

Estimating 3D Motion and Forces of Person-Object Interactions from Monocular Video

Zongmian Li^{1,2}

Ivan Laptev^{1,2}

¹DI ENS PSL

²Inria

Jiri Sedlar³

Nicolas Mansard⁴

³CIIRC, CTU in Prague

Justin Carpentier^{1,2}

Josef Sivic^{1,2,3}

⁴LAAS-CNRS

Abstract

In this paper, we introduce a method to automatically reconstruct the 3D motion of a person interacting with an object from a single RGB video. Our method estimates the 3D poses of the person and the object, contact positions, and forces and torques actuated by the human limbs. The main contributions of this work are three-fold. First, we introduce an approach to jointly estimate the motion and the actuation forces of the person on the manipulated object by modeling contacts and the dynamics of their interactions. This is cast as a large-scale trajectory optimization problem. Second, we develop a method to automatically recognize from the input video the position and timing of contacts between the person and the object or the ground, thereby significantly simplifying the complexity of the optimization. Third, we validate our approach on a recent MoCap dataset with ground truth contact forces and demonstrate its performance on a new dataset of Internet videos showing people manipulating a variety of tools in unconstrained environments.

1. Introduction

People can easily learn how to break concrete with a sledgehammer or cut hay using a scythe by observing other people performing such tasks in instructional videos, for example. They can also easily perform the same task in a different context. This involves advanced visual intelligence capabilities such as recognizing and interpreting complex person-object interactions that achieve a specific goal. Understanding such complex interactions is a key to building autonomous machines that learn how to interact with the physical world by observing people.

*Please see our project webpage [2] for trained models, data and code.

¹Département d'informatique de l'ENS, École normale supérieure, CNRS, PSL Research University, 75005 Paris, France.

³Czech Institute of Informatics, Robotics and Cybernetics at the Czech Technical University in Prague, Czech Republic.

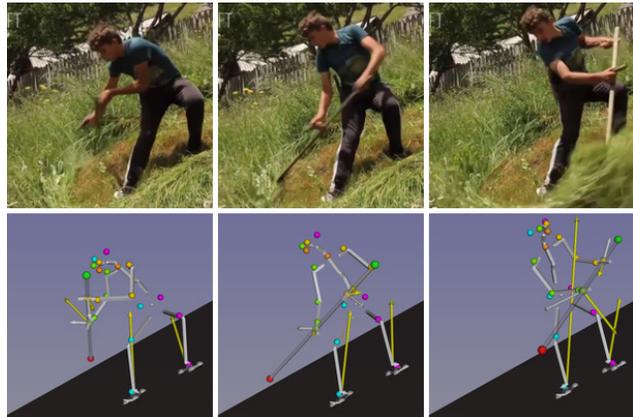


Figure 1: Our method automatically reconstructs in 3D the object manipulation action captured by a single RGB video. **Top:** Frames from the input video. **Bottom:** the output human and object 3D motion including the recovered contact forces (yellow arrows) and moments (white arrows).

This work makes a step in this direction and describes a method to estimate the 3D motion and actuation forces of a person manipulating an object given a single unconstrained video as input, as shown in Figure 1. This is an extremely challenging task. First, there are inherent ambiguities in the 2D-to-3D mapping from a single view: multiple 3D human poses correspond to the same 2D input. Second, human-object interactions often involve contacts, resulting in discontinuities in the motion of the object and the human body part in contact. For example, one must place a hand on the hammer handle before picking the hammer up. The contact motion strongly depends on the physical quantities such as the mass of the object and the contact forces exerted by the hand, which renders modeling of contacts a very difficult task. Finally, the tools we consider in this work, such as hammer, scythe, or spade, are particularly difficult to recognize due to their thin structure, lack of texture, and frequent occlusions by hands and other human parts.

To address these challenges, we propose a method to jointly estimate the 3D trajectory of both the person and the object by visually recognizing contacts in the video

and modeling the dynamics of the interactions. We focus on rigid stick-like hand tools (e.g. hammer, barbell, spade, scythe) with no articulation and approximate them as 3D line segments. Our key idea is that, when a human joint is in contact with an object, the object can be integrated as a constraint on the movement of the human limb. For example, the hammer in Figure 1 provides a constraint on the relative depth between the person’s two hands. Conversely, 3D positions of the hands in contact with the hammer provide a constraint on the hammer’s depth and 3D rotation. To deal with contact forces, we integrate physics in the estimation by modeling dynamics of the person and the object. Inspired by recent progress in humanoid locomotion research [16], we formulate person-object trajectory estimation as an optimal control problem given the contact state of each human joint. We show that contact states can be automatically recognized from the input video using a deep neural network.

2. Related work

Here we review the key areas of related work in both computer vision and robotics literature.

Single-view 3D pose estimation aims to recover the 3D joint configuration of the person from the input image. Recent human 3D pose estimators either attempt to build a *direct mapping* from image pixels to the 3D joints of the human body or break down the task into *two stages*: estimating pixel coordinates of the joints in the input image and then lifting the 2D skeleton to 3D. Existing direct approaches either rely on generative models to search the state space for a plausible 3D skeleton that aligns with the image evidence [59, 26, 25] or, more recently, extract deep features from images and learn a discriminative regressor from the 2D image to the 3D pose [37, 49, 53, 62]. Building on the recent progress in 2D human pose estimation [51, 50, 34, 14], two-stage methods have been shown to be very effective [5, 71, 9, 19] and achieve state-of-the-art results [47] on 3D human pose benchmarks [35]. To deal with depth ambiguities, these estimators rely on good pose priors, which are either hand-crafted or learnt from large-scale MoCap data [71, 9, 37]. However, unlike our work, these methods do not consider explicit models for 3D person-object interactions with contacts.

Understanding human-object interactions involves both recognition of actions and modeling of interactions. In action recognition, most existing approaches that model human-object interactions do not consider 3D, instead model interactions and contacts in the 2D image space [29, 20, 70, 55]. Recent works in scene understanding [36, 24] consider interactions in 3D but have focused on static scene elements rather than manipulated objects as we do in this work. Tracking 3D poses of people interacting with the environment has been demonstrated for bipedal walk-

ing [12, 13] or in sports scenarios [66]. However, these works do not consider interactions with objects. Furthermore, [66] requires manual annotation of the input video.

There is also related work on modeling person-object interactions in robotics [60] and computer animation [10]. Similarly to people, humanoid robots interact with the environment by creating and breaking contacts [32], for example, during walking. Typically, generating artificial motion is formulated as an optimal control problem, transcribed into a high-dimensional numerical optimization problem, seeking to minimize an objective function under contact and feasibility constraints [21, 58]. A known difficulty is handling the non-smoothness of the resulting optimization problem introduced by the creation and breaking of contacts [67]. Due to this difficulty, the sequence of contacts is often computed separately and not treated as a decision variable in the optimizer [38, 64]. Recent work has shown that it may be possible to decide both the continuous movement and the contact sequence together, either by implicitly formulating the contact constraints [54] or by using invariances to smooth the resulting optimization problem [48, 68].

In this paper, we take advantage of rigid-body models introduced in robotics and formulate the problem of estimating 3D person-object interactions from monocular video as an optimal control problem under contact constraints. We overcome the difficulty of contact irregularity by first identifying the contact states from the visual input, and then localizing the contact points in 3D via our trajectory estimator. This allows us to treat multi-contact sequences (like walking) without manually annotating the contact phases.

Object 3D pose estimation methods often require depth or RGB-D data as input [61, 22, 33], which is restrictive since depth information is not always available (e.g. for outdoor scenes or specular objects), as is the case of our instructional videos. Recent work has also attempted to recover object pose from RGB input only [11, 56, 69, 39, 52, 28, 57]. However, we found that the performance of these methods is limited for the stick-like objects we consider in this work. Instead, we recover the 3D pose of the object via localizing and segmenting the object in 2D, and then jointly recovering the 3D trajectory of both the human limbs and the object. As a result, both the object and the human pose help each other to improve their joint 3D trajectory by leveraging the contact constraints.

Instructional videos. Our work is also related to recent efforts in learning from Internet instructional videos [46, 6, 6] that aim to segment input videos into clips containing consistent actions. In contrast, we focus on extracting a detailed representation of the object manipulation in the form of a 3D person-object trajectory with contacts and underlying manipulation forces.

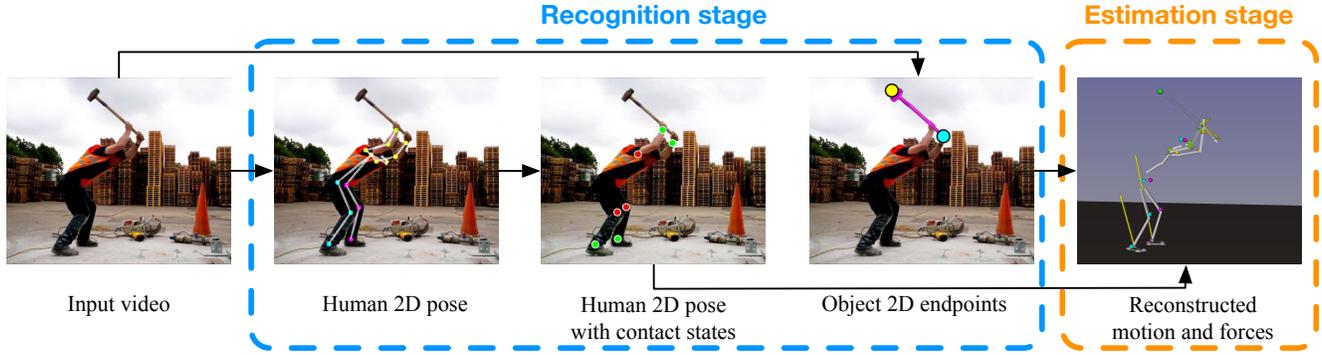


Figure 2: Overview of the proposed method. In the recognition stage, the system estimates from the input video the person’s 2D joints, the hammer’s 2D endpoints and the contact states of the individual joints. The human joints and the object endpoints are visualized as colored dots in the image. Human joints recognized as in contact are shown in green, joints not in contact in red. In the estimation stage, these image measurements are fused in a trajectory estimator to recover the human and object 3D motion together with the contact positions and forces.

3. Approach overview

We are given a video clip of a person manipulating an object or in another way interacting with the scene. Our approach, illustrated in Figure 2, receives as input a sequence of frames and automatically outputs the 3D trajectories of the human body, the manipulated object, and the ground plane. At the same time, it localizes the contact points and recovers the contact forces that actuate the motion of the person and the object. Our approach proceeds along two stages. In the first, *recognition stage*, we extract 2D measurements from the input video. These consist of 2D locations of human joints, 2D locations of a small number of predefined object endpoints, and contact states of selected joints over the course of the video. In the second, *estimation stage*, these image measurements are then fused in order to estimate the 3D motion, 3D contacts, and the controlling forces of both the person and the object. The person and object trajectories, contact positions, and contact forces are constrained jointly by our carefully designed contact motion model, force model, and dynamics equations. Finally, the reconstructed object manipulation sequence can be applied to control a humanoid robot via behavior cloning.

In the following, we start in Section 4 by describing the estimation stage giving details of the formulation as an optimal control problem. Then, in Section 5 we give details of the recognition stage including 2D human pose estimation, contact recognition, and object 2D endpoint estimation. Finally, we describe results in Section 6.

4. Estimating person-object trajectory under contact and dynamics constraints

We assume that we are provided with a video clip of duration T depicting a human subject manipulating an object. We encode the 3D poses of the human and the object, including joint translations and rotations, in the configuration

vectors q^h and q^o , for the human and the object respectively. We define a constant set of K contact points between the human body and the object (or the ground plane). Each contact point corresponds to a human joint, and is activated whenever that human joint is recognized as in contact. At each contact point, we define a contact force f_k , whose value is non-zero whenever the contact point k is active. The state of the complete dynamical system is then obtained by concatenating the human and the object joint configurations q and velocities \dot{q} as $x := (q^h, q^o, \dot{q}^h, \dot{q}^o)$. Let τ_m^h be the joint torque vector describing the actuation by human muscles. This is a $n_q - 6$ dimensional vector where n_q is the dimension of the human body configuration vector. We define the control variable u as the combination of the joint torque vector together with contact forces at the K contact joints, $u := (\tau_m^h, f_k, k = 1, \dots, K)$. To deal with sliding contacts, we further define a contact state c that consists of the relative positions of all the contact points with respect to the object (or ground) in the 3D space.

Our goal is two-fold. We wish to (i) estimate smooth and consistent human-object and contact trajectories \underline{x} and \underline{c} , while (ii) recovering the control \underline{u} which gives rise to the observed motion¹. This is achieved by jointly optimizing the 3D trajectory \underline{x} , contacts \underline{c} , and control \underline{u} given the measurements (2D positions of human joints and object endpoints together with contact states of human joints) obtained from the input video. The intuition is that the human and the object’s 3D poses should match their respective projections in the image while their 3D motion is linked together by the recognized contact points and the corresponding contact forces. In detail, we formulate person-object interaction estimation as an optimal control problem with contact and

¹In this paper, trajectories are denoted as underlined variables, e.g. \underline{x} , \underline{u} or \underline{c} .

dynamics constraints:

$$\underset{\underline{x}, \underline{u}, \underline{c}}{\text{minimize}} \quad \sum_{e \in \{\text{h}, \text{o}\}} \int_0^T l^e(x, u, c) dt, \quad (1)$$

$$\text{subject to} \quad \kappa(x, c) = 0 \quad (\text{contact motion model}), \quad (2)$$

$$\dot{x} = f(x, c, u) \quad (\text{full-body dynamics}), \quad (3)$$

$$u \in \mathcal{U} \quad (\text{force model}), \quad (4)$$

where e denotes either ‘h’ (human) or ‘o’ (object), and the constraints (2)-(4) must hold for all $t \in [0, T]$. The loss function l^e is a weighted sum of multiple costs capturing (i) the data term measuring the discrepancy between the observed and re-projected 2D joint and object endpoint positions, (ii) the prior on the estimated 3D poses, (iii) the physical plausibility of the motion and (iv) the smoothness of the trajectory. Next, we in turn describe these cost terms as well as the insights leading to their design choices. For simplicity, we ignore the superscript e when introducing a cost term that exists for both the human l^h and the object l^o component of the loss. We describe the individual terms using continuous time notation as used in the overall problem formulation (1). A discrete version of the problem as well as the optimization and implementation details are relegated to Section 4.5.

4.1. Data term: 2D re-projection error

We wish to minimize the re-projection error of the estimated 3D human joints and 3D object endpoints with respect to the 2D measurements obtained in each video frame. In detail, let $j = 1, \dots, N$ be human joints or object endpoints and p_j^{2D} their 2D position observed in the image. We aim to minimize the following *data term*

$$l_{\text{data}} = \sum_j \rho(p_j^{2D} - P_{\text{cam}}(p_j(q))), \quad (5)$$

where P_{cam} is the camera projection matrix and p_j the 3D position of joint or object endpoint j induced by the person-object configuration vector q . To deal with outliers, we use the robust Huber loss, denoted by ρ .

4.2. Prior on 3D human poses

A single 2D skeleton can be a projection of multiple 3D poses, many of which are unnatural or impossible exceeding the human joint limits. To resolve this, we incorporate into the human loss function l^h a pose prior similar to [9]. The pose prior is obtained by fitting the SMPL human model [43] to the CMU MoCap data [1] using MoSh [44] and fitting a Gaussian Mixture Model (GMM) to the resulting SMPL 3D poses. We map our human configuration vector q^h to a SMPL pose vector θ and compute the likelihood under the pre-trained GMM

$$l_{\text{pose}}^h = -\log(p(q^h; \text{GMM})). \quad (6)$$

During optimization, l_{pose}^h is minimized in order to favor more plausible human poses against rare or impossible ones.

4.3. Physical plausibility of the motion

Human-object interactions involve contacts coupled with interaction forces, which are not included in the data-driven cost terms (5) and (6). Modeling contacts and physics is thus important to reconstruct object manipulation actions from the input video. Next, we outline models for describing the motion of the contacts and the forces at the contact points. Finally, the contact motions and forces, together with the system state \underline{x} , are linked by the laws of mechanics via the dynamics equations, which constrain the estimated person-object interaction. This full body dynamics constraint is detailed at the end of this subsection.

Contact motions. In the recognition stage, our contact recognizer predicts, given a human joint (for example, left hand, denoted by j), a sequence of contact states $\delta_j : t \rightarrow \{1, 0\}$. Similarly to [16], we call a *contact phase* any time segment in which j is in contact, i.e., $\delta_j = 1$. Our key idea is that the 3D distance between human joint j and the active contact point on the object (denoted by k) should remain zero during a contact phase:

$$\|p_j^h(q^h) - p_k^c(x, c)\| = 0 \quad (\text{point contact}), \quad (7)$$

where p_j^h and p_k^c are the 3D positions of joint j and object contact point k , respectively. Note that position of the object contact point $p_k^c(x, c)$ depends on the state vector x describing the human-object configuration and the relative position c of the contact along the object. The position of contact p_k^c is subject to a feasible range denoted by \mathcal{C} . For stick-like objects such as hammer, \mathcal{C} is approximately the 3D line segment representing the handle. For the ground, the feasible range \mathcal{C} is a 3D plane. In practice, we implement $p_k^c \in \mathcal{C}$ by putting a constraint on the trajectory of relative contact positions \underline{c} .

Equation (7) applies to most common cases where the contact area can be modeled as a point. Examples include the hand-handle contact and the knee-ground contact. To model the *planar contact* between the human sole and ground, we approximate each sole surface as a planar polygon with four vertices, and apply the point contact model at each vertex. In our human model, each sole is attached to its parent ankle joint, and therefore the four vertex contact points of the sole are active when $\delta_{\text{ankle}} = 1$.

The resulting overall contact motion function κ in problem (1) is obtained by unifying the point and the planar contact models:

$$\kappa(x, c) = \sum_j \sum_{k \in \phi(j)} \delta_j \left\| T^{(kj)}(p_j^h(q^h)) - p_k^c(x, c) \right\|, \quad (8)$$

where the external sum is over all human joints. The internal sum is over the set of active object contact points mapped to their corresponding human joint j by mapping $\phi(j)$. The mapping $T^{(kj)}$ translates the position of an ankle joint j to its corresponding k -th sole vertex; it is an identity mapping for non-ankle joints.

Contact forces. During a contact phase of the human joint j , the environment exerts a contact force f_k on each of the active contact points in $\phi(j)$. f_k is always expressed in contact point k 's local coordinate frame. We distinguish two types of contact forces: (i) 6D spatial forces exerted by objects and (ii) 3D linear forces due to ground friction. In the case of object contact, f_k is an unconstrained 6D spatial force with 3D linear force and 3D moment. In the case of ground friction, f_k is constrained to lie inside a 3D friction cone \mathcal{K}^3 (also known as the quadratic Lorentz ‘‘ice-cream’’ cone [16]) characterized by a positive friction coefficient μ . In practice, we approximate \mathcal{K}^3 by a 3D pyramid spanned by a basis of $N = 4$ generators, which allows us to represent f_k as the convex combination $f_k = \sum_{n=1}^N \lambda_{kn} g_n^{(3)}$, where $\lambda_{kn} \geq 0$ and $g_n^{(3)}$ with $n = 1, 2, 3, 4$ are the 3D generators of the contact force. We sum the contact forces induced by the four sole-ground contact points and express a unified contact force in the ankle's frame:

$$f_j = \sum_{k=1}^4 \begin{pmatrix} f_k \\ p_k \times f_k \end{pmatrix} = \sum_{k=1}^4 \sum_{n=1}^N \lambda_{jkn} g_{kn}^{(6)}, \quad (9)$$

where p_k is the position of contact point k expressed in joint j 's (left/right ankle) frame, \times is the cross product operator, $\lambda_{jkn} \geq 0$, and $g_{kn}^{(6)}$ are the 6D generators of f_j . Please see the appendix in the extended version of this paper available at [40] for additional details including the expressions of $g_n^{(3)}$ and $g_{kn}^{(6)}$.

Full body dynamics. The full-body movement of the person and the manipulated object is described by the Lagrange dynamics equation:

$$M(q)\ddot{q} + b(q, \dot{q}) = g(q) + \tau, \quad (10)$$

where M is the generalized mass matrix, b covers the centrifugal and Coriolis effects, g is the generalized gravity vector and τ represents the joint torque contributions. \dot{q} and \ddot{q} are the joint velocities and joint accelerations, respectively. Note that (10) is a unified equation which applies to both human and object dynamics, hence we drop the superscript e here. Only the expression of the joint torque τ differs between the human and the object and we give the two expressions next.

For human, it is the sum of two contributions: the first one corresponds to the internal joint torques (exerted by the

muscles for instance) and the second one comes from the contact forces:

$$\tau^h = \begin{pmatrix} \mathbf{0}_6 \\ \tau_m^h \end{pmatrix} + \sum_{k=1}^K (J_k^h)^T f_k, \quad (11)$$

where τ_m^h is the human joint torque exerted by muscles, f_k is the contact force at contact point k and J_k^h is the Jacobian mapping human joint velocities \dot{q}^h to the Cartesian velocity of contact point k expressed in k 's local frame. Let n_q^h denote the dimension of q^h , \dot{q}^h and \ddot{q}^h , then τ_m^h and J_k^h are of dimension $n_q^h - 6$ and $3 \times n_q^h$, respectively. We model the human body and the object as free-floating base systems. In the case of human body, the six first entries in the configuration vector q correspond to the 6D pose of the free-floating base (translation + orientation), which is not actuated by any internal actuators such as human muscles. This constraint is taken into consideration by adding the zeros in Eq. (11).

In the case of the manipulated object, there is no actuation other than the contact forces exerted by the human. Therefore, the object torque is expressed as

$$\tau^o = - \sum_{\text{object contact } k} (J_k^o)^T f_k, \quad (12)$$

where the sum is over the object contact points, f_k is the contact force, and J_k^o denotes the object Jacobian, which maps from the object joint velocities \dot{q}^o to the Cartesian velocity of the object contact point k expressed in k 's local frame. J_k^o is a $3 \times n_q^o$ matrix where n_q^o is the dimension of object configuration vectors q^o , \dot{q}^o and \ddot{q}^o .

We concatenate the dynamics equations of both human and object to form the overall dynamics in Eq. (3) in problem (1), and include a *muscle torque* term $l_{\text{torque}}^h = \|\tau_m^h\|^2$ in the overall cost. Minimizing the muscle torque acts as a regularization over the energy consumption of the human body.

4.4. Enforcing the trajectory smoothness

Regularizing human and object motion. Taking advantage of the temporal continuity of video, we minimize the sum of squared 3D joint velocities and accelerations to improve the smoothness of the person and object motion and to remove incorrect 2D poses. We include the following *motion smoothing* term to the human and object loss in (1):

$$l_{\text{smooth}} = \sum_j \left(\|\nu_j(q, \dot{q})\|^2 + \|\alpha_j(q, \dot{q}, \ddot{q})\|^2 \right), \quad (13)$$

where ν_j and α_j are the spatial velocity and the spatial acceleration² of joint j , respectively. In the case of object, j represents an endpoint on the object. By minimizing l_{smooth} , both the linear and angular movements of each joint/endpoint are smoothed simultaneously.

²Spatial velocities (accelerations) are minimal and unified representations of linear and angular velocities (accelerations) of a rigid body [23]. They are of dimension 6.

Regularizing contact motion and forces. In addition to regularizing the motion of the joints, we also regularize the contact states and control by minimizing the velocity of the contact points and the temporal variation of the contact force. This is implemented by including the following *contact smoothing* term in the cost function in problem (1):

$$l_{\text{smooth}}^c = \sum_j \sum_{k \in \phi(j)} \delta_j \left(\omega_k \|\dot{c}_k\|^2 + \gamma_k \|\dot{f}_k\|^2 \right) dt, \quad (14)$$

where \dot{c}_k and \dot{f}_k represent respectively the temporal variation of the position and the contact force at contact point k . ω_k and γ_k are scalar weights of the regularization terms \dot{c}_k and \dot{f}_k . Note that some contact points, for example the four contact points of the human sole during the sole-ground contact, should remain fixed with respect to the object or the ground during the contact phase. To tackle this, we adjust ω_k to prevent contact point k from sliding while being in contact.

4.5. Optimization

Conversion to a numerical optimization problem. We convert the continuous problem (1) into a discrete nonlinear optimization problem using the collocation approach [8]. All trajectories are discretized and constraints (2), (3), (4) are only enforced on the “collocation” nodes of a time grid matching the discrete sequence of video frames. The optimization variables are the sequence of human and object poses $[x_0 \dots x_T]$, torque and force controls $[u_1 \dots u_T]$, contact locations $[c_0 \dots c_T]$, and the scene parameters (ground plane and camera matrix). The resulting problem is nonlinear, constrained and sparse (due to the sequential structure of trajectory optimization). We rely on the Ceres solver [4], which is dedicated to solving sparse estimation problems (e.g. bundle adjustment [65]), and on the Pinocchio software [17, 18] for the efficient computation of kinematic and dynamic quantities and their derivatives [15]. Additional details are given in the appendix available in the extended version of this paper [40].

Initialization. Correctly initializing the solver is key to escape from poor local minima. We warm-start the optimization by inferring the initial configuration vector q_k at each frame using the human body estimator HMR [37] that estimates the 3D joint angles from a single RGB image.

5. Extracting 2D measurements from video

In this section, we describe how 2D measurements are extracted from the input video frames during the first, recognition stage of our system. In particular, we extract the 2D human joint positions, the 2D object endpoint positions and the contact states of human joints.

Estimating 2D positions of human joints. We use the state-of-the-art Openpose [14] human 2D pose estimator,

which achieved excellent performance on the MPII Multi-Person benchmark [7]. Taking a pretrained Openpose model, we do a forward pass on the input video in a frame-by-frame manner to obtain an estimate of the 2D trajectory of human joints, $p_j^{\text{h},2\text{D}}$.

Recognizing contacts. We wish to recognize and localize contact points between the person and the manipulated object or the ground. This is a challenging task due to the large appearance variation of the contact events in the video. However, we demonstrate here that a good performance can be achieved by training a contact recognition CNN module from manually annotated contact data that combine both still images and videos harvested from the Internet. In detail, the contact recognizer operates on the 2D human joints predicted by Openpose. Given 2D joints at video frame i , we crop fixed-size image patches around a set of joints of interest, which may be in contact with an object or ground. Based on the type of human joint, we feed each image patch to the corresponding CNN to predict whether the joint appearing in the patch is in contact or not. The output of the contact recognizer is a sequence δ_{ji} encoding the contact states of human joint j at video frame i , i.e. $\delta_{ji} = 1$ if joint j is in contact at frame i and zero otherwise. Note that δ_{ji} is the discretized version of the contact state trajectory δ_j presented in Sec. 4.

Our contact recognition CNNs are built by replacing the last layer of an ImageNet pre-trained Resnet model [31] with a fully connected layer that has a binary output. We have trained separate models for five types of joints: hands, knees, foot soles, toes, and neck. To construct the training data, we collect still images of people manipulating tools using Google image search. We also collect short video clips of people manipulating tools from Youtube in order to also have non-contact examples. We run Openpose pose estimator on this data, crop patches around the 2D joints, and annotate the resulting dataset with contact states.

Estimating 2D object pose. The objective is to estimate the 2D position of the manipulated object in each video frame. To achieve this, we build on instance segmentation obtained by Mask R-CNN [30]. We train the network on shapes of object models from different viewpoints and apply the trained network on the test videos. The output masks and bounding boxes are used to estimate object endpoints in each frame. The resulting 2D endpoints are used as input to the trajectory optimizer. Details are given next.

In the case of barbell, hammer and scythe, we created a single 3D model for each tool, roughly approximating the shapes of the instances in the videos, and rendered it from multiple viewpoints using a perspective camera. For spade, we annotated 2D masks of various instances of the tool in thirteen different still images. The shapes of the rendered 3D models or 2D masks are used to train Mask R-CNN for instance segmentation of each tool. The training set is aug-

mented by 2D geometric transformations (translation, rotation, scale) to handle the changes in shapes of tool instances in the videos. In addition, domain randomization [42, 63] is applied to handle the variance of instances and changes in appearance in the videos caused by illumination: the geometrically transformed shape is filled with pixels from a random image (foreground) and pasted on another random image (background). We utilized random images from the MS COCO dataset [41] for this purpose. We use a Mask R-CNN (implementation [3]) model pre-trained on the MS COCO dataset and re-train the head layers for each tool.

At test time, masks and bounding boxes obtained by the re-trained Mask R-CNN are used to estimate the coordinates of tool endpoints. Proximity to coordinates of estimated wrist joints is used to select the mask and bounding box in case multiple candidates are available in the frame. To estimate the main axis of the object, a line is fitted through the output binary mask. The endpoints are calculated as the intersection of the fitted line and boundaries of the bounding box. Using the combination of the output mask and the bounding box compensates for errors in the segmentation mask caused by occlusions. The relative orientation of the tool (i.e. the head vs. the handle of the tool) is determined by spatial location of endpoints in the video frames as well as by their proximity to the estimated wrist joints.

6. Experiments

In this section we present quantitative and qualitative evaluation of the reconstructed 3D person-object interactions. Since we recover not only human poses but also object poses and contact forces, evaluating our results is difficult due to the lack of ground truth forces and 3D object poses in standard 3D pose benchmarks such as [35]. Consequently, we evaluate our motion and force estimation quantitatively on a recent Biomechanics video/MoCap dataset capturing challenging dynamic parkour motions [45]. In addition, we report joint errors and show qualitative results on our newly collected dataset of videos depicting handtool manipulation actions.

6.1. Parkour dataset

This dataset contains videos capturing human subjects performing four typical parkour actions: two-hand jump, moving-up, pull-up and a single-hand hop. These are highly dynamic motions with rich contact interactions with the environment. The ground truth 3D motion and contact forces are captured with a Vicon motion capture system and several force plates. The 3D motion and forces are reconstructed with frame rates of 400Hz and 2200Hz, respectively, whereas the RGB videos are captured in a relatively lower rate of 25Hz, making this dataset a challenge for computer vision algorithms due to motion blur.

Method	Jump	Move-up	Pull-up	Hop	Avg
SMPLify [9]	121.75	147.41	120.48	169.36	139.69
HMR [37]	111.36	140.16	132.44	149.64	135.65
Ours	98.42	125.21	119.92	138.45	122.11

Table 1: Mean per joint position error (in mm) of the recovered 3D motion for each action on the Parkour dataset.

	L. Sole	R. Sole	L. Hand	R. Hand
Force (N)	144.23	138.21	107.91	113.42
Moment (N-m)	23.71	22.32	131.13	134.21

Table 2: Estimation errors of the contact forces exerted on soles and hands on the Parkour dataset.

Evaluation set-up. We evaluate both the estimated human 3D motion and the contact forces. For evaluating the accuracy of the recovered 3D human poses, we follow the common approach of computing the mean per joint position error (MPJPE) of the estimated 3D pose with respect to the ground truth after rigid alignment [27]. We evaluate contact forces without any alignment: we express both the estimated and the ground truth 6D forces at the position of the contact aligned with the world coordinate frame provided in the dataset. We split the 6D forces into linear and moment components and report the average Euclidean distance of the linear force and the moment with respect to the ground truth.

Results. We report joint errors for different actions in Table 1 and compare the results with the HMR 3D human pose estimator [37], which is used to warm-start our method. To make it a fair comparison, we use the same Openpose 2D joints as input. In addition, we evaluate the recent SMPLify [9] 3D pose estimation method. Our method outperforms both baselines by more than 10mm on average on this challenging data. Finally, Table 2 summarizes the force estimation results. To estimate the forces we assume a generic human physical model of mass 74.6 kg for all the subjects. Despite the systematic error due to the generic human mass assumption, the results in Table 2 validate the quality of our force estimation at the soles and the hands during walking and jumping. We observe higher errors of the estimated moments at hands, which we believe is due to the challenging nature of the Parkour sequences where the entire person’s body is often supported by hands. In this case, the hand may exert significant force and torque to support the body, and a minor shift in the force direction may lead to significant errors.

6.2. Handtool dataset

In addition to the Parkour data captured in a controlled set-up, we would like to demonstrate generalization of our approach to the “in the wild” Internet instructional videos. For this purpose, we have collected a new dataset of object manipulation videos, which we refer to as the Handtool dataset. The dataset contains videos of people manipulating four types of tools: *barbell*, *hammer*, *scythe*, and *spade*. For each type of tool, we chose among the top videos re-

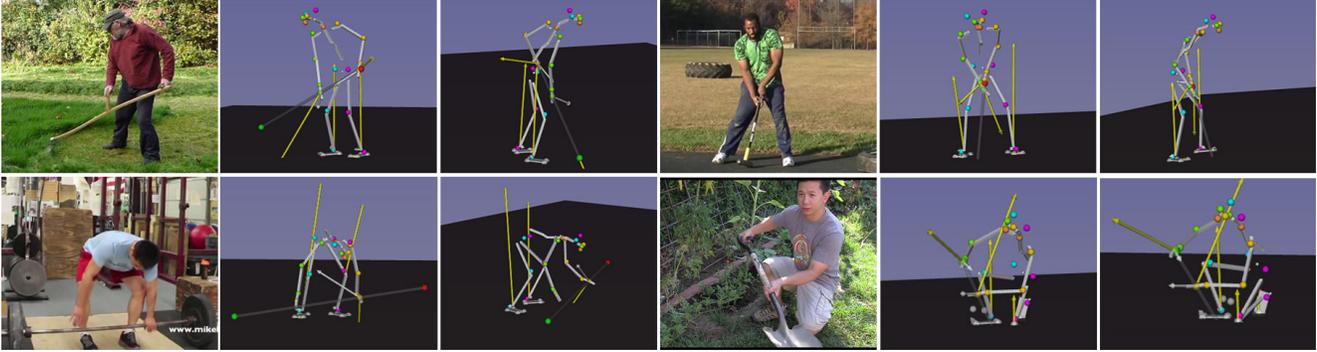


Figure 3: Example qualitative results on the Handtool dataset. Each example shows the input frame (left) and two different views of the output 3D pose of the person and the object (middle, right). The yellow and the white arrows in the output show the contact forces and moments, respectively. Note how the proposed approach recovers from these challenging unconstrained videos the 3D configuration of the person-object interaction together with the contact forces and moments. **For additional video results please see the project webpage [2].**

turned by YouTube five videos covering a range of actions. We then cropped short clips from each video showing the whole human body and the tool.

Evaluation of 3D human poses. For each video in the Handtool dataset, we have annotated the 3D positions of the person’s left and right shoulders, elbows, wrist, hips, knees, and ankles, for the first, the middle, and the last frame. We evaluate the accuracy of the recovered 3D human poses by computing their MPJPE after rigid alignment. Quantitative evaluation of the recovered 3D poses is shown table 3. On average, our method outperforms the strong HMR [37] and SMPLify [9] baselines. However, the differences between the methods are reaching the limits of the accuracy of the manually provided 3D human pose annotations on this dataset. Videos available at project website [2] demonstrate that our model produces smooth 3D motion, respecting person-object contacts and capturing the interaction of the person with the tool. This is not the case for the HMR [37] and SMPLify [9] baselines that are applied to individual frames and do not model the interaction between the person and the tool. Example results for our method are shown in Figure 3. For additional results including examples of the main failure modes please see the appendix in the extended version of this paper available at [40].

Evaluation of 2D object poses. To evaluate the quality of estimated object poses, we manually annotated 2D object endpoints in every 5th frame of each video in the Handtool dataset and calculated the 2D Euclidean distance (in pixels) between each manually annotated endpoint and its estimated 2D location provided by our method. The 2D location is obtained by projecting the estimated 3D tool position back to the image plane. We compare our results to the output of the Mask R-CNN instance segmentation baseline [30] (which provides initialization for our person-object interaction model). In Table 4 we report for both methods the percentage of endpoints for which the estimated endpoint location lies within 25, 50, and 100 pixels from the anno-

Method	Barbell	Spade	Hammer	Scythe	Avg
SMPLify [9]	130.69	135.03	93.43	112.93	118.02
HMR [37]	105.04	97.18	96.34	115.42	103.49
Ours	104.23	95.21	95.87	114.22	102.02

Table 3: Mean per joint position error (in mm) of the recovered 3D human poses for each tool type on the Handtool dataset.

Method	Barbell	Spade	Hammer	Scythe
Mask R-CNN [30]	33/42/54	54/79/93	35/44/45	63/72/76
Ours	38/71/98	57/86/99	61/91/99	69/88/98

Table 4: The percentage of endpoints for which the estimated 2D location lies within 25/50/100 pixels (in 600×400 pixel image) from the manually annotated ground truth location.

tated ground truth endpoint location. The results demonstrate that our approach provides more accurate and stable object endpoint locations compared to the Mask R-CNN baseline thanks to modeling the interaction between the object and the person.

7. Conclusion

We have developed a visual recognition system that takes as input video frames together with a simple object model, and outputs a 3D motion of the person and the object including contact forces and torques actuated by the human limbs. We have validated our approach on a recent MoCap dataset with ground truth contact forces. Finally, we have collected a new dataset of unconstrained instructional videos depicting people manipulating different objects and have demonstrated benefits of our approach on this data. Our work opens up the possibility of large-scale learning of human-object interactions from Internet instructional videos [6].

Acknowledgments. We warmly thank Bruno Watier (Université Paul Sabatier and LAAS-CNRS) and Galo Maldonado (ENSAM Paris-Tech) for setting up the Parkour dataset. This work was partly supported by the ERC grant LEAP (No.336845), the H2020 Memmo project, CIFAR Learning in Machines&Brains program, and the European Regional Development Fund under the project IMPACT (reg. no. CZ.02.1.01/0.0/0.0/15_003/0000468).

References

- [1] CMU Graphics Lab Motion Capture Database. <http://mocap.cs.cmu.edu>.
- [2] Project webpage (code/dataset). <https://www.di.ens.fr/willow/research/motionforcesfromvideo/>.
- [3] W. Abdulla. Mask R-CNN for object detection and instance segmentation on Keras and TensorFlow. https://github.com/matterport/Mask_RCNN, 2017.
- [4] S. Agarwal, K. Mierle, and Others. Ceres solver. <http://ceres-solver.org>.
- [5] I. Akhter and M. J. Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, 2015.
- [6] J.-B. Alayrac, P. Bojanowski, N. Agrawal, I. Laptev, J. Sivic, and S. Lacoste-Julien. Unsupervised learning from narrated instruction videos. In *CVPR*, 2016.
- [7] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.
- [8] L. T. Biegler. *Nonlinear programming: concepts, algorithms, and applications to chemical processes*, volume 10, chapter 10. Siam, 2010.
- [9] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it simple: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*, 2016.
- [10] R. Boulic, N. M. Thalmann, and D. Thalmann. A global human walking model with real-time kinematic personification. *The Visual Computer*, 6(6):344–358, Nov 1990.
- [11] E. Brachmann, F. Michel, A. Krull, M. Ying Yang, S. Gumhold, et al. Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image. In *CVPR*, 2016.
- [12] M. A. Brubaker, D. J. Fleet, and A. Hertzmann. Physics-based person tracking using simplified lower-body dynamics. In *CVPR*, 2007.
- [13] M. A. Brubaker, L. Sigal, and D. J. Fleet. Estimating contact dynamics. In *CVPR*, 2009.
- [14] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017.
- [15] J. Carpentier and N. Mansard. Analytical derivatives of rigid body dynamics algorithms. In *Robotics: Science and Systems (RSS 2018)*, 2018.
- [16] J. Carpentier and N. Mansard. Multi-contact locomotion of legged robots. *IEEE Transactions on Robotics*, 2018.
- [17] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiraux, O. Stasse, and N. Mansard. The Pinocchio C++ library—A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives. In *International Symposium on System Integrations*, 2019.
- [18] J. Carpentier, F. Valenza, N. Mansard, et al. Pinocchio: fast forward and inverse dynamics for poly-articulated systems. <https://stack-of-tasks.github.io/pinocchio>, 2015–2017.
- [19] C.-H. Chen and D. Ramanan. 3d human pose estimation= 2d pose estimation+ matching. In *CVPR*, 2017.
- [20] V. Delaitre, J. Sivic, and I. Laptev. Learning person-object interactions for action recognition in still images. In *NIPS*, 2011.
- [21] M. Diehl, H. Bock, H. Diedam, and P.-B. Wieber. Fast Direct Multiple Shooting Algorithms for Optimal Robot Control. In *Fast Motions in Biomechanics and Robotics*. 2006.
- [22] A. Doumanoglou, R. Kouskouridas, S. Malassiotis, and T.-K. Kim. 6d object detection and next-best-view prediction in the crowd. In *CVPR*, 2016.
- [23] R. Featherstone. *Rigid body dynamics algorithms*. Springer, 2008.
- [24] D. F. Fouhey, V. Delaitre, A. Gupta, A. A. Efros, I. Laptev, and J. Sivic. People watching: Human actions as a cue for single view geometry. *IJCV*, 110(3):259–274, 2014.
- [25] J. Gall, B. Rosenhahn, T. Brox, and H.-P. Seidel. Optimization and filtering for human motion capture. *IJCV*, 87(1-2):75, 2010.
- [26] S. Gammeter, A. Ess, T. Jäggl, K. Schindler, B. Leibe, and L. Van Gool. Articulated multi-body tracking under egomotion. In *ECCV*, 2008.
- [27] J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975.
- [28] A. Grabner, P. M. Roth, and V. Lepetit. 3D Pose Estimation and 3D Model Retrieval for Objects in the Wild. In *CVPR*, 2018.
- [29] A. Gupta, A. Kembhavi, and L. S. Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *PAMI*, 31(10):1775–1789, 2009.
- [30] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick. Mask R-CNN. *CoRR*, abs/1703.06870, 2017.
- [31] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [32] A. Herdt, N. Perrin, and P.-B. Wieber. Walking without thinking about it. In *International Conference on Intelligent Robots and Systems (IROS)*, 2010.
- [33] S. Hinterstoisser, V. Lepetit, N. Rajkumar, and K. Konolige. Going further with point pair features. In *ECCV*, 2016.
- [34] E. Insafutdinov, L. Pishchulin, B. Andres, M. Andriluka, and B. Schiele. Deepcruc: A deeper, stronger, and faster multi-person pose estimation model. In *ECCV*, 2016.
- [35] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *PAMI*, 36(7):1325–1339, jul 2014.
- [36] Y. Jiang, H. Koppula, and A. Saxena. Hallucinated humans as the hidden context for labeling 3d scenes. In *CVPR*, 2013.
- [37] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018.
- [38] J. Kuffner, K. Nishiwaki, S. Kagami, M. Inaba, and H. Inoue. Motion planning for humanoid robots. In *Robotics Research. The Eleventh International Symposium*, 2005.
- [39] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox. DeepIM: Deep Iterative Matching for 6D Pose Estimation. In *ECCV*, 2018.
- [40] Z. Li, J. Sedlar, J. Carpentier, I. Laptev, N. Mansard, and J. Sivic. Estimating 3D motion and forces of person-object interactions from monocular video. *arXiv preprint arXiv:1904.02683*, 2019.

- [41] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014.
- [42] V. Loing, R. Marlet, and M. Aubry. Virtual training for a real application: Accurate object-robot relative localization without calibration. *IJCV*, Jun 2018.
- [43] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6):248, 2015.
- [44] M. M. Loper, N. Mahmood, and M. J. Black. MoSh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 33(6):220:1–220:13, Nov. 2014.
- [45] G. Maldonado, F. Bailly, P. Souères, and B. Watier. Angular momentum regulation strategies for highly dynamic landing in Parkour. *Computer Methods in Biomechanics and Biomedical Engineering*, 20(sup1):123–124, 2017.
- [46] J. Malmaud, J. Huang, V. Rathod, N. Johnston, A. Rabinovich, and K. Murphy. What’s cookin’? interpreting cooking videos using text, speech and vision. *arXiv preprint arXiv:1503.01558*, 2015.
- [47] J. Martinez, R. Hossain, J. Romero, and J. J. Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017.
- [48] I. Mordatch, E. Todorov, and Z. Popović. Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (TOG)*, 31(4):43, 2012.
- [49] F. Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *CVPR*, 2017.
- [50] A. Newell, Z. Huang, and J. Deng. Associative embedding: End-to-end learning for joint detection and grouping. In *NIPS*, 2017.
- [51] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [52] M. Oberweger, M. Rad, and V. Lepetit. Making Deep Heatmaps Robust to Partial Occlusions for 3D Object Pose Estimation. In *ECCV*, 2018.
- [53] G. Pavlakos, X. Zhou, K. G. Derpanis, and K. Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *CVPR*, 2017.
- [54] M. Posa, C. Cantu, and R. Tedrake. A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research*, 33(1):69–81, 2014.
- [55] A. Prest, V. Ferrari, and C. Schmid. Explicit modeling of human-object interactions in realistic videos. *PAMI*, 35(4):835–848, 2013.
- [56] M. Rad and V. Lepetit. Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In *ICCV*, 2017.
- [57] M. Rad, M. Oberweger, and V. Lepetit. Feature Mapping for Learning Fast and Accurate 3D Pose Inference from Synthetic Images. In *CVPR*, 2018.
- [58] G. Schultz and K. Mombaur. Modeling and optimal control of human-like running. *IEEE/ASME Transactions on mechatronics*, 15(5):783–792, 2010.
- [59] H. Sidenbladh, M. J. Black, and D. J. Fleet. Stochastic tracking of 3d human figures using 2d image motion. In *ECCV*, 2000.
- [60] Y. Tassa, T. Erez, and E. Todorov. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [61] A. Tejani, D. Tang, R. Kouskouridas, and T.-K. Kim. Latent-class hough forests for 3d object detection and pose estimation. In *ECCV*, 2014.
- [62] B. Tekin, A. Rozantsev, V. Lepetit, and P. Fua. Direct prediction of 3d body poses from motion compensated sequences. In *CVPR*, 2016.
- [63] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *CoRR*, abs/1703.06907, 2017.
- [64] S. Tonneau, A. Del Prete, J. Pettré, C. Park, D. Manocha, and N. Mansard. An Efficient Acyclic Contact Planner for Multipled Robots. *IEEE Transactions on Robotics (TRO)*, 2018.
- [65] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustment modern synthesis. In *International workshop on vision algorithms*, 1999.
- [66] X. Wei and J. Chai. Videomocap: Modeling physically realistic human motion from monocular video sequences. *ACM Trans. Graph.*, 29(4):42:1–42:10, July 2010.
- [67] E. R. Westervelt, J. W. Grizzle, and D. E. Koditschek. Hybrid zero dynamics of planar biped walkers. 2003.
- [68] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli. Gait and trajectory optimization for legged systems through phase-based end-effector parameterization. *IEEE Robotics and Automation Letters*, 3(3):1560–1567, 2018.
- [69] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *CoRR*, abs/1711.00199, 2017.
- [70] B. Yao and L. Fei-Fei. Recognizing human-object interactions in still images by modeling the mutual context of objects and human poses. *PAMI*, 34(9):1691–1703, 2012.
- [71] X. Zhou, M. Zhu, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *CVPR*, 2016.