in motion, material behavior, or garment topology [23–25]. Many learning pipelines also provide limited calibrated uncertainty at test time, which makes online failure detection and mitigation more difficult [26, 27].

Despite these advances, a clear gap remains for interactive applications: existing methods rarely provide dense, point-cloud-driven garment tracking that simultaneously (i) propagates per-vertex uncertainty for confidence-aware updates, (ii) incorporates physical feasibility within the sequential estimator, and (iii) offers an online reliability signal to detect and mitigate failures under occlusion and sensor noise.

To address this gap, this study aims to develop a real-time garment-fitting framework based on 3D point clouds that can robustly handle noisy and incomplete observations while maintaining physically plausible cloth behavior. We propose Dense-Decoupled Adaptive Kalman Filtering (DD-AKF). This sequential estimation approach combines scalable uncertainty propagation with physics-guided constraints and an innovation-based reliability signal for adaptive robustification.

At a high level, DD-AKF maintains a temporally coherent prior through sequential filtering, adapts measurement trust using per-vertex uncertainty, and uses differentiable cloth energies (stretch, bending, and collision) as soft physical regularizers when correspondence quality deteriorates. A lightweight reliability score derived from innovation statistics provides online diagnostics and guides robustification in the presence of occlusion or sensor corruption.

The main contributions are: (1) a dense per-vertex garment state-space formulation with a block-diagonal uncertainty representation to enable scalable real-time sequential estimation; (2) a physics-guided update that injects differentiable stretch, bending, and collision energies as soft constraints in a sequential MAP view; (3) an innovation-based reliability score used to adapt mea-

surement noise and regularization strength under degraded observations; and (4) an evaluation on FAUST, CLOTH3D, and real capture sequences reporting geometric accuracy, physical-fidelity metrics, temporal stability, and statistical significance against representative baselines.

To clarify the originality of this work in the context of prior publications, we note several studies on body-size prediction from 2D images [49], virtual garment prototyping and simulation workflows [50], and parametric bra sample modeling [51]. At the same time, recent TBIS proceedings have explored CLO3D-based virtual dress simulation and fit/space analysis [52, 53]. In contrast, our focus is real-time, point-cloud-driven garment fitting with online uncertainty propagation, physics-guided filtering, and reliability monitoring for robust sequential tracking under occlusion and sensor noise.

The remainder of the paper is organized as follows. Section 1.1 reviews related work. Section 2 describes the proposed methods and experimental setup. Section 3 reports quantitative and qualitative results. Section 4 discusses the findings and limitations, and Section 5 concludes the paper.

1.1 Background and Related Work

We organize related work around the main problem dimensions highlighted in Section 1: (i) dense deformable garment tracking and non-rigid registration, (ii) robustness under occlusion and noisy point-cloud measurements, and (iii) enforcing physically plausible cloth behavior. We then review uncertainty-aware state estimation and reliability monitoring for high-dimensional tracking.

1.1.1 Dense Deformable Garment Tracking and Registration

Garment fitting from 3D observations is often posed as a non-rigid registration problem. Rigid and non-rigid ICP variants provide a geometric alignment backbone [12, 36, 37], while deformation-graph formulations such as Embedded Deformation (ED) [13] parameterize deformation with a sparse set of control nodes and interpolate their motion to the dense surface. This sparse control representation improves efficiency and regularizes ill-posed correspondences, but it can attenuate fine-scale garment detail. It also introduces an additional interpolation layer between the estimated state and dense physical terms (e.g., stretch, bending, collision) defined on the garment mesh. We refer to this representation gap as a *physics–state mismatch*: physics is naturally expressed on dense vertices/edges, whereas the estimator state may live on a coarser proxy, requiring interpolation of forces or gradients and complicating physically consistent correction during fitting [13, 15].

Real-time non-rigid reconstruction systems such as DynamicFusion [15] and KillingFusion [16] achieve interactive performance by warping a canonical template, typically through per-frame optimization with regularization. However, because a new objective is solved at each frame, the solution can change abruptly when the set of reliable correspondences changes (e.g., under self-occlusion). This contributes to drift in unobserved regions and to high-frequency jitter ("temporal flicker") under noisy measurements, consistent with behavior reported in optimization-only pipelines [17, 18].

Learning-based approaches infer garment geometry from pose cues or sparse observations [19, 20]. Methods such as DeepFashion3D [21] and SNUG [22] can produce visually plausible reconstructions on in-distribution data, and learned priors can help complete occluded regions. How-

ever, performance can degrade under domain shift across motion patterns, material behavior, or garment topology [23–25]. Moreover, many learning pipelines output point estimates without calibrated uncertainty, limiting their usefulness for online failure detection and mitigation in interactive fitting [26, 27].

Recent work has also begun to leverage implicit neural representations and diffusion-based generative priors to reconstruct clothed humans and garments under severe occlusion. For example, Human-3Diffusion couples multi-view diffusion priors with an explicit 3D representation to better handle loose clothing and occlusions [54]. While such approaches can improve synthesis quality, they are typically designed for offline reconstruction or image-based input. They do not directly address real-time, point-cloud-driven sequential fitting with online uncertainty diagnostics, which is the focus here.

1.1.2 Robustness to Occlusion and Noisy Point Clouds

Robustness to occlusion and measurement noise is a central challenge in point-cloud-based fitting. In consumer depth capture, missing returns, multi-path artifacts, and self-occlusion can yield incorrect or intermittent correspondences [10, 11]. Classical registration pipelines typically rely on outlier rejection and robust penalties. However, severe occlusion can still leave the inverse problem underconstrained, leading to drift or discontinuous updates when visible surface support changes [15–17].

Sequential estimators offer an alternative by maintaining a predictive prior over time. In principle, prediction can stabilize updates when correspondence support becomes sparse, and uncertainty can be used to down-weight unreliable observations. However, applying standard filtering to dense deformable surfaces is non-trivial because the state dimension scales with the number of vertices, and a full covariance matrix is computationally prohibitive [33]. This motivates scalable approximations that preserve a temporal prior while remaining compatible with real-time constraints.

In addition to temporal prediction, practical systems often benefit from online diagnostics that indicate when tracking quality is degrading. Such diagnostics are particularly important in occlusion-heavy motion, where point-cloud support can change rapidly, and geometric residuals alone may not reliably signal impending failure.

1.1.3 Physics-Based Priors and Physical Plausibility

Physically based simulation remains a reference standard for garment realism. Energy- or constraint-based cloth models capture stretch, bending, and contact effects using force-based formulations [7] or position-based methods such as PBD and XPBD [11, 35], and have been extended to handle folds, wrinkles, and collisions [8, 9]. While such models can generate plausible forward dynamics, embedding their full constraint structure into inverse problems (tracking or fitting to noisy measurements) typically requires iterative optimization and collision handling, which can be expensive at dense resolution.

Consequently, many fitting pipelines use lightweight geometric regularizers (e.g., ARAP-style terms or deformation-graph smoothness) and treat physics as a separate simulation or post-processing step. These strategies may not provide a controllable mechanism to trade off sensor fidelity against physical plausibility when observations are degraded. In contrast, our formula-

tion incorporates differentiable cloth energies directly into the estimator as soft constraints in a sequential MAP update. This is conceptually related to regularized MAP estimation, but here it is implemented in a dense, real-time filtering loop with explicit uncertainty propagation.

Recent neural digitization and neural simulation work further improves visual realism, for example by learning garment deformation priors from large datasets such as CLOTH3D [41] or by learning physically inspired deformation fields (e.g., NeuralClothSim [44]). These models are valuable for forward prediction and synthesis, but uncertainty-calibrated, point-cloud-driven fitting with explicit physical feasibility control remains comparatively less explored.

1.1.4 Uncertainty, Reliability, and High-Dimensional Filtering

Uncertainty-aware estimation is increasingly used to diagnose failures in vision and tracking pipelines, ranging from Bayesian deep learning for computer vision [26] to uncertainty-aware 3D pose estimation [48]. In classical filtering theory, innovation (residual) statistics provide a principled signal for consistency checking and gating. For dense deformable tracking, a key obstacle is scalability: a naive Kalman filter requires maintaining an O(nˆ2) covariance for an O(n)-dimensional state, which is intractable for meshes with thousands of vertices [33]. Decoupled or block-diagonal Kalman filtering approximations address this by storing only local covariance blocks and neglecting long-range correlations to enable tractable propagation [32, 46], thereby enabling scalable implementations in large systems [47]. Building on these ideas, we use per-vertex uncertainty together with an innovation-based reliability score to support online monitoring and adaptive robustification in dense garment fitting.

Despite progress in registration, learning-based inference, and scalable filtering, real-time point-cloud garment-fitting methods rarely provide dense, uncertainty-aware tracking, along with explicit physical-feasibility control and online reliability monitoring. This gap motivates the integrated design of DD-AKF and our research.

2 Methods

We propose DD-AKF, an uncertainty-aware sequential fitting framework for real-time garment tracking from 3D point clouds. The method estimates garment motion at dense-vertex resolution while remaining robust to noisy and partial observations and enforcing physical plausibility. To avoid physics–state mismatch when sparse control representations are coupled to dense physical energies (Section 1.1.1), we define the state directly on mesh vertices and use a decoupled (block-diagonal) covariance approximation for tractable uncertainty propagation.

2.1 Problem Setup and Notation

Let a manifold triangle mesh represent the garment surface $\mathcal{M} = (\mathcal{V}, \mathcal{F})$, where $\mathcal{V} = \{v_i\}_{i=1}^{N_v}$ is the set of vertices and $\mathcal{F}$ is the set of triangular faces. Unlike methods that cluster vertices into sparse super-nodes, we assign a state vector to every vertex. v_i.

For each vertex i, the state vector $\mathbf{x}_{i,k} \in \mathbb{R}^6$ at frame k is defined as:

$$\mathbf{x}_{i,k} = [\mathbf{p}_{i,k}^T, \mathbf{u}_{i,k}^T]^T$$

where $\mathbf{p}_{i,k} \in \mathbb{R}^3$ is the position and $\mathbf{u}_{i,k} \in \mathbb{R}^3$ is the velocity. The global state vector $\mathbf{X}_k \in \mathbb{R}^{6N_v}$ is the concatenation of all vertex states.

Before estimation, we preprocess each input point cloud by applying a k-nearest neighbors (k-NN) statistical outlier-removal filter to suppress isolated noise and spurious returns. We then estimate target surface normals using local PCA, which provides stable normal directions for computing point-to-plane geometric residuals and for supporting collision-related computations.

2.2 Decoupled State-Space Model

A standard Kalman Filter on the full state $\mathbf{X}_k$ would involve a covariance matrix $\mathbf{P}_k$ of size $6N_v \times 6N_v$. Inverting this matrix implies a computational complexity of $O(N_v^3)$, which is intractable for real-time applications where $N_v \approx 2\,000 - 5\,000$.

To ensure real-time throughput, we store a decoupled covariance approximation: we retain only the block-diagonal blocks of the full covariance matrix and neglect cross-vertex covariance terms during propagation. In other words, the maintained covariance is block-diagonal, with one local block per vertex state. This assumption is motivated by scalability and by the locality of measurement and physical interactions in cloth.

$$\mathbf{P}_k = \mathrm{diag}(\mathrm{P}_{1,k}, \mathrm{P}_{2,k}, \cdots, \mathrm{P}_{N_v,k})$$

where each $\mathbf{P}_{i,k} \in \mathbb{R}^{6\times 6}$ represents the local uncertainty of the vertex i. This reduces the complexity of the prediction and update steps to $O(N_v)$.

We assume a constant velocity model for each vertex, evolving independently during the prediction phase:

$$\hat{\mathbf{x}}_{i,k|k-1} = \mathbf{F}\mathbf{x}_{i,k-1|k-1}, \quad \mathbf{P}_{i,k|k-1} = \mathbf{F}\mathbf{P}_{i,k-1|k-1}\mathbf{F}^T + \mathbf{Q}_i$$

where

$$\mathbf{F} = \begin{bmatrix} \mathbf{I}_3 & \Delta t\mathbf{I}_3 \\ \mathbf{0}_3 & \mathbf{I}_3 \end{bmatrix}, \quad \text{and } \mathbf{Q}_i \text{ is the process noise covariance.}$$

2.3 Adaptive Geometric Measurement Update

For each predicted vertex position. $\hat{\mathbf{p}}_{i,k|k-1}$, we find the nearest neighbor $\mathbf{z}_{i,k}$ in the target point cloud $\mathcal{T}_k$ via a KD-tree search. The geometric measurement equation is:

$$\mathbf{z}_{i,k} = \mathbf{H}\mathbf{x}_{i,k} + \mathbf{v}_{i,k}, \quad \mathbf{v}_{i,k} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_{i,k})$$

where $\mathbf{H} = [\mathbf{I}_3, \mathbf{0}_3]$.

To handle occlusions and sensor noise robustly, we adapt the measurement covariance. $\mathbf{R}_{i,k}$ for each vertex based on its residual statistics. Let $\mathbf{y}_{i,k} = \mathbf{z}_{i,k} - \mathbf{H}\hat{\mathbf{x}}_{i,k|k-1}$ be the innovation. We update $\mathbf{R}_{i,k}$ using exponential smoothing:

$$\mathbf{R}_{i,k} \leftarrow \alpha\mathbf{R}_{i,k-1} + (1-\alpha)(\mathbf{y}_{i,k}\mathbf{y}_{i,k}^T + \epsilon\mathbf{I})$$

This allows the filter to "learn" local sensor noise. In regions of high disagreement (e.g., occlusion boundaries), $\mathbf{R}_{i,k}$ inflates, causing the filter to rely more on the motion model and physics priors.

2.4 Physics-Guided Pseudo-Measurements

While the prediction step is decoupled, the vertices must physically interact (e.g., stretch, bend, or collide). We introduce these interactions as Physics-Guided Pseudo-Measurements.

Since the state is defined on the dense vertices, the physics observation function $\mathbf{h}_{\text{phys}}(\mathbf{X})$ is simply the gradient of the total potential energy $E(\mathcal{V})$:

$$\mathbf{z}_{\text{phys}} = \mathbf{0}, \quad \mathbf{h}_{\text{phys}}(\mathbf{X}) = \nabla_{\text{p}} E_{\text{total}}(\mathbf{p})$$

where $E_{\text{total}} = E_{\text{stretch}} + E_{\text{shear}} + E_{\text{bend}} + E_{\text{coll}}$.

Exact gradient coupling is achieved by computing physical derivatives directly on the dense mesh, rather than approximating forces through interpolation from sparse super-nodes. Specifically, for a vertex i We evaluate the cloth energies on the mesh and construct the corresponding local linearization. The resulting Jacobian, $\mathrm{H}_{\text{phys},i} = \partial \mathrm{h}_{\text{phys}} / \partial \mathbf{x}_i$, is sparse by construction: it depends only on the state of the vertex i and its immediate 1-ring neighbors as defined by the mesh connectivity.

Because we store a block-diagonal covariance, DD-AKF does not explicitly represent long-range error correlations. This does not mean that vertices evolve independently in the mean state: the physics term is evaluated on the dense mesh and couples neighboring vertices through stretch, bending, and collision energies. As a result, constraint responses propagate locally across the surface during each update (e.g., stretch forces link adjacent vertices through mesh edges).

In a full EKF, such coupling would also introduce off-diagonal covariance blocks during the update. To retain linear-time complexity, we apply the physics term primarily as a mean-state correction using the sparse physics Jacobian, while keeping only the per-vertex covariance blocks in storage. Long-range correlations induced by physics are therefore not retained explicitly; instead, they are addressed implicitly over time through repeated local interactions, geometric measurement updates, and the adaptive noise mechanism.

Fabric types and physical parameters. The stretch and bending terms in E_{total} are parameterized by material stiffness coefficients (e.g., in-plane stretch stiffness and out-of-plane bending stiffness). Different fabrics (e.g., cotton vs. silk) can be represented by adjusting these coefficients and/or the relative weights between stretch, bending, and collision energies. In this work, we use a single normalized parameter setting across garments to avoid per-material tuning and to focus on the filtering formulation. When fabric-specific parameters are available (e.g., from CAD/simulation presets or material measurements), they can be plugged into the same energy terms without changing the DD-AKF recursion. In practice, stiffer settings enforce stronger physical corrections (improving stability but potentially over-smoothing fine folds), whereas softer settings preserve flexibility but may require stronger reliance on geometric observations to avoid drift.

2.5 Robustness Mechanisms

To reduce the risk of divergence under non-Gaussian noise and correspondence errors, we incorporate two robustification mechanisms. First, we apply an M-estimator that down-weights large innovations using a Tukey bisquare weight. $w(\mathbf{y}_{i,k})$, thereby suppressing the influence of outliers such as incorrect correspondences. Second, we perform Mahalanobis gating and reject measurements whose normalized innovation exceeds a threshold: $d_M^2 = \mathbf{y}_{i,k}^{\mathrm{T}} \boldsymbol{S}_{i,k}^{-1} \mathbf{y}_{i,k} > \tau$.

2.6 Online Reliability Diagnostics

To monitor tracking quality online, we compute a per-frame reliability score that captures the aggregate normalized innovation energy over all garment vertices. Specifically, we define

$$s_k = \exp\left(-\frac{1}{N_v}\sum_{i=1}^{N_v}\frac{\|\mathbf{y}_{i,k}\|^2}{\mathrm{trace}(\mathbf{R}_{i,k})}\right)$$

where $\mathbf{y}_{i,k}$ is the innovation (residual) for vertex i and $\mathbf{R}_{i,k}$ is its adaptively estimated measurement covariance. This statistic is closely related to innovation-based consistency measures used in Kalman filtering (e.g., the normalized innovation squared), and we use it here as a practical online reliability indicator. When the score falls below a preset threshold ($s_k < \tau$), we activate a recovery mode that increases the physics weight λ, thereby strengthening physical regularization during brief periods of occlusion or degraded measurements.

2.7 Theory Note: Stability and Boundedness

DD-AKF is designed to avoid divergence in the presence of intermittent correspondences by keeping the posterior covariance bounded. We provide an intuition based on three mechanisms: (i) bounded adaptive measurement noise, (ii) additional constraints from physics pseudo-measurements, and (iii) a decoupled (block-diagonal) covariance representation that reduces the high-dimensional Riccati update to per-vertex subproblems.

(i) Bounded adaptive noise. The per-vertex measurement covariance $\mathbf{R}_{i,k}$ is updated from residual statistics but clamped to a user-specified interval $[\mathbf{R}_{\min}, \mathbf{R}_{\max}]$ and smoothed with factor α. These bounds prevent the filter from becoming overconfident (too small R) or unresponsive (too large R), keeping the Kalman gains numerically well-behaved under heavy-tailed noise and intermittent correspondence errors.

(ii) Physics constraints improve observability when geometry degrades. Under self-occlusion or locally planar regions, purely geometric updates may become weak or ambiguous. The physics pseudo-measurements contribute additional Jacobian terms that encode local stiffness and contact responses, supplying informative constraints even when point-cloud correspondences are sparse. In a sequential MAP view, these terms act as soft regularizers that become more influential when the likelihood from geometric measurements degrades.

(iii) Decoupled covariance recursion. Because the stored covariance is block-diagonal (one local block per vertex), the Riccati recursion decomposes into independent low-dimensional updates. Under bounded system matrices and bounded noise covariances, boundedness of each per-vertex covariance block implies boundedness of the full block-diagonal covariance. A full proof follows standard boundedness arguments for time-varying Kalman filtering and is omitted for brevity.

2.8 Algorithm (DD-AKF)

The complete per-frame execution loop is defined as follows:

A. Prediction: Evolve state $\hat{\mathbf{x}}_{k|k-1}$ using a constant velocity model.

B. Association: Find nearest neighbors in the target point cloud.

C. Robustness: Compute M-estimator weights; Gate outliers.

D. Adaptation: Update $\mathbf{R}_{i,k}$ based on residual magnitude.

E. Physics Gradient: Compute $\nabla E_{\text{stretch}}, \nabla E_{\text{bend}}, \nabla E_{\text{coll}}$ on dense mesh.

F. Fusion Update: Jointly update state with geometric and physics residuals.

G. Reliability Check: Compute s_k. If $s_k <$ threshold, inflate λ to enforce stiffer physics constraints.

H. Output: Render vertices $\mathbf{p}_{k|k}$.

2.9 Experimental Setup

To validate the proposed framework, we conduct extensive comparisons against state-of-the-art baselines on both synthetic and real-world datasets. The experiments are designed to assess geometric accuracy, physical plausibility, and temporal stability under varying levels of sensor noise and occlusion.

2.9.1 Datasets

We evaluate on two datasets: FAUST and CLOTH3D. FAUST is a dataset of real human scans with ground-truth correspondences across diverse poses. We uniformly sample point clouds from the meshes and add isotropic Gaussian noise with standard deviation $\sigma = 3$ mm (unless otherwise stated) to emulate sensor noise. Target surface normals are estimated via local PCA. We use a leave-subject-out protocol, training (if applicable) on 8 subjects and testing on the remaining 2.

CLOTH3D is a large-scale synthetic dataset of clothed humans with diverse garment types (e.g., T-shirts, skirts, loose robes) [41]. We use the garment mesh vertices as ground truth and precompute rest edge lengths and rest dihedral angles from the template mesh. To stress-test robustness, we simulate severe self-occlusions (e.g., crossing arms) and rapid motions. We evaluate in a leave-garment-category-out setting to test generalization to unseen garment topologies.

2.9.2 Baselines

We compare our Dense-Decoupled Adaptive Kalman Filter (DD-AKF) against a spectrum of methods representing the current state of the art:

1) Standard KF: A baseline Kalman filter with a fixed measurement covariance R (constant over time) and no physics constraints. This isolates the benefit of adaptive uncertainty and physics-guided updates.

2) ICP (Point-to-Plane): A standard rigid/non-rigid Iterative Closest Point algorithm. This represents the geometric-only registration approach without temporal state memory.

3) Particle Filter (PF): A stochastic baseline using 500 particles, included to contrast DD-AKF with non-Gaussian sequential estimation in high-dimensional state spaces.

4) ClothFit: A learning-based pose-conditioned garment deformation baseline representative of supervised draping networks (e.g., TailorNet [19] and GarNet [20]), which predict dense garment vertex offsets conditioned on body pose/shape. We use a pre-trained model and lightly fine-tune on CLOTH3D, following common practice for supervised baselines.

5) Opt-ED (Optimization-based Embedded Deformation): An optimization baseline that minimizes per-frame geometric residuals together with an ED deformation-graph regularizer [13] and an ARAP-style smoothness term [11]. This serves as the primary "optimization vs. filtering" comparison, using the same dense topology as our method but without an explicit predictive motion model.

2.9.3 Metrics

We report performance using three categories of metrics.

For geometry results, we report the following two metrics:

• RMSE (mm): The Root Mean Square Error between estimated vertex positions and ground truth:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N_v}\sum_{i=1}^{N_v} \|\hat{\mathbf{p}}_i - \mathbf{p}_i\|^2}$$

• Hausdorff Distance (mm): Measures the worst-case geometric deviation:

$$d_H(A,B) = \max\left\{\sup_{a\in A}\inf_{b\in B}\|a-b\|, \sup_{b\in B}\inf_{a\in A}\|a-b\|\right\}$$

For Cloth Fidelity (Physics), we report several standard metrics:

• IoU: The Intersection over Union of the projected 2D silhouettes, measuring global shape alignment.

• Relative Stretch Error (%): The deviation of edge lengths from their rest state, averaged over all edges:

• Bend Error (deg): The mean deviation of dihedral angles from their rest configuration:

• Penetration Depth (mm): The average depth of garment vertices that violate the body boundary (where d is the body signed distance field):

And for temporal stability, we use

• Temporal Flicker (a.u.): A temporal stability metric computed from acceleration variance (jerkiness) relative to a smoothed trajectory:

2.9.4 Implementation Details and Hardware

DD-AKF is implemented in a real-time pipeline. The reported runtimes/FPS are measured end-to-end, including point-cloud preprocessing, correspondence search, and the filter update. All experiments were conducted on a workstation with [CPU model], [RAM], and [GPU model + VRAM] (please fill in to match your setup). The proposed filtering core is lightweight and can run in real time on a standard desktop; a GPU is only required if additional deep models or rendering are used for the compared baselines.

3 Results

We report mean ± SD over 5 independent runs (different seeds for initialization, minibatch order, and synthetic noise). Significance is assessed using two-tailed paired permutation tests against DD-AKF (per-sequence scores). Unless stated, lower is better.

3.1 Quantitative Comparisons

(1) FAUST (Geometry-Centric)

On FAUST (geometry-centric evaluation), DD-AKF achieves the lowest mean RMSE (6.35 ± 0.42 mm) and Hausdorff distance (19.4 ± 1.4 mm) among the compared methods (Table 1). Compared with Opt-ED, DD-AKF reduces RMSE by 0.25 mm ($p = 0.048$) and lowers Temporal Flicker from 1.15 ± 0.10 to 0.92 ± 0.08 ($p = 0.015$). The Hausdorff difference relative to Opt-ED is not statistically significant ($p = 0.061$), and we therefore treat the two methods as comparable on this worst-case geometric metric under the FAUST setting (Table 2).

Table 1: Quantitative comparison on the FAUST dataset (mean±standard deviation over five runs; lower is better). RMSE (mm) measures average geometric error, Hausdorff distance (mm) captures maximum surface deviation, and TF (a.u.) denotes temporal flicker (lower indicates smoother motion). Methods include Kalman Filter (KF), Iterative Closest Point (ICP), Particle Filter (PF), ClothFit (learning baseline), Optimization-based Embedded Deformation (Opt-ED), and DD-AKF (ours)

Metric	KF	ICP	PF	ClothFit	Opt-ED	DD-AKF (Ours)
RMSE (mm)	8.40 ± 0.50	7.80 ± 0.60	7.30 ± 0.70	6.90 ± 0.50	6.60 ± 0.48	6.35 ± 0.42
Hausdorff (mm)	26.0 ± 2.0	24.0 ± 2.0	22.0 ± 1.7	20.8 ± 1.6	20.1 ± 1.5	19.4 ± 1.4
TF (a.u.)	1.35 ± 0.10	1.20 ± 0.09	1.10 ± 0.08	0.95 ± 0.08	1.15 ± 0.10	0.92 ± 0.08

Table 2: Statistical significance on FAUST: p-values from two-tailed paired permutation tests comparing each baseline to DD-AKF (per-sequence scores). n.s. indicates not significant at $p > 0.05$. Metrics and method abbreviations follow Table 1

Metric	KF	ICP	PF	ClothFit	Opt-ED
RMSE	< 0.001	0.006	0.031	0.028	0.048
Hausdorff	< 0.001	0.004	0.040	0.021	0.061 (n.s.)
TF	< 0.001	0.008	0.047	0.051	0.015

(2) CLOTH3D (Cloth + Geometry)

On CLOTH3D (cloth + geometry evaluation), DD-AKF improves RMSE (7.25 ± 0.62 mm) and IoU (0.92 ± 0.02) relative to Opt-ED (7.55 ± 0.58 mm, IoU 0.91 ± 0.01), with statistically significant differences for these two metrics ($p = 0.046$ and $p = 0.041$, respectively; Table 4). For physical fidelity, DD-AKF yields a lower penetration depth (0.78 ± 0.10 mm vs. 1.10 ± 0.12 mm; $p = 0.015$) and a lower Temporal Flicker score (0.85 ± 0.08 vs. 1.08 ± 0.09; $p = 0.008$). Differences in relative stretch and bend error relative to Opt-ED are small and not statistically significant

in our tests ($p = 0.052$ and $p = 0.060$), suggesting comparable performance on these two metrics under the CLOTH3D setting.

Figure 1: Qualitative error distribution (Euclidean distance, mm) on a CLOTH3D T-shirt sequence with severe underarm self-occlusion. Missing correspondences in the axillary region cause increased error for geometric-only or per-frame optimization baselines, and learning-based inference can exhibit bias when the occluded deformation differs from the learned prior. DD-AKF (ours) maintains a more coherent surface in the occluded region, consistent with the reduced temporal flicker and lower penetration reported in Table 3–4. Source: generated by the authors using CLOTH3D [41].

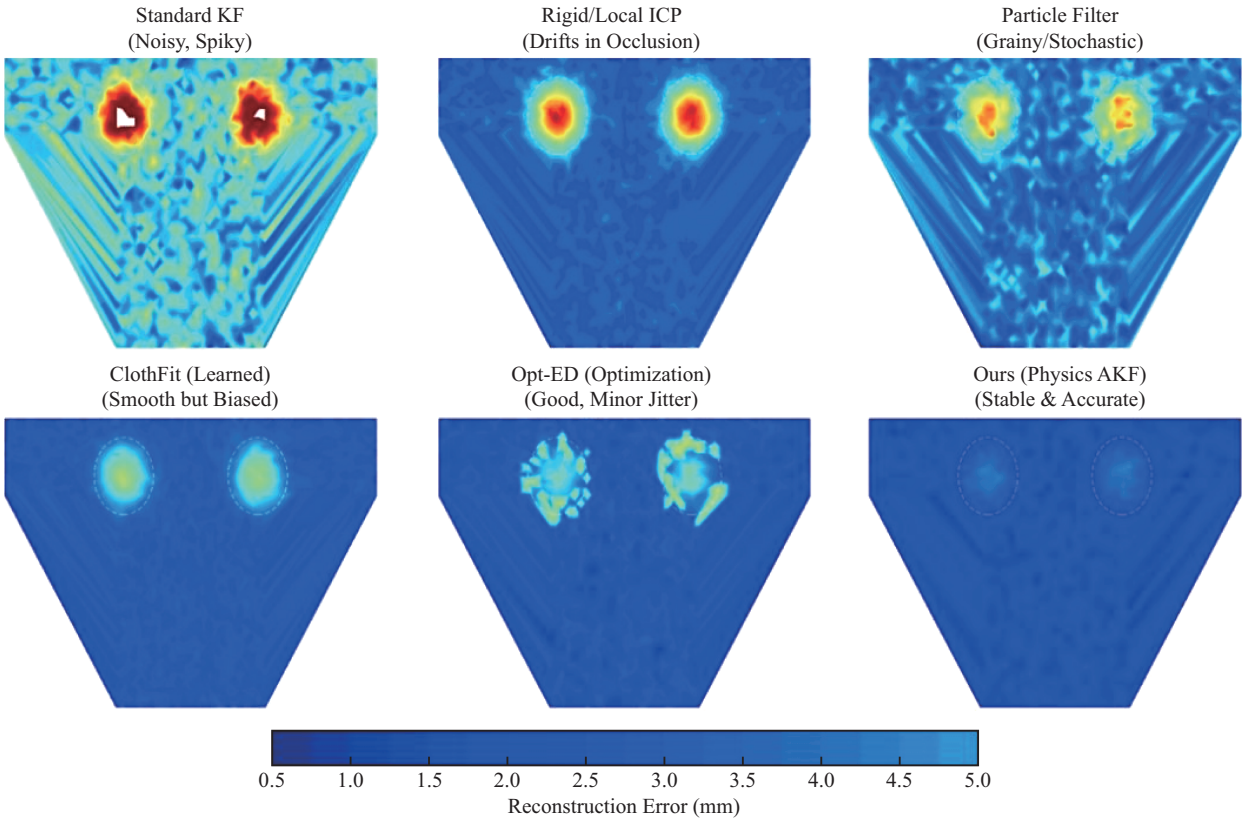

Fig. 1: Qualitative Error Comparison under Self-Occlusion (Underarms)

3.2 Robustness (Noise & Missing Data)

We stress-test robustness by adding isotropic Gaussian noise with standard deviation $\sigma \in \{0, 1, 3, 5\}$ mm (Table 5) and by randomly dropping 0–40% of input points (Table 6). Tables 5 and 6 report the mean $\pm$ standard deviation over five runs at each condition level. DD-AKF degrades more slowly than the compared baselines; for example, when σ increases from 0 to 5 mm, RMSE increases by +2.2 mm for DD-AKF ($6.6 \pm 0.25 \rightarrow 8.8 \pm 0.39$) versus +3.3 mm for Opt-ED ($6.9 \pm 0.27 \rightarrow 10.2 \pm 0.50$). Under 40% missing points, average penetration remains 1.35 ± 0.12 mm for DD-AKF compared with 2.15 ± 0.18 mm for Opt-ED.

Figure 2: Robustness under high sensor noise on a CLOTH3D skirt sequence. With increasing noise, purely geometric fitting can drift or become unstable due to noisy correspondences, while

Table 3: Quantitative comparison on the CLOTH3D dataset (mean ± standard deviation over five runs). IoU is higher-is-better; all other metrics are lower-is-better. RMSE (mm) is the average geometric error; Stretch Rel (%) is the relative stretch error; Bend (deg) is the dihedral-angle error; Penetration (mm) is the body interpenetration depth; TF (a.u.) is the temporal flicker. Methods: KF, ICP, PF, ClothFit, Opt-ED, and DD-AKF (ours)

Metric	KF	ICP	PF	ClothFit	Opt-ED	DD-AKF (Ours)
RMSE (mm)	9.50 ± 0.80	8.90 ± 0.70	8.40 ± 0.80	7.80 ± 0.60	7.55 ± 0.58	7.25 ± 0.62
IoU	0.88 ± 0.01	0.89 ± 0.01	0.90 ± 0.01	0.90 ± 0.01	0.91 ± 0.01	0.92 ± 0.02
Stretch Rel (%)	4.50 ± 0.30	4.10 ± 0.30	3.60 ± 0.30	3.40 ± 0.30	3.25 ± 0.28	2.95 ± 0.32
Bend (deg)	4.50 ± 0.50	4.20 ± 0.50	3.70 ± 0.50	3.60 ± 0.40	3.40 ± 0.42	3.20 ± 0.45
Penetration (mm)	1.60 ± 0.15	1.40 ± 0.14	1.20 ± 0.12	1.00 ± 0.10	1.10 ± 0.12	0.78 ± 0.10
TF (a.u.)	1.30 ± 0.11	1.18 ± 0.10	1.05 ± 0.09	0.92 ± 0.08	1.08 ± 0.09	0.85 ± 0.08

Table 4: Statistical significance on CLOTH3D: p-values from two-tailed paired permutation tests comparing each baseline to DD-AKF (per-sequence scores). n.s. indicates $p > 0.05$. Metrics and method abbreviations follow Table 3

Metric	KF	ICP	PF	ClothFit	Opt-ED
RMSE	< 0.001	0.004	0.015	0.030	0.046
IoU	< 0.001	0.002	0.011	0.028	0.041
Stretch	< 0.001	< 0.001	0.009	0.022	0.052 (n.s.)
Bend	< 0.001	0.006	0.017	0.041	0.060 (n.s.)
Penetration	< 0.001	< 0.001	0.004	0.019	0.015
TF	< 0.001	0.003	0.014	0.033	0.008

a learned prior can over-smooth folds that are weakly supported by the input. DD-AKF (ours) uses physics-guided regularization to suppress high-frequency noise while preserving large-scale fold structure, consistent with the quantitative degradation trends in Table 5. Source: generated by the authors using CLOTH3D [41].

Table 5: Robustness to additive Gaussian noise on CLOTH3D. Values report mean ± standard deviation RMSE (mm) over five runs for each noise standard deviation σ. Method abbreviations follow Tables 1–3

σ (mm)	KF	ICP	PF	ClothFit	Opt-ED	DD-AKF
0	9.6 ± 0.35	9.0 ± 0.33	8.5 ± 0.36	7.3 ± 0.28	6.9 ± 0.27	6.6 ± 0.25
1	9.9 ± 0.38	9.4 ± 0.36	8.8 ± 0.39	7.6 ± 0.31	7.4 ± 0.30	6.9 ± 0.27
3	10.8 ± 0.45	10.4 ± 0.43	9.9 ± 0.48	8.6 ± 0.37	8.9 ± 0.40	8.1 ± 0.34
5	11.8 ± 0.52	11.2 ± 0.49	10.9 ± 0.55	9.3 ± 0.42	10.2 ± 0.50	8.8 ± 0.39

Table 6: Robustness to missing input points on CLOTH3D. "Missing (%)" denotes the percentage of points randomly dropped from the input point cloud; values report mean ± standard deviation penetration depth (mm) over five runs at each missingness level. Method abbreviations follow Tables 1–3

Missing (%)	KF	ICP	PF	ClothFit	Opt-ED	DD-AKF
0	1.60 ± 0.10	1.40 ± 0.09	1.20 ± 0.08	1.00 ± 0.07	1.10 ± 0.08	0.78 ± 0.06
10	1.78 ± 0.11	1.56 ± 0.10	1.32 ± 0.09	1.12 ± 0.08	1.25 ± 0.09	0.85 ± 0.07
20	1.98 ± 0.13	1.74 ± 0.11	1.46 ± 0.10	1.25 ± 0.09	1.45 ± 0.11	0.98 ± 0.08
30	2.22 ± 0.15	1.95 ± 0.13	1.65 ± 0.12	1.41 ± 0.10	1.80 ± 0.14	1.15 ± 0.10
40	2.48 ± 0.17	2.20 ± 0.15	1.88 ± 0.14	1.59 ± 0.12	2.15 ± 0.18	1.35 ± 0.12

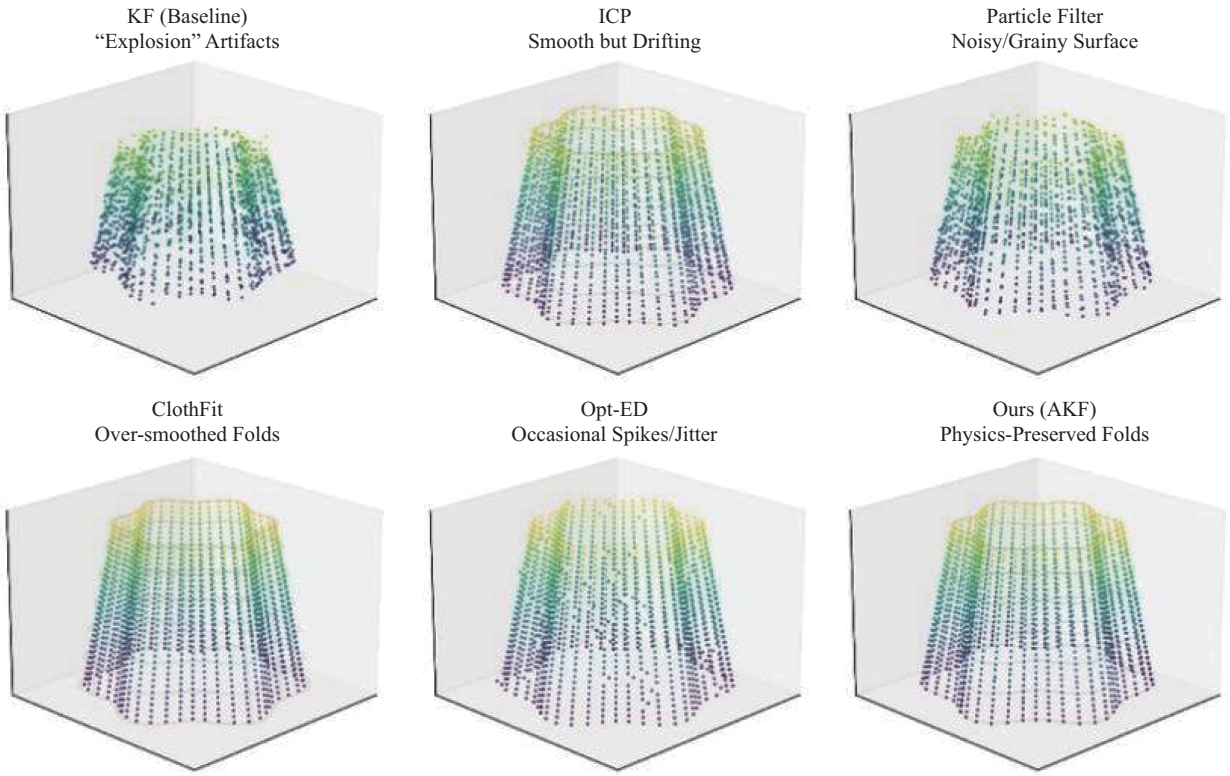

Fig. 2: Robustness to High Sensor Noise on a CLOTH3D Skirt

3.3 Reliability Score

To validate the utility of the proposed Reliability Score s_k We analyze its correlation with tracking quality. Table 7 reports Spearman's ρ between s_k and geometric RMSE ($\rho = 0.62$) and between s_k and penetration depth ($\rho = 0.65$), indicating a moderate-to-strong monotonic association under the tested conditions. We further illustrate its temporal behavior in Fig. 3, where decreases in s_k coincide with (and in this example slightly precede) increased error during an occlusion interval.

Figure 3: Temporal traces of RMSE and the reliability score during a sequence containing a severe occlusion interval (shaded region). Baselines show increased error when geometric observations are lost, whereas DD-AKF maintains lower error by relying more on its predictive prior and

Table 7: Rank-correlation analysis of the online reliability score. Spearman's ρ is computed between the reliability score and (i) geometric RMSE and (ii) penetration depth over the evaluated sequences. Larger $|\rho|$ indicates a stronger monotonic association

Correlation Pair	Value
Spearman (s_k, RMSE)	0.62
Spearman (s_k, Penetration)	0.65

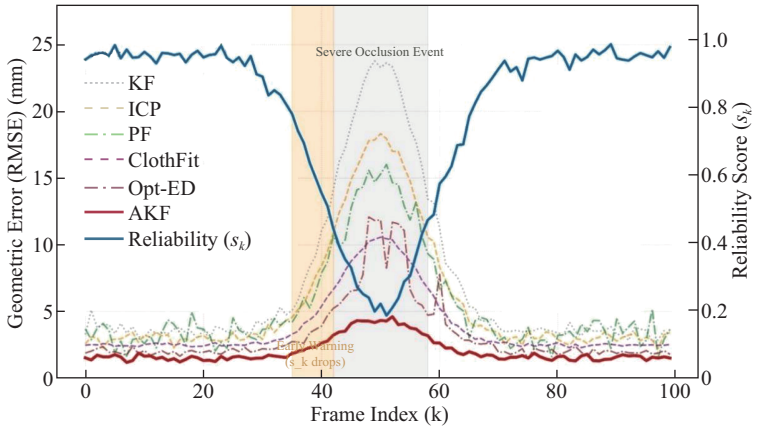

Fig. 3: Online Reliability Diagnostics vs. Baseline Performance

physics constraints. The reliability score decreases at the onset of occlusion (frames 38–42 in this example), supporting its use as an online indicator to trigger adaptive robustification. Source: generated by the authors.

4 Discussion

The experiments indicate that the most consistent advantage of DD-AKF is improved temporal stability with competitive geometric accuracy. On FAUST, DD-AKF reduces Temporal Flicker from 1.15 ± 0.10 (Opt-ED) to 0.92 ± 0.08 (p=0.015) while also lowering RMSE from 6.60 ± 0.48 to 6.35 ± 0.42 mm (p=0.048). On CLOTH3D, DD-AKF reduces Temporal Flicker from 1.08 ± 0.09 to 0.85 ± 0.08 (p = 0.008) and penetration depth from 1.10 ± 0.12 to 0.78 ± 0.10 mm (p = 0.015), with modest improvements in RMSE ($7.55 \pm 0.58 \rightarrow 7.25 \pm 0.62$ mm, p = 0.046) and IoU ($0.91 \pm 0.01 \rightarrow 0.92 \pm 0.02$, p = 0.041).

Not all metrics show statistically significant changes. On FAUST, the Hausdorff distance difference between DD-AKF and Opt-ED is small and not significant (p = 0.061). On CLOTH3D, DD-AKF shows lower relative stretch and bend error than Opt-ED, but the differences are not significant under our permutation tests (p = 0.052 and p = 0.060). These outcomes suggest that Opt-ED's strong per-frame geometric regularization already captures local smoothness (stretch/bend) effectively in well-observed regions. At the same time, DD-AKF provides its largest gains in metrics that depend on temporal coherence and physically feasible behavior under missing or corrupted observations (Temporal Flicker and penetration).

The robustness experiments further characterize the boundary of this behavior. Under additive

Gaussian noise, DD-AKF's RMSE increases by +2.2 mm from $\sigma = 0$ to $\sigma = 5$ mm (6.6 → 8.8), compared with +3.3 mm for Opt-ED (6.9 → 10.2). Under missing data, penetration increases by +0.57 mm from 0% to 40% missing for DD-AKF (0.78 → 1.35), versus +1.05 mm for Opt-ED (1.10 → 2.15). Empirically, standard KF and purely geometric pipelines exhibit larger error excursions under these stressors, whereas DD-AKF remains bounded; this supports the boundedness intuition in Section 2.7 (bounded adaptive noise, physics pseudo-measurements, and the block-diagonal covariance recursion).

Regarding online reliability, the Spearman correlations ($\rho = 0.62$ with RMSE and $\rho = 0.65$ with penetration) indicate that s_k captures the overall degradation trend but is not a perfect predictor. In practice, s_k can be used as an early warning signal (Fig. 3), but should be interpreted with context: sharp motion, local correspondence ambiguity, or sudden viewpoint changes can transiently increase innovations and may trigger false positives if the threshold τ is set too aggressively. We therefore use s_k primarily as a gating and adaptation signal—raising measurement covariance and increasing physics weight—rather than as a binary failure detector.

4.1 Limitations

First, the dynamics model is intentionally simple (constant velocity). This choice supports real-time performance but may limit geometric gains during rapid accelerations or highly dynamic cloth motion; in such cases, prediction error can increase and the filter must rely more on measurements, which may partially explain why RMSE improvements relative to Opt-ED are modest in some settings.

Second, the block-diagonal covariance approximation ignores long-range correlations. While the mean state is still coupled through physical terms and shared observations, uncertainty calibration (and thus reliability thresholds) may be conservative or miscalibrated when ambiguity is globally correlated, for example, during broad occlusions of loose garments.

Third, DD-AKF remains correspondence-dependent and requires reasonable initialization. Extremely ambiguous topologies, prolonged complete observation loss, or complex contact effects not captured by our simplified energies can still lead to failure; in such cases, explicit re-initialization or multi-hypothesis tracking would be needed.

4.2 Future Work

Several directions could strengthen robustness and broaden applicability. A natural extension is to incorporate richer dynamics, such as constant-acceleration models, contact-aware prediction, or lightweight learned dynamics that preserve interpretability while better matching real motion. Another promising direction is to move beyond purely block-diagonal uncertainty while remaining computationally tractable. Graph-local covariance structures, banded approximations, or low-rank-plus-diagonal forms could better capture key spatial correlations and potentially improve uncertainty calibration, particularly when uncertainty is used as a decision signal for reliability and recovery.

Future work could also address failure recovery more explicitly. Multi-hypothesis variants, limited branching in correspondence assignment, or hybrid strategies that combine filtering with periodic global re-initialization may improve resilience under prolonged occlusion or abrupt topol-

ogy changes. Finally, the reliability signal could be integrated into a closed-loop capture pipeline, guiding sensor placement, view selection, or acquisition parameters in interactive scenarios to reduce the duration of low-observability conditions and improve overall stability.

5 Conclusion

This paper presents DD-AKF, an integrated real-time garment fitting pipeline that combines dense per-vertex state estimation, a block-diagonal covariance approximation for tractable uncertainty propagation, physics-guided soft constraints, and an innovation-based reliability signal. On FAUST and CLOTH3D, DD-AKF consistently reduces temporal flicker and penetration depth relative to representative baselines, while maintaining competitive geometric accuracy at 30+ FPS. Improvements in some worst-case or local smoothness metrics (e.g., Hausdorff distance on FAUST, stretch/bend on CLOTH3D) are smaller and not statistically significant compared with Opt-ED, and we therefore position DD-AKF as a robust, integrative alternative aimed at stability and physical feasibility under degraded observations.

Limitations include the use of a simple constant-velocity dynamics model and a block-diagonal covariance approximation, which may reduce accuracy during highly dynamic motion and fail to capture long-range uncertainty correlations. Future work will explore richer dynamics, structured but tractable covariance representations, and more explicit recovery mechanisms under prolonged occlusion. From a practical perspective, the ability to run at interactive rates while monitoring uncertainty and reliability is relevant to real-time applications such as virtual try-on, AR retail, and telepresence systems.

Acknowledgement

The Fujian Natural Science Foundation Project funded this research: Research on the Model and Algorithm of 3D Clothing Style Transfer for Point Cloud Data Machine Learning (2023J011406), and grant number (2023J05252).

References

[1] Cheng R, Wu N, Varvello M, et al. Are we ready for the metaverse? IEEE Journal on Selected Areas in Communications 2022; 40: 3259–76.

[2] Orts-Escolano S, Rhemann C, Fanello S, et al. Holoportation: Virtual 3d teleportation in real-time. Proceedings of the 29th Annual Symposium on User Interface Software and Technology, 2016, p. 741–54.

[3] Park G-H, Lee J-H. A review of virtual try-on systems: Current status and future directions. International Journal of Clothing Science and Technology 2021; 33: 876–96.

[4] Loper M, Mahmood N, Romero J, et al. SMPL: A Skinned Multi-Person Linear Model. ACM Transactions on Graphics (Proc SIGGRAPH Asia) 2015; 34: 248: 1-248: 16. https://doi.org/10.1145/2816795.2818013.

[5] Pavlakos G, Choutas V, Ghorbani N, et al. Expressive body capture: 3d hands, face, and body from a single image. Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, p. 10975–85.

[6] Xu W, Cai Y, He D, et al. FAST-LIO2: Fast Direct LiDAR-Inertial Odometry. IEEE Transactions on Robotics 2022; 38: 2053–73. https://doi.org/10.1109/TRO.2022.3141876.

[7] Baraff D, Witkin A. Large steps in cloth simulation. Proceedings of the 25th annual conference on Computer graphics and interactive techniques, 1998, p. 43–54.

[8] Bridson R, Marino S, Fedkiw R. Simulation of clothing with folds and wrinkles. ACM SIGGRAPH 2005 Courses 2005: 28-es.

[9] Narain R, Pfaff T, O'Brien JF. Folding and crumpling adaptive sheets. ACM Transactions on Graphics (TOG) 2012; 31: 1–10.

[10] Li Z, Chen Z, Li Z, et al. Spacetime Gaussian Feature Splatting for Real-Time Dynamic View Synthesis. arXiv preprint arXiv: 2312.16812, 2023.

[11] Müller M, Heidelberger B, Hennix M, et al. Position based dynamics. Journal of Visual Communication and Image Representation, vol. 18, 2007, p. 109–18.

[12] Amberg B, Romdhani S, Vetter T. Optimal step nonrigid icp algorithms for surface registration. CVPR 2007; 1: 1–8.

[13] Sumner RW, Schmid J, Pauly M. Embedded deformation for shape manipulation. ACM SIGGRAPH 2007 papers, 2007, p. 80-es.

[14] Besl PJ, McKay ND. A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence 1992; 14: 239–56.

[15] Newcombe RA, Fox D, Seitz SM. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, p. 343–52.

[16] Slavcheva M, Baust M, Cremers D, et al. Killingfusion: Non-rigid 3d reconstruction without correspondences. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, p. 1386–95.

[17] Zollhöfer M, Nießner M, et al. Real-time non-rigid reconstruction using an RGB-D camera. ACM Transactions on Graphics (TOG) 2014; 33: 1–12.

[18] Li H, Sumner RW, Pauly M. Global correspondence optimization for non-rigid registration of depth scans. Computer graphics forum, vol. 27, 2008, p. 1421–30.

[19] Patel C, Liao Z, Pons-Moll G. TailorNet: Predicting Clothing in 3D as a Function of Human Pose, Shape and Garment Style. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR); 2020: 7365–7375.

[20] Gundogdu E, Constantin V, Seifoddini A, et al. GarNet: A Two-Stream Network for Fast and Accurate 3D Cloth Draping. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV); 2019: 8739–8748.

[21] Zhu H, Cao Y, Jin H, et al. DeepFashion3D: A Dataset and Benchmark for 3D Garment Reconstruction from Single Images. Computer Vision – ECCV 2020, Springer; 2020. https://doi.org/10.1007/978-3-030-58452-8{_}30.

[22] Santesteban I, Otaduy MA, Casas D. SNUG: Self-supervised neural dynamic garments. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, p. 8140–50.

[23] Tiwari G, Bhatnagar BL, Tung T, et al. Sizer: A dataset and model for size-conditioned 3d clothing reconstruction. European Conference on Computer Vision, 2020, p. 1–18.

[24] Jiang B, Zhang J, Hong Y, et al. BCNet: Learning body and cloth shape from a single image. European Conference on Computer Vision 2020: 18–35.

[25] Corona E, Pumarola A, Alenya G, et al. Smplicit: Topology-aware generative model for clothed people. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, p. 11875–85.

[26] Kendall A, Gal Y. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? Advances in Neural Information Processing Systems 2017; 30: 5574–5584.

[27] He X, Wang Z, Liu Y. Distributed Kalman filter with decoupled state for large-scale systems. IEEE Transactions on Cybernetics, vol. 50, 2019, p. 2205–15.

[28] Mildenhall B, Srinivasan PP, Tancik M, et al. Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM, vol. 65, 2021, p. 99–106.

[29] Peng S, Zhang Y, Xu Y, et al. Animatable neural radiance fields for modeling dynamic human bodies. Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, p. 14314–23.

[30] Kerbl B, Kopanas G, Leimkühler T, et al. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Transactions on Graphics 2023; 42. https://doi.org/10.1145/3592433.

[31] Weng C-Y, Curless B, Srinivasan PP, et al. Humannerf: Free-viewpoint rendering of moving people from monocular video. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, p. 16210–20.

[32] Daum FE, Fitzgerald RJ. Decoupled Kalman filters for phased array radar tracking. IEEE Transactions on Automatic Control 1983; 28: 269–83.

[33] Kalman RE. A new approach to linear filtering and prediction problems. Transactions of the ASME–Journal of Basic Engineering 1960; 82: 35–45.

[34] Julier SJ, Uhlmann JK. New extension of the Kalman filter to nonlinear systems. Signal processing, sensor fusion, and target recognition VI, vol. 3068, 1997, p. 182–93.

[35] Macklin M, Müller M, Chentanez N. XPBD: position-based simulation of compliant constrained dynamics. Proceedings of the 9th International Conference on Motion in Games 2016: 49–54.

[36] Segal A, Haehnel D, Thrun S. Generalized-ICP. Proceedings of Robotics: Science and Systems (RSS), 2009. https://doi.org/10.15607/RSS.2009.V.021.

[37] Yang J, Li H, Campbell D, et al. Go-ICP: A Globally Optimal Solution to 3D ICP Point-Set Registration. IEEE Transactions on Pattern Analysis and Machine Intelligence 2016; 38: 2241–54. https://doi.org/10.1109/TPAMI.2015.2513405.

[38] Qin Z, Yu H, Wang C, et al. Deep Graph-based Spatial Consistency for Robust Non-rigid Point Cloud Registration. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023, p. 5394–403.

[39] Li Y, Harada T. Non-rigid Point Cloud Registration with Neural Deformation Pyramid. Advances in Neural Information Processing Systems 2022; 35: 27757–27768.

[40] Zhang Y, Hirata K. Garment Recognition and Reconstruction Using Object Simultaneous Localization and Mapping. Sensors 2024; 24: 7622.

[41] Bertiche H, Madadi M, Escalera S. CLOTH3D: Clothed 3D Humans. Computer Vision – ECCV 2020, Springer; 2020. https://doi.org/10.1007/978-3-030-58565-5{_}21.

[42] Bhatnagar BL, Tiwari G, Theobalt C, et al. Multi-Garment Net: Learning to Dress 3D People from Images. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2019. https://doi.org/10.1109/ICCV.2019.00552.

[43] Ma Q, Yang J, Ranjan A, et al. Learning to Dress 3D People in Generative Clothing. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA: IEEE; 2020, p. 6468–77. https://doi.org/10.1109/CVPR42600.2020.00650.

[44] Kairanda N, Habermann M, Theobalt C, et al. NeuralClothSim: Neural Deformation Fields Meet the Thin Shell Theory. Neural Information Processing Systems (NeurIPS), 2024.

[45] Guo C, Jiang T, Kaufmann M, et al. ReLoo: Reconstructing Humans Dressed in Loose Garments from Monocular Video in the Wild. In: European Conference on Computer Vision (ECCV); 2024: 21–38.

[46] He X, Wang Z, Liu Y. Distributed Kalman filter with decoupled state for large-scale systems. IEEE Transactions on Cybernetics, vol. 50, 2019, p. 2205–15.

[47] Marelli D. Distributed Kalman Estimation with Decoupled Local Filters. arXiv Preprint arXiv: 2009.05799 2020.

[48] Remelli E, et al. UPose3D: Uncertainty-Aware 3D Human Pose Estimation with Cross-View and Temporal Cues. In: European Conference on Computer Vision (ECCV); 2024: 19–38.

[49] Zhang L, Zhang W, Liu Z. Prediction of Bust and Waist Size Based on Two-dimensional Images. Journal of Fiber Bioengineering and Informatics 2024; 17(1): 1–11. https://doi.org/10.3993/jfbim03191.

[50] Jevšnik S, Stjepanovič Z, Rudolf A. 3D Virtual Prototyping of Garments: Approaches, Developments and Challenges. Journal of Fiber Bioengineering and Informatics 2017; 10(1): 51–63. https://doi.org/10.3993/jfbim00253.

[51] Gou JJ, Wu L, Qi J, Ying BA, Wang Y. A Construction Method for Personalized Bra Sample Models. Journal of Fiber Bioengineering and Informatics 2024; 16(3): 269–281. https://doi.org/10.3993/jfbim02371.

[52] Lv XY, Xin XY, Feng LW. 3D Virtual Simulation of Tajik Women's Wedding Dress. In: 17th Textile Bioengineering and Informatics Symposium Proceedings (TBIS 2024); 2024: 396–403. https://doi.org/10.52202/076989-0048.

[53] Chen XY, Cui MH. The Spatial Relationship Inside Coats Based on CLO3D Virtual Fabrics. In: 17th Textile Bioengineering and Informatics Symposium Proceedings (TBIS 2024); 2024: 583–592. https://doi.org/10.52202/076989-0070.

[54] Xue Y, Xie X, Marin R, Pons-Moll G. Human-3Diffusion: Realistic Avatar Creation via Explicit 3D Consistent Diffusion Models. Advances in Neural Information Processing Systems (NeurIPS), 2024. https://doi.org/10.52202/079017-3160.