

Single Image Human Proxemics Estimation for Visual Social Distancing

Maya Aghaei¹, Matteo Bustreo^{1,2}, Yiming Wang¹, Gianluca Bailo¹, Pietro Morerio¹, and Alessio Del Bue¹

¹Pattern Analysis & Computer Vision (PAVIS), Istituto Italiano di Tecnologia, Genova, Italy

²Dipartimento di Ingegneria Navale, Elettrica, Elettronica e delle Telecomunicazioni, University of Genova, Italy

Abstract

In this work, we address the problem of estimating the so-called “Social Distancing” given a single uncalibrated image in unconstrained scenarios. Our approach proposes a semi-automatic solution to approximate the homography matrix between the scene ground and image plane. With the estimated homography, we then leverage an off-the-shelf pose detector to detect body poses on the image and to reason upon their inter-personal distances using the length of their body-parts. Inter-personal distances are further locally inspected to detect possible violations of the social distancing rules. We validate our proposed method quantitatively and qualitatively against baselines on public domain datasets for which we provided groundtruth on inter-personal distances. Besides, we demonstrate the application of our method deployed in a real testing scenario where statistics on the inter-personal distances are currently used to improve the safety in a critical environment.

1. Introduction

The recent worldwide pandemic emergency raised attention on daily-life circumstances which previously were not a cause of concern. Among them, there are constraints on the physical distance between people as an effective measure to reduce the virus spread. However, beyond the enforcement of such rules, a critical issue for safety is to verify and quantify the actual compliance of people with these restrictions which indeed have a substantial impact on our social life [12]. To this end, a lot of solutions have been proposed [27]. However, cameras in video surveillance settings offer arguably a more viable infrastructure to control the so-called *Social Distancing* (SD).

Visual Social Distancing (VSD) [17] is a particular case of the SD estimation problem, in which the inter-personal

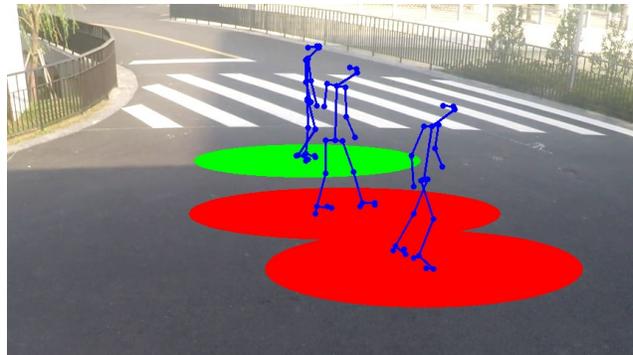


Figure 1. An example of the output of our proposed pipeline. The pedestrians are modeled as skeletal figures and each disc on the ground plane represents the minimum allowed inter-personal distance. The violation of social distancing is shown in red.

distance is estimated from a single uncalibrated image or video. VSD solutions use camera networks that are often deployed in pre-existing video surveillance settings. This allows fast integration of VSD to increase the safety of the population, by detecting recurrent SD violations or by generating statistic analysis. This information can be used to identify areas that are subject to crowding in order to redefine and improve safety in public and private places.

The study in [17], although inspiring, does not propose a methodology for solving VSD, but rather provides an overview of related issues as well as a few potential guidelines to implement solutions. Differently we propose a practical method that estimates a safe space for each detected person (i.e. person proxemics) on uncalibrated images, and detects violations of SD rules by measuring overlapping of safe spaces across people in the image scene (see Fig. 1).

Estimating VSD requires the computation of metric inter-personal distances (e.g. $\leq 2m$) between detected people in the scene. A convenient region to estimate distances is the ground plane where people stand, strengthened by

the fact that locally it can be considered as planar. We can therefore estimate a local planar homography between the ground plane and the image plane and relate them to a known metric reference.

Automatic homography estimation or camera calibration previously appeared as stand-alone problems in the computer vision literature [24, 49]. However, to the best of our knowledge, there is no prior work that proposes a solution for inter-personal distance estimation from a single uncalibrated image. Our method is a practical solution that takes advantage of the common installation scenario of surveillance cameras to perform both homography estimation and *local* metric inference. We exploit the standard surveillance settings, where people appear up-right, relying on the fact that surveillance cameras are typically placed at a certain height with a certain tilt angle from the ground which leads to approximately a uniform projection of lengths across cameras. We estimate the homography matrix between the bird’s eye view ground plane and image plane by intuitively selecting two auxiliary perspective ratios. Upon estimation of homography, people are detected and a local reference is inferred from them to estimate pairwise distances. With subjective evaluation, we demonstrate that auxiliary ratio parameters, even selected by human users, offer close proximity with the groundtruth. Moreover, we choose to use and validate body pose detectors since they offer richer information about the 2D body joints, whose length and relative position can be leveraged to provide a metric reference.¹

We have evaluated our method on VSD-adapted public datasets against baselines in order to validate the design choices of the approach. Additionally, as the proposed framework is currently being deployed in real-world environments, we show real statistics given by the system being deployed 24/7 in one of these scenarios to facilitate faster and better adaptation of spaces to new SD regulations. The main contributions of this work are:

- We propose the first² VSD method to identify the local human proxemics for automatic SD rule violation detection on single uncalibrated images.
- We propose a semi-automatic approach for homography approximation between ground plane and image plane by intuitively selecting two auxiliary ratios.
- We provide an analysis over human body joints to select the best body parts to be used as reference for scene metric inference.
- We validate various stages of our method over three public datasets that we labeled with pairwise human distance measures to adapt them for VSD analysis. Code is available at <https://github.com/>

¹Note that human body detection is not the focus of this work, and we rely on off-the-shelf methods (e.g. OpenPose [13]).

²To the authors knowledge, at the time of submission there were no previous peer-reviewed works addressing VSD with available implementations.

IIT-PAVIS/Social-Distancing.

- We present the advantages and challenges of deploying our proposed method deployed in real scenarios.

2. Related work

According to [17], estimating VSD, requires the solution of a set of computer vision tasks, namely, scene geometry understanding and person detection/body pose estimation. Indeed, the geometry of the scene is the first important step to define a local metric reference for measuring inter-personal distances. The second yet equally important task, is the detection of people in the scene in a possibly crowded environment. Once the target people are correctly localised in the scene, their distance can be locally estimated to understand if their mutual distance is lower than the predefined threshold (e.g. $\sim 2m$). In this work, our focus is on scene geometry understanding from monocular images, instead of using camera calibration, or other sensory inputs such as stereo images [62] or depth [60]. In this regard, the most similar work to us is Monoloco [9], in which authors propose a solution for 3D human localization from a monocular camera. The model behind Monoloco is a feed-forward neural network that predicts human 3D localization. Eventually, the VSD can be calculated from the 3D position of the people in the scene. Despite similarities between Monoloco and our proposed algorithm in simplicity and objectives, in contrast to Monoloco our proposed model does not require any camera calibration information (not even a default value) nor a training procedure, and it does not intend to solve the VSD problem globally through 3D localization of pedestrians but locally for each pair of people. In the following, as there is no further previous work on estimating VSD from a single uncalibrated image, we will describe the related work to the two main sub-tasks involved in the solutions for this problem.

Scene Geometry Understanding The task of measuring social distancing from images requires the definition of a metric reference. This problem is strongly related to the single view metrology topic [16] as we consider the most common case of a fixed camera. An initial solution for estimating inter-personal distances requires the identification of the ground plane where people stand [28, 25, 36, 1, 58]. Such ground plane serves in many video surveillance systems to visualise the scene as a bird’s eye view for ease of statistics representation. Most works impose the assumption that the ground plane is planar. Then, the problem is to estimate a homography given some reference elements (e.g., known objects or manual measurements) extracted from the scene or using the information of detected vanishing points in the image [16, 45, 40, 59, 5, 4, 63, 34]. Another common approach is to calibrate fixed cameras by observing the motion of dynamic objects such as pedestrians [30, 35, 50, 61]. As a unique approach, AutoRect [14] targets the restoration of

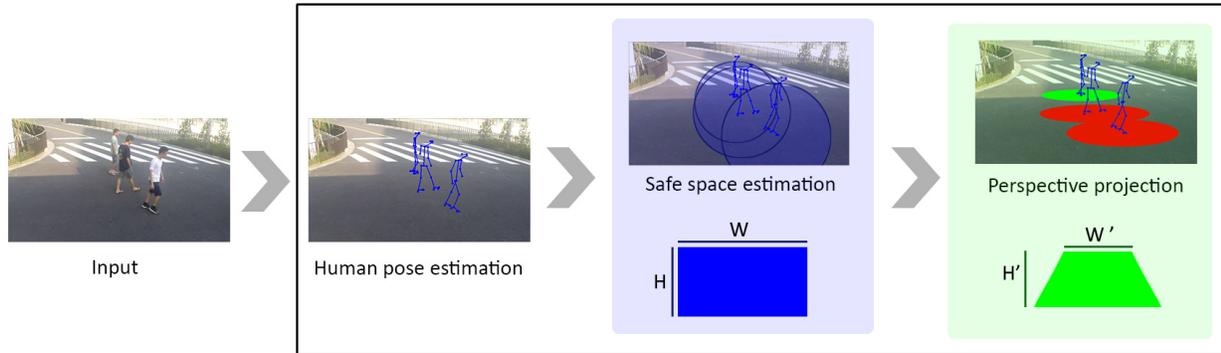


Figure 2. A schematic representation of our proposed pipeline for estimation of safe space and detection of its violation, comprised of three main computational stages: body-joint detection, body height estimation for metric inference, and ground-plane geometry estimation.

bird’s view homography from the perspective scene image by restoring the parallelism of lines by computing 2D homography from only two estimated vanishing points, allowing the algorithm to be used in more general scenarios. Furthermore, approaches based on deep learning also attempt at estimating directly camera pose and intrinsic parameters on a single image [26, 33]. In this work, we propose a simple yet effective approach that does not require exhaustive training procedures but only two auxiliary perspective ratios for approximating the homography matrix.

Nonetheless, VSD estimation in addition to camera intrinsic/extrinsic parameters estimation or ground plane detection requires a metric reference. Such information can be coarsely computed in the scene given objects of known dimension or by using a standardized height of pedestrians as a rule of thumb [55, 6, 54]. Similarly, in this work, we rely on a set of human body joints to infer the metric reference. It is important to note that a local solution for VSD is a simpler problem than global solutions which entail estimation of every metric distance among people at any position in the image. While the global solutions can be hard to estimate, for example, when the single ground plane assumption is violated, local solutions only require distance measure calculation when two or more pedestrians get close enough for triggering the necessity of a measure.

Person Detection and Pose Estimation Person detection has reached impressive performance in the last decade given the interest in the automotive industry and other application fields [7]. Real-time approaches can now estimate human pose in complex scenarios [13] and even reconstruct a 3D mesh of the person body [22]. The majority of the approaches estimate not only people’s location as a bounding box but also 2D stick-like figures, thus conveying a schematic representation of the pose. Recently, several methods augment 2D poses in 3D or infer directly a 3D pose in a normalised reference system [56, 41, 52, 39, 37, 10, 38, 44, 64].

Specific pedestrian detection techniques have been de-

signed to work in crowded scenes [53, 57, 32, 20], where skeleton-based representations are often dropped in favor of saliency-based masks, especially focusing on heads. When the image resolution becomes too low to spot single people, regression-based approaches are employed [11, 31, 46, 51], providing in some cases density measures [48, 43, 47]. This information, merged with a geometric model of the scene, can lead to a solution for measuring the average SD in the field of view. Recently, new efforts tend towards solving body detection and pose estimation in crowded environments [21, 29], the very same scenario that SD is trying to impede. Yet, finding the location of people in such cases is of relevant importance for buzzing alerts or creating statistics of overcrowded areas. To this end, the people detection module has to be robust to severe self and other occlusions, different image scales, and indoor/outdoor scenarios.

It is worth to highlight that, person detection in bounding box format does not account for body pose variations (*e.g.*, sitting, riding) and might negatively impact the estimation of height and cause wrong SD analysis. In this case, detecting body joints has certain advantages. This is due to the fact that for obtaining an approximation of the metric reference, or even calibrating the cameras, usually the person height is used as a coarse proxy as computed from a bounding box or using more precise techniques [55, 54, 18, 23]. Given a metric reference from scene geometry and detected people pose in the scene, the VSD can be solved as a distance on the ground plane among the detected pedestrians. As previously discussed, this information can be estimated locally or pairwise in order to reduce the complexity of estimating a global reference system for the whole image.

3. Method

The goal of our proposal is to detect SD violations by identifying the overlapping of the circular safe space around the pedestrians using images from any uncalibrated cameras. To achieve this goal, our pipeline is comprised of

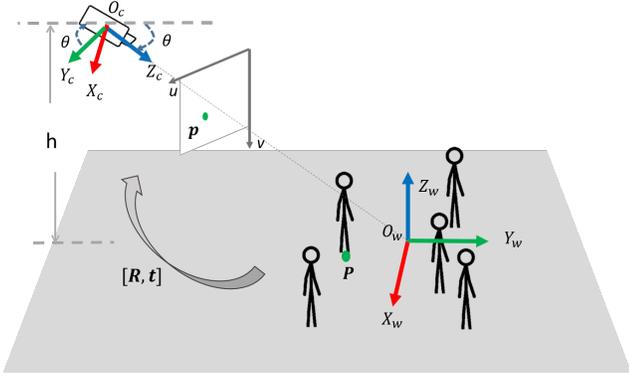


Figure 3. Illustration of the world and camera coordinate system under our assumptions that i) the camera only has a tilt angle θ with zero roll and pan angle, and ii) the camera has the height h from the ground and the world origin locates at where the camera principal axis intersect on the ground plane. The \mathbf{P} on the ground plane can be mapped to the pixel \mathbf{p} on the image through the homography matrix.

three main steps: the ground-plane homography estimation, person detection and localisation, and joint-based metric inference (see Fig. 2). Given any image captured by an uncalibrated camera, we first estimate the homography matrix between the ground plane and image plane (up to a scale) in a semi-automatic manner. We then detect people in the image using a human body-joint detector, e.g. OpenPose [13], to exploit the rich information embedded in the human pose. Finally, for each detected person, we analyse the body joints to robustly localise its ground-plane centroid on the image, and then to infer the safe space and detect the SD violation. In the following, we detail each of the three modules of our proposed method.

3.1. Homography Estimation

In this section, we cover the approach to estimate the homography matrix of any given image plane to the ground plane. Let us consider a pinhole camera model with the camera intrinsic matrix $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ and extrinsic matrix $[\mathbf{R}|\mathbf{t}]$. The intrinsic matrix \mathbf{K} contains the camera parameters that are related to the focal length, skew, and the principal point. The extrinsic matrix $[\mathbf{R}|\mathbf{t}]$ denotes the 3D coordinate transform from the world to the camera reference, where $\mathbf{R} \in SO(3)$ is the rotation matrix and $\mathbf{t} \in \mathbb{R}^{3 \times 1}$ is the translation vector. Any 3D point in the world coordinate \mathbf{P} can be projected to the image plane at pixel position \mathbf{p} as:

$$s\mathbf{p} = \mathbf{K}[\mathbf{R}|\mathbf{t}]\mathbf{P} = \mathbf{K} \begin{bmatrix} \mathbf{r}_1 & \mathbf{r}_2 & \mathbf{r}_3 & \mathbf{t} \end{bmatrix} \mathbf{P}, \quad (1)$$

where s is the scale factor and both $\mathbf{P} = [X, Y, Z, 1]^T$ and $\mathbf{p} = [u, v, 1]^T$ are in its homogeneous coordinate format.

In the case of projecting 3D points onto the ground plane,

i.e. $Z = 0$, to the image plane, Eq. 1 can be simplified as:

$$s\mathbf{p} = \mathbf{K} \begin{bmatrix} \mathbf{r}_1 & \mathbf{r}_2 & \mathbf{t} \end{bmatrix} \mathbf{P}' = \mathbf{H}\mathbf{P}', \quad (2)$$

where \mathbf{H} is the *homography matrix* to project points on the ground plane with metrics to the image plane and $\mathbf{P}' = [X, Y, 1]^T$. In our problem setting, both \mathbf{K} and $[\mathbf{R}|\mathbf{t}]$ are not available for the computation of homography matrix \mathbf{H} , we therefore approximate \mathbf{H} (up to a scale) with reasonable assumptions based on the standard scenario observed from surveillance cameras. One key observation is that people are appearing (almost) upright in the image plane, which means we can assume the roll angle to be 0° . Moreover, we can set the camera pan angle as 0° and the world origin at the intersection point of the camera principal axis on the ground plane with the height as h (as shown in Fig. 3). Let θ be the camera tilt angle, we will have $\mathbf{r}_1 = [1, 0, 0]^T$ and $\mathbf{r}_2 = [0, -\cos(\theta), -\sin(\theta)]^T$ and $\mathbf{t} = [0, -\frac{h}{\tan(\theta)}, h]^T$. Under such assumptions, lines that are parallel to the X-axis in the ground plane remain horizontal in the projected image plane, and lines parallel to the Y-axis converge to the vanishing point along Y-axis in the image plane. Therefore, a rectangular area on the ground plane with a width W and a height H would be projected into an isosceles trapezoidal shape with short-based width W' and height H' (see Fig. 2), where the horizontal ratio $\rho_h = \frac{W'}{W}$ and vertical ratio $\rho_v = \frac{H'}{H}$ are purely related to the camera tilt angle θ and the camera height h .

By manually tuning ρ_h and ρ_v , one can locate the corresponding four corners between the rectangular shape and the projected isosceles trapezoidal shape, in order to estimate the homography matrix $\tilde{\mathbf{H}}$ up to a scale. Note that $\tilde{\mathbf{H}}$ does not reflect the real metric mapping between pixel and length as \mathbf{K} is not known (for details regarding metric inference, please refer to Sec. 3.3). Approximated homography $\tilde{\mathbf{H}}$ however, can be used to project the circular safe space around each person, to form an approximation of an ellipse on the image plane in perspective. With subjective evaluation, we prove that human choice is reliable and can approach the optimal solution with grid search using the groundtruth inter-personal distances. Using these two ratios to facilitate the manual choice of the correspondence points, we can achieve a reasonable homography estimation when the camera calibration is not available.

3.2. Image-Plane Person Detection and Localisation

We detect people on the image using a human body detector, e.g. OpenPose [13], that provides the person detection as the set of detailed body joints. Let $\{\mathbf{D}_i\}_{i \in [1, N]}$ be the set of person detections, where N denote the total number of detected people on the image. Each detection \mathbf{D}_i is composed of a set of body joints, i.e. $\mathbf{D}_i = \{\mathbf{j}_i^k\}_{k \in [1, K]}$, and each joint corresponds to a position on the image plane

$\mathbf{j}_i^k = [u_i^k, v_i^k]$, where K is the number of joints that are supposed to be detected ($K = 25$ for OpenPose). Out of the detected joints \mathbf{D}_i , we obtain the bounding box $\mathbf{B}_i = [u_i^{min}, v_i^{min}, u_i^{max}, v_i^{max}]$ with $[u_i^{min}, v_i^{min}]$ being the minimum and $[u_i^{max}, v_i^{max}]$ being the maximum among all detected joints.

By localising a detected person, we mean to compute the pixel position $\mathbf{p}_i = [u_i, v_i]$ based on \mathbf{D}_i that represents the person on the ground plane. It is necessary to guarantee that \mathbf{p}_i locates on the ground plane as the homography matrix \mathbf{H} is estimated w.r.t it. To this end, we only consider person detections as valid if at least one joint of the feet is detected. For each valid detection, we decouple the estimation of u_i and v_i and make use of the rich information provided by the detected joints for a more robust localisation.

Regarding u_i , we exploit the fact that, for a standing-up or sitting person, certain part of human body, i.e. *head*, *neck*, or *torso* are often located on the middle line of the body. When any of the above mentioned joints are detected, we represent u_i by their averaged horizontal pixel position. If none of the aforementioned parts are detected, u_i is represented by the horizontal pixel position corresponding to the centroid of the bounding box \mathbf{B}_i , i.e. $u_i = \frac{u_i^{min} + u_i^{max}}{2}$. For v_i , as the valid detection inevitably contains joints corresponding to feet, we can represent it as the vertical pixel position averaged over all the detected joints of feet.

3.3. Metric Inference

The specified *safe space* around the detected person \mathbf{D}_i in the scene, is a circular area centered at the person with radius r , which is set according to the SD rules. Theoretically, the mapping from the length unit to the pixel unit on the ground plane can be obtained by calibrating the camera. However, when the calibration is not available, metric inference from pixel becomes a challenging problem to tackle. We, therefore, propose a practical metric inference solution that can be applicable to many surveillance scenarios, using the structural priors of the human body.

Studies on the average human body height over the years demonstrated that human body height has only marginal differences (a few centimeters) among different races [3]. In addition to body height, the length of body parts also maintain relatively stable attributes [42]. Thus, the human body potentially can be used to infer approximate length metrics for the image pixels. In this regard, given that we have estimated the homography between the ground plane and image plane, it would be a natural choice to use a set of human joints parallel to the ground plane, such as *shoulders* or *hips* for inferring the length metric. However, these joints are only reliable if the person appears perpendicular to the camera principal axis and any rotation of the body from this angle leads to a large variation onto the image plane.

Other joints such as *legs*, *arms*, or *torso*, are more robust

in maintaining a proper metric reference, but they lie mostly on the vertical plane that is orthogonal to the ground plane. This means that such metric references cannot be directly applied to measure distances on the ground plane. However, under most surveillance settings, where the camera is installed at a height and a tilt angle of a certain range, we can use the body parts on the vertical plane to practically approximate the metric on the ground plane. More specifically, onto the image plane, we consider that the safe space of a detected person is within a circle centered at \mathbf{p}_i with a radius r that can be set based on the pixel length of those vertical body parts. As validated in the experimental section, *torso* proves to be the most robust option for metric reference estimation among different body parts. The circle is then projected with the homography matrix into an approximated ellipse shape. The violation of SD is detected if the ellipse of a detected person overlaps with any ellipse of other detected people, leading to an $SD \leq 2m$.

4. Experiments

4.1. Evaluation Datasets

For quantitative evaluation, we use real recordings of people in both indoor and outdoor scenes from publicly available datasets. In these datasets, the 3D location of people on the ground plane can be estimated using either the homography matrix or the given camera intrinsic and extrinsic. In particular, we prepared three VSD-related datasets, *Epfl-Mpv-VSD*, *Epfl-Wildtrack-VSD*, and *OxTown-VSD*, out of the publicly available Epfl-Mpv [19], Epfl-Wildtrack [15], and OxTown [8], respectively. In all these datasets, for each frame we compute the groundtruth as the pairwise inter-person distance across people using their 3D location in the scene as provided by the benchmark.

We make use of the Epfl-MPV dataset as a dataset that covers an indoor environment. Our Epfl-Mpv-VSD uses only the lab sequences of Epfl-MPV where there are 4 sequences recorded by 4 cameras that cover a lab area with up to 6 people moving freely inside. Each sequence has 2955 frames. Epfl-Wildtrack-VSD is composed of seven synchronised sequences from the Epfl Wildtrack HD dataset served for multi-camera multi-people tracking. All sequences have a large overlapping field of view covering an open outdoor area, and each of them is composed of 401 frames. OxTown-VSD makes use of the OxTown sequence of 7498 frames, that covers a daily street in Oxford from a single static camera. Although the two datasets are featuring outdoor environments, the camera settings are similar to indoor surveillance cameras, therefore applicable to our method evaluation.

We applied OpenPose [13] to all frames within the above-mentioned datasets to detect people in the format of 25 skeletal joints. A detection is considered to be valid only



Figure 4. Image samples from the three datasets for the method evaluation with VSD results using the best parameters overlaid. Samples from left-to-right are from Epfl-Mpv-VSD, Epfl-Wildtrack-VSD, and OxTown-VSD.

if at least one joint from the feet is detected and at least a total of 13 joints of the person are detected. For each frame, the groundtruth pairwise inter-person distance is computed among valid detections using the homography matrix. For both Epfl-Wildtrack-VSD and OxTown-VSD, we make use of the camera intrinsic and extrinsic to compute the homography matrix as in Eq. 2. For Epfl-Mpv-VSD, the homography matrix is provided by the original dataset up to a scale, where the scale is inferred using a generic dimension of the objects in the scene, such as door width, and rug size. All distances measured are given in meters.

4.2. Evaluation Protocol

In order to benchmark the proposed VSD method, we only consider the detected people by the pose detector in the groundtruth, disregarding the missed detections by the detector from the entire analysis and thus disentangling the VSD performance from the performance of the detector³

We measure the performance of the algorithm by evaluating its ability to successfully report violations of the SD, namely, considering VSD as a binary classification problem. We report the results of all experiments within the standard metrics used in binary classification evaluation as $Precision = \frac{TP}{TP+FP}$, $Recall = \frac{TP}{TP+FN}$, and $F1 - score = 2 * \frac{Precision * Recall}{Precision + Recall}$. In this experimental evaluation, we considered a violation, if the distance between two people in the groundtruth has a value of less than $2m$. In the case, where our VSD algorithm reports that the safe space of a person collides with the safe space of another person in the image, then they are considered as TP cases. Note that FN here means the missed violation of safe distance by the VSD algorithm, rather than by the detector⁴.

Subjective evaluation on homography estimation: We have mentioned earlier that our proposed solution is most useful when camera calibration is not available. However, the question might be “*how humans would perform for tuning ρ_v and ρ_h , relying purely on their intuitive geometrical*

³For the sake of future comparisons, we will release body pose detection results on our dataset as part of the groundtruth.

⁴In the main paper we only report the F1-score results, and show the corresponding Precision and Recall values in the Supplementary Material.

understanding of the scene?”. To answer this question, we formulated a subjective experiment in which we asked 10 human subjects to tune ρ_v and ρ_h parameters till the visual feedback, in the form of projected circles centered at each detected pedestrian onto the image, is reasonable for them. Each subject was asked to tune these parameters once for every single sequence available in each dataset. For the pair of ρ_v and ρ_h selected by each subject for each sequence, we computed the $F1 - score$ measure for that sequence and report it in the form of *average \pm std* among all the subjects, in the Table 1, ‘Human estimated’ column.

Homography parameters tuning with grid search: This experiment aims to answer “*what is the best achievable performance using our proposed algorithm?*”. To this end, we tune the optimal ρ_v and ρ_h values by grid searching separately over each sequence. We discretise the search space of ρ_v and $\rho_h \in [0, 1]$ with a step 0.1. For each sequence, the combination of ρ_v and ρ_h that produces the best performance is then selected as the optimal parameters. In Table 1, ‘Grid search’ column, we report the quantitative results per each sequence using the optimal parameters.

Comparison with automatic homography estimation: Last, we benchmark the proposed VSD pipeline against two methods. AutoRect [14] provides a model for single image automatic homography estimation. Monoloco [9] is a method for monocular 3D human localization from a single image. To compare our proposed model with AutoRect, the circular safe space around each detected subject on the scene is projected onto the ground plane through an estimate of \mathbf{H} , computed by evaluating vanishing points and projecting them to infinity, while the \mathbf{H} is estimated real time for each frame of the sequence. The results of this experiment are listed in Table 1, column ‘AutoRect \mathbf{H} ’. In the last column of Table 1, the results of applying Monoloco on each sequence of the datasets are reported. We show results using the ‘default’ camera intrinsic matrix \mathbf{K} provided by the Monoloco algorithm, and the groundtruth \mathbf{K} when available (Epfl-Wildtrack-VSD and OxTown-VSD). Note that all the experiments in Table 1, except Monoloco, the metric reference is set to the ‘Torso’ for coherence.

Ablation study on metric reference: As discussed in

Table 1. Investigating the estimation of the Homography matrix \mathbf{H}

Dataset	Seq.	Proposed VSD				AutoRect \mathbf{H} [14]	Monoloco [9] (+ camera intrinsics)
		Human estimated		Grid search			
		F1-score	ρ_h	ρ_v	F1-score	F1-score	F1-score
EPFL-MPV	C0	77.43 \pm 1.12	0.7	0.5	77.90	73.47	73.46 (N.A.)
	C1	71.24 \pm 4.92	0.5	0.6	75.39	61.17	70.19 (N.A.)
	C2	75.69 \pm 3.87	0.6	0.6	78.67	78.14	77.20 (N.A.)
	C3	73.34 \pm 3.51	0.5	0.6	75.86	58.36	72.60 (N.A.)
	Avg.	74.42 \pm 3.35	-	-	76.94	67.75	73.36 (N.A.)
EPFL-wildtrack	C1	84.78 \pm 1.11	0.8	0.7	86.31	61.80	70.07 (70.50)
	C2	83.84 \pm 1.19	0.8	0.6	85.57	57.27	68.31 (68.11)
	C3	85.52 \pm 2.08	0.8	0.7	87.96	45.21	66.73 (67.10)
	C4	81.10 \pm 4.03	0.6	0.8	85.54	35.06	64.29 (65.13)
	C5	67.63 \pm 3.63	0.8	0.8	69.91	50.99	54.80 (54.69)
	C6	62.87 \pm 2.23	0.8	0.8	65.27	39.54	49.64 (49.58)
	C7	84.46 \pm 3.51	0.5	0.7	86.96	55.63	68.44 (68.74)
	Avg.	78.60 \pm 2.54	-	-	81.07	49.77	63.18 (63.40)
OxTown	-	76.94 \pm 4.52	0.5	0.8	81.04	51.78	54.57 (54.54)

Table 2. Investigating metric reference choice, using the fixed homography-related parameters (the best choice among subjects).

Dataset	Seq.	Ratios		Body parts			BBX height (1.7 m)
				Leg length (0.85 m)	Arm length (0.70 m)	Torso length (0.60 m)	
		ρ_h	ρ_v	F1-score	F1-score	F1-score	F1-score
PFL-MPV	C0	0.6	0.5	77.14	74.68	77.64	76.71
	C1	0.5	0.6	74.80	71.19	75.38	73.90
	C2	0.8	0.5	78.43	70.02	77.12	78.63
	C3	0.8	0.5	73.41	66.38	71.19	74.69
	Avg.	-	-	75.94	70.56	75.98	75.96
EPFL-wildtrack	C1	0.5	0.8	83.93	81.46	84.10	79.18
	C2	0.7	0.5	81.07	80.47	82.64	81.18
	C3	0.5	0.8	83.57	81.00	85.37	79.13
	C4	0.6	0.8	85.95	73.30	86.68	81.52
	C5	0.8	0.8	68.40	72.58	69.82	64.65
	C6	0.8	0.7	62.98	63.53	63.72	59.73
	C7	0.6	0.8	87.20	76.95	84.03	86.62
	Avg.	-	-	79.01	75.61	79.48	76.00
OxTown	-	0.5	0.8	82.59	73.03	81.04	72.38

Sec. 3.3, pose detectors facilitate a better estimation of metric reference with regards to human body detection as a bounding box. In Table 2, ‘BBX height’ column we report the results of our proposed algorithm using the height of the bounding box (obtained using OpenPose joints as detailed in Sec. 3.2). As inferring the length of any of the body parts from a single bounding box is implausible, we consider the height of the detected bounding box as the average height of a person (1.70m) and from there we estimate the radius r of the safe space. Moreover, we also validate metric references using various body parts, namely, ‘Leg’, ‘Arm’, and ‘Torso’. More details on the actual OpenPose joints used to infer such pixel lengths can be found in the Supplementary Material. The obtained results from this experiment can be found in Table 2 in the respective columns. Each body part

is associated with a metric reference according to population statistics [42, 2], as detailed in the table. For the sake of fair comparisons, ρ_v and ρ_h are kept the same for this ablation study according to the best subjective estimation.

4.3. Experimental Results Discussion

As it can be extrapolated from Table 1, the best average obtained results for all datasets is achieved by grid-searching the best parameters. Our original proposal to rely on human geometrical understanding to estimate ρ_v and ρ_h (Table 1, ‘Human estimated’ column) performs only slightly lower than the best choice of parameters (Table 1, ‘Grid search’ column). Moreover, the relatively low *std* of the results in most of the sequences is an indication of the inter-raters agreement on the *correctness* measures. Besides, the

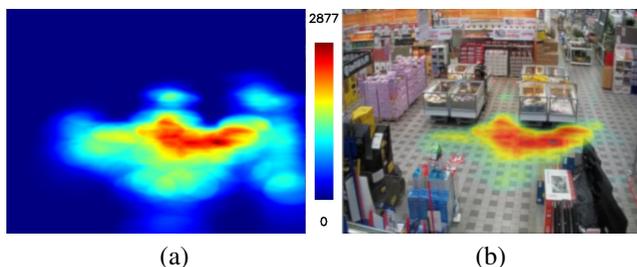


Figure 5. The heatmap in the images represents the areas on which SD violations mostly occur. (a) the maximum number of observed violations is depicted in red and overlaid in the image frame in (b). The heatmap is produced using observations collected in 1-hour from a monitored retail shop. The area with the most observed violations is due to an exhibit placed in the main entrance of the shop that causes increasing SD violations.

proposed VSD algorithm, either requiring manual tuning of parameters or validating them, on average per dataset, performs better than applying automatic estimation of \mathbf{H} using AutoRect. Monoloco provides instead the 3D localization starting from pose detections and the matrix for the intrinsic camera parameters \mathbf{K} . If \mathbf{K} Monoloco has a default value. It can be seen that using such a default \mathbf{K} or the groundtruth one does not lead to relevant differences, indicating that Monoloco is not very sensitive to it in the evaluated datasets. Instead, the Monoloco network is probably very biased towards its training data since our proposed algorithm outperforms it in all the sequences (although it outperforms AutoRect).

Concerning Table 2, it can be seen that on average using bounding boxes as the metric reference is not a reliable option. This observation emphasizes the importance of having information regarding human body joints instead of a mere bounding box, as already discussed. The last observation regards the choice of different body parts, namely 'Leg length', 'Arm length', or 'Torso length' for inferring distances in the image. We observe that torso provides a more robust reference, in general performing better than arm and leg. This is reasonable since, in normal conditions in which people are walking in the scene, legs and arms can undergo various deformations due to the multiple numbers of joints in them, while the torso is merely one straight line which connects the neck to the hips. For OxTown only, the leg reference performs slightly better than the torso one. Arm is not a good reference body part as it often undergoes occlusions with torso, preventing from a correct measurement of the metric reference. Additional qualitative results in the form of videos are provided in the supplementary material.

4.4. Real-world Scenario Challenges and Next Steps

Our proposed setup is currently able to monitor up to 18 cameras with 640x480 resolution, running on a sin-

gle Alienware Aurora R8 PC equipped with an NVidia RTX2080 GPU at 1 fps. The system has been successfully applied for monitoring social distancing 24/7 in challenging scenarios such as corporate offices, airports, and shopping malls, where more than 100,000 observations per hour are being evaluated. The collected statistics are used for both generating real time alarms finalised to break apart big groups of persons and for generating statistics in order to achieve effective space redesigns (see Fig. 5).

The collection of this huge stream of data on real application scenarios allows us to highlight the most important open issues that future developments should focus on:

- System alarms for the violations of safe spaces should not be based only on single frame analysis, but rather encompass temporal information and people tracking.
- Currently many of the SD laws allow family members to stay close to each other. A robust system ideally should be able to recognise family member formations and discard their violations from the computed ones.
- Our approach is frame-based and it can be applied to any kind of camera under certain setting constraints. In the case of fixed cameras, a statistical approach can be used for inferring metric references. This statistical approach would allow the system to be more robust to outliers (such as kids on the scene).

5. Conclusion

We provided the first² VSD method for estimating SD from single uncalibrated images together with a augmented VSD annotations for three public datasets for further VSD evaluation purposes. The proposed method performs favorably against competing baselines, showing the importance of using a pose estimator instead of a bounding box detector. This allows exploiting human body parts as robust metric references; in particular, we find the length of the torso as the most reliable. In addition, we demonstrated the competence of the proposed semi-automatic homography estimation against automatic estimation using state-of-the-art. Our proposed method might fail when the underlying pose detector fails to properly detect the body joints due to its strong dependence on the joint positional information. The promising performance of the proposed method is also guaranteed under certain camera positioning which limits its applicability to surveillance settings in particular. In the same line of reasoning, our semi-automatic method is only capable of estimating one homography per scene which leads to deteriorations in its performance when multiple homographies are present in the scene. Future work will provide a more in-depth analysis of the causes of SD violations (e.g. presence of families, use of personal protective equipment) and the development of forecasting methods to avoid the formation of gatherings.

References

- [1] S. A. Abbas and A. Zisserman. A geometric approach to obtain a bird's eye view from an image. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 4095–4104, 2019.
- [2] National Aeronautics and Space Administration. Man systems integration standards. Technical report, 1987.
- [3] Joerg Baten and Matthias Blum. Human height since 1820. 2014.
- [4] Jean-Charles Bazin and Marc Pollefeys. 3-line ransac for orthogonal vanishing point detection. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4282–4287. IEEE, 2012.
- [5] Jean-Charles Bazin, Yongduek Seo, Cédric Demonceaux, Pascal Vasseur, Katsushi Ikeuchi, Inso Kweon, and Marc Pollefeys. Globally optimal line clustering and vanishing point estimation in manhattan world. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 638–645. IEEE, 2012.
- [6] Chiraz BenAbdelkader and Yaser Yacoob. Statistical body height estimation from a single image. In *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition*, pages 1–7. IEEE, 2008.
- [7] Rodrigo Benenson, Mohamed Omran, Jan Hosang, and Bernt Schiele. Ten years of pedestrian detection, what have we learned? In *European Conference on Computer Vision*, pages 613–627. Springer, 2014.
- [8] Ben Benfold and Ian Reid. Stable multi-target tracking in real-time surveillance video. In *CVPR*, pages 3457–3464, June 2011.
- [9] Lorenzo Bertoni, Sven Kreiss, and Alexandre Alahi. Monoloco: Monocular 3d pedestrian localization and uncertainty estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6861–6871, 2019.
- [10] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016.
- [11] Lokesh Boominathan, Srinivas SS Kruthiventi, and R Venkatesh Babu. Crowdnet: A deep convolutional network for dense crowd counting. In *ACM International Conference on Multimedia*, pages 640–644, 2016.
- [12] Samantha K Brooks, Rebecca K Webster, Louise E Smith, Lisa Woodland, Simon Wessely, Neil Greenberg, and Gideon James Rubin. The psychological impact of quarantine and how to reduce it: rapid review of the evidence. *The Lancet*, 2020.
- [13] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [14] Krishnendu Chaudhury, Stephen DiVerdi, and Sergey Ioffe. Auto-rectification of user photos. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 3479–3483. IEEE, 2014.
- [15] Tatjana Chavdarova, Pierre Baqué, Andrii Maksai, Stéphane Bouquet, Cijo Jose, Louis Lettry, Francois Fleuret, Pascal Fua, and Luc Van Gool. Wildtrack: A multi-camera hd dataset for dense unscripted pedestrian detection. In *Proc. of IEEE/CVF Conference On Computer Vision And Pattern Recognition*, pages 5030–5039, New York, 2018. IEEE.
- [16] Antonio Criminisi, Ian Reid, and Andrew Zisserman. Single view metrology. *International Journal of Computer Vision*, 40(2):123–148, 2000.
- [17] M. Cristani, A. D. Bue, V. Murino, F. Setti, and A. Vinciarelli. The visual social distancing problem. *IEEE Access*, 8:126876–126886, 2020.
- [18] Ratan Dey, Madhurya Nangia, Keith W Ross, and Yong Liu. Estimating heights from photo collections: A data-driven approach. In *Proceedings of the second ACM conference on Online social networks*, pages 227–238, 2014.
- [19] Francois Fleuret, Jerome Berclaz, Richard Lengagne, and Pascal Fua. Multicamera people tracking with a probabilistic occupancy map. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):267–282, Feb. 2008.
- [20] Zheng Ge, Zequn Jie, Xin Huang, Rong Xu, and Osamu Yoshie. Ps-rcnn: Detecting secondary human instances in a crowd via primary object suppression. *arXiv preprint arXiv:2003.07080*, 2020.
- [21] Thomas Golda, Tobias Kalb, Arne Schumann, and Jürgen Beyerer. Human pose estimation for real-world crowded scenarios. In *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–8. IEEE, 2019.
- [22] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018.
- [23] Semih Gunel, Helge Rhodin, and Pascal Fua. What face and body shapes can tell us about height. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- [24] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [25] Derek Hoiem, Alexei A Efros, and Martial Hebert. Recovering surface layout from an image. *International Journal of Computer Vision*, 75(1):151–172, 2007.
- [26] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gambaretto, Sunil Hadap, and Jean-François Lalonde. A perceptual measure for deep single image camera calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2354–2363, 2018.
- [27] Joseph Johnson Jr, Shiblee Hasan, David Lee, Chris Hluchan, and Nazia Ahmed. Social-distancing monitoring using portable electronic devices. 2020.
- [28] Lily Lee, Raquel Romano, and Gideon Stein. Monitoring activities from multiple video streams: Establishing a common coordinate frame. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):758–767, 2000.
- [29] Jiefeng Li, Can Wang, Hao Zhu, Yihuan Mao, Hao-Shu Fang, and Cewu Lu. Crowdpose: Efficient crowded scenes

- pose estimation and a new benchmark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10863–10872, 2019.
- [30] Jingchen Liu, Robert T Collins, and Yanxi Liu. Surveillance camera autocalibration based on pedestrian height distributions. In *Proceedings of the British Machine Vision Conference*, page 144, 2011.
- [31] Jiang Liu, Chenqiang Gao, Deyu Meng, and Alexander G Hauptmann. Decidenet: Counting varying density crowds through attention guided detection and density estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5197–5206, 2018.
- [32] Songtao Liu, Di Huang, and Yunhong Wang. Adaptive nms: Refining pedestrian detection in a crowd. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [33] Manuel López-Antequera, Roger Marí, Pau Gargallo, Yubin Kuang, Javier Gonzalez-Jimenez, and Gloria Haro. Deep single image camera calibration with radial distortion. In *Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [34] Xiaohu Lu, Jian Yaoy, Haoang Li, and Yahui Liu. 2-line exhaustive searching for real-time vanishing point estimation in manhattan world. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 345–353. IEEE, 2017.
- [35] Fengjun Lv, Tao Zhao, and Ramakant Nevatia. Camera calibration from video of a walking human. *IEEE transactions on pattern analysis and machine intelligence*, 28(9):1513–1518, 2006.
- [36] Yunze Man, Xinshuo Weng, Xi Li, and Kris Kitani. Groundnet: Monocular ground plane normal estimation with geometric consistency. In *Proceedings of the 27th ACM International Conference on Multimedia*, page 2170–2178, 2019.
- [37] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2640–2649, 2017.
- [38] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Mohamed Elgharib, Pascal Fua, Hans-Peter Seidel, Helge Rhodin, Gerard Pons-Moll, and Christian Theobalt. Xnect: Real-time multi-person 3d human pose estimation with a single rgb camera. *arXiv preprint arXiv:1907.00837*, 2019.
- [39] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 36(4):1–14, 2017.
- [40] Faraz M Mirzaei and Stergios I Roumeliotis. Optimal estimation of vanishing points in a manhattan world. In *2011 International Conference on Computer Vision*, pages 2454–2461. IEEE, 2011.
- [41] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2823–2832, 2017.
- [42] Vitor Murtinho. Leonardo’s vitruvian man drawing: a new interpretation looking at leonardo’s geometric constructions. *Nexus Network Journal*, 17(2):507–524, 2015.
- [43] A Rangel-Huerta, AL Ballinas-Hernández, and A Muñoz-Meléndez. An entropy model to measure heterogeneity of pedestrian crowds using self-propelled agents. *Physica A: Statistical Mechanics and its Applications*, 473:213–224, 2017.
- [44] Grégory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net++: Multi-person 2d and 3d pose detection in natural images. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [45] Carsten Rother. A new approach to vanishing point detection in architectural environments. *Image and Vision Computing*, 20(9-10):647–655, 2002.
- [46] Francesco Setti, Davide Conigliaro, Michele Tobanelli, and Marco Cristani. Count on me: learning to count on a single image. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(8):1798–1806, 2018.
- [47] Vishwanath A Sindagi and Vishal M Patel. Cnn-based cascaded multi-task learning of high-level prior and density estimation for crowd counting. In *IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2017.
- [48] Vishwanath A Sindagi and Vishal M Patel. A survey of recent advances in cnn-based single image crowd counting and density estimation. *Pattern Recognition Letters*, 107:3–16, 2018.
- [49] R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer International Publishing, 2011.
- [50] Zheng Tang, Yen-Shuo Lin, Kuan-Hui Lee, Jenq-Neng Hwang, and Jen-Hui Chuang. Esther: Joint camera self-calibration and automatic radial distortion correction from tracking of walking humans. *IEEE Access*, 7:10754–10766, 2019.
- [51] Y. Tian, Y. Lei, J. Zhang, and J. Z. Wang. Padnet: Pandensity crowd counting. *IEEE Transactions on Image Processing*, 29:2714–2727, 2020.
- [52] Denis Tome, Chris Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2500–2509, 2017.
- [53] Jennifer Vandoni, Emanuel Aldea, and Sylvie Le Hégarat-Masclé. Evidential query-by-committee active learning for pedestrian detection in high-density crowds. *International Journal of Approximate Reasoning*, 104:166–184, 2019.
- [54] Johan Vester. Estimating the height of an unknown object in a 2d image. *KTH CSIC*, 2012.
- [55] Robert Wagner, Daniel Crispell, Patrick Feeney, and Joe Mundy. 4-d scene alignment in surveillance video. *arXiv preprint arXiv:1906.01675*, 2019.
- [56] Chunyu Wang, Yizhou Wang, Zhouchen Lin, Alan L Yuille, and Wen Gao. Robust estimation of 3d human poses from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2361–2368, 2014.

- [57] Xinlong Wang, Tete Xiao, Yuning Jiang, Shuai Shao, Jian Sun, and Chunhua Shen. Repulsion loss: Detecting pedestrians in a crowd. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [58] Jamie Watson, Michael Firman, Aron Monszpart, and Gabriel J. Brostow. Footprints and free space from a single color image. In *Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [59] Horst Wildenauer and Allan Hanbury. Robust camera self-calibration from monocular images of manhattan worlds. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2831–2838. IEEE, 2012.
- [60] L. Xia, C. Chen, and J. K. Aggarwal. Human detection using depth information by kinect. In *Proc. of Computer Vision and Pattern Recognition Workshops*, pages 15–22, 2011.
- [61] Yan Xu, Vivek Roy, and Kris Kitani. Estimating 3d camera pose from 2d pedestrian trajectories. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2568–2577. IEEE, 2020.
- [62] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. In *Proc. of International Conference on Learning Representations*, 2020.
- [63] Lilian Zhang, Huimin Lu, Xiaoping Hu, and Reinhard Koch. Vanishing point estimation and line classification in a manhattan world with a unifying camera model. *International Journal of Computer Vision*, 117(2):111–130, 2016.
- [64] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 398–407, 2017.