

Gravity as a Reference for Estimating a Person’s Height from Video

Didier Bieler¹ Semih Günel¹ Pascal Fua¹ Helge Rhodin^{1,2}
¹EPFL, Lausanne, Switzerland ²UBC, Vancouver, Canada
 firstname.lastname@epfl.ch

Abstract

Estimating the metric height of a person from monocular imagery without additional assumptions is ill-posed. Existing solutions either require manual calibration of ground plane and camera geometry, special cameras, or reference objects of known size. We focus on motion cues and exploit gravity on earth as an omnipresent reference ‘object’ to translate acceleration, and subsequently height, measured in image-pixels to values in meters. We require videos of motion as input, where gravity is the only external force. This limitation is different to those of existing solutions that recover a person’s height and, therefore, our method opens up new application fields. We show theoretically and empirically that a simple motion trajectory analysis suffices to translate from pixel measurements to the person’s metric height, reaching a MAE of up to 3.9 cm on jumping motions, and that this works without camera and ground plane calibration.

1. Introduction

Estimating metric scale from a monocular image or video recordings is a fundamental problem in computer vision [5, 14] and important for determining distances in forensics, autonomous driving, person re-identification, and structure-from-motion (SfM). In general, object size and distance cancel in perspective projection—which makes the problem ill-posed. However, solutions exist if cameras and ground floor are manually calibrated [12, 31, 18], special cameras for depth-of-field sweeping are available [10], or reference objects of known scale are present [30, 14, 30, 20]. Some solutions studied relationships that are particular for persons, such as height, appearance, facial features, body proportions [3, 13]. However, only uncertain and biased predictions could be obtained.

This paper aims at a new approach to estimating a person’s height using motion cues in videos. The main idea is to use the omnipresent gravity on earth as a reference ‘object’. Newton’s second equation of motion dictates that the trajectory of an object is a parabola, a function of time, its

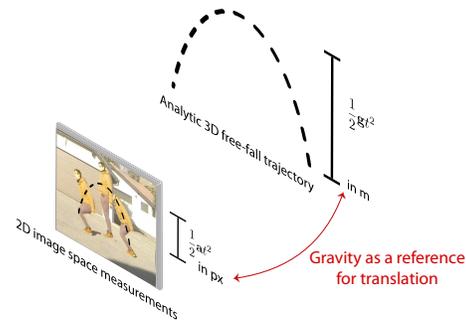


Figure 1. **Idea.** We exploit gravity as a reference *object* for mapping image height measurements in pixel [px] to metric height [m].

initial speed and position, with the curvature determined by the acceleration induced by constant external forces [22]. To this end, we restrict ourselves to cases where gravitation is the only source of external acceleration and the camera is static, so that acceleration in the image can be uniquely attributed to gravity. By relating acceleration in the image and to gravity on earth, we can then translate measurements in pixels to metric height in meters, as sketched in Figure 1.

Although this strategy restricts the application to scenarios entailing free-fall motions, it applies to any video of a person jumping or running, where air friction and other sources of acceleration are negligible. Because this limitation is orthogonal to those of existing solutions, our method opens up new application fields, such as metric monocular SfM, automatic person re-identification in videos with unknown camera geometry, and metric human pose estimation from a single camera, as demonstrated in Figure 2.

Our method is inspired by approaches that estimate the 3D trajectory of rigid objects in free fall [16, 23, 24, 17, 29]. All these methods assume a calibrated camera, known gravity direction, rigid objects and focus on object position instead of scale estimation. In contrast, we show that our strategy for height estimation is applicable without knowing or constraining the initial object speed or position, that it can be generalized to account for the internal forces present in articulated person motion, and does not require knowledge of any camera parameters nor ground plane position.

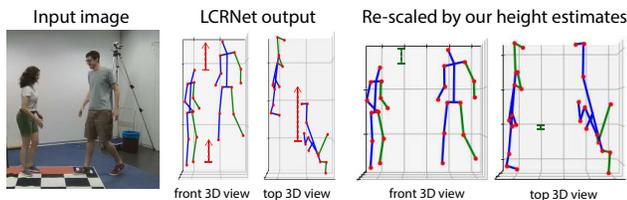


Figure 2. **Application to metric 3D pose estimation.** **Left:** Input and output of LCRNet, a 3D human pose estimation method [25], which does not recover the correct scales (marked in red). **Right:** Our height estimate from directed on-spot jumping recover the scale and relative depth of the hip locations (marked in green).

We validate our findings on a new test set containing 12 persons performing seven different jumping and running motions at two distances. Our obtained mean absolute error (MAE) of 3.9 cm, is nearly half of the 6.5 cm MAE by [13], which demonstrates the success of our method and the importance of geometric and physical constraints. Our code and dataset is available at <https://cvlab.epfl.ch/articulated-free-fall-dataset>.

2. Related Work

Multiple approaches for estimating the size of a person or object have been proposed in the literature. We list the most important ones, focusing on the existing conditions used to make monocular height estimation well-posed.

Height from camera geometry. Without external scale information, object size and distance is ambiguous according to the basic pinhole camera model. In practice, lenses have a limited depth of field, which shape-from-defocus techniques exploit [21, 27]. While it can be used to guess depth orderings in a single image, a focal sweep across multiple images or a specialized camera [10] is required for metric scale reconstruction. These constraints preclude applications to monocular height estimation from single images and videos recorded with consumer cameras.

Scale from scene and camera geometry. One method for estimating height from images is to place several reference objects nearby the target object [30, 20]. Placing multiple references is cumbersome and has been overcome by relating a single reference to other points in the scene through vanishing points of the ground plane [5, 14, 30]. An alternative is to manually calibrate the camera height and orientation in relation to the ground plane [12, 31, 18]. Thereby, the height of standing people can be inferred by locating the head and foot position in a single image, triangulating the 3D foot position on the known ground plane, and scaling the head-foot distance to this 3D position. These methods do not require to alter the scene but they are not applicable to legacy videos and require an expert to calibrate.

Height from statistics. Dey et al. [7] propose a unique solution that exploits statistics across a group of people by measuring relative heights in group pictures and connecting these in large image collections. Absolute height is then estimated from the resulting network of relative heights by enforcing consistency with the average human. Medical studies determined a correlation of human height and ratios of limb proportions [1] and the ratio of head to shoulders [28, 15], but these are difficult to estimate from images. Recent methods have attempted to capture such statistical relations directly from anthropometric measurements [3, 11] and collection of images [13, 6] though black-box regression using linear regression and deep learning, respectively. We show that these statistical methods suffer from errors and bias towards the average human height, as monocular scale estimation is ill-posed without geometric constraints.

Physics-based trajectory modeling. Physical constraints are widely used for refining and estimating trajectories of rigid objects. Kim et al. [16] recover the 3D positions of a soccer ball in relation to the known height of a player or goalpost by exploiting that the ball trajectory follows a parabola in free flight. Ohno et al. [23] model gravity and air drag explicitly to recover metric 3D football trajectories from a calibrated camera. This strategy has also been shown to generalize to other projectiles [24]. Kumar et al. [17] analyze tennis and use physics to fill-in frames for which no multi-view triangulation of the ball is available. All of these methods assume a calibrated camera, rigid object, focus on object position, and, thereby, do not address our aim of recovering a person’s height from uncalibrated video.

3. Theory

Gravitation is omnipresent and roughly constant on the earth’s surface, with an acceleration of $g \approx 9.81m/s^2$ towards the center of the planet and the variation across the surface below four percent.¹

A rigid object in free-fall, where gravitation is the only acceleration, is explained by Newton’s equation of motion:

$$\mathbf{p}(t) = \frac{1}{2}\mathbf{g}t^2 + \mathbf{v}_0t + \mathbf{p}_0, \quad (1)$$

where \mathbf{p}_0 and \mathbf{v}_0 are 3D-vectors capturing the initial position and velocity, respectively, t the elapsed time, and $\mathbf{g} = n\mathbf{g}$ the acceleration in direction \mathbf{n} (vertically, depending on the chosen coordinate system).

The underlying idea is to measure motion in video and relate the quantities estimated in pixel units to metric meters through Eq. 1. Figure 1 sketches this relation graphically. For example, in the special case of $\mathbf{v}_0 = 0$, Newton’s

¹Extremes are, e.g., $9.78m/s^2$ in Singapore and $9.83m/s^2$ in Oslo.

equation dictates that an object moves $4.9m$ after one second of free fall. Hence, estimating a motion of $|\mathbf{p}^{\text{px}}|$ pixels (px) over a video of length $t = 1s$ yields $1\text{px} = \frac{1}{|\mathbf{p}^{\text{px}}|}4.9m$, uniquely determining the px to meter ratio, which in turn allows to translate height measurements of the moving object in the image to metric units in 3D world coordinates. Next, we derive this relation formally with respect to the camera position and orientation. We first derive them for the case of rigid objects and then for articulated motion.

3.1. Computing a Rigid Object's Height

In this section we derive a linear relation, the factor q , that maps the distances measured in the image, e.g., height h^{px} measured from upper to lower object extend, to corresponding height in meters, $h = h^{\text{px}}q$. Notably, this factor can be computed for unknown distance d , focal length f , and gravity direction, by relating the gravitational constant g to the measured image acceleration \mathbf{a}^{px} .

Formally, the observed 2D motion $(\mathbf{p}_1^{\text{px}}, \dots, \mathbf{p}_T^{\text{px}})$ on the image plane across a video of T frames are samples of the projected 3D motion $\mathbf{p}^{\text{px}}(t) = \Pi(\mathbf{p}(t))$. In this work, we will approximate the projection process with scaled orthographic projection, Π . In terms of camera coordinates, with optical axis pointing towards the third coordinate, the projected free-fall motion (Eq. 1) is then explained by

$$\begin{aligned} \mathbf{p}^{\text{px}}(t) &= \frac{f}{d} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \mathbf{p}(t) \\ &= \frac{f}{d} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \left(\frac{t^2}{2} \mathbf{g} + \mathbf{v}_0 t + \mathbf{p}_0 \right), \end{aligned} \quad (2)$$

here d is the object distance to the camera. This projection formula implies that the image motion is also a parabola,

$$\mathbf{p}^{\text{px}}(t) = \frac{1}{2} \mathbf{a}^{\text{px}} t^2 + \mathbf{v}_0^{\text{px}} t + \mathbf{p}_0^{\text{px}}. \quad (3)$$

The estimation of \mathbf{a}^{px} , \mathbf{v}_0^{px} and \mathbf{p}_0^{px} from the input video is explained in Section 4.3. By relating the quadratic terms in Eq. 2 and Eq. 3, we obtain the following relation between the measured 2D and predicted 3D acceleration,

$$\begin{aligned} \frac{1}{2} \mathbf{a}^{\text{px}} t^2 &= \frac{f}{d} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \frac{t^2}{2} \mathbf{g} \\ \Leftrightarrow \quad \mathbf{a}^{\text{px}} &= \frac{f}{d} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \mathbf{n} g. \end{aligned} \quad (4)$$

This relation is sketched in Fig. 1. Although \mathbf{n} , f , and d are unknown, they are constants. For an object at d , their combined effect is determined by the acceleration quotient,

$$\mathbf{q} = \frac{d}{f \mathbf{n}^{\text{px}}} = \frac{g}{\mathbf{a}^{\text{px}}} \quad \text{with} \quad \mathbf{n}^{\text{px}} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \mathbf{n}, \quad (5)$$

where g is known, \mathbf{a}^{px} is the observed acceleration in Eq. 3. Note that, vector-scalar operations are here element-wise.

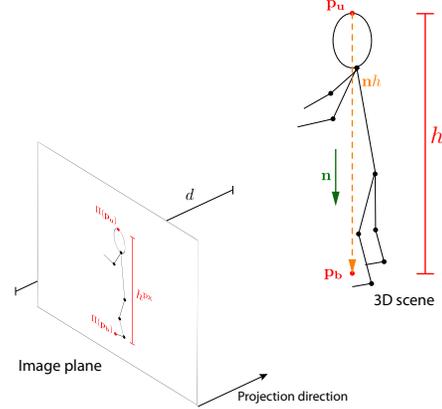


Figure 3. **Projection of height.** The projected height, h^{px} is the distance between two projected points $\Pi(\mathbf{p}_u)$ and $\Pi(\mathbf{p}_b)$ that are aligned with the direction of gravity \mathbf{n} in 3D and span height h .

We now turn to the height estimation using \mathbf{q} . We define height as the distance between two 3D points \mathbf{p}_u and \mathbf{p}_b that are in a line with the direction of gravity \mathbf{n} , see Figure 3. The sought function is derived by applying Π on \mathbf{p}_u and \mathbf{p}_b , using the linearity of Π , and substituting \mathbf{q} from Eq. 5,

$$\begin{aligned} \mathbf{h}^{\text{px}} &= \Pi(\mathbf{p}_u) - \Pi(\mathbf{p}_b) \\ &= \frac{f}{d} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} (\mathbf{p}_u - \mathbf{p}_b) \\ &= \frac{f}{d} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \mathbf{n} h = \frac{h}{\mathbf{q}}, \end{aligned} \quad (6)$$

where, \mathbf{h}^{px} is the difference between the projected reference points used for height estimation—in our experiments it is the head-to-heel vector measured in the image. This relation provides two solutions for h , respectively for the vertical and horizontal components of \mathbf{q} and \mathbf{a}^{px} . This system of equations could be solved in the least-squares sense.

However, in practice, videos are predominantly captured upright and the vertical direction will be dominant. We therefore use the vertical solution directly. In the following, we denote the vertical components of \mathbf{q} , \mathbf{a}^{px} and \mathbf{h}^{px} respectively with scalars q , a^{px} and h^{px} . The translation rule between h^{px} in px and metric height is then

$$h = h^{\text{px}} q. \quad (7)$$

Note that the gravity direction and camera projection matrix are subsumed in q . Hence, equations hold for any affine camera model and no ground plane nor camera calibration is needed. Also the simplified Eq. 7 applies to unknown camera orientation, so long vertical acceleration is non-zero.

Only if n and f are known, i.e., when the direction of gravity, camera intrinsic, and extrinsic parameters are calibrated, it is true that q can be further decomposed to compute the object's distance d and extend in all directions, which was the focus of previous studies [16, 23, 24, 17, 29].

3.2. Computing a Person’s Height

A jumping or otherwise moving person is likely to articulate the arms and other body parts actively. Instead of trying to model the complex interactions between body parts, we model the human body as a closed system and switch to Center of Mass (COM) computations. Conveniently, by the law of conservation of momentum, the center of mass of a closed system will move at a constant speed even if internal parts move in other ways, and it also follows Newton’s equation of motion (Eq. 1) when accelerated by g . Moreover, the (scaled orthographic) projection of the 3D-COM is the 2D-COM of the individually projected body parts, so that Eq. 4 holds as well.

To sum up, to estimate the height of a person we propose to compute the COM trajectory in 2D, solve Eq. 5 for q , and apply q in Eq. 7 to translate the height h^{px} from head to heel measured in pixel units to absolute height h in m.

In the following, we will explain how the needed quantities, such as COM trajectory can be inferred automatically from videos without requiring camera calibration.

4. Algorithms

The algorithm splits into four consecutive steps: computing object position and COM per frame, detecting free fall events, estimating image acceleration, and measuring height in the image for total height computation.

4.1. Measuring the COM

For algorithmic validation, we analyze the ideal projectile trajectory of a rigid and uniformly colored ball. In this toy example, the image position measurements can easily be automated. We segment the ball from the background through color thresholding and determine the object center, $\mathbf{p}^{px}(t)$, by fitting a circle to the segmentation contour.

For articulated human motion, we detect the person’s body parts with AlphaPose [8], a neural network that has been trained to predict person keypoints in the input image. Given a set of T RGB images $I \in \mathbb{R}^{3 \times W \times H}$ representing the T frames of the original video, AlphaPose outputs $J = 17$ person 2D joint positions $(\mathbf{p}_t^{px,j})_{j=1}^J$ for each frame t . Examples are shown in Fig. 4. If multiple persons are detected, we select the largest one, the one in the foreground.

To compute the needed COM position, $\mathbf{p}_t^{px} = \sum_{j=0}^J r_j \mathbf{p}_t^{px,j}$, of the J parts at positions $\mathbf{p}_t^{px,j}$, one needs to know the ratio r_j between the weight of body part j in relation to the total weight. For persons, the absolute weight varies significantly and it appears to be at least as hard to estimate weight as to recover the sought height. However, the relative weight of bodyparts varies little across individuals. We took the mean mass distribution estimated in a large-scale study in [4], which has been widely used for COM computation in the past.

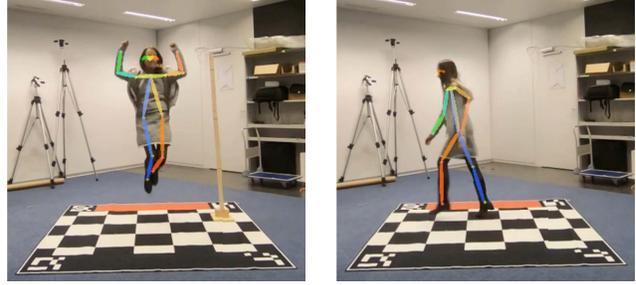


Figure 4. **Bodypart keypoint detections** using AlphaPose on frontal and lateral motion, displayed as colored skeleton overlay.

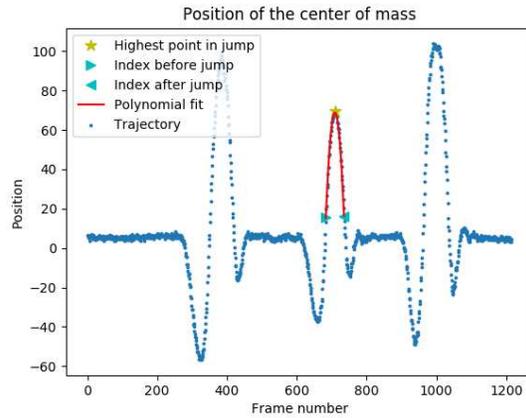


Figure 5. **COM trajectory for on-spot jumping**, with the polynomial fit, detected start, maximum, and end points marked.

4.2. Detecting Free-Fall

We assume that gravity is the only external force applied on the person (neglecting air friction). However, this is only true during jumps and other free-fall phases. Thus, we derive a simple yet effective algorithm for distinguishing flight and ground contact phases. First, we localize maxima in the trajectory, the highest points $M = \{m | p_m^{px} > p_t^{px} \text{ for } t \in [m - 10, \dots, m + 10]\}$ in the trajectory of vertical pixel locations p_t^{px} . Second, we compute the position of contact, p_{floor}^{px} as the median across the first 100 frames, expecting that the motion starts in a standing pose. We then select the frame interval $S_m = [p_s^{px}, \dots, p_e^{px}]$ such that all $p_t^{px} \in S_m$ are at least 15 % above the ground in relation to the jump peak p_m^{px} . This procedure ensures the points belonging to the jump initiation are excluded, those where the body is accelerated and in contact with the ground. A selection example is visualized in Figure 5.

For lateral motion the ground contact point varies dependent on the position in the image, lense distortion, and camera orientation. For these, we switch to an alternative approach and determine the distance between a maxima and the two neighboring minimum values and select only those points that are in the upper half.

4.3. Estimating Acceleration

A key challenge is the estimation of a^{px} from the COM trajectory. Without loss of generality, we assume that the camera roll is small and, hence, acceleration is predominantly vertical. This can be ensured by rotating the image in the direction of maximal acceleration. Moreover, we read the camera frame rate from the video meta data, to measure acceleration in pixel per seconds instead of number of frames. We propose two ways of estimating the image acceleration:

Baseline-distance-based. At the highest point of a jump the velocity is zero. Hence, the linear term in Eq. 1 must be zero. Therefore, motion in the video after this turning point till inception with the ground can be uniquely attributed to gravitational acceleration. Taking the highest point p_m and the last point before ground contact, p_e^{px} , we can easily solve $p_e^{px} - p_m^{px} = \frac{1}{2}a^{px}(e - m)^2$ for a^{px} since v_0^{px} is zero and p_0^{px} constant. Here $(e - m)$ is the time difference in seconds. However, these point estimates are prone to error and the highest point might happen to be in-between two samples.

Ours. Under the free-fall assumption, the COM trajectory must follow a second order polynomial, with the quadratic term representing the acceleration and the linear term the velocity in pixels. Therefore, we fit a polynomial of degree two on the curve in the least squares sense. Since the COM trajectory is estimated through AlphaPose, wrong and inaccurate detections occur. We test two measures to counteract. First, we utilize the confidence output of AlphaPose and exclude points where the score drops below 2. Second, we apply random sample consensus (RANSAC) [9] on top.

The fitting process and the effect of outliers rejection is exemplified in Figure 6. The acceleration a^{px} is then simply two-times the quadratic coefficient of the curve fit. Note that the image acceleration needs to be measured in pixels per seconds squared using the video frame rate information.

4.4. Converting Length Estimates to Human Height

To estimate the total height of a person, we need to know the pixel location of the top of the head to the heel in a standing pose. Unfortunately, no such off-the-shelf detectors exist. The most utmost locations that AlphaPose returns are the nose and ankle points. As a stop-gap solution, we propose to infer a correction factor to go from nose-to-ankle to total height. We compute the mean ratio, $c = 1.17 \pm 0.03$, between the person’s pixel height and ankle-nose distance, determined by AlphaPose, over 29 images taken from the web. The low standard deviation suggests that this linear approximation is accurate.

In practice, we measure the pixel height at the first frames of the video, assuming an upright stance. The final height is subsequently recovered by translating that image

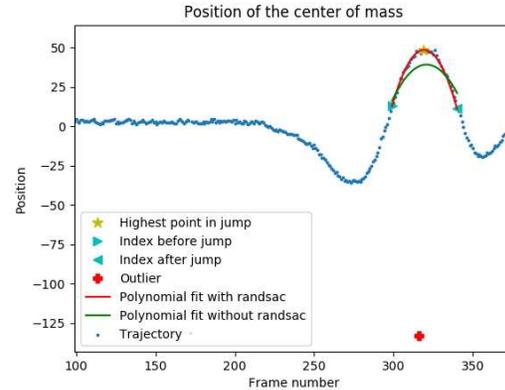


Figure 6. **Outlier removal example.** The red diamond marks the outlier and the red curve is the one that excludes it during fitting, leading to significantly improved fits.

height measurements to meters with Eq. 7 and multiplying the result with c .

5. Experimental Evaluation

We first study the attainable accuracy of the gravity-based height estimation at hand of rigid objects, where automatic detection is easy and reliable. Subsequently, we analyze the feasibility of estimating a person’s height during jumps and runs of varying complexity and compare it to ground truth measurements and to the appearance-based solution proposed in [13]. We introduce new tests sets for both setups:

- **FallingBall.** We drop a tennisball at several distances to the camera, spaced 50cm apart, as sketched in Figure 7. Two bounces of the ball are recorded at 120 Hz, the second one is about 50% lower due to the absorbed energy at inception with the ground. The introduced ball detection algorithm is used to calculate the ball diameter in pixels. The ground truth diameter is 7.3 cm.
- **ArticulatedFreeFall.** We recorded 12 subjects, at 30 Hz, located at four and seven meter distance to the camera. Examples are shown in Figure 8. Seven motions are tested: low, high, jumping jack, and funny on-spot jumps, as well as lateral running, exaggerated running, and jumping with swinging arms. Funny jumps are undirected and participants chose asymmetric, articulated poses, see Figure 9 for examples. The dataset is available at <https://cvlab.epfl.ch/articulated-free-fall-dataset>.

On top of this, we test human height estimation on community videos downloaded from YouTube. We chose a parour clip with movements of varying complexity in different outdoor scenes. Qualitative results of these, including limitations, are shown in the supplemental video.

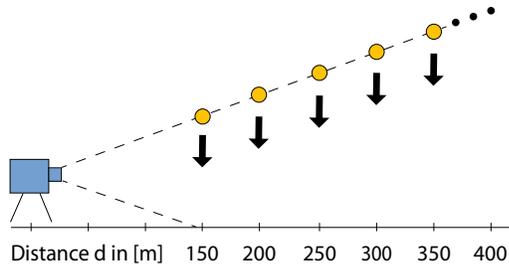


Figure 7. **Illustration of the tennisball experiment.** The ball is dropped in front of the camera so that it is still in the field of view.

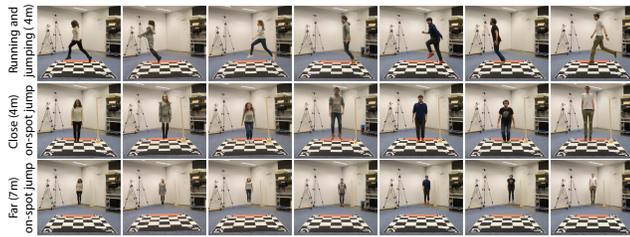


Figure 8. **Examples from our new ArticulatedFreeFall dataset,** that is captured at 4 and 7 m distance s and features lateral running, jumping and on-spot jumping.

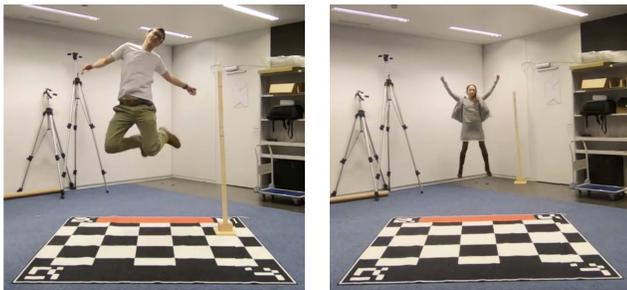


Figure 9. **Example of funny and jumping jack poses.** The wooden reference object is used to estimate relative errors of q .

For quantitative evaluation of estimated height \hat{h} and conversion factor \hat{q} we compute absolute and signed errors to report:

- **accuracy** across trials as the Mean Absolute Error (MAE) between predicted and estimated quantities.
- **bias** across trials as the Mean signed Error (ME).

Both measures are computed for absolute and relative errors since relative errors are easier to grasp for the small values of q . For consistency analysis, we report the Standard Deviation (SD) for absolute and signed errors using the \pm notation.

The ground truth ratio q is computed from a reference object of known height and its height in the images.

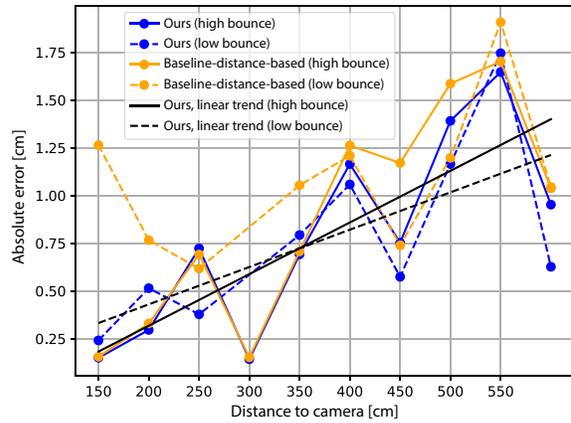


Figure 10. **Absolute size estimation error** as a function of the distance to the camera. The error is small overall, below one centimeter for distances up to 4m and follows a linear trend with respect to the distance. The varied bounce heights do not influence the accuracy significantly.

5.1. Estimating a Rigid Object’s Height

This experiment analyzes height estimation accuracy on a tennis ball being thrown from varying heights and distances to the camera. Figure 10 plots the absolute error with respect to the distance to the camera. Overall, the error is small and below one centimeter for distances up to four meters. Acceleration estimation with curve fitting, **Ours**, slightly outperforms the simpler distance-based strategy, **Baseline-distance-based**. This marginal difference can be explained with the high accuracy of the ball detection, larger differences can be seen in the subsequent evaluation on persons.

To analyze the overall behavior we show the linear fit of errors from **Ours** as black lines in Figure 10. The error increases with respect to the distance from the camera. This is expected, as the projected object size and motion decrease in scale proportionally to the distance while the image estimation error stays constant due to the pixel discretization. We analyzed two different bounce heights. Although the lower bounce is roughly 50% smaller, we found no significant influence on the scale estimation accuracy. This can be seen on the linear fits in Figure 10.

5.2. Estimating a Person’s Height

Qualitative results and exemplar videos from the new **ArticulatedFreeFall** dataset are shown in the supplemental video. Here we analyze the accuracy quantitatively. All videos are processed with the described COM computation and automatic flight-phase detection using per-frame AlphaPose estimates.

Table 1. **Error analysis per subject for on-spot jumps.** We compute the median prediction across the four jumps of each subject and report the relative errors, absolute errors, as well as the bias (ME) and accuracy (MAE) across all subjects. Estimates are accurate and, in contrast to [13], unbiased with 3.9 MAE and 0.8 ME.

Detailed results using Ours:

Subjects			4m distance				7m distance				Combined 4m and 7m	
ID	Height [cm]	Sex [f or m]	Error [cm]	Rel. [%]	Error [cm/px]	Rel. [%]	Error [cm]	Rel. [%]	Error [cm/px]	Rel. [%]	Error [cm]	Rel. [%]
S1	163.0	m	-2.4	-1.5	-0.015	-3.7	3.0	1.8	0.035	4.8	1.3	0.8
S2	178.0	m	-1.0	-0.6	0.000	0.0	-3.3	-1.9	0.030	4.1	-2.4	-1.4
S3	175.0	m	0.6	0.3	0.000	0.0	2.6	1.5	-0.025	-3.4	1.3	0.7
S4	159.0	f	-7.4	-4.7	-0.030	-7.3	2.3	1.5	0.005	0.7	-2.6	-1.6
S5	175.0	m	0.3	0.2	0.010	2.4	-9.1	-5.2	-0.040	-5.5	-2.8	-1.6
S6	157.0	f	0.8	0.5	0.000	0.0	1.1	0.7	0.005	0.7	1.1	0.7
S7	188.0	m	-5.4	-2.9	-0.005	-1.2	-2.2	-1.2	0.025	3.4	-4.7	-2.5
S8	183.0	m	-3.2	-1.8	0.010	2.4	-11.2	-6.1	-0.015	-2.1	-6.1	-3.3
S9	163.0	f	23.4	14.4	0.050	12.2	-0.4	-0.3	-0.010	-1.4	13.2	8.1
S10	170.0	f	4.9	2.9	0.000	0.0	10.1	6.0	0.015	2.1	7.9	4.7
S11	173.0	m	-3.9	-2.3	-0.025	-6.1	2.6	1.5	0.025	3.4	0.0	0.0
S12	173.0	m	0.2	0.1	0.005	1.2	7.9	4.5	0.025	3.4	3.4	2.0

Bias (ME±STD) across subjects:

Ours	0.6±7.9	0.4±4.8	0.00±0.02	0±4.9	0.3±6.1	0.2±3.5	0.01±0.02	0.9±3.3	0.8±5.4	0.6±3.2
Baseline-distance-based	-7±21.3	4.3±12.5	0.02±0.05	4±13.1	-5.7±21.1	3.3±12	0.02±0.09	2.2±12.2	-7.7±16.5	4.5±9.5
Baseline-population-mean	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	-2.5±9.5	n/a
Günel et al. [13]	3.7±6.8	2.3±4	n/a	n/a	3.6±6.9	2.3±4.1	n/a	n/a	3.7±6.7	2.3±4

Accuracy (MAE±STD) across subjects:

Ours	4.5±6.4	2.7±3.9	0.01±0.02	3±3.8	4.6±3.8	2.7±2.1	0.02±0.01	2.9±1.6	3.9±3.7	2.3±3.2
Baseline-distance-based	18.3±11.9	10.8±7	0.05±0.03	11.7±6.4	13.6±16.7	7.8±9.5	0.05±0.07	7.2±9.9	15.5±8.8	9±5
Baseline-population-mean	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	8.1±5.0	n/a
Günel [13]	6.5±4	3.8±2.5	n/a	n/a	6.5±3.9	3.8±2.5	n/a	n/a	6.5±3.9	3.8±2.4

Table 2. **Error analysis per on-spot jump type.** We compute the accuracy (MAE) and bias (ME) across all 12 subjects. Estimation accuracy is roughly the same for all jump heights, with outliers in each class leading to high mean errors and standard deviations. This highlights the importance of taking multiple measurements.

	4m distance				7m distance			
	Error in h		Relative error in q		Error in h		Relative error in q	
	Bias [cm] (ME±STD)	Accuracy [cm] (MAE±STD)	Bias [%] (ME±STD)	Accuracy [%] (MAE±STD)	Bias [cm] (ME±STD)	Accuracy [cm] (MAE±STD)	Bias [%] (ME±STD)	Accuracy [%] (MAE±STD)
J1 (low)	-0.3±18.5	13.1±12.4	0±11.2	8±7.5	-14.8±37.3	20.7±34	-7.2±21	11.1±19
J2 (high)	-0.6±12.4	8.0±9.1	-0.1±8.1	5.8±5.4	1.6±7.9	6.5±4.2	2±4.8	4.4±2.4
J3 (jack)	1.1±8	5.6±5.6	1.1±5.1	3.5±3.8	0.2±7	5.3±4.3	1.3±4.3	3.1±3.1
J4 (funny)	0.4±11.9	9.4±6.8	0.8±6.8	5.2±4.2	4±14	11.6±8.1	3.6±8.1	7.7±3.9

Accuracy analysis. The introduced error metrics are evaluated in Table 1. The predictions from all four on-spot jumps are accumulated taking the median, independently for each of the 12 subjects. The overall height estimation accuracy of 3.9 cm (MAE) is quite good given the large distance of 4 to 7 meters. Notably is also the low bias of 0.8 cm. Accuracies vary across subjects, but without apparent correlation to subject gender and height. For instance, S4 and S9 have the same height (159 and 163 cm) and gender but largely different errors. Moreover, the largest errors are distributed across all heights, e.g., S7 (188 cm) and S9 (163 cm) have both an error above 4.5 cm. As for the tennis ball experiment, results are slightly more accurate and have lower standard deviation for jumps closer to the camera.

Table 2 sheds light on the performance in terms of jump complexity. Low jumps have the highest error, presumably due to their short duration and resulting low numbers of samples. High and simple jumps are most accurate, followed by the symmetric jumping jack. In general, occasional AlphaPose failures occur during all jump types, resulting in relatively high MAE and SD. However, taking the median prediction across all jumps, including the difficult funny poses, increases prediction accuracy and SD significantly, as analyzed before.

Lateral jumping and running motions are more difficult to capture due to the occluded body side. Therefore, AlphaPose exhibits many false joint detections, leading to temporal jitter in the derived COM trajectories. Table 3 reports the accuracy and bias across all subjects and independently for slow run, fast run, and jumping. While Jumping is still acceptable with a MAE of 6.6 cm, the reconstructions of runs are noisy. Part of the error comes from the fact that flight phases during running are very short and thereby provide little data points. The difference between complex and simple jumps can be seen on the example trajectory in Figure 11. This problem could be mitigated by recording with frame rates larger than the 30 Hz used here. Another reason could be the athletic articulation leading to errors in the COM computation. We discuss alternative counter measures in the subsequent limitation and future work section.

Quantitative comparison. We compare our approach to the only existing method for estimating height from monocular and uncalibrated footage [13]. Since that method regresses height with a neural network from single images, we apply it to all video frames and compute error values on the median prediction across all frames, to be comparable. [13] shows large errors for tall and short persons, high

Table 3. **Error analysis per subject for lateral motion.** Dynamic running and jumping motions lateral to the camera are hard to capture due to self-occlusions of half of the body, leading to unreliable COM trajectories and increased errors compared to the tested frontal jumps. Still, jumping is adequate with 6.6 cm MAE.

	Jumping, 4m				Slow Run, 4m				Fast Run, 4m			
	h		q (= 0.0041)		h		q (= 0.0041)		h		q (GT: = 0.0041)	
	Error [cm]	Rel. [%]	Error [cm/px]	Rel. [%]	Error [cm]	Rel. [%]	Error [cm/px]	Rel. [%]	Error [cm]	Rel. [%]	Error [cm/px]	Rel. [%]
Bias (ME+STD)	0.4±7.5	0.2±4.4	-0.01±0.02	-1.2±4.2	11.7±19.7	6.8±11.8	-0.02±0.04	-5.4±9.7	11.4±8.5	6.6±5	-0.01±0.02	-3.2±4.5
Accuracy (MAE+STD)	6.6±3	3.9±1.8	0.01±0.01	3.3±2.8	18.9±12.2	11.2±7.3	0.03±0.03	7.2±8.3	12.3±7	7.1±4.1	0.02±0.02	4±3.7

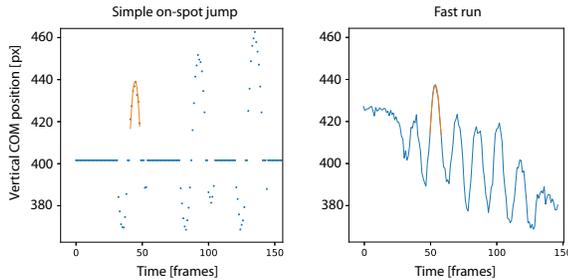


Figure 11. **COM trajectory comparison**, for 3 on-spot jumps (left) and 6 running steps (right). Running has shorter flight phases and contains temporal jitter around the peaks.

MAE of 6.5 cm, and strong bias of 3.7 cm, see bottom of Table 1 [13]. This can be explained with a tendency to the mean person height due to the uncertainty of visual cues. The naive **Baseline-population-mean**, predicting the population mean height of 168.9 m, is largely outperformed.

Ablation study. We now turn to analyzing independent model components and alternatives. The negligible bias of 0.5 cm (in Table 1) towards overestimating height shows that the simple linear translation from ankle-nose to total height works well. This is further evidenced by the similar relative errors of q , which depends only on the accuracy of a^{px} , compared to that of h , that includes both steps.

To validate the importance of outlier rejection, we disable the confidence-based outlier removal. Without any rejection, strong outliers in the pose estimation distort the curve fitting, shown for one jump in Figure 6. Using RANSAC on top helps to reduce some errors, but worsens others. In total it increases the MAE of h to 7.4 cm.

Finally, we repeat the comparison of the curve fitting, **Ours**, to the simpler **Baseline-distance-based** approach. Unlike for the ball experiment, **Baseline-distance-based** performs significantly worse, see the bottom part of Table 1. This is because **Ours** can tolerate moderate AlphaPose errors by fitting to the entire jump, while **Baseline-distance-based** is sensitive to the peak and end frame pair.

The distance dependent error due to the assumed scaled-orthographic projection model is analyzed in the supplemental document using a simulation. The document further contains an analysis on the resilience to camera rotation.

6. Limitations and Future Work

Although our results are up to two times better than the one of Günel et al., there are multiple technical constraints. Besides requiring free-flight motion, which restricts the application scenarios, the assumption of a static camera is the strongest limitation. However, we believe that existing video stabilization [19], video panorama [2], and SfM techniques [26] could be applied to resolve this constraint by registering moving videos to a static frame. In fact, our strategy could be integrated into SfM methods and provide the scale information missing in monocular reconstruction.

Detecting free-flight, is a challenge that we only touched in this study. Currently the flight and ground contact phases can only be distinguished within a hand-selected snippet that must contain flight phases. While sufficient for many tasks such as forensics and analyzing legacy videos, developing fully-automatic techniques, perhaps by analyzing foot-ground contact visually, poses an important future research direction for domains that require automation.

It is important to reinforce that our method recovers metric height, the absolute extend of an object in the direction of gravity, but no other extends. Computing scales in arbitrary directions, such as object width, would require knowledge of the 3D gravity direction and camera calibration to disentangle the effects of foreshortening and pixel extend.

7. Conclusion

We have explored a new approach for estimating the height of a person jumping and running. It is applicable to monocular videos of persons and objects alike, whenever free-fall with negligible air friction is present. We hope that this approach will enable new applications, as its limitations are complementary to those of existing monocular height estimation approaches. The precision is high, attaining up to 3.9 cm MAE on our new **ArticulatedFreeFall** dataset. In the future, we expect that advances in person keypoint detection will further improve the proposed gravity-based height estimation strategy.

Acknowledgement. We thank artist *Michael Dick* for his outstanding parkour performances. This work was supported in part by the Swiss National Science Foundation.

References

- [1] Donald Adjeroh, Deng Cao, Marco Piccirilli, and Arun Ross. Predictability and Correlation in Human Metrology. In *IEEE International Workshop on Information Forensics and Security*, 2010.
- [2] Aseem Agarwala, Ke Colin Zheng, Chris Pal, Maneesh Agrawala, Michael Cohen, Brian Curless, David Salesin, and Richard Szeliski. Panoramic video textures. In *ACM Transactions on Graphics*, volume 24, pages 821–827. ACM, 2005.
- [3] Chiraz BenAbdelkader and Yaser Yacoob. Statistical Body Height Estimation from a Single Image. In *Automated Face and Gesture Recognition*, pages 1–7, 2008.
- [4] Carles E. Clouser, John T. McConville, and J.W. Young. Weight, Volume, and Center of Mass Segments of the Human Body. *Journal of Occupational and Environmental Medicine*, 13(5):270, 1971.
- [5] Antonio Criminisi, Ian Reid, and Andrew Zisserman. Single view metrology. *International Journal of Computer Vision*, 40(2):123–148, 2000.
- [6] Antitza Dantcheva, Francois Bremond, and Piotr Bilinski. Show me your face and i will tell you your height, weight and body mass index. pages 3555–3560, 08 2018.
- [7] Ratan Dey, Madhurya Nangia, Keith W. Ross, and Yong Liu. Estimating Heights from Photo Collections: A Data-Driven Approach. In *ACM conference on Online social network*, 2014.
- [8] H.-S. Fang, S. Xie, Y.-W. Tai, and C. Lu. RMPE: Regional multi-person pose estimation. In *ICCV*, 2017.
- [9] Martin A. Fischler and Robert C. Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Communications ACM*, 24(6):381–395, 1981.
- [10] Todor Georgiev, Zhan Yu, Andrew Lumsdaine, and Sergio Goma. Lytro Camera Technology: Theory, Algorithms, Performance Analysis. In *Multimedia Content and Mobile Devices*, 2013.
- [11] Claire C. Gordon, Thomas Churchil and Charles E. Clouser, Bruce Bradtmiller, and Jhon T. McConville. Anthropometric Survey of US Army Personnel: Methods and Summary Statistics 1988. Technical report, Anthropology Research Project Inc Yellow Springs OH, 1989.
- [12] Ye Peng Guan. Unsupervised Human Height Estimation from a Single Image. *Journal of Biomedical Science and Engineering*, 2009.
- [13] Semih Günel, Helge Rhodin, and Pascal Fua. What Face and Body Shapes Can Tell About Height. *arXiv*, 2018.
- [14] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [15] Kinuyo Kato and Atsuki Higashiyama. Estimation of Height for Persons in Pictures. In *Perception & psychophysics*, 1998.
- [16] Taeone Kim, Yongduek Seo, and Ki-Sang Hong. Physics-based 3d position analysis of a soccer ball from monocular image sequences. In *International Conference on Computer Vision*, pages 721–726. IEEE, 1998.
- [17] Anil Kumar, P. Shashidhar Chavan, V.K. Sharatchandra, David S. Sumam, Phillip Kelly, and Noel E. O’Connor. 3d estimation and visualization of motion in a multicamera network for sports. In *Irish Machine Vision and Image Processing Conference*, pages 15–19. IEEE, 2011.
- [18] Shengzhe Li, Van Huan Nguyen, Mingjie Ma, Cheng-Bin Jin, Trung Dung Do, and Hakil Kim. A Simplified Nonlinear Regression Method for Human Height Estimation in Video Surveillance. In *EURASIP Journal on Image and Video Processing*, 2011.
- [19] Shuaicheng Liu, Lu Yuan, Ping Tan, and Jian Sun. Steadyflow: Spatially smooth optical flow for video stabilization. In *Conference on Computer Vision and Pattern Recognition*, pages 4209–4216, 2014.
- [20] Jenny Ljungberg and Johanna Sönerstam. Estimation of Human Height from Surveillance Camera Footage -A Reliability Study. Master’s thesis, KTH, 2008.
- [21] George Mather. Image Blur as a Pictorial Depth Cue. In *Proc. R. Soc. Lond. B*, 1996.
- [22] Isaac Newton. *The Mathematical Principles of Natural Philosophy*. 1729.
- [23] Y Ohno, Jun Miura, and Yoshiaki Shirai. Tracking Players and Estimation of the 3D Position of a Ball in Soccer Games. In *International Conference on Pattern Recognition*, 2000.
- [24] Evan Ribnick, Stefan Atev, and Nikolaos P. Papanikolopoulos. Estimating 3d positions and velocities of projectiles from monocular views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):938–944, 2009.
- [25] Gregori Rogez, Philippe Weinzapfel, and Cordelia Schmid. Lcr-Net++: Multi-Person 2D and 3D Pose Detection in Natural Images. In *arXiv preprint arXiv:1803.00455*, 2018.
- [26] Johannes L. Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016.
- [27] Jianping Shi, Xin Tao, Li Xu, and Jiaya Jia. Break ames room illusion: depth from general single images. *ACM Transactions on Graphics*, 34(6):225, 2015.
- [28] Tzzy-Yuang Shiang. A Statistical Approach to Data Analysis and 3D Geometric Description of the Human Head and Face. In *Proceedings of the National Science Council, Republic of China. Part B, Life sciences*, 1999.
- [29] Jonas Sköld. Estimating 3d-trajectories from monocular video sequences, 2015.
- [30] Johan Vester. Estimating the Height of an Unknown Object in a 2D Image. Master’s thesis, KTH, 2012.
- [31] Xiao Zhou, Peilin Jiang, Xuetao Zhang, Bin Zhang, and Fei Wang. The Measurement of Human Height Based on Coordinate Transformation. In *ICIC*. 2016.