

Optical Non-Line-of-Sight Physics-based 3D Human Pose Estimation

Mariko Isogawa, Ye Yuan, Matthew O’Toole, Kris Kitani
Carnegie Mellon University




Figure 1. **Non-Line-of-Sight (NLOS) Physics-Based 3D Human Pose Estimation.** A pulsed laser and a transient sensor record the time light takes to travel from a point on a wall to the person hidden from sight (left). Given a sequence of 3D transient measurements from an optical NLOS system (top right), our method estimates a 3D pose sequence using a physics-based policy (bottom right).

Abstract

We describe a method for 3D human pose estimation from transient images (i.e., a 3D spatio-temporal histogram of photons) acquired by an optical non-line-of-sight (NLOS) imaging system. Our method can perceive 3D human pose by ‘looking around corners’ through the use of light indirectly reflected by the environment. We bring together a diverse set of technologies from NLOS imaging, human pose estimation and deep reinforcement learning to construct an end-to-end data processing pipeline that converts a raw stream of photon measurements into a full 3D human pose sequence estimate. Our contributions are the design of data representation process which includes (1) a learnable inverse point spread function (PSF) to convert raw transient images into a deep feature vector; (2) a neural humanoid control policy conditioned on the transient image feature and learned from interactions with a physics simulator; and (3) a data synthesis and augmentation strategy based on depth data that can be transferred to a real-world NLOS imaging system. Our preliminary experiments suggest that our method is able to generalize to real-world NLOS measurement to estimate physically-valid 3D human poses.¹

1. Introduction

Our goal is to develop a vision-based human pose estimation technique capable of inferring the 3D pose sequence

of a person without direct line of sight from the sensing device (see Fig. 1). To this end, this work brings together, for the first time, state-of-the-art techniques in computational imaging, human pose estimation, humanoid control, and deep reinforcement learning – all under one system architecture. From computational imaging, we use volumetric reconstruction algorithms that efficiently process transient measurements (a 3D spatio-temporal histogram of photons) captured from a non-line-of-sight (NLOS) imaging system. From vision-based pose estimation and humanoid control, we make use of deep reinforcement learning in a physics simulator to learn an image feature-conditioned 3D human pose policy that abides by the laws of physics. All components are integrated into a single end-to-end learnable architecture which takes the raw transient images, converts them to a deep feature representation using a learnable inverted point spread function (PSF), and uses a neural humanoid control policy conditioned on the deep feature to estimate physically valid 3D human poses. In short, our proposed end-to-end system models the complete flow of information from the photon-level all the way to the high-level cognitive task of 3D human pose estimation.

In order to capture the actions of a person around a corner, we employ a NLOS imaging technique using a pulsed laser and a time-of-flight sensor [30]. The imaging process involves sending light pulses towards a visible surface, and measuring the travel time of light reflected back in response; this represents a “transient image” [31, 4, 16]. The transient image used here is a 3D volume that encodes the ar-

¹Project page: https://marikoisogawa.github.io/project/nlos_pose

rival of individual photons over a spatial region (first two dimensions) and over time (third dimension); see Fig. 1 for a sequence of transient images. Optical NLOS imaging at visible or near-IR wavelengths offers several advantages, such as the ability to image at long-ranges [3], reconstruct surfaces at micron-scales [32], and robustly perform imaging for multiple objects with different BRDFs [13].

While NLOS imaging gives us the ability to look around corners, the raw transient measurements acquired by a NLOS system have several unique properties that make 3D human pose estimation difficult. In order to attain fast image acquisition necessary for 3D human pose understanding, existing NLOS solutions limit the resolution of the transient image in both space and time [13]. This makes it difficult to capture small shape details and fast motion, both of which are important for estimating a human pose sequence. Furthermore, due to the light lost after multiple scattering events, very few photons reach the sensor and the acquired transient image can therefore be very noisy. All of these characteristics make it very challenging to estimate 3D human pose directly from the transient image.

Due to the noisy nature of transient images, it is important to make use of prior information about the body and its dynamics to help solve the 3D human pose estimation problem. In particular, we have prior knowledge about the structure of the human body (*e.g.*, arms, head, and legs are connected to the torso) and we also know that human pose must abide by the laws of physics. Based on this prior knowledge, we make use of a parametrized humanoid model and a physics-based deep reinforcement learning approach to estimate 3D human pose from a sequence of transient images. Inspired by [35], we first define a Markov Decision process (MDP) where the state is defined as the current pose of a humanoid model and a sequence of transient images, the action is defined as the target joint angles of a humanoid, the environment is a physics simulator (*i.e.*, MuJoCo [28]) and the reward function is an imitation-based objective function [18]. By solving the MDP via reinforcement learning, we are able to learn an optimal policy that maps states to actions. In our scenario, the policy function encodes how the humanoid should move such that it best explains the 3D pose information captured in the transient image sequence.

The deep neural network design used to implement the policy of the MDP is a critical representational element of the system, as it must have the computational components necessary to map a 3D transient image to a vector of humanoid joint angles. Inspired by recent NLOS imaging techniques, we model our network based on existing reconstruction algorithms used to process transient images, and introduce P2PSF Net: a neural network that learns “corrections” that improve the NLOS image reconstruction process. A deep feature computed by the P2PSF Net is then passed to a bi-directional recurrent neural architecture to

ensure that the output of the policy (the next joint position) follows the laws of motion and matches demonstrated human pose dynamics. The ambitious goal of the policy is to estimate fine-scale 3D human pose from transient images and overcome the resolution limits and noisy characteristics of NLOS imaging.

As mentioned above, the MDP formulation for human pose estimation makes use of an imitation-based reward function which requires a large amount of annotated data in the form of transient images and a corresponding 3D human pose sequence. As a practical matter, current NLOS imaging systems are quite costly to build and non-trivial to calibrate. As such, the need for a large amount of annotated data and the cost of collecting such data are at odds with each other. While we do believe that these practical hurdles will be resolved through new innovations over time, for this work, we address them by synthesizing pseudo-transient images using the 3D body volume recovered with a depth camera. To learn our MDP policy, we rely solely on this synthetic data synchronized with ground truth 3D human pose acquired by a motion capture system.

Synthesizing transient images from depth maps for training can create a large domain gap between the training and test data distributions. To ensure that we minimize the size of this domain gap, it is critical to implement principled data augmentation techniques based on the physics of the imaging process, *e.g.*, reproducing the noise characteristics associated with real transient images. We introduce several augmentation techniques to make the 3D human pose estimation policy able to handle real-world NLOS data.

To summarize, our contributions are as follows:

- (1) We are the first to propose an end-to-end data processing pipeline that models information flow from optical NLOS transient measurements to 3D human pose estimation.
- (2) We present a novel policy network architecture that integrates an inverse point spread function (PSF), transient images, and human pose as one deep neural network.
- (3) We describe a method for generating pseudo (*i.e.*, synthetic) transient images from depth images. We propose several critical processes for augmenting pseudo-transient images to help the learned 3D human pose policy to generalize better to real-world data.
- (4) We provide extensive experimentation with both synthetic and real transient images and show that our model outperforms other baseline methods.

2. Related Work

NLOS Imaging and Its Applications. NLOS imaging has received significant attention recently [1, 2, 3, 6, 8, 9, 10, 12, 13, 17, 30, 14, 21, 27, 29, 32, 36]. This includes a wide variety of solutions, in terms of hardware systems and reconstruction algorithms, that operate in different parts of the electromagnetic spectrum. Most NLOS solutions

work with visible (380-740 nm) or near-IR (740-1500 nm) light [3, 13, 32]; we refer to these as *optical* NLOS systems. Others choose to operate at longer wavelengths, including long-wave IR (8-15 μm) [14], terahertz (millimeter scale) [21], WiFi and radio frequency (RF) (centimeter scale) [10], or even sound waves [12].

Operating in a specific part of the spectrum affects the nature of the NLOS signal that can be used for pose estimation. For example, RF or WiFi signals enable through-the-wall pose estimation [1, 10, 36], since longer electromagnetic waves tend to pass through objects; however, as highlighted in Fig. 2, this also comes at a number of fundamental limitations when estimating pose dynamics: the spatial and angular resolution must be limited, making it hard to distinguish small objects or motions; objects tend to behave specularly at long wavelengths, presenting a challenge when attempting to form images; RF signals can be highly attenuated when passing through thick objects; and RF signals are blocked entirely when passing through other surfaces, like metal or water. As a result, they cannot always be applied to scenes shown in Fig. 2 (a), (c) and (d).

Rather than using RF signals to attempt imaging through walls, most NLOS works choose to operate in the optical domain and image *around* corners. This is done by reflecting light off of walls to indirectly illuminate a hidden scene, and measuring the time light takes to return to the wall with transient sensors, *e.g.*, single photon avalanche diodes (SPADs). Amongst these works, many attempt to reconstruct 3D volumes [6, 13, 17, 30] where each voxel represents the scene’s reflectivity at a 3D location; or compute a surface representation [29, 32], offering the potential for more accurate reconstructions. Even regular cameras can be used to track a person walking around in a room under either passive [2] or active [8, 9, 27] illumination.

Data-driven methods for optical NLOS imaging have also received some attention, *e.g.*, to discretely classify the pose of a human from raw transient measurements [24] or form NLOS images with a regular camera [27]. Network architectures have also been proposed for denoising transient measurements from SPADs, in the context of line-of-sight imaging [11]. However, there has been no prior work addressing the task of learning full physics-based 3D human pose estimation in a NLOS context.

Physics Based 3D Human Pose Estimation. Enforcing physics-based constraints on human pose dynamics is used commonly for simulated humanoid control [18, 19, 20, 34, 35]. These methods use deep reinforcement learning (DeepRL) to learn control policies that can reproduce humanoid motion inside a physics simulator. Among these methods, [34, 35] use the optical flow of egocentric videos as additional inputs to the motion policy to estimate physically-valid human poses. We use a similar RL-based pose estimation framework proposed in [35] for its ability to




Figure 2. Overview of WiFi/RF (*i.e.*, through-the-wall) and visible/near-IR (*i.e.*, around-the-corner) NLOS human pose estimation. (a) The person is hidden from view in a nearby corridor, but seen indirectly by reflecting light off of a visible surface. (b) The person is hidden in an enclosed room, which can only be accessed with WiFi/RF that pass through walls. (c) The person is within a partially-enclosed space that can be accessed through either method. (d) The walls are either too thick or made from a material that prevents WiFi and RF signals from passing through.

estimate physically-accurate human poses with only limited amounts of visually information. However, the transient images that we use as input present many new challenges when learning the policy, which we address in this paper.

3. Methodology

There are three key ideas behind our method. First, to model physically-valid human pose, we use a humanoid model inside a physics simulator and model hidden 3D human pose as the result of a humanoid control policy conditioned on the input transient images (Sec. 3.1). Second, to make our model applicable to real captured data, we synthesize “pseudo-transient” images as training data which replicates the visual characteristics of real transient measurements (Sec. 3.2). Third, to enhance our pose estimation, we further introduce P2PSF Net that improves the NLOS imaging process (Sec. 3.3).

3.1. Physics-based 3D Pose Estimation Framework

Given a sequence of transient images $\tau_{1:T}$, the goal of our physics-based pose estimation pipeline is to predict a physically-valid pose sequence $p_{1:T}$. As shown in Fig. 3, the pipeline contains two main parts: (1) We use a feature extractor consisting of the NLOS imaging pipeline and P2PSF Net to extract transient features $\psi_{1:T}$ from the transient images $\tau_{1:T}$; (2) We use a humanoid policy conditioned on the transient features $\psi_{1:T}$ to control the humanoid inside a physics simulator and generate the pose sequence $p_{1:T}$ underlying the transient images. In this section, we will focus on part 2, *i.e.*, humanoid control, and leave details on how to extract transient features (part 1) to Sec. 3.2 and 3.3.

Following [35], we formalize the task of estimating a

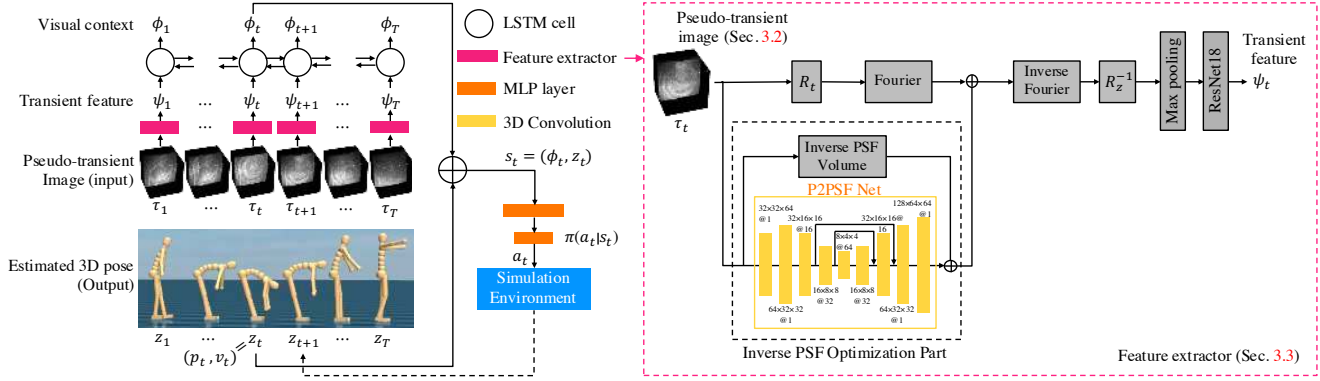


Figure 3. Our DeepRL based photons to 3D human pose estimation framework under the laws of physics.

pose sequence $p_{1:T}$ from a transient image sequence $\tau_{1:T}$ with a Markov Decision process (MDP). The MDP is defined by a tuple $\mathcal{M} = (S, A, P, R, \gamma)$ of states, actions, transition dynamics, a reward function, and a discount factor. As Fig. 3 shows, at each time step, the humanoid agent samples an action a_t from a policy $\pi(a_t|s_t)$ whose input state s_t contains both the visual context ϕ_t (defined later) and the humanoid state z_t (*i.e.*, position and velocity of the human model joints). Next, the environment generates the next state s_{t+1} through physics simulation and gives the agent a reward r_t based on how well the humanoid’s 3D pose aligns with the ground-truth. To solve this MDP, we apply the PPO policy gradient method [25] to obtain the optimal policy π^* that maximizes the expected discounted return $\mathbb{E}[\sum_{t=1}^T \gamma^{t-1} r_t]$. At test time, starting from some initial state s_1 , we roll out the policy π^* to generate state sequence $s_{1:T}$, from which we extract the output pose sequence $p_{1:T}$. In the following, we will discuss the details of each component.

State s_t consists of the state of the humanoid z_t and the visual context ϕ_t . z_t consists of the pose p_t (root position/orientation and joint angles) and velocity v_t (root linear/angular velocities and joint velocities). The visual context ϕ_t is extracted from the transient feature $\psi_{1:T}$ with a bi-directional LSTM (BiLSTM) as shown in Fig. 3. During training, we set the starting state z_1 to the ground-truth \hat{z}_1 . Since we have no access to the ground truth at test time, we learn a regressor \mathcal{F} that maps the visual context ϕ_t to its corresponding state z_t .

Action a_t specifies the target joint angles of the Proportional-Derivative (PD) controller for each degree of freedom (DoF) of the humanoid joints except for the root. The joint torques are computed based on the stable PD controllers in [26] using the specified target joint angles a_t .

Policy $\pi_\theta(a_t|s_t)$ is represented by a Gaussian distribution $\mathcal{N}(\mu; \Sigma)$ with a fixed diagonal covariance matrix treated as hyperparameters. We use a θ -parametrized multilayer perceptron (MLP) with two hidden layers (300, 200) and ReLU

activations to model the mapping from s_t to the mean μ_t .

Reward function. To encourage the policy to output a pose sequence $p_{1:T}$ that matches the ground-truth $\hat{p}_{1:T}$, we use the reward function proposed in [35]. The specific design of the reward function is given in the supplementary materials.

3.2. Pseudo-Transient Image

Before introducing our transient data synthesis from depth maps (Sec. 3.2.2), we provide background on the NLOS imaging procedure used in this work in Sec. 3.2.1.

3.2.1 Background: Confocal NLOS Imaging

Confocal NLOS imaging refers to a specialized raster-scanning procedure for capturing transient measurements [17]. Measurements are the result of illuminating and imaging a common point (x', y') on a visible surface. After illuminating this point with a pulse laser, light scatters to hidden regions of an environment and returns back to the wall at a later instant of time. A SPAD measures the transient response at this same point, represented as a histogram of photon arrival times [16]. This procedure is repeated for a uniform and planar 2D grid of points across the surface; the collection of transients is stored as a 3D transient image $\tau(x, y, t)$.

The objective of NLOS imaging is to convert this 3D transient $\tau(x', y', t)$ into a discretized reconstruction volume $\rho(x, y, z)$, which represents the reflectivity at every point (x, y, z) in space. The presence of an object in voxel (x, y, z) produces a non-zero reflectance value $\rho(x, y, z)$, where the wall is located at $z = 0$.

After resampling the transient measurements τ along the time dimension ($\tilde{\tau} = R_t\{\tau\}$) and resampling the reconstruction volume ρ along the depth dimension ($\tilde{\rho} = R_z\{\rho\}$), the forward image formation model for confocal NLOS images becomes a simple 3D convolution operation [17]:

$$\tilde{\tau} = h * \tilde{\rho}, \quad (1)$$

where the hidden volume $\tilde{\rho}$ is convolved with a known 3D point spread function (PSF) represented by h . This PSF h describes the transient response of a single scatterer.

Equivalently, we can rewrite this convolution in matrix-vector form as follows:

$$\tau = R_t^{-1} F^{-1} \hat{H} F R_z \rho, \quad (2)$$

where we vectorize volumes τ and ρ . The matrix F represents a 3D discrete Fourier transform and \hat{H} is a diagonal matrix representing the Fourier transform of the PSF h .

The NLOS imaging process reconstructs a 3D volume ρ_* from a transient image τ by inverting Eq. 2 and solving a 3D deconvolution procedure (*e.g.*, using the Wiener filter):

$$\rho_* = R_z^{-1} F^{-1} \underbrace{\left[\frac{\hat{H}^*}{|\hat{H}|^2 + \frac{1}{\alpha}} \right]}_{\text{the inverse PSF}} F R_t \tau, \quad (3)$$

where a user-defined parameter α controls how sensitive the inverse PSF is to noise.

3.2.2 Pseudo-Transient Image Synthesis

We propose a procedure for synthesizing pseudo-transient images to generate annotated training data for our pose estimator. This involves using a motion capture (MoCap) system to obtain 3D ground truth human pose, which is synchronized with a depth camera to capture depth images used for synthesizing pseudo-transient images.

Given a depth map $d(x, y)$, we compute a corresponding synthesized reflectance volume for our scene: $\rho_s(x, y, z) = a$ when $z = d(x, y)$, and zero otherwise. The scalar $a = 100$ is a constant value representing the amount of light reflected by voxels in the volume.² We then convert the synthesized volume ρ_s into transient measurements τ by using the image formation model described in Eq. 2.

Note, however, that the image formation model is currently incomplete, since it does not model all characteristics of real-world transient measurements. For dynamic scenes (*e.g.*, a person walking around in a room), acquisition times must be short. Raster scanning transient measurements, therefore, is affected by motion, and the number of sampled points on the wall may be limited. The sensors used for NLOS imaging (*e.g.*, SPADs) are also affected by different types of sensor noise.

To close the domain gap that exists between the synthesized data and the real transient measurements, we introduce Poisson noise, temporal blur, temporal shifts, and temporal re-sampling into the image formation model.

Poisson noise. At low light levels, the number of photons detected by a SPAD approximately follow Poisson noise

²Though we assume uniform reflectance in this paper, an intensity image could be used to encode non-uniform reflectance information.

characteristics [16]. We therefore apply Poisson noise to our synthetic transient measurements, as a post-process operation on Eq. 2.

Temporal blur. The temporal profile of a transient measurement is affected by jitter and the shape of the laser pulse. The result is that the transient measurements become blurred in the time domain. Before the Poisson noise step, we introduce temporal blur by convolving the transients measurements with a Gaussian [16]. The standard deviation of the Gaussian is characterized by its full width at half maximum (FWHM), which is reported to be 70 picoseconds for the real-world NLOS data used in this work [13].

Temporal shift. The arrival times of photons are correlated with the distance of the hidden person from the wall. We augment our data by biasing the depth values by a constant amount d prior to synthesizing pseudo-transient images, which temporally shifts the transient images. In our experiments, we make use of five levels of shifts, augmenting the training data up to a factor of five.

Temporal re-sampling. Confocal NLOS imaging requires raster scanning a visible surface point by point. The mechanics of current systems impose limits on the speed of this raster scanning procedure. For example, we make use of NLOS data sampled at 32×32 locations at a frame rate of only 4 Hz [13]. Since points and corresponding transients are scanned sequentially, performance of pose estimation can be affected by fast moving body parts.

We propose a procedure to temporally re-sample transient data. First, to simulate the raster scanning procedure, we capture depth maps at 30 Hz, convert them to pseudo-transient measurements, and simulate raster-scanned measurements by down-sampling the result to 4 Hz. Second, we up-sample the pseudo-transient data from 4 Hz back to 30 Hz. Each transient image is a collection of transients scanned from time t_k to $t_k + \frac{1}{4}$, where $t_k = \frac{k}{4}$ represents the start time of the k^{th} frame (in seconds). To generate a 30 Hz transient sequence, we simply assemble transients captured within the same time range, but set the start time of the k^{th} frame to $t_k = \frac{k}{30}$. For real transient images used at test time, we also perform above procedure to up-sample them to 30 Hz to be compatible with the humanoid policy.

3.3. P2PSF Net: Photon to Inverse PSF Network

To predict 3D human pose from transient images, as discussed in Sec. 3.1, Our DeepRL based framework needs a feature extractor to obtain transient feature ψ_t from transient image τ_t . We model the feature extractor by incorporating aspects of the confocal NLOS imaging process (described in Sec. 3.2.1) into our feature extractor network. Note that learning to estimate 3D human pose from transient images is challenging for two key reasons: (1) the transient images are noisy, have low spatial resolution, and are recorded at slow frame rates; (2) the inverse PSF, labeled

in Eq. 3, makes several simplifying assumptions about the NLOS imaging process.

To overcome these challenges, we propose P2PSF Net: the photon-to-inverse-PSF network. Its objective is to adjust the inverse PSF volume used in the confocal NLOS reconstruction process (*e.g.*, to help handle realistic sensor noise, calibration errors, motion during the acquisition period). P2PSF Net is a 3D volume-to-volume network with nine 3D convolution layers (see Fig. 3). It also has two residual connections, which are inspired by past works [5, 22]. Given a transient image of resolution $(x, y, t) = 32 \times 32 \times 64$, the network outputs a volume $(x, y, d) = 128 \times 64 \times 64$, matching the size of the inverse PSF volume used in confocal NLOS imaging.

To extract transient features ψ_t , we first reconstruct the reflectance volume ρ_* following the procedure used in confocal NLOS imaging (Eq. 3) with the modification that we add the output of P2PSF Net as a corrective volume to the inverted PSF volume. Given the reconstructed reflectance volume ρ_* , we then generate a 2D heat map by applying a single max pooling layer across the depth axis of the volume. The 2D heat map is then passed to a ResNet-18 [5] pretrained on ImageNet [23] to extract ψ_t .

4. Experimental Settings

4.1. Datasets

Pseudo-Transient Image Dataset. We generated a large set of synthetic pseudo-transient images from MoCap data. The dataset is 1 hour long and has 103200 frames with 30 FPS. It consists of five subjects performing various complex motions: walking, jumping, turning, bending backward/forward, rotating, and transitioning between all of these motions. Our method does not require segmenting the pose sequences or labeling the actions.

Real Captured Transient Image Dataset. To further showcase our methods applicability, we also test our method on actual transient images [13]. This is a challenging dataset since it contains a different subject and the data acquisition conditions are different compared to the training data. This test data also contains significant noise, samples a sparse grid (32×32), and records measurements at a low frame-rate (*i.e.*, 4 Hz). The person is wearing a retroreflective tracksuit to increase the light signal used in NLOS imaging. Through our temporal re-sampling procedure (Sec. 3.2.2), we up-sampled this data to 30 FPS and 1000 frames. Since this dataset has no pose ground truth, we extract 2D keypoint ground truth with AlphaPose [33] using a third-person RGB camera.

4.2. Baseline Methods

1. **V2V-PoseNet [15]:** Since there is no existing work on optical NLOS human pose estimation from tran-

sient images, we compare against the state-of-the-art method, V2V-PoseNet, for depth volume-based pose estimation [15]. We train V2V-PoseNet using the reconstructed depth volume ρ_* computed from our synthetic pseudo-transient images using the regular NLOS imaging process (without P2PSF Net) defined in Eq. 3.

2. **PoseReg:** To investigate the effect of physics inside our physics-based pose estimation framework, we compare our method a regression-based method, which directly maps the visual context ϕ_t to the humanoid state z_t without using any physics. We integrate root linear/angular velocities to generate global positions and orientations of the pose sequence.
3. **PoseReg w/o P2PSF:** a variant of PoseReg that does not use proposed P2PSF Net in the NLOS imaging process.
4. **Ours w/o P2PSF:** a variant of our method that does not use proposed P2PSF Net.

4.3. Evaluation Metrics

We use the following metrics to evaluate both the accuracy and physical correctness of each method:

1. **Mean Per-Joint Position Error (MPJPE):** a pose-based metric that measures the Euclidean distance (measured in millimeters) between the ground truth and prediction for a joint defined as $\frac{1}{TJ} \sum_{t=1}^T \sum_{j=1}^J \| (x_t^j - x_t^{root}) - (\hat{x}_t^j - \hat{x}_t^{root}) \|_2$, where x_t^j is the j^{th} joint position of estimated pose and \hat{x}_t^j is the ground truth. x_t^{root} and \hat{x}_t^{root} represent root joint positions of the estimate and the ground truth.
2. **2D Keypoint Error (E_{key}):** A pose-based metric used for real captured dataset, calculated as $\frac{1}{TJ} \sum_{t=1}^T \sum_{j=1}^J \| y_t^j - \hat{y}_t^j \|_2$. Here, y_t^j is the j^{th} 2D keypoint of estimated pose and \hat{y}_t^j is the ground truth. For both estimated and ground-truth, we set the hip keypoint as the origin and scale the coordinates to make the height between shoulder and hip equal to 0.5.
3. **Velocity Error (E_{vel}):** A physics-based metric that measures the Euclidean distance between the generated velocity sequence $v_{1:T}$ and the ground-truth $\hat{v}_{1:T}$, calculated as $\frac{1}{T} \sum_{t=1}^T \| v_t - \hat{v}_t \|_2$.
4. **Average Acceleration (A_{acc}):** A physics-based metric that uses the average magnitude of joint accelerations to measure the smoothness of the pose sequence, calculated as $\frac{1}{TN} \sum_{t=1}^T \| \dot{v}^t \|_1$ where \dot{v}_t denotes joint accelerations and N is the number of actuated DoFs.

4.4. Implementation Details

Simulation Environment. We use MuJoCo [28] as the physics simulator for a humanoid that consists of 58 DoFs and 21 rigid bodies. We use stable PD controllers [26] to compute joint torques. The gains k_p ranges from 50 to 500,




Figure 4. Qualitative results for (a) single-subject and (b) cross-subject pose estimation.

where joints such as legs and spine have larger gains while arms and head have smaller gains; k_d is set to $0.1k_p$.

Networks and Training. We set the reward weights (w_q, w_e, w_p, w_v) to (0.5, 0.3, 0.1, 0.1). We use PPO [25] with a clipping epsilon of 0.2 for policy optimization. The discount factor γ is 0.95. We use Adam [7] to optimize the policy and value function with learning rate $5e-5$. The policy typically converges after 3k iterations, which takes about a day on a GeForce RTX 2080 Ti.

5. Experiments and Results

We evaluate our method against the baselines in three different experiment settings and one ablation study: (1) single subject with pseudo-transient data, using the same subject in both training and testing; (2) cross-subject with pseudo-transient data, using different subjects for training and testing; (3) cross-subject with real captured data, where we use real transient images even though our model is trained solely on synthetic data. We further conduct an ablation study to show the importance of each augmentation in our proposed pseudo-transient image synthesis procedure.

Single-Subject Evaluation. We train our model for each subject. We use a 80-20 train-test data split. As shown in the quantitative results (Table 1 (left)), we can see our approach outperforms other baselines in terms of both pose-based

metric (MPJPE) and physics-based metrics (\mathbf{E}_{vel} , \mathbf{A}_{accel}). We also present qualitative results in Fig. 4 (a). As the white dotted rectangles show, V2V-PoseNet and our method without P2PSF Net are not able to reproduce finer motion like walking. Also, non-physics-based methods (*i.e.*, all baselines except Ours w/o P2PSF) often cause the humanoid’s foot to sink into the ground (highlighted by the yellow circles). In contrast, our method produces 3D human pose closer to the ground-truth than any other baseline.

Cross-Subject Evaluation. To further evaluate the robustness and generalization of our method, we perform cross-subject experiments where we train our models on four subjects and test on the other subject. This is challenging since different people have unique action characteristics. As shown in Table 1(middle), our method once again outperforms other baselines in terms of MPJPE and \mathbf{E}_{vel} . For the smoothness metric \mathbf{A}_{accel} , V2V-PoseNet is quantitatively the best but this is because it mainly outputs a single pose only (indicated by the high MPJPE) and cannot estimate finer movement such as walking or raising hands, as shown in Fig. 4 (b). As the qualitative results in Fig. 4 (b) shows, our method with P2PSF Net produces poses closest to the ground truth, although it still fails during difficult sequences. Green arrows in Fig. 4 (b) highlight large differences in the estimated pose compared to the ground truth.

Table 1. Results on single (left), cross-subject (middle), and real captured data (right).

Method	Single Subject			Cross Subject			Real Data
	MPJPE ↓	E_{vel} ↓	A_{accl} ↓	MPJPE ↓	E_{vel} ↓	A_{accl} ↓	E_{key} ↓
V2V-PoseNet [15]	123.9	5.12	4.61	137.7	4.85	3.67	0.185
PoseReg w/o P2PSF	109.0	5.96	9.92	114.8	5.53	8.93	0.178
PoseReg w/ P2PSF	100.8	5.99	9.98	108.7	5.54	8.74	0.175
Ours w/o P2PSF	98.0	4.96	4.61	110.8	4.83	4.39	0.176
Ours	96.1	4.92	4.33	108.6	4.77	<u>4.16</u>	0.173

Table 2. Ablation study results.

Method	E_{key} ↓
(a) w/o Poisson	0.174
(b) w/o temp. blur	0.179
(c) w/o temp. shift	0.197
(d) w/o temp. resampling	0.185
(e) All noise types	0.173

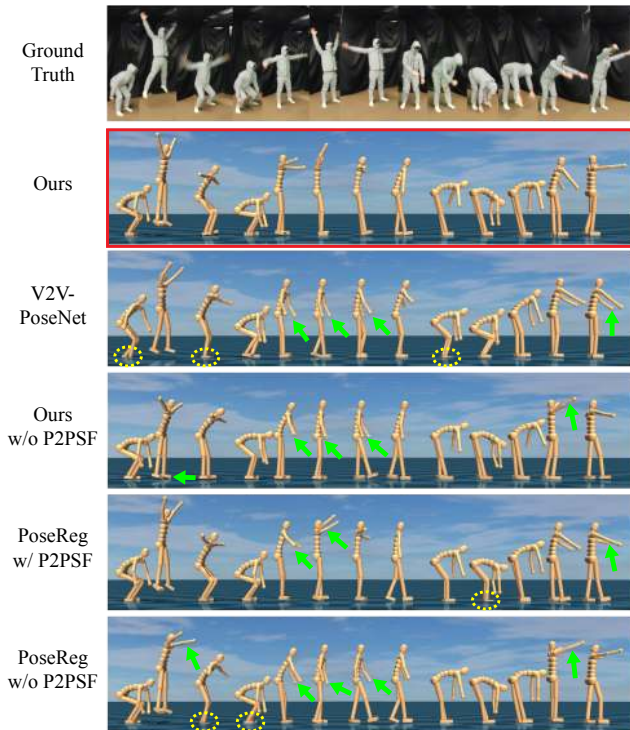


Figure 5. Qualitative results for real captured transient.

Real Captured Cross-Subject Evaluation. To show our approach’s ability to work with real-world data, we further test our method on the real captured transient image dataset described in Sec. 4.1, with our model trained with the same dataset with single subject experiment. As shown in Table 1 (Right), our approach produces more accurate poses than the baselines. We also present qualitative results in Fig. 5. As the figure shows, our method with P2PSF Net is better than the baselines in estimating human poses, including jumping (second from the left) and hand raising (sixth from the left, first from the right). Again, non-physics-based methods cause the humanoid’s foot to sink into the ground (highlighted by the yellow circle). The green arrows once again identify large discrepancies with the ground truth.

Ablative Analysis. As described in Sec. 3.2, our pseudo-transient images model the characteristics of real captured transient measurements. This ablation test investigates the effect of the four operations used to minimize the domain

gap between the pseudo and real transient images: (a) Poisson noise, (b) temporal blur, (c) temporal shift, and (d) temporal re-sampling. To investigate the importance of each of them, we train our model four times under the real captured cross-subject setting, and each time exclude one of the four operations. As indicated in Table 2, the temporal shifts are of critical importance for performance. Temporal re-sampling also improves accuracy. Removing Poisson noise and temporal blur independently does not severely affect performance. The combination of all four operations result in the best pose accuracy.

6. Conclusion

This work brings together a diverse set of sub-areas of computer vision, including state-of-the-art in computational imaging, physics-based vision, human pose estimation and physics-based deep reinforcement learning. As a result of this integration, we have shown for the first time that it is possible to take noisy real-world photon-level measurements of human motion and transform that information into a high-level understanding of human body dynamics aided by the power of data-driven machine learning.

While the primary technical focus of this work is to better understand how visual information should be represented and processed to enable 3D pose estimation from NLOS imaging, the technology described in this work also has some practical applications for next generation autonomous systems. In the context of autonomous driving, the ability to detect and track people outside of the line of sight of its sensors can be instrumental in informing planning algorithms and preventing accidents. In the context of domestic robots, the ability to see around walls could help robots make more informed decisions when entering a room or avoiding collisions. Though more research is necessary to lower the financial cost and computational complexity of the NLOS imaging system described in this work, we believe that this preliminary work shows the remarkable potential for higher-level reasoning using NLOS imaging in the real-world.

Acknowledgements. We thank Ioannis Gkioulekas for many helpful suggestions. M. Isogawa is supported by NTT Corporation. M. O’Toole is supported by the DARPA REVEAL program.

References

- [1] Fadel Adib, Chen-Yu Hsu, Hongzi Mao, Dina Katabi, and Frédo Durand. Capturing the human figure through a wall. *ACM Transactions on Graphics (TOG)*, 34(6):1–13, 2015. [2](#), [3](#)
- [2] Katherine L. Bouman, Vickie Ye, Adam B. Yedidia, Frédo Durand, Gregory W. Wornell, Antonio Torralba, and William T. Freeman. Turning corners into cameras: Principles and methods. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2289–2297, Oct 2017. [2](#), [3](#)
- [3] Susan Chan, Ryan E. Warburton, Genevieve Garipey, Jonathan Leach, and Daniele Faccio. Non-line-of-sight tracking of people at long range. *Optics Express*, 25(9):10109–10117, 2017. [2](#), [3](#)
- [4] Ioannis Gkioulekas, Anat Levin, Frédo Durand, and Todd Zickler. Micron-scale Light Transport Decomposition Using Interferometry. *ACM Transactions on Graphics (TOG)*, 34(4):37:1—37:14, 2015. [1](#)
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [6](#)
- [6] Felix Heide, Wolfgang Heidrich, and Matthias B Hullin. Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sens. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3222–3229, 2014. [2](#), [3](#)
- [7] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. [7](#)
- [8] Jonathan Klein, Christoph Peters, Martin Laurenzis, and Matthias Hullin. Tracking objects outside the line of sight using 2d intensity images. *Scientific Reports*, 6(32491):32491:1–32491:9, 2016. [2](#), [3](#)
- [9] Jonathan Klein, Christoph Peters, Martin Laurenzis, and Matthias Hullin. Non-line-of-sight mocap. In *ACM SIGGRAPH 2017 Emerging Technologies*, SIGGRAPH '17, pages 18:1–18:2. ACM, 2017. [2](#), [3](#)
- [10] Tianhong Li, Lijie Fan, Mingmin Zhao, Yingcheng Liu, and Dina Katabi. Making the invisible visible: Action recognition through walls and occlusions. In *arXiv preprint*, 2019. [2](#), [3](#)
- [11] David B. Lindell, Matthew O’Toole, and Gordon Wetzstein. Single-Photon 3D Imaging with Deep Sensor Fusion. *ACM Transactions on Graphics (TOG, SIGGRAPH)*, (4), 2018. [3](#)
- [12] David B. Lindell, Gordon Wetzstein, and Vladlen Koltun. Acoustic non-line-of-sight imaging. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6780–6789, 2019. [2](#), [3](#)
- [13] David B. Lindell, Gordon Wetzstein, and Matthew O’Toole. Wave-Based Non-Line-of-Sight Imaging using Fast f_k Migration. *ACM Transactions on Graphics (TOG)*, 38(4):116, 2019. [2](#), [3](#), [5](#), [6](#)
- [14] Tomohiro Maeda, Yiqin Wang, Ramesh Raskar, and Achuta Kadambi. Thermal non-line-of-sight imaging. In *IEEE International Conference on Computational Photography (ICCP)*, pages 1–11, 2019. [2](#), [3](#)
- [15] Gyeongsik Moon, Juyong Chang, and Kyoung Mu Lee. V2v-posenet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [6](#), [8](#)
- [16] Matthew O’Toole, Felix Heide, David B. Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. Reconstructing Transient Images from Single-Photon Sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2289–2297, 2017. [1](#), [4](#), [5](#)
- [17] Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338, 2018. [2](#), [3](#), [4](#)
- [18] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)*, 37(4):143:1–143:14, July 2018. [2](#), [3](#)
- [19] Xue Bin Peng, Glen Berseth, KangKang Yin, and Michiel van de Panne. Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Transactions on Graphics (TOG, SIGGRAPH 2017)*, 36(4), 2017. [3](#)
- [20] Xue Bin Peng, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. Sfv: Reinforcement learning of physical skills from videos. *ACM Transactions on Graphics (TOG)*, 37(6), Nov. 2018. [3](#)
- [21] Albert Redo-Sanchez, Barmak Heshmat, Alireza Aghasi, Salman Naqvi, Mingjie Zhang, Justin Romberg, and Ramesh Raskar. Terahertz time-gated spectral imaging for content extraction through layered structures. *Nature Communications*, 7:12665, 2016. [2](#), [3](#)
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, 2015. [6](#)
- [23] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. [6](#)
- [24] Guy Satat, Matthew Tancik, Otkrist Gupta, Barmak Heshmat, and Ramesh Raskar. Object classification through scattering media with deep learning on time resolved measurement. *Optics Express*, 25(15):17466–17479, Jul 2017. [3](#)
- [25] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. In *arXiv preprint*, volume abs/1707.06347, 2017. [4](#), [7](#)
- [26] Jie Tan, Karen Liu, and Greg Turk. Stable proportional-derivative controllers. *IEEE Computer Graphics and Applications*, 31(4):34–44, 2011. [4](#), [6](#)
- [27] Matthew Tancik, Guy Satat, and Ramesh Raskar. Flash photography for data-driven hidden scene recovery. In *arXiv preprint*, 2018. [2](#), [3](#)
- [28] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS)*, pages 5026–5033. IEEE, 2012. [2](#), [6](#)

- [29] Chia-Yin Tsai, Aswin C. Sankaranarayanan, and Ioannis Gkioulekas. Beyond volumetric albedo – a surface optimization framework for non-line-of-sight imaging. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#), [3](#)
- [30] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounqi G Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3:745, 2012. [1](#), [2](#), [3](#)
- [31] Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mounqi G Bawendi, Diego Gutierrez, and Ramesh Raskar. Femto-Photography: Capturing and Visualizing the Propagation of Light. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 32(4), 2013. [1](#)
- [32] Shumian Xin, Sotiris Nousias, Kiriakos N Kutulakos, Aswin C Sankaranarayanan, Srinivasa G Narasimhan, and Ioannis Gkioulekas. A Theory of Fermat Paths for Non-Line-of-Sight Shape Reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#), [3](#)
- [33] Yuliang Xiu, Jiefeng Li, Haoyu Wang, Yinghong Fang, and Cewu Lu. Pose Flow: Efficient online pose tracking. In *British Machine Vision Conference (BMVC)*, 2018. [6](#)
- [34] Ye Yuan and Kris Kitani. 3D Ego-Pose Estimation via Imitation Learning. In *European Conference on Computer Vision (ECCV)*, pages 763–778, 2018. [3](#)
- [35] Ye Yuan and Kris Kitani. Ego-pose estimation and forecasting as real-time pd control. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. [2](#), [3](#), [4](#)
- [36] Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, and Dina Katabi. Through-Wall Human Pose Estimation Using Radio Signals. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7356–7365, 2018. [2](#), [3](#)