

Walking on Thin Air: Environment-Free Physics-based Markerless Motion Capture

Micha Livne*, Leonid Sigal[†], Marcus A. Brubaker[‡] and David J. Fleet*

**Department of Computer Science
University of Toronto
Toronto, Canada
{mlivne, fleet}@cs.toronto.edu*

[†]*Department of Computer Science
University of British Columbia
Vancouver, Canada
lsigal@cs.ubc.ca*

[‡]*Lassonde School of Engineering
York University
Toronto, Canada
mab@eecs.yorku.ca*

Abstract—We propose a generative approach to physics-based motion capture. Unlike prior attempts to incorporate physics into tracking that assume the subject and scene geometry are calibrated and known a priori, our approach is automatic and online. This distinction is important since calibration of the environment is often difficult, especially for motions with props, uneven surfaces, or outdoor scenes. The use of physics in this context provides a natural framework to reason about contact and the plausibility of recovered motions. We propose a fast data-driven parametric body model, based on linear-blend skinning, which decouples deformations due to pose, anthropometrics and body shape. Pose (and shape) parameters are estimated using robust ICP optimization with physics-based dynamic priors that incorporate contact. Contact is estimated from torque trajectories and predictions of which contact points were active. To our knowledge, this is the first approach to take physics into account without explicit *a priori* knowledge of the environment or body dimensions. We demonstrate effective tracking from a noisy single depth camera, improving on state-of-the-art results quantitatively and producing better qualitative results, reducing visual artifacts like foot-skate and jitter.

Keywords—Computer Graphics; Computer Vision; Physics; 3D Human Pose Tracking;

I. INTRODUCTION

Markerless motion capture methods enable reconstruction of detailed motion and dynamic geometry of the body (and sometimes garments) from multiple streams of video [1] or depth data [2], [3]. Recent human tracking methods are able to handle video captured in the wild, but still suffer from visually significant artifacts (jittering, feet/contact skating). This issue is significant as people are sensitive to such artifacts (*e.g.*, foot-skate is perceptible at levels less than 21 mm [4]).

To address these challenges we propose a generative 3D human tracking approach that takes physics-based prior knowledge into account when estimating pose over time. The use of physics in this context is compelling as it provides a natural framework to reason about contact and the plausibility of the recovered motions online. Prior attempts to use physics for tracking assume that the subject and scene geometry are known a priori and calibrated [5], [6], that contact states are annotated by a user [7], or that

optimization can be performed off-line (*i.e.*, in batch) [8]. In contrast, our approach is online, without manual input. Beginning with the first frame, the subject and the contact state(s) are estimated online during tracking, without a priori knowledge of the environment. This is an important distinction, as calibration of the environment can be difficult, especially when capturing motions with props, on uneven surfaces or outdoors.

Our main contribution is the use of a physics-based prior without an explicit model of the environment. To our knowledge, it is the first tracking approach to incorporate physics without any explicit *a priori* knowledge of the environment or body dimensions. We demonstrate that the approach is effective in tracking from a single depth camera, improving on state-of-the-art results quantitatively and qualitatively, greatly reducing visually unpleasant artifacts such as foot-skate and jitter.

II. RELATED WORK

3D Human Tracking: Markerless motion capture, estimating the skeletal motion of a subject, has a rich history in vision and graphics (for an extensive survey see [9]). Methods can be broken into two classes: model-based and regression-based (or generative vs. discriminative). Regression-based methods estimate pose directly by regressing pose from image feature descriptors (*e.g.* [3], [10]–[14]). Model-based approaches exploit a generative model for the body and image, and optimize for generative parameters that explain the image observations (*e.g.* [15], [16]). The former is faster, but generally less accurate (unless the problem domain is highly constrained). Model-based approaches may be more accurate, but tend to be slower as they require iterative or stochastic optimization, and suffer from local optima.

Use of Physics in Tracking: Physics-based tracking has been proposed as a way to regularize pose under the assumption that physics is a universal prior that requires no assumptions about one’s motion (given a physical model). Early work dates back to [17] and [18], however they focused on simple motions in absence of contact. More recently Brubaker *et al.* [5] proposed a low-dimensional



Figure 1: Illustration of results produced by our physics-based tracker.

model of the lower-body to track walking subjects from monocular video. A more general data-driven physics-based filter, applicable to variety of motions, was proposed in [6]. In [19] a controller-based approach is proposed where a physics-based full body controller, instead of sequence of poses, is estimated. In all cases the body proportions and the scene geometry were assumed to be known. In [7] a physics-based tracking approach is formulated as a batch optimization problem with known contact points and ground geometry. In [8] the parameters of the planar ground model are estimated from data, however, the method assumes a parametric structure of ground geometry (a plane) and reasonably accurate 3D input, obtained using a binocular system. We build on this work with one notable distinction: we assume no knowledge of ground geometry or subject proportions. A notable distinction from prior work is [20] which attempted to encode ground constraints directly in the kinematics, without a physics model. Like other methods, however, it required *a priori* knowledge of the ground plane.

Contact Estimation: We briefly note that contact estimation and sampling has been used in other domains of graphics as well. One example is hand manipulation [21], where a randomized search over the hand-object contacts is proposed as the strategy for finding pose of the hand manipulating an object over time. Contact invariant optimization [22] attempts to sidestep the problem of explicit contact estimation by searching over the space of contacts at the same time as behaviour of the character. Such approaches, while interesting, require batch processing and long compute times, making them inapplicable for real-time capable full body tracking.

III. METHOD

Our tracking pipeline is depicted in Fig. 2. We describe in this section a fast body mesh model (Sec. III-A), discrete formulation of a physical engine and physics-based motion prior (Sec. III-B), and a tracking framework that utilizes those in order to facilitate physics-based 3D human tracking (Sec. III-C). We also describe a method to pre-process input point cloud data that allows us to automatically initialize tracking, is fast, and simple to implement (Sec. III-D).

A. Fast Data-driven Parametric Body Model

In what follows we exploit a new SCAPE-like model (see [23]) for tracking. With an explicit skeleton, anthropometrics

(bone length) and body shape parameters, our model is easy to manipulate and control. The anthropometrics parameters offers direct control over deformations due to bone lengths. The body shape parameters allow for control over the shape, independent of anthropometrics. To the best of our knowledge, explicit control over anthropometrics and shape is not straightforward with other existing body models.

The body is modelled as a 3D triangulated mesh and comprise 69 DOF and 26 body parts

$$\mathcal{M}(\Theta) = \{\mathbf{p}(\Theta; \mathbf{B}_\ell, \mathbf{B}_\beta), E\} \quad (1)$$

where \mathbf{B}_ℓ and \mathbf{B}_β are basis matrices, which capture variations in the mesh due to anthropometrics and body shape respectively, and $\Theta = (\mathbf{q}, \ell, \beta)$, where \mathbf{q} denotes articulated pose, and ℓ and β denote coordinate vectors within the two subspaces. The N mesh vertices of a canonical pose (called the template pose $\tilde{\mathbf{q}}^s$) for a given subject s are given by a vector $\tilde{\mathbf{p}}^s \in \mathbb{R}^{3N \times 1}$

$$\tilde{\mathbf{p}}^s(\ell^s, \beta^s) = \mathbf{B}_\ell(\ell^s - \hat{\ell}) + \mathbf{B}_\beta\beta^s \quad (2)$$

where $\hat{\ell}$ denotes the mean anthropometrics within the subspace. The anthropometrics basis \mathbf{B}_ℓ represents a linear mapping from bone lengths (relative to the mean) to a base template mesh. The basis \mathbf{B}_β provides a linear mapping of body shape coefficients into a deformation from a base template mesh. We enforce orthogonality of the two subspaces during the basis learning stage. We discuss how we learn the basis in the supplementary material. The final mesh is calculated using Linear Mesh Blending (LMB)

$$\mathbf{p}_i^s(\Theta) = \sum_{b \in \mathcal{B}_i} w_{ib} \mathbf{M}_b(\ell, \mathbf{q}) \mathbf{M}_b^{-1}(\ell, \tilde{\mathbf{q}}^s) \tilde{\mathbf{p}}_i^s(\ell, \beta) \quad (3)$$

where \mathcal{B}_i is the set of bones (*i.e.*, rigid body parts) that influence the position of vertex i , w_{ib} is the influence of bone b on vertex i (assumed to be constant for all poses), and $\mathbf{p}^s \in \mathbb{R}^{3N \times 1}$ is the final mesh. We trained our model on the Hasler dataset [24]. Fig. 3 and Fig. 4 depict our LMB model.

B. Environment-Free Physics-Based Priors

Priors are used in optimization to regularize the loss, pushing the optimal solution to a more desired manifold. As such, we would like our priors to be as generic as possible in




Figure 2: An overview of the tracking pipeline. We extract connected components and geodesic extrema from point cloud. Followed by human detection and pose initialization. We initialize the pose with the nearest-neighbour pose from a database, in geodesic descriptor space. Next, we register body anthropometrics, shape, and pose. Tracking is performed by updating pose with ICP optimization, while estimating contact state through forces pattern.



Figure 3: Visualization of linear mesh blending. From left to right: target pose, template pose, weights, template mesh, and target mesh. The final vertices position vector $\mathbf{p}(\mathbf{q})$ is calculated by weighing vertex position w.r.t. transformations of a few joints applied to the template vertex position.




Figure 4: A template is composed of anthropometrics ℓ , and body shape β . Top row: anthropometrics variations with constant body shape. Bottom row: body shape variations with constant anthropometrics.

order to generalize well. Physics-based priors exploit physical dynamics as an informative but general prior on motion, to help ensure that tracking yields a plausible motion. To that end we formulate our model of articulated dynamics using discrete mechanics [25]. This has many desirable properties such as direct mapping to discrete observations, conservation of energy, and computational efficiency (see [26]).

Variational Integrator: In the variational formulation of Lagrangian mechanics, the motion of a system is described by a function known as a *discrete Lagrangian*, $\mathcal{L}^d(\mathbf{q}_{k-1}, \mathbf{q}_k)$ where \mathbf{q}_k denotes the generalized coordinates of the system (e.g., a stick figure) at time step k . The discrete Lagrangian is an approximation of the continuous Lagrangian, and is used in a discrete formulation of the principal of least action to derive discrete mechanics (see [25] for more details). The evolution of the system is then given by discrete Euler-

Lagrange equations

$$\underbrace{D_1 \mathcal{L}^d(\mathbf{q}_k, \mathbf{q}_{k+1}) + D_2 \mathcal{L}^d(\mathbf{q}_{k-1}, \mathbf{q}_k)}_{\mathcal{E}(\mathbf{q}_{k-1}, \mathbf{q}_k, \mathbf{q}_{k+1})} = \mathbf{f}_{k+1} \quad (4)$$

where \mathbf{f} is the vector of net generalized forces applied to the system, and D_i is a partial derivative operator with respect to the i th function parameter.

Contact: Contact is one of the greatest challenges, both computationally and theoretically, with physical dynamics. Such problems are less severe if contact times and locations are known, or provided by a user (e.g., [7]), but in most real-world tracking problems contact is unknown at inference time. Despite the added complexity, contact represents a strong constraint on motion (e.g., feet skate should not happened during contact), and as such is a desirable element

of the prior. To avoid dependency on prior knowledge of the environment or manual intervention, we infer contact states as part of a generative model. This reduces the computational challenge of handling inequality constraints into enforcing holonomic constraints, wherein one adds a constraint term, \mathcal{L}_c , to the Lagrangian \mathcal{L}^d . For a set of constraints, given by a function equation $\mathbf{g}(\mathbf{q}) = 0$, \mathcal{L}_c is constructed as

$$\mathcal{L}_c(\mathbf{q}) = \mathbf{g}(\mathbf{q})^T \boldsymbol{\lambda} \quad (5)$$

where $\boldsymbol{\lambda}$ is a vector of Lagrange multipliers.

Root Forces: We refer to the forces applied to the root node of the kinematic tree as root forces. The root node represents the global translation/rotation, and forces applied to it represents external forces applied to our physical model. Newton’s 2nd law states that changes (over time) to the total momentum of a physical system are equal to external forces applied to the system. Our model represents a system that has no external forces but contact. As a result, we use the existence of root forces as an indicator that our contact model is incomplete (following [8]). By choosing a contact configuration that minimizes the external forces, we enhance our model with contact in order to enforce the assumption that contact is the only option for our model to change its momentum, and propel itself. Alternatively, applying a direct force to the root node of a kinematic tree can be thought of as a human wearing a jet-pack. By minimizing the root forces, we discourage that option.

Contact Estimation: To determine contact, we trained an independent binary classifier (per possible contact point) on the forces of tracked subjects, assuming contact-free motions. Effectively, we learn to infer contact from the forces that drive our model in the absence of contact. Currently we use four possible contact points at the heel and toe of each foot. A logistic regressor is trained to estimate the probability of contact given these forces:

$$p(c_k^i | \mathbf{f}_k) = \sigma_i(\mathbf{f}_k) \quad (6)$$

where c_k^i is a binary variable indicating the contact state of point i in frame k , \mathbf{f}_k is the vector of net generalized forces at frame k (in the absence of contact), and $\sigma_i(\mathbf{x}) = (1 + \exp(-\alpha_0^i - \mathbf{x}^T \boldsymbol{\alpha}_1^i))^{-1}$ is a sigmoid function. Parameters $\alpha_0^i, \boldsymbol{\alpha}_1^i$ were learned for each contact point independently.

Physics-based Prior: With contact state determined, we can estimate the contact forces by minimizing root forces (following [8]). We can then write Eq. 4 with the additional contact constraints as

$$\mathcal{E}(\mathbf{q}_{k-1}, \mathbf{q}_k, \mathbf{q}_{k+1}) + \frac{\partial \mathbf{g}^T}{\partial \mathbf{q}_k} \boldsymbol{\lambda} = \mathbf{f}_{k+1} \quad (7)$$

where $\boldsymbol{\lambda}$ is a vector of Lagrange multipliers for the holonomic contact constraints \mathbf{g} . Given the selected contact configuration (*i.e.*, active contact points), we can estimate \mathbf{f} such that the forces on the root node are minimized. We

achieve that by minimizing the squared norm of the root forces

$$\left\| \mathbf{I}_{root} \left(\boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}^T}{\partial \mathbf{q}_k} \boldsymbol{\lambda} \right) \right\|_2^2 \quad (8)$$

where \mathbf{I}_{root} is a square selection matrix in the size of \mathbf{q} , with ones on the diagonal to select the six degrees of freedom of the root node. The regularized LS solution (by adding a small constant to the diagonal of \mathbf{I}_{root}) yields the contact forces required to minimize root forces, *i.e.*,

$$\boldsymbol{\lambda}^* = \left(\frac{\partial \mathbf{g}^T}{\partial \mathbf{q}_k} \mathbf{I}_{root} \frac{\partial \mathbf{g}}{\partial \mathbf{q}_k} \right)^{-1} \frac{\partial \mathbf{g}^T}{\partial \mathbf{q}_k} \mathbf{I}_{root} \boldsymbol{\varepsilon} \quad (9)$$

We can then calculate the final forces

$$\mathbf{f}_{k+1}^* = \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}^T}{\partial \mathbf{q}_k} \boldsymbol{\lambda}^* \quad (10)$$

that are used as a prior in tracking. The prior over the forces \mathbf{f}_{k+1}^* is defined twofold. We would like to minimize the root forces as a generic prior for a plausible motion, and we would like to minimize the internal torques to reduce jitter. Notice that given a contact configuration, both $\boldsymbol{\lambda}^*$ and \mathbf{f}_{k+1}^* are functions of $(\mathbf{q}_{k-1}, \mathbf{q}_k, \mathbf{q}_{k+1})$, thus our physics-based prior over the forces is applied directly over poses \mathbf{q} .

C. Registration and Tracking

Tracking is accomplished in an online fashion, by maximizing the posterior distribution over state parameters at each frame. As is common in online filtering, we assume conditional observation independence, and a second-order Markov model to account for acceleration in the physics-based prior. Accordingly, the posterior over state parameters at time k is proportional to the data likelihood and the conditional distribution over state parameters given those at previous time steps

$$p(\Theta_k | \mathcal{D}_k, \Theta_{k-1:k-2}) \propto p(\mathcal{D}_k | \Theta_k) p(\Theta_k | \Theta_{k-1:k-2}) \quad (11)$$

where \mathcal{D}_k is an input 3D point cloud at time k . By assuming Gaussian noise in observations, the negative log likelihood of the data term in Eq. 11 becomes

$$-\log p(\mathcal{D}_k | \Theta_k) = \sum_{(\mathbf{p}', \mathbf{d}') \in \Psi_k} \|\mathbf{p}' - \mathbf{d}'\|_2^2 \quad (12)$$

where Ψ_k holds all matching body model points $\mathbf{p}' \in \mathbf{p}(\Theta_k)$ and data points $\mathbf{d}' \in \mathcal{D}_k$ at time k . The matching was done in a standard Iterative Closest Point (ICP [27]) manner, and matched closest body and data points with a maximal distance threshold and pruning of back facing vertices. The data term captures the discrepancy between the model surface of the body, encoded by mesh vertices $\mathbf{p}_i(\Theta_k)$, and the observed depth data points.

The negative log likelihood of the conditional state probability is based on the physics-based priors, as described above, and takes the following form:

$$-\log p(\Theta_k | \Theta_{k-1:k-2}) = \gamma_1 \|\mathbf{f}_{root}^k\|^2 + \gamma_2 \|\mathbf{f}_{-root}^k\|^2 \quad (13)$$

where γ_1, γ_2 are prior weights, \mathbf{f}_{-root}^k comprises all but the root forces and accounts for smoothness in torques, and \mathbf{f}_{root}^k are root forces which account to physical plausibility.

Tracking is formulated as the optimization of a global objective \mathcal{F} , to find the parameters Θ_k at each frame that minimize errors between the body model, denoted $\mathcal{M}(\Theta_k)$, and an input 3D point cloud \mathcal{D}_k at time k . The objective is the negative log likelihood of Eq. 11, *i.e.*,

$$\mathcal{F}(\Theta_{k-2:k}, \mathcal{D}_k) = -\log p(\mathcal{D}_k | \Theta_k) - \log p(\Theta_k | \Theta_{k-1:k-2}) \quad (14)$$

as defined in Eq. 12, and Eq. 13 above. A natural way to optimize this objective function is to use a variant of ICP, *i.e.*, by alternating between correspondence and parameter optimization. Empirically, ICP tends to be both fast and accurate. We register our model in the first frame by optimizing Eq. 14 w.r.t. all parameters $(\mathbf{q}, \ell, \beta)$, and in the following frames update the pose \mathbf{q} only, while holding ℓ and β fixed.

D. Pre-processing

The proposed ICP algorithm requires initialization of the body model parameters Θ . In what follows we describe a fast and simple initialization method. Following [28], we exploit the observation that the geodesic distances between human end-effectors (*i.e.*, head, hands, feet) are both large and relatively independent of body pose, in order to automatically initialize tracking.

Geodesic Extrema as Scale/Rotation Invariant Mesh Features: Given an input point cloud \mathcal{D} , we generate a mesh \mathcal{D}_{mesh} (*i.e.*, connecting vertices with edges) using a greedy projection method for fast triangulation of unordered point clouds [29]. In case of grid-based depth input data, we use a method similar to [28], with a cut-off distance between nearby vertices (*i.e.*, threshold over maximal distance).

Given a connected component, we extract the first five geodesic extrema $\{\mathbf{g}_i | \mathbf{g}_i \in \mathcal{D}_{mesh}\}_{i=1}^5$ from the geodesic centroid of the mesh, $\bar{\mathbf{g}}$, as in [28]. We order the geodesic extrema by geodesic distance, so that $d(\mathbf{g}_i, \bar{\mathbf{g}}) \leq d(\mathbf{g}_{i+1}, \bar{\mathbf{g}})$, where $d(\cdot, \cdot)$ is the geodesic distance between two points. We define two features from which we detect human-like meshes, for labelling end-effectors and for finding initial poses for tracking. The first is a 4D vector that encodes the ratio of the ordered geodesic distances:

$$\phi_{ratio} = (r_1, \dots, r_4)^T, \quad r_j \equiv \frac{d(\mathbf{g}_{j+1}, \bar{\mathbf{g}})}{d(\mathbf{g}_j, \bar{\mathbf{g}})} \quad (15)$$

These features act like moments to describe the geodesic eccentricity of the point cloud. The second feature encodes geometric shape:

$$\phi_{pos} = \left\{ \begin{array}{l} \text{angles between all triplets} \\ \text{angles between all orientations} \end{array} \right\} \quad (16)$$

where orientation is defined as the vector between a geodesic extrema and the point 30cm along the geodesic path to the geodesic centroid.




Figure 5: Three examples of pose query based on ϕ_{pos} . Since, by design, the feature is scale and rotation invariant, we fetch different orientations of similar poses.

Detecting human-like components: We use the distance-ratio features to detect connected-components that might be people in the scene. To that end we learn a 4D Gaussian distribution over ϕ_{ratio} of human meshes. This distribution then provides a probability that a connected component is a plausible person. A threshold of 0.1 on that probability is used to cull non-human components. Even with this simple method we accurately detect about 90% of the human components with minimal false positives, which is sufficient for our application.

Pose Initialization: Given the extrema feature descriptors, we can register an unregistered point cloud \mathcal{D} by finding poses from a database of labelled point clouds whose features ϕ_{pos} are most similar to those of the point cloud (see examples in Fig. 5). In more details: we fetch a pose from a pose data-base (based on ϕ_{pos} L2 distance), we align the database mesh (with mean Θ) and the data point-cloud (based on fitted Ellipsoids to vertices), we estimate Θ (ICP).

E. Method Summary

To summarize, the tracking pipeline is as follow:

- 1: Divide \mathcal{D} into connected components
- 2: Remove non-human connected components by using ϕ_{ratio}
- 3: Initialize pose and register first frame (ICP)
- 4: **for all frames do**
- 5: Initialize pose with previous pose
- 6: Execute ICP
- 7: **end for**

IV. EXPERIMENTS

A. Execution Speed

The tracking system was implemented in Python, with the core physical components in C++. It was evaluated on a desktop running OS X, with Intel Core i7/2.3GHz and 8 GB RAM. At present it runs at 0.2[fps] with physical priors, and 1.96[fps] with body model only.

B. Quantitative Comparison

We used SMMC-10 dataset [30] for quantitative comparison. It comprises synchronized Vicon mocap marker data

Contact Point	No Kalman Filter	With Kalman Filter
Left Toe	94.6	98.9
Right Toe	95.6	98.5
Left Heel	77.4	84.0
Right Heel	87.9	95.9

Table I: A comparison between contact predictor performance, with and without Kalman filter. The values represent the percentage of predicting ground truth contact state.

	MJPE	MJIE
Proposed Method (Data Only)	1.7[cm]	0.975
Proposed Method (Physics)	1.89[cm]	0.973
Ganapathi	N/A	0.971
Baak	~ 5 [cm]	N/A

Table II: Quantitative results. The metrics MJPE and MJIE are defined in Sec. IV-B. While accurate numbers for Baak were not available, it is at best 5[cm], as shown in his work from 2013.

and Mesa SwissRanger ToF depth data. The depth data have significant amounts of noise, as can be seen in Fig. 7, and the accompanying video. We compare our method to [28] and [31] on the same dataset. We achieve state-of-the-art tracking accuracy (see Table II). We show the results of two metrics: Mean Joint Prediction Error (MJPE) - the RMSE of predicting the mocap markers from our skeleton joints, and Mean Joint Indicator Error (MJIE) - the average of all joint predictions that were within a range of 10[cm] from the target joint. Interestingly, when based solely on MSE metrics (as defined above), the physics-based prior does not appear to significantly affect performance (*e.g.*, see Table II). On the contrary, the accompanying video demonstrates how MSE metric does not reflect the physical plausibility of a motion. That is, there are many motions for a given MSE metric, most of which are not physically plausible per se.

Due to high SNR in SMMC-10 dataset’s depth scans, and as a result in our force estimation, we used an online Kalman filter as a noise filtering technique. We have found that this improves performance, reducing prediction error, by roughly 50% (see Table I).

We used a threshold of 0.8 contact probability to reduce sticky contact (*i.e.*, false positive contact prediction). This is the result of enforcing holonomic (equality) contact constraints instead of inequality constraints. Despite simplifying the joint distribution over all contact points into an independent probabilistic model per contact point, our model was accurate enough to allow tracking with 4 possible contact points. The contact prediction model predicted the correct contact configuration with more than 95% mean accuracy (see Table I).




Figure 6: Registration results of an accurate laser scan. From left to right: initial alignment, ICP correspondences, gradient descent step, and final registration.




Figure 7: Registration results of noisy depth sensor (SMMC-10 dataset). In yellow is the noisy depth input data. Notice that the amplitude of noise is larger than the thickness of arms and legs. Despite that, we get a reasonable registration, when visually inspected.

C. Qualitative Comparison

Fig. 6 depicts the process of registering an accurate laser scan with our model. Notice how we learn anthropometrics, body shape and pose. Similarly, in Fig. 7 we register our model to the first frame in a tracking sequence. Due to the low dimensionality of the model we are still able to register a plausible model to very noisy data.

We tested our registration technique (Sec. III-C) on Hasler [24], and SMMC-10 [31] datasets with promising initial results. When inspecting the tracking results in the accompanying video, there are some visible artifacts in the mesh model. Those, however, are mainly due to different pose distributions in training the mesh model and during tracking, rather than due to fundamental limitation of the model (excluding known artifacts of linear mesh blending such as volume collapse). Another interesting property of our mesh model is volume prediction per registered mesh. We used the mesh volume to calculate the inertial description needed for the physical priors, treating the volume as water. When compared with ground truth, we had an average weight prediction error of 1.5% of 115 subjects.

The true power of our approach is with the reduction of visual artifacts. While we do not remove jitter entirely, it is attenuated, when compared with data-only tracking. A more dramatic result is how foot-slide is removed in cases where contact is correctly detected. Despite the fact that our false positive contact estimation (wrong contact




Figure 8: An example of contact prediction. In the top graph, blue is the ground truth contact, while green is the estimated probability. The bottom graph depicts the three body parts with the largest weights in predicting contact.

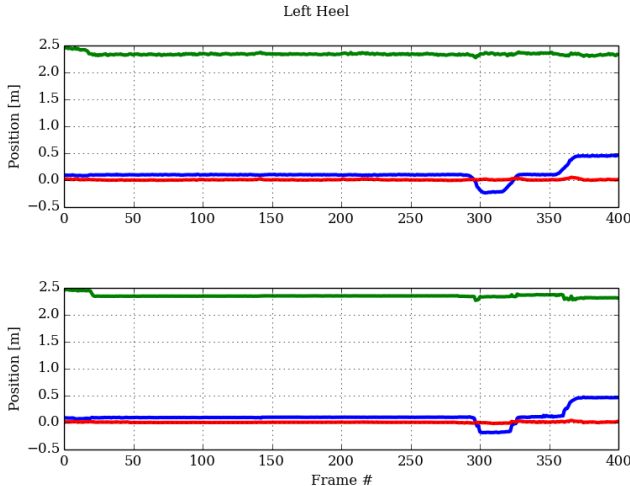


Figure 9: An example of the effects of contact on the foot. Top plot: RGB for XYZ of left heel position, with no contact constraints. Bottom: with contact constraints. Notice how the contact reduced jitter and foot slide, while still preserving the global pattern.

prediction) caused occasional visual artifacts, the value of removing foot-slide is much more noticeable, as evident in the accompanying video. To better understand how the force predictor works, Fig. 8 plots ground truth contacts, along with the corresponding forces. We considered the lowest marker, along with all markers up to $5[cm]$ away as in contact, due to lack of contact ground truth. Despite having noisy ground truth, our simple predictor was able to perform well on most frames.

Fig. 9 demonstrates how adding the contact constraint acts as a strong motion prior. While smoothing pose will reduce jitter, it will also reduce discontinuities in motion due to contact. On the other hand, applying physical contact constraints can smooth jitter while allowing abrupt changes in motion.

V. DISCUSSION

We propose an online physics-based 3D human tracking approach that incorporates physics-based priors into tracking without the need for subject calibration or knowledge about the environment. The use of physics in this context is compelling as it allows us to minimize visual artifacts, most notably jitter and foot-skate which results from noise and occlusions. We demonstrate that we can infer contact from joint torque trajectories computed by inverse dynamics. We show that our method is effective at tracking from a noisy single depth sensor and produces quantitative results that are on par or better than current state-of-the-art, while at the same time qualitatively reducing visual artifacts.

Our contact prediction model, while conceptually compelling, is relatively simple. For example, we predict all of the contact points independently, despite the fact that contact patterns (especially for contact points on connected segments) are clearly correlated. We believe the prediction model can be further improved by structured prediction that incorporate these correlations in contact state. While our method predicts the environment online, it currently does not aggregate these predictions, which may be important for longer sequences.

We also note that in our method the ground can both push and pull on the body when contact is established. This can sometimes be seen in the video. While this behaviour is not realistic in terms of underlying physical behaviour, we nevertheless believe it allows us to overcome vast amounts of noise in the observations (especially where feet can easily be confused with occluders and the ground plane).

Finally, we note how MSE metrics do not capture the dynamics of a motion. For example, the same MJPE can represent a motion that is the ground truth with a constant added to it, or the ground truth with a Gaussian noise with the same constant as a standard deviation, and a zero mean. This exposes the limitation of relying on a MSE metric to assess the quality of tracking results, and highlight the advantage of using physics-based prior. In a sense, the physics-based prior shapes the results to be qualitatively superior, despite not improving the MSE metric itself.

REFERENCES

- [1] E. de Aguiar, C. Stoll, C. Theobalt, N. Ahmed, H.-P. Seidel, and S. Thrun, "Performance capture from sparse multi-view video," *ACM Trans. Graph.*, vol. 27, no. 3, pp. 98:1–98:10, Aug. 2008.

- [2] X. Wei, P. Zhang, and J. Chai, "Accurate realtime full-body motion capture using a single depth camera," *ACM Trans. Graph.*, vol. 31, no. 6, pp. 188:1–188:12, Nov. 2012.
- [3] A. Haque, B. Peng, Z. Luo, A. Alahi, S. Yeung, and F. Li, "Viewpoint invariant 3d human pose estimation with recurrent error feedback," *CoRR*, vol. abs/1603.07076, 2016.
- [4] M. Prazak, L. Hoyet, and C. O'Sullivan, "Perceptual evaluation of footskate cleanup," *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2011.
- [5] M. A. Brubaker, D. J. Fleet, and A. Hertzmann, "Physics-based person tracking using the anthropomorphic walker," *International Journal of Computer Vision*, vol. 87(1), pp. 140–155, 2010.
- [6] M. Vondrak, L. Sigal, and O. Jenkins, "Dynamical simulation priors for human motion tracking," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, no. 1, pp. 52–65, 2013.
- [7] X. Wei and J. Chai, "Videomocap: Modeling physically realistic human motion from monocular video sequences," *ACM Trans. Graphics (SIGGRAPH)*, vol. 29(4), 2010.
- [8] M. A. Brubaker, L. Sigal, and D. J. Fleet, "Estimating contact dynamics," in *Proc. IEEE ICCV*, 2009.
- [9] T. Moeslund, A. Hilton, and V. Kruger, "A survey of advances in vision-based human motion capture and analysis," *CVIU*, vol. 104, no. 2, 2006.
- [10] L. Ren, G. Shakhnarovich, J. K. Hodgins, H. Pfister, and P. Viola, "Learning silhouette features for control of human motion," *ACM Transactions on Graphics*, 2005.
- [11] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, "Real-time human pose recognition in parts from a single depth image," in *CVPR*, 2011.
- [12] R. Y. Wang and J. Popovic, "Real-time hand-tracking with a color glove," *ACM Transactions on Graphics*, vol. 28, no. 3, 2009.
- [13] X. Zhou, M. Zhu, K. Derpanis, and K. Daniilidis, "Sparseness meets deepness: 3d human pose estimation from monocular video," in *CVPR*, 2016.
- [14] B. Tekin, I. Katircioglu, M. Salzmann, V. Lepetit, and P. Fua, "Structured Prediction of 3D Human Pose with Deep Neural Networks," *Bmvc*, pp. 1–11, 2016.
- [15] J. Deutscher and I. Reid, "Articulated body motion capture by stochastic search," *International Journal of Computer Vision*, vol. 61, no. 2, pp. 185–205, 2004.
- [16] P. F. Felzenszwalb and D. P. Huttenlocher, "Pictorial structures for object recognition," *Int. J. Comput. Vision*, vol. 61, no. 1, pp. 55–79, Jan. 2005.
- [17] D. Metaxas and D. Terzopoulos, "Shape and nonrigid motion estimation through physics-based synthesis," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 15, no. 6, pp. 580–591, 1993.
- [18] C. R. Wren and A. Pentland, "Dynamic models of human motion," in *IEEE International Conference on Automatic Face and Gesture Recognition*, 1998.
- [19] M. Vondrak, L. Sigal, J. Hodgins, and O. Jenkins, "Video-based 3d motion capture through biped control," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 27:1–27:12, Jul. 2012.
- [20] B. Rosenhahn, C. Schmaltz, and T. Brox, "Staying well grounded in markerless motion capture," *Pattern . . .*, pp. 385–395, 2008.
- [21] Y. Ye and C. K. Liu, "Synthesis of detailed hand manipulations using contact sampling," *ACM Trans. Graphics (SIGGRAPH)*, vol. 31, no. 4, 2012.
- [22] I. Mordatch, E. Todorov, and Z. Popovic, "Discovery of complex behaviors through contact-invariant optimization," *ACM Trans. Graphics (SIGGRAPH)*, vol. 31, no. 4, 2012.
- [23] D. Anguelov, P. Srinivasan, and D. Koller, "Scape: shape completion and animation of people," *ACM Transactions on Graphics*, 2005.
- [24] N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel, "A statistical model of human pose and body shape," *Computer Graphics Forum*, vol. 28, no. 2, pp. 337–346, 2009.
- [25] J. E. Marsden and M. West, "Discrete mechanics and variational integrators," *Acta Numerica*, vol. 10, pp. 357–514, 2001.
- [26] E. Johnson and T. Murphey, "Scalable variational integrators for constrained mechanical systems in generalized coordinates," *Robotics, IEEE Transactions on*, vol. 25, no. 6, pp. 1249–1261, 2009.
- [27] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, Feb 1992.
- [28] A. Baak, M. Müller, G. Bharaj, H.-P. Seidel, and C. Theobalt, "A data-driven approach for real-time full body pose reconstruction from a depth camera," in *Consumer Depth Cameras for Computer Vision*, ser. Advances in Computer Vision and Pattern Recognition, A. Fossati, J. Gall, H. Grabner, X. Ren, and K. Konolige, Eds. Springer London, 2013, pp. 71–98.
- [29] Z. C. Marton, R. B. Rusu, and M. Beetz, "On Fast Surface Reconstruction Methods for Large and Noisy Datasets," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Kobe, Japan, May 12-17 2009.
- [30] V. Ganapathi and C. Plagemann. (2010, March) SMMC-10 dataset. <http://ai.stanford.edu/varung/cvpr10>.
- [31] Ganapathi, Varun, Plagemann, and Christian, "Real-time human pose tracking from range data," *Computer Vision—ECCV*, pp. 1–14, 2012.