

# NurtureNet: A Multi-task Video-based Approach for Newborn Anthropometry

Yash Khandelwal<sup>1†</sup>    Mayur Arvind<sup>1</sup>    Sriram Kumar<sup>1</sup>    Ashish Gupta<sup>1</sup>  
 Sachin Kumar Danisetty<sup>1</sup>    Piyush Bagad<sup>2</sup>    Anish Madan<sup>2</sup>    Mayank Lunayach<sup>2</sup>  
 Aditya Annavajjala<sup>2</sup>    Abhishek Maiti<sup>2</sup>    Sansiddh Jain<sup>2</sup>    Aman Dalmia<sup>2</sup>  
 Namrata Deka<sup>2</sup>    Jerome White<sup>2</sup>    Jigar Doshi<sup>2</sup>    Angjoo Kanazawa<sup>3</sup>  
 Rahul Panicker<sup>2</sup>    Alpan Raval<sup>1</sup>    Srinivas Rana<sup>1</sup>    Makarand Tapaswi<sup>1†</sup>

Wadhvani Institute for Artificial Intelligence (WIAI)

<sup>1</sup>currently at WIAI, <sup>2</sup>work done while at WIAI; <sup>3</sup>UC Berkeley

† {yash, makarand}@wadhvaniai.org

## Abstract

Malnutrition among newborns is a top public health concern in developing countries. Identification and subsequent growth monitoring are key to successful interventions. However, this is challenging in rural communities where health systems tend to be inaccessible and under-equipped, with poor adherence to protocol. Our goal is to equip health workers and public health systems with a solution for contactless newborn anthropometry in the community.

We propose NurtureNet, a multi-task model that fuses visual information (a video taken with a low-cost smartphone) with tabular inputs to regress multiple anthropometry estimates including weight, length, head circumference, and chest circumference. We show that visual proxy tasks of segmentation and keypoint prediction further improve performance. We establish the efficacy of the model through several experiments and achieve a relative error of 3.9% and mean absolute error of 114.3 g for weight estimation. Model compression to 15 MB also allows offline deployment to low-cost smartphones.

## 1. Introduction

The first 4 weeks of life are critical for a newborn’s physiological and neurological development. Conditions such as malnutrition and malabsorption during this phase lead to neonatal morbidities and in extreme cases even mortality. Thus, tracking a newborn’s growth over the first few weeks is an important public health responsibility [43].

The weight of a newborn is an important statistic that captures its overall health and well-being [4, 8, 13, 20]. Other measurements such as length, head circumference, and chest circumference are also useful for assessing growth or related developmental disorders [3, 63]. However, there




Figure 1. Illustration contrasting traditional approaches (a-c) for newborn anthropometry to what our proposed solution (d) enables. (a) A measuring tape is used to measure head and chest circumference. (b) An infantometer is used to capture length. (c) The newborn is suspended from a cloth and hooked up to a spring balance to measure weight. (d) Our proposed solution replaces all the above tasks and only requires the data collector to take a short video with a low-cost smartphone.

are several challenges in accurately capturing it in low- and middle-income countries (LMICs).

As seen in Fig. 1(c), traditional methods for measuring weight in community settings use a spring balance (least count 100 g), from which the newborn is suspended. This results in two main sources of error: (i) Human factors: the panicked mother supporting her baby from the bottom; motion of the spring balance as the newborn moves; difficulty in ascertaining the reading due to parallax; cultural challenges such as reluctance towards “outsiders” handling babies; and data handling malpractices leading to reporting

challenges. (ii) Instrument factors: old machines whose springs are no longer taut result in positive errors (over-prediction); poorly calibrated or uncertified instruments; and even unavailability of the instrument due to supply chain issues. Similar challenges also apply to other anthropometric measurements such as newborn length or head and chest circumference.

There are also logistical factors at play. Rural communities may be several miles away from health facilities with poor mobility options and limited inter-connectivity. Poverty further affects their ability to avail health facilities. Geographical barriers like rough terrain or rivers, and seasonal challenges such as extreme heat and heavy rain make it challenging for both, families (with newborns) to reach health centers and for health workers (carrying heavy instruments) to visit rural communities.

Our goal is to develop a contactless, geo-tagged, easy-to-use solution that provides accurate anthropometry estimates for a newborn (age 0-42 days). We wish to leverage the proliferation of mobile phone adoption in rural areas of LMICs enabling AI technologies to improve the daily activities of frontline health workers while ensuring automatic reporting for timely public health response and policy formulation. Our technology is suitable primarily in rural community settings over health facilities that may have good instrumentation and well-trained staff. In such rural settings, there are about 1 million health workers in our country, making this solution amenable for large-scale impact.

To facilitate widespread adoption, we make several design choices. (i) We restrict ourselves to RGB videos captured on low-cost mobile devices and forego complex depth sensors that may curtail adoption in rural areas. (ii) A reference object is needed to provide a sense of metric scale and we use easily available wooden rulers instead of chessboards. (iii) We develop a simple protocol for capturing the video that enables viewing the newborn from multiple angles without the need for a dedicated video capture setup or specialized hardware. (iv) We restrict modeling to simple architectures that can be compressed and deployed on a low-cost smartphone to enable offline inference.

We present a CNN-based model that can ingest video frames, aggregate their information, augment it with relevant tabular inputs, and estimate the newborn’s weight. This weight estimation model is extended to estimate several anthropometric measurements through a multi-task setup (NurtureNet). We also propose proxy tasks such as baby segmentation and keypoint estimation that can assist the model in focusing on the baby’s shape and pose respectively, resulting in improved performance. Notably, building on state-of-the-art segmentation and keypoint estimation methods, we show that segmentation masks and keypoints need not be annotated for each frame, and pseudo-labels can be used instead.

Overall, our contributions can be summarized as follows: (i) We present a vision application for newborn anthropometry based on a standard RGB video captured with a low-cost smartphone that enables widespread adoption and impact in rural community settings. (ii) We propose and train a multi-task model, NurtureNet, that ingests the video and is assisted by tabular inputs to estimate several anthropometric measurements simultaneously. (iii) We show the benefits of modeling auxiliary vision tasks (*e.g.* segmentation and keypoints) with pseudo-labels for training anthropometry models; and (iv) We present thorough experiments to show the impact of various modeling ideas, including evaluation of compressed models for offline phone deployment.

## 2. Related Work

Vision techniques have been used for various applications in newborn and infant healthcare, specially for anthropometry.

**Computer vision for newborns** is used in applications such as heart rate monitoring [46], General Movement Assessment (GMA) for early detection of cerebral palsy [39, 47], postnatal age estimation [60], and even identification based on footprints [32]. Related to anthropometry, there is work on estimating birth weight using ultrasound videos prior to birth [44]. An infant’s length is estimated using easy-to-detect stickers or markers [58], children’s (aged 2-5 years) height based on point clouds [61], and even height for adults using multiple images and a large reference object [33]. However, the above methods require specialized hardware or equipment and are therefore not applicable in low resource areas. A different approach, Baby Naapp [11] aims to use vision tools to eliminate the need for manual transcription by capturing and analyzing images/videos of devices - spring balances for weight and measuring tapes for circumference. Closest to our work, single image based weight and height estimation is performed using CNN-based regression [50]. However, their goal is different from ours as they aim to clinically estimate birth weight by capturing images in controlled settings (hospital). We show that such an approach performs poorly in community settings where protocol adherence may be difficult.

**Pose estimation** for newborns or infants is a popular task. Tracking body movements for newborns is critical, and early detection of abnormalities can prevent long-term health effects [47]. Specifically, tracking baby pose over time is useful to perform GMA [51], indicative of conditions such as cerebral palsy, autism spectrum disorder, and Rett syndrome. Approaches for infant body pose estimation include handcrafted features such as histograms of 3D joints [38] or Random Ferns on depth images [15]. Depth-only videos have also been used along with CNNs to directly regress pose [41, 64]. Towards GMA, CNN-based pose regression models have been developed that work with

RGB or RGB-D data [1, 12, 29, 42]. Transformer models are making inroads in infant pose estimation with RGB images [6] or depth and pressure images [30]. As a proxy to pose estimation, body part segmentation may also be used to understand infant movement [65]. 3D parametric models and pose estimates are also used to estimate the height and weight of adults [59]. For anthropometry, we show that segmentation and keypoint detection are good proxy tasks that help the model focus on the newborn.

**3D parametric models for adults and infants.** The Skinned Multi-Person Linear (SMPL) [35] model has been wildly popular in modeling the 3D shape of a human and has stemmed a flurry of methods [7, 19, 22, 26, 27, 31, 62]. However, adult human models cannot be used directly for newborns, predominantly due to changes in body shape proportions [18, 48]. Hence, a Skinned Multi-Infant Linear (SMIL) model was proposed [16]. Unfortunately, the model does not fit well to our case as it is trained on European infants in the 2-4 months age range with a significantly higher weight distribution. In contrast, we are interested in anthropometry for newborns up to 42 days of age in LMICs.

**Tabular methods.** Contrary to vision-based approaches, tabular data in the form of electronic health records (EHR) are used for infant [23] and fetal weight [37] prediction. These methods use maternal attributes, economic factors, and other aspects related to the gestation period as predictive features [21, 28, 36, 49]. However, it can be hard to train or deploy models in regions where such records are inaccessible or not carefully curated. In our work, we use tabular data (birth weight and age) to augment and assist visual features. Importantly, we show that our vision model augmented with tabular inputs is robust to errors in the tabular data that may be common due to misreporting.

### 3. Method

We formulate anthropometry as a regression problem and introduce an end-to-end pipeline to estimate the weight  $w$ , length  $l$ , head circumference  $h$ , and chest circumference  $c$  of an infant with age  $a \in [0, 42]$  days. Our model ingests a video  $\mathcal{V}$  and is augmented by tabular information such as the birth weight  $w^0$  and the current age to regress:

$$[w, l, h, c] = f_{\theta}(\mathcal{V}, w^0, a), \quad (1)$$

where  $\theta$  are the model’s learnable parameters.

The visual component of the model learns an implicit shape representation through the fusion of multiple frames that capture the newborn from several angles (Sec. 3.1). Furthermore, we encourage the model to focus on the newborn by asking it to predict a segmentation mask and keypoints in a bootstrapped multi-task setting (Sec. 3.2). Finally, we show how the visual features can be augmented

with tabular data, resulting in significant improvements (Sec. 3.3). Fig. 2 illustrates the overall approach.

#### 3.1. Video-based Anthropometry

**How to record a video?** Before addressing modeling, we briefly talk about how we record the video  $\mathcal{V}$ . Predicting the metric shape of an entity using a monocular camera often requires a reference object. However, considering the large-scale rural use-case of our solution, we switch from the chessboard (♠) a classic and accurate reference object used for camera calibration [57]) to a wooden ruler (📏 length 30 cm) that is easily available to health workers. The newborn is laid on a bedsheet spread on a flat surface with a 📏 placed below the newborn (in the same plane). While we remove all clothes for the newborn, no specific instructions are provided for the bedsheet. The data collectors (or health workers) are trained to capture a video by starting from the top of the baby and making a smooth arc as illustrated in the supplement. We filter videos by quality to ensure the newborn and reference object are clearly visible for a majority of the video (details in the supplement).

**Video-based weight estimation.** Consider a video  $\mathcal{V} = [f_1, \dots, f_T]$  with  $T$  frames. We sample  $N < T$  frames from the video and pass them through a CNN backbone  $\phi(\cdot)$  to obtain frame-level representations  $\mathbf{x}_i = \phi(f_i)$  for each selected frame  $f_i$ , where  $\mathbf{x}_i \in \mathbb{R}^d$ . Our approach combines the individual frame-level representations via a pooling function  $\mathbf{z} = \rho(\{\mathbf{x}_i\}_{i=1}^N)$ , and is followed by a Multi-Layer Perceptron (MLP) with one hidden layer to estimate the weight  $w \in \mathbb{R}^1$  in kg.

We explore several pooling functions  $\rho(\cdot)$  ranging from average and max-pooling to complex ones involving vanilla attention [2]. Interestingly, given the nature of the problem, we find that a permutation invariant pooling (such as max-pooling) provides good results, as suggested by previous works on 3D shape [54]. We estimate the weight:

$$w = \text{MLP}_w(\max(\mathbf{x}_1, \dots, \mathbf{x}_N)) = \text{MLP}_w(\mathbf{z}). \quad (2)$$

Similar to work in action recognition [66], during training, we randomly select  $N$  frames from the video as a form of data augmentation, while during inference, we pick linearly spaced frames. The model parameters including the CNN backbone  $\phi(\cdot)$  are trained using the L1 loss:

$$L_w = |w - w_{\text{gt}}|, \quad (3)$$

where  $w_{\text{gt}}$  is the ground-truth weight of the newborn (in kg).

#### 3.2. Multi-task Learning

Multi-task learning generally leads to performance improvements when the tasks are related to each other [53].

**Anthropometric measurements.** Along with weight, we also predict other measurements such as length, head circumference, and chest circumference. Naturally, a taller

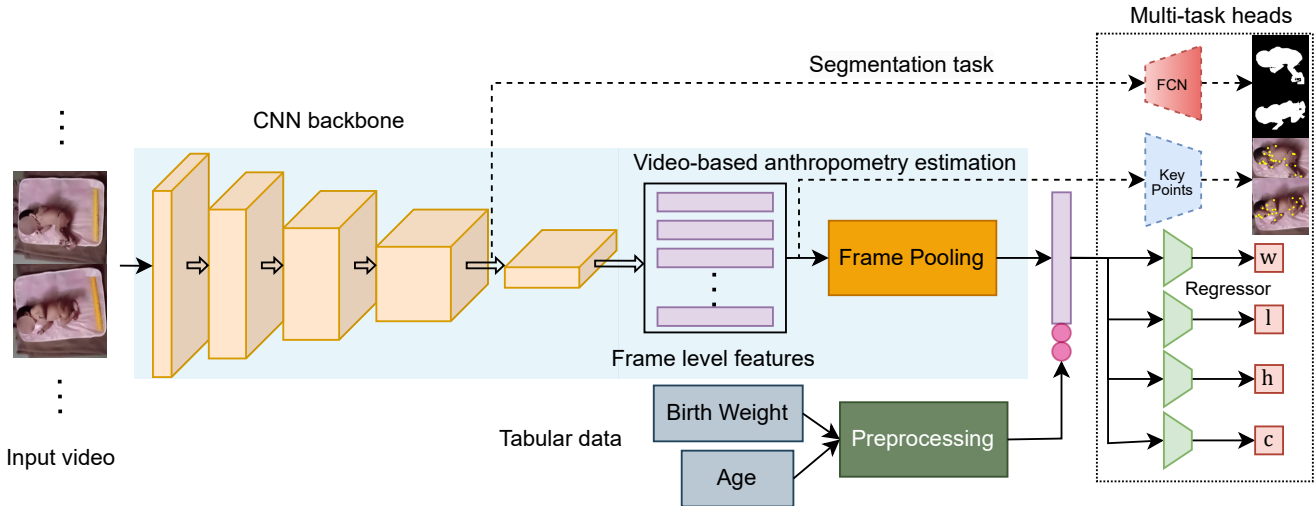


Figure 2. Overview of the proposed approach. Input video frames are sub-sampled and processed using a CNN and fused using a pooling module. Tabular data is normalized between  $[0, 1]$  and concatenated to this video representation. We use independent MLP regressors to predict anthropometry measures: weight, length, head circumference, and chest circumference. Additionally, we introduce two proxy tasks only used during training: newborn pixel segmentation predicted through an FCN head and keypoint estimation through a simple MLP.

$r(w, l)$	$r(w, h)$	$r(w, c)$	$r(l, h)$	$r(l, c)$	$r(h, c)$
0.7574	0.7176	0.7743	0.5747	0.5901	0.7340

Table 1. We observe high correlation (Pearson correlation coefficient) across anthropometric measurements on the training set: weight  $w$ , length  $l$ , head circumference  $h$ , chest circumference  $c$ .

baby is likely to be heavier, or a baby with a bigger chest may be better off with respect to nutrition. We compute the Pearson correlation coefficients across pairs of anthropometric measurements on our training set. As seen in Table 1, weight is strongly correlated with length, head circumference, and chest circumference.

We attach additional task heads, similar to the MLP used for weight estimation, to the pooled video representation. Specifically, we create three new task heads to predict each measurement in cm: length  $l = \text{MLP}_l(\mathbf{z})$ , head circumference  $h = \text{MLP}_h(\mathbf{z})$ , and chest circumference  $c = \text{MLP}_c(\mathbf{z})$ . The model is trained jointly to optimize:

$$L_{\text{anthro}} = \lambda_w L_w + \lambda_l L_l + \lambda_h L_h + \lambda_c L_c, \quad (4)$$

where  $L_l, L_h, L_c$  are L1 losses applied to length, head circumference, and chest circumference respectively; and  $\lambda_{(\cdot)}$  are loss weight coefficients.

**Visual prediction tasks.** While all the above tasks require ground-truth measurements to be collected at the time of video capture, we now present visual tasks that can be annotated post data collection. In particular, we consider pixel-level newborn segmentation and keypoint estimation, with the intent to encourage the model to learn representations that focus on the newborn.

While annotating segmentation masks or keypoints for

each frame of each video is possible, it is an expensive and time-consuming affair. We circumvent this through a bootstrapped approach. For baby segmentation masks, we finetune a PointRender segmentation model [25] on  $\sim 500$  videos with 10 linearly spaced frames from each. Similarly, for keypoints, we finetune HRNet [56] on  $\sim 1500$  videos with 20 linearly spaced frames from each. We apply both models to all video frames of the training set and use the predictions as pseudo-labels during multi-task training.

Our complete multi-task model has a segmentation head, a keypoint estimation head, and all the other anthropometric regression heads (see Fig. 2). We use a Fully Convolutional Network (FCN) head [34] to perform segmentation since we do not need fine precision. For keypoint estimation, we use a simple 2-layer MLP that regresses the spatial coordinates of keypoints from each frame embedding  $\mathbf{x}_l$ . The model is trained end-to-end through a combination of all losses:

$$L_{\text{total}} = L_{\text{anthro}} + \lambda_m \sum_l L_m(m'_l, \hat{m}_l) + \lambda_k \sum_l L_k(k'_l, \hat{k}_l), \quad (5)$$

where  $m'_l$  and  $k'_l$  are the frame-level segmentation mask and keypoints generated by our multi-task model,  $\hat{m}_l$  and  $\hat{k}_l$  are pseudo-labels for the mask and keypoints,  $L_m$  is Dice loss [55],  $L_k$  is L1 loss used for keypoints, and  $\lambda_m$  and  $\lambda_k$  are loss weights for masks and keypoints. During inference, we drop both the proxy heads.

### 3.3. Augmenting with Tabular Information

The weight of a baby reduces immediately after birth, recovers around days 7-10, and then follows a mostly linear growth trend [10]. Hence, knowing the birth weight  $w^0$  and current age  $a$  is often useful. We incorporate this meta in-

Source	Visits			Newborns		
	Train	Val	Test	Train	Val	Test
Region 1	8735	1096	1075	2304	293	280
Region 2	1590	185	220	447	51	64
Total	10325	1281	1295	2751	344	344

Table 2. Number of visits and newborns from rural home settings.

formation by normalizing them in the  $[0, 1]$  range and concatenating them to the output of the pooling layer. Our final weight regression head (similar to other measurements) is:

$$w = \text{MLP}_w([\mathbf{z}, w^0, a]). \quad (6)$$

We refer to this visual model augmented with tabular features as *W-NurtureNet* when used to only estimate weight and *NurtureNet* when used to estimate all anthropometric measures in a multi-task setting.

## 4. Experiments

We now present empirical validation of our approach on a large dataset collected in community deployment settings.

### 4.1. Setup

**Dataset collection.** The data has been collected by 28 trained personnel from rural home settings across 2 geographically diverse regions. They typically use Android smartphones with a 2-5 MP camera, with cost under \$150.

Our dataset consists of 3439 newborns that are visited on average 3.75 times in the first 42 days of life. At each visit, we capture three videos with different reference objects (only one is used for one experiment). While each visit is treated independently in the context of training and evaluation, all visits of a newborn are in the same split. Ethics committee approvals were obtained prior to data collection and the process itself involves taking informed consent, capturing videos, measuring the weight and anthropometry for the newborn, and providing home-based care recommendations if the newborn is not doing well. See supplement for more details.

We split the dataset into train (80%), validation (10%) and test (10%), while ensuring all visits of a baby are included in the same split. Table 2 shares the demographics of the dataset. The weight distribution across splits is matched to the overall dataset distribution (Fig. 3 (Left)).

**Ground-truth.** We obtain accurate ground-truth weight readings by using a calibrated digital weighing machine (least count 10 g). The machine is robust to newborn movement as it stabilizes and locks the recorded value. A calibrated infantometer is used to measure the length of the newborn, and tape measures are used for measuring the head and chest circumference.










Input	Backbone	RefObj	SS	Pooling	MAE (g)	BMAE (g)
Frame	IN		–	–	224.6	437.2
Video	IN		–	Average	170.0	290.6
	IN		–	Vanilla SA	167.1	291.9
	IN		–	Max	158.5	223.7
Video	CLIP		–	Max	145.4	207.2
Video	CLIP		✓	Max	139.9	211.3
	CLIP		✓	Max	129.0	189.0

Table 3. Impact of input, reference object (chessboard  and wooden ruler ) , frame subsampling (SS), and pooling methods. MAE and BMAE are reported for weight estimation on the validation set. The backbone is ResNet-50 pretrained either on ImageNet (IN) [14] or CLIP [45]. We see consistent performance improvement across all modifications – video-based models, CLIP as backbone, and including sub-sampling.

**Evaluation metrics.** We adopt the Mean Absolute Error (MAE) as our primary evaluation metric. We also report Balanced MAE (BMAE) that averages MAE across each bin of interest as the weight distribution is non-uniform. Finally, as an error of 200 g for a 2 kg newborn is worse than that for a 4 kg baby, we report relative errors  $|w - w_{\text{gt}}|/w_{\text{gt}}$  to highlight errors for newborns with lower weights.

**Implementation details.** We fine-tune a ResNet-50 CNN [14] to produce  $d=2048$  dimensional representations  $\mathbf{x}_l$  for each frame  $f_l$ . We choose  $N=25$  frames from a video with an average duration of 12 s. The selected frames are padded to form a square and resized to  $224 \times 224 \times 3$ . Augmentations such as vertical / horizontal flips, translations, and color jitter are applied during training. We use the Adam optimizer [24] and train the model for 200 epochs. When not mentioned otherwise, the initial learning rate is  $10^{-5}$ . We use a StepLR scheduler wherein the learning rate steps down by a factor of 2 every 50 epochs.

### 4.2. Model Ablations

**Frame- vs. Video-based weight estimation.** In the video-based weight estimation approach (Sec. 3.1), we fuse  $N$  representations early on in the network. An alternative frame-based approach would be to use individual frames to obtain weight estimates and average over multiple frames during inference (similar to [50]). A constant learning rate of  $10^{-5}$  works best for frame-based models.

Table 3 shows that video-based models outperform frame-based by a large margin: reduce MAE by 66 g. In particular, we also observe that max-pooling outperforms average pooling and vanilla self-attention (SA) [2].

A key hyperparameter for video-based models is the number of sampled frames  $N$ . Fig. 3 (Middle) shows that the weight MAE reduces dramatically with increasing  $N$ , reaches the minimum around  $N=25$  and slightly increases








Figure 3. **Left:** Weight distribution for the training set. **Middle:** Impact of varying the number of frames  $N$  during evaluation on the validation set. For training, we use  $N=25$ . The model used here is NurtureNet, that augments video information with tabular data and uses proxy tasks of baby segmentation mask and keypoints. **Right:** Effect on weight MAE on the validation set when adding noise sampled from a uniform distribution to the birth weight for NurtureNet models.

thereafter. We choose  $N=25$  for the rest of the experiments.


**Impact of ResNet-50 parameters.** We experiment with a ResNet-50 encoder pretrained on the ImageNet (IN) dataset and one pretrained using the Contrastive Language-Image Pretraining (CLIP) technique [45]. For IN models, an initial learning rate of  $10^{-4}$  stepped down by a factor of 2 every 30 epochs works best. As seen in Table 3, CLIP-based CNN initialization results in better performance, and this encoder is used in all further experiments.

**Impact of frame sampling.** The frame selection process influences the representation and the weight estimate. To reduce this dependency, we introduce subsampling as an augmentation during training. Specifically, we randomly pick  $N'=40$  frames, and subsample 10 subsets of  $N=25$  frames with replacement. By requiring the model to produce the same estimate (Eq. 2) across these subsets, we make the model less sensitive to frame selection. During inference, we do not use subsampling. Table 3 shows that subsampling typically results in a modest improvement of 5.5 g. We employ this technique for further experiments.

**Impact of reference object.** Table 3 shows that using , a standard object for camera calibration, as a reference object over the  gives a 10.9 g improvement on MAE. However, for wider adoption and given 's availability in resource constrained areas, we restrict our solution to using a .

**Impact of multi-task approaches.** Table 4 shows the results of combining various tasks, and the corresponding MAE. As a simple baseline, row 0 displays performance for using the mean value of the training set as the prediction. In all experiments, the loss coefficients are set to  $\lambda_w=5.0$ ,  $\lambda_l=0.1$ ,  $\lambda_h=0.1$ ,  $\lambda_c=0.1$ ,  $\lambda_m=3.0$ , and  $\lambda_k=100.0$ , to scale the importance of various losses. Rows 1-4 show the results when performing each anthropometry estimation task independently, indicating that they fare much better to not using any model. Row 5 shows the effect of including proxy visual tasks, which leads to a performance gain of 16.5 g. Row 6, compared to rows 1-4,

	Tasks					MAE				
	W	L	H	C	M	K	W (g)	L (cm)	H (cm)	C (cm)
0	Train set mean						495.8	2.36	1.80	2.38
1	✓	-	-	-	-	-	139.9	-	-	-
2	-	✓	-	-	-	-	-	1.53	-	-
3	-	-	✓	-	-	-	-	-	1.15	-
4	-	-	-	✓	-	-	-	-	-	1.37
5	✓	-	-	-	✓	✓	123.4	-	-	-
6	✓	✓	✓	✓	-	-	138.2	1.51	1.13	1.37
7	✓	✓	✓	✓	✓	✓	124.4	1.39	1.08	1.27

Table 4. Video-based model with CLIP backbone, , Max pooling, and frame subsampling, under different multi-task configurations. W: weight, L: length, H: head circumference, C: chest circumference, M: segmentation mask, K: keypoints. Performance on the validation set improves as we incorporate all tasks.

	Model	Tasks	Weight estimation	
			MAE (g)	BMAE (g)
0	Video-only	W	139.9	211.3
1	Tab-only	W	202.5	322.6
2	W-NurtureNet	W	127.7	196.6
3	NurtureNet	W M K	113.7	171.8
4	NurtureNet	W L H C M K	115.6	181.3

Table 5. W-NurtureNet concatenates the visual representation to the tabular inputs to regress weight to show improved performance (validation set). NurtureNet is the multi-task equivalent. The Tasks column shows the set of tasks on which the model is trained.

shows negligible change. This indicates that we can use one multi-task model that estimates all anthropometric measures with one video. Finally, row 7, combines all tasks showing marked improvement across all measurements.

**Tabular only model** Tab-Only is posed as a simple linear regression that takes in 2 inputs (birth weight  $w^0$  and age  $a$ ) and predicts the current weight  $w$ . Row 1 in Ta-

	Train set mean	Hu + SVR	HOG + SVR	Regionprops + SVR	Hu + Regionprops + SVR
W (g)	495.8	470.3	466.7	399.4	393.0

Table 6. Hand-crafted models perform poorly on weight regression (validation set) compared to proposed models.

ble 5 shows that this model achieves a competitive MAE of 202.5 g. However, this model is not useful in practice as it predicts the same weight for all babies with a given birth weight after  $a$  days, *i.e.* this model cannot predict deviations from the mean growth for the population (training set).

**Augmenting visual information with tabular inputs.** We augment our visual representation  $\mathbf{z}$  by concatenating them with  $[0, 1]$  normalized tabular inputs and send them forward to regress weight (W-NurtureNet, Sec. 3.3). Row 2 of Table 5 shows that W-NurtureNet achieves an impressive 12.2 g improvement in MAE from 139.9 g to 127.7 g. Rows 3 and 4 show results on training the multi-task NurtureNet. While row 3 shows a 14 g improvement in MAE on using auxiliary tasks, row 4 is a unified model that can estimate all anthropometry measurements and shows a 12 g improvement in weight MAE.

**Effect of errors in recorded birth weight.** To deploy models in rural settings, an important factor to consider is the erroneous nature of tabular inputs. We simulate these errors as  $\tilde{w}^0 = w^0 + \epsilon$ , where the noise  $\epsilon$  is sampled from a uniform distribution  $\mathcal{U}(-q, q)$  and  $q$  corresponds to the maximum deviation in kg. Fig. 3 (Right) shows that our models are quite robust to noisy inputs. In fact, when  $q=0.5$  kg, W-NurtureNet achieves an MAE of 136.5 g (worse than when  $q=0$  by 9 g), still better than our video-based model at 139.9 g. NurtureNet is more robust to noise than W-NurtureNet and results in a  $\sim 5$  g increase in MAE. Finally, the Tab-Only model is highly sensitive to errors in birth weight and results in an MAE of 307.4 g (up by 105 g).

### 4.3. Baselines, Results Summary

We now present and evaluate a few baselines for weight estimation: (i) A naïve approach is to predict the mean of the training set. This acts like the upper bound of the error for any model. (ii) A second approach uses structural information that can be extracted from predicted segmentation masks of the newborn and the ruler. We extract hand-crafted representations in the form of Hu image moments [17] or region features [5]. (iii) We also evaluate the Histogram of Oriented Gradients (HOG) features [9] that are popular in classical computer vision literature. RBF-kernel Support Vector Regressors (SVR) [52] are used to obtain anthropometry estimates from all three representations.

**Baseline ablations.** Table 6 shows the MAE for weight estimation for all three feature representations and a combination. Regionprops features, together with Hu moments

Method	MAE (g)	BMAE (g)
Train set mean	483.5	1091.8
Best hand-crafted approach	390.1	716.7
Frame-based method	214.0	409.7
Best video-only model	139.0	222.5
W-NurtureNet	126.4	207.1
NurtureNet (W L H C M K)	<b>114.3</b>	<b>157.5</b>

Table 7. Weight estimation performance on the test set.

Weight Bin (kg)	Test Set Count	Video-based model			NurtureNet		
		MAE (g)	E80 (g)	% Rel	MAE (g)	E80 (g)	% Rel
1 - 2	40	207.5	305.1	12.2	127.3	180.9	7.3
2 - 2.5	226	135.2	218.4	5.9	108.1	170.5	4.7
2.5 - 3	420	111.5	177.7	4.1	97.7	160.0	3.6
3 - 3.5	353	137.2	224.4	4.2	115.1	181.8	3.6
3.5 - 4	179	148.0	229.3	4.0	123.5	197.6	3.3
4 - 5	73	245.4	425.8	5.7	187.6	288.5	4.4
All	1291	139.0	223.2	4.8	114.3	179.8	3.9

Table 8. Results sliced by weight bins on the test set. Metrics: E80 corresponds to the 80<sup>th</sup> percentile absolute error and indicates that 80% of the samples have an error less than this value. % Rel corresponds to mean absolute relative error. NurtureNet improves over the video-based model across all weight bins.

show best performance for weight estimation (393.0 g). Here as well, the reference object is useful and computing features from both the baby and ruler regions improves performance. Supplement presents additional details on features and ablations. However, performance of all baselines is far from proposed deep-learning based approaches.

**Comparing best approaches.** We summarize the key methods on the test set in Table 7. In particular, we observe that video-based models (139.0 g) achieve a large improvement over the best hand-crafted representations (390.1 g) and single frame based approach [50] (214.0 g). Augmenting the video-only models with tabular data improves MAE to 126.4 g. Finally, NurtureNet results in best performance, achieving an MAE of 114.3 g, while presenting a unified model for newborn anthropometry. Table 9 (Uncompressed) presents results on all measures of NurtureNet.

### 4.4. Analysis and Discussion

**Results sliced by weight bins** are presented in Table 8. NurtureNet brings large improvements in MAE from 15 to 80 g across all bins, but particularly for the low 1 to 2 kg and high 4 to 5 kg bins. Inclusion of the multi-task approach and the tabular inputs also reduces relative and 80<sup>th</sup> percentile errors across all bins. Low errors in the 1 to 2.5 kg bins are especially important to identify underweight newborns.

**Predictions vs. Ground-truth** Another way to assess the performance of our proposed model is a scatter plot of the model’s predictions against the ground-truth. Fig. 4 shows the scatter plot for NurtureNet on the test set. While we




Figure 4. Scatter plot showing predicted weight vs. ground-truth weight for NurtureNet on the test set. The best fit line is obtained using least squares and lies very close to the  $y=x$  diagonal, indicating the goodness of our model.  $R^2$  is the coefficient of determination and  $PCC$  is the Pearson correlation coefficient.

Model	Size (MB)	GFLOPs	MAE (g, cm, cm, cm)			
			W	L	H	C
Uncompressed	121.6	5.38	114.3	1.33	1.04	1.25
Pruned	30.4	1.35	116.0	1.26	1.02	1.21
Quantized	15.0	1.35	117.8	1.53	1.05	1.26

Table 9. Performance of the uncompressed NurtureNet, pruned, and quantized models on the test set. We are able to reduce models by  $8\times$  with minimal loss in weight estimation performance.

observe that the best fit line is close to the  $y=x$  diagonal, our model tends to slightly over-predict for low weight and under-predict for higher weight newborns. This can be attributed to the dataset imbalance (Fig. 3 (Left)).

**Model compression.** Training our model requires one V100 GPU with 32 GB memory. However, our goal is to deploy the NurtureNet model on a low-cost smartphone that can be used by health workers in underserved geographies. Furthermore, the lack of internet coverage necessitates a drastic reduction in the memory and computational footprint of the model to enable on-device and offline inference. We prune NurtureNet using the NNI library [40]. Specifically, we use the L1NormPruner twice to discard half the output channels having the smallest L1 norm of weights in each iteration, effectively reducing the size of the compute requirement by 75%. Further, we perform static quantization, converting the FP32 weights and activations to INT8. The result is a model that is  $8\times$  smaller and  $4\times$  faster with an acceptable deterioration in performance up to 3.5 g MAE. Table 9 shows that compression leads to a negligible performance drop for weight (W), head circumference (H), and chest circumference (C), but an acceptable increase in length (L) error.

**NurtureNet vs. Conventional practice.** We conduct a preliminary and independent field study to analyze the errors in weight measurements made through conventional practices and compare them against NurtureNet. Weight readings taken by health workers using spring balances are compared against calibrated digital weighing machines used for ground-truth weight measurements. We observe an MAE of 183 g (N=92) for conventional methods indicating the challenges of recording such data in rural community settings. Note, that this result is biased towards being lower as the health workers knew that they were being monitored and can only be expected to be worse in real scenarios. NurtureNet achieves a lower MAE at 114.3 g (N=1295), indicating the field-readiness of our approach.

**Limitations.** While AI models can provide meaningful accuracy in many cases, they cannot be perfect on all samples, particularly when it comes to complex problems like estimating the weight of a baby. This may be due to various factors such as newborn clothing, lighting conditions, environmental conditions, camera angles, the position of the baby relative to the camera, or the baby’s movements. The model may also struggle to accurately estimate the weight of babies with certain physical characteristics (e.g. missing limbs) or rare medical conditions that affect growth.

## 5. Conclusion

We present a vision system for newborn anthropometry from a short video taken with a low-cost smartphone. Our proposed approach computes a video representation and augments it with tabular data to obtain weight estimates. We extend this model through multi-task training to simultaneously estimate other anthropometric measurements such as the length, head circumference, and chest circumference (NurtureNet). A proxy task of predicting baby segmentation masks and keypoints further improves the weight estimation performance. Using pruning and quantization, we compress NurtureNet to 15 MB, allowing offline inference and deployment on low-cost smartphones.

This solution is envisioned as a public health screening tool and is currently not intended for diagnostic or clinical settings where good anthropometric instruments and trained personnel are available. Such a tool provides a convenient, geo-tagged, and contactless way for health workers and public health systems to monitor the growth and development of newborns, enabling targeted interventions to drive better health outcomes.

**Acknowledgments.** We thank the Bill & Melinda Gates Foundation (BMGF) and the Fondation Botnar for supporting this work. We thank Niloufer Hospital at Hyderabad for initial data collection. We are grateful to SEWA Rural at Gujarat and PGIMER at Chandigarh for collecting data in community settings. We thank all members of WIAI who were involved for their contributions.



## References

- [1] Hamid Abbasi, Sarah Mollet, Sian Williams, Lilian Lim, Malcolm Battin, Thor F Besier, and Angus McMorland. Deep-Learning for Automated Markerless Tracking of Infants General Movements. *bioRxiv*, 2022. 3
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. In *International Conference on Learning Representations (ICLR)*, 2015. 3, 5
- [3] S K Bhargava, S Ramji, Arun Kumar, Man Mohan, Jasbir Marwah, and H P Sachdev. Mid-arm and chest circumferences at birth as predictors of low birth weight and neonatal mortality in the community. *British Medical Journal*, 291: 1617–1619, 1985. 1
- [4] Hannah Blencowe, Julia Krasevec, Mercedes De Onis, Robert E Black, Xiaoyi An, Gretchen A Stevens, Elaine Borghi, Chika Hayashi, Diana Estevez, Luca Cegolon, et al. National, regional, and worldwide estimates of low birth-weight in 2015, with trends from 2000: a systematic analysis. *The Lancet Global Health*, 7(7):e849–e860, 2019. 1
- [5] Wilhelm Burger, Mark James Burge, Mark James Burge, and Mark James Burge. *Principles of digital image processing*. Springer, 2009. 7
- [6] Xu Cao, Xiaoye Li, Liya Ma, Yi Huang, Xuan Feng, Zening Chen, Hongwu Zeng, and Jianguo Cao. AggPose: Deep Aggregation Vision Transformer for Infant Pose Estimation. In *International Joint Conference on Artificial Intelligence*, 2022. 3
- [7] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *European Conference on Computer Vision (ECCV)*. Springer, 2020. 3
- [8] Parul Christian, Sun Eun Lee, Moira Donahue Angel, Linda S Adair, Shams E Arifeen, Per Ashorn, Fernando C Barros, Caroline HD Fall, Wafaie W Fawzi, Wei Hao, et al. Risk of childhood undernutrition related to small-for-gestational age and preterm birth in low-and middle-income countries. *International journal of epidemiology*, 42(5):1340–1355, 2013. 1
- [9] Navneet Dalal and Bill Triggs. Histograms of Oriented Gradients for Human Detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 7
- [10] Diane DiTomasso and Mary Cloud. Systematic review of expected weight changes after birth for full-term, breastfed newborns. *Journal of Obstetric, Gynecologic & Neonatal Nursing*, 48(6):593–603, 2019. 4
- [11] R Fletcher, X Soriano Díaz, H Bajaj, and Suparna Ghosh-Jerath. Development of smart phone-based child health screening tools for community health workers. In *2017 IEEE Global Humanitarian Technology Conference (GHTC)*. IEEE, 2017. 2
- [12] Daniel Groos, Lars Adde, Ragnhild Støen, Heri Ramampiaro, and Espen AF Ihlen. Towards human-level performance on automatic pose estimation of infant spontaneous movements. *Computerized Medical Imaging and Graphics*, 95, 2022. 3
- [13] Huaiting Gu, Lixia Wang, Lingfei Liu, Xiu Luo, Jia Wang, Fang Hou, Pauline Denis Nkomola, Jing Li, Genyi Liu, Heng Meng, et al. A gradient relationship between low birth weight and IQ: A meta-analysis. *Scientific reports*, 7(1): 18035, 2017. 1
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [15] Nikolas Hesse, Gregor Stachowiak, Timo Breuer, and Michael Arens. Estimating Body Pose of Infants in Depth Images Using Random Ferns. In *International Conference on Computer Vision Workshops (ICCVW)*, 2015. 2
- [16] Nikolas Hesse, Sergi Pujades, Javier Romero, Michael J. Black, Christoph Bodensteiner, Michael Arens, Ulrich G. Hofmann, Uta Tacke, Mijna Hadders-Algra, Raphael Weinberger, Wolfgang Müller-Felber, and A. Sebastian Schroeder. Learning an Infant Body Model from RGB-D Data for Accurate Full Body Motion Analysis. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*. Springer, 2018. 3
- [17] Ming-Kuei Hu. Visual pattern recognition by moment invariants. *IRE transactions on information theory*, 8(2):179–187, 1962. 7
- [18] Xiaofei Huang, Nihang Fu, Shuangjun Liu, and Sarah Ostadabbas. Invariant representation learning for infant pose estimation with small data. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, 2021. 3
- [19] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [20] François R Jornayvaz, Peter Vollenweider, Murielle Bochud, Vincent Mooser, Gérard Waeber, and Pedro Marques-Vidal. Low birth weight leads to obesity, diabetes and increased leptin levels in adults: the CoLaus study. *Cardiovascular diabetology*, 15(1):1–10, 2016. 1
- [21] Manzur Kader and Nirmala KP Perera Perera. Socio-economic and nutritional determinants of low birth weight in India. *North American journal of medical sciences*, 6(7), 2014. 3
- [22] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end Recovery of Human Shape and Pose. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [23] Wasif Khan, Nazar Zaki, Mohammad M Masud, Amir Ahmad, Luqman Ali, Nasloon Ali, and Luai A Ahmed. Infant birth weight estimation and low birth weight classification in United Arab Emirates using machine learning algorithms. *Scientific Reports*, 12(1), 2022. 3
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 5
- [25] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 4

- [26] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [27] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [28] Stefan Kuhle, Bryan Maguire, Hongqun Zhang, David Hamilton, Alexander C Allen, KS Joseph, and Victoria M Allen. Comparison of logistic regression with machine learning methods for the prediction of fetal growth abnormalities: a retrospective cohort study. *BMC pregnancy and childbirth*, 18(1), 2018. 3
- [29] Daniel G Kyrollos, Kim Greenwood, JoAnn Harrold, and James R Green. Transfer Learning Approaches for Neonate Head Localization from Pressure Images. In *2022 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, 2022. 3
- [30] Daniel G. Kyrollos, Anthony Fuller, Kim Greenwood, JoAnn Harrold, and James R. Green. Under the Cover Infant Pose Estimation using Multimodal Data. *IEEE Transactions on Instrumentation and Measurement*, 2023. 3
- [31] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [32] Eryun Liu. Infant footprint recognition. In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [33] Yingying Liu, Arcot Sowmya, and Heba Khamis. Single camera multi-view anthropometric measurement of human height and mid-upper arm circumference using linear regression. *PloS one*, 13(4):e0195600, 2018. 2
- [34] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 4
- [35] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6), 2015. 3
- [36] Patrícia Loreto, Hugo Peixoto, António Abelha, and José Machado. Predicting low birth weight babies through data mining. In *New Knowledge in Information Systems and Technologies: Volume 3*. Springer, 2019. 3
- [37] Yu Lu, Xi Zhang, Xianghua Fu, Fangxiong Chen, and Kelvin KL Wong. Ensemble machine learning for estimating fetal weight at varying gestational age. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2019. 3
- [38] Kevin D McCay, Edmond SL Ho, Claire Marcroft, and Nicholas D Embleton. Establishing pose based features using histograms for the detection of abnormal infant movements. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2019. 2
- [39] L Meinecke, N Breitbach-Faller, C Bartz, R Damen, G Rau, and C Disselhorst-Klug. Movement analysis in the early detection of newborns at risk for developing spasticity due to infantile cerebral palsy. *Human Movement Science*, 25(2), 2006. 2
- [40] Microsoft. Neural Network Intelligence, 2021. 8
- [41] Sara Moccia, Lucia Migliorelli, Virgilio Carnielli, and Emanuele Frontoni. Preterm infants’ pose estimation with spatio-temporal features. *IEEE Transactions on Biomedical Engineering*, 67(8), 2019. 2
- [42] Haomiao Ni, Yuan Xue, Liya Ma, Qian Zhang, Xiaoye Li, and Sharon X Huang. Semi-supervised body parsing and pose estimation for enhancing infant general movement assessment. *Medical Image Analysis*, 83, 2023. 3
- [43] World Health Organization and World Health Organization. Nutrition for Health. *WHO child growth standards: growth velocity based on weight, length and head circumference: methods and development*. World Health Organization, 2009. 1
- [44] Szymon Płotka, Michał K Grzeszczyk, Robert Brawura-Biskupski-Samaha, Paweł Gutaj, Michał Lipa, Tomasz Trzcziński, and Arkadiusz Sitek. BabyNet: Residual Transformer Module for Birth Weight Prediction on Fetal Ultrasound Video. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2022. 2
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, 2021. 5, 6
- [46] Lorenzo Scalise, Natascia Bernacchia, Ilaria Ercoli, and Paolo Marchionni. Heart rate measurement in neonatal patients using a webcam. In *2012 IEEE International Symposium on Medical Measurements and Applications Proceedings*. IEEE, 2012. 2
- [47] A. Sebastian Schroeder, Raphael Weinberger Nikolas Hesse, Uta Tacke, Lucia Gerstl, Anne Hilgendorff, Florian Heinen, Michael Arens, Linze J. Dijkstra, Sergi Pujades Rocamora, Michael J. Black, Christoph Bodensteiner, and Mijna Hadders-Algra. General Movement Assessment from videos of computed 3D infant body models is equally effective compared to conventional RGB video rating. *Early Human Development*, 144:104967, 2020. 2
- [48] Giuseppa Sciortino, Giovanni Maria Farinella, Sebastiano Battiato, Marco Leo, and Cosimo Distante. On the estimation of children’s poses. In *Image Analysis and Processing-ICIAP 2017: 19th International Conference*. Springer, 2017. 3
- [49] D Senthilkumar and S Paulraj. Prediction of low birth weight infants and its risk factors using data mining techniques. In *Proceedings of the 2015 international conference on industrial engineering and operations management*, 2015. 3
- [50] Huaijing Shu, Lirong Ren, Liping Pan, Dongmin Huang, Hongzhou Lu, and Wenjin Wang. Single image based infant body height and weight estimation. In *CVPRW*, 2023. 2, 5, 7
- [51] Nelson Silva, Dajie Zhang, Tomas Kulvicius, Alexander Gail, Carla Barreiros, Stefanie Lindstaedt, Marc Kraft, Sven Bölte, Luise Poustka, Karin Nielsen-Saines, Florentin

- Wörgötter, Christa Einspieler, and Peter B. Marschik. The future of General Movement Assessment: The role of computer vision and machine learning - A scoping review. *Research in Developmental Disabilities*, 110:103854, 2021. 2
- [52] Alex J Smola and Bernhard Schölkopf. A Tutorial on Support Vector Regression. *Statistics and Computing*, 14, 2004. 7
- [53] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *International Conference on Machine Learning (ICML)*. PMLR, 2020. 3
- [54] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, 2015. 3
- [55] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso. Generalised Dice overlap as a deep learning loss function for highly unbalanced segmentations. In *International Workshop on Deep Learning in Medical Image Analysis*, 2017. 4
- [56] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 4
- [57] Richard Szeliski. *Computer Vision: Algorithms and Applications*. Springer Nature, 2022. 3
- [58] Maolong Tang, Ming-Ting Sun, Leonardo Seda, James Swanson, and Zhengyou Zhang. Measuring Infant’s Length with an Image. In *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2018. 2
- [59] Neerja Thakkar, Georgios Pavlakos, and Hany Farid. The reliability of forensic body-shape identification. In *PCVPRW*, 2022. 3
- [60] Mercedes Torres Torres, Michel Valstar, Caroline Henry, Carole Ward, and Don Sharkey. Postnatal gestational age estimation of newborns using small sample deep learning. *Image and vision computing*, 83, 2019. 2
- [61] Anusua Trivedi, Mohit Jain, Nikhil Kumar Gupta, Markus Hinsche, Prashant Singh, Markus Matiaschek, Tristan Behrens, Mirco Militeri, Cameron Birge, Shivangi Kaushik, et al. Height estimation of children under five years using depth images. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2021. 2
- [62] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [63] C Geoffrey Woods and Alasdair Parker. Investigating microcephaly. *Archives of Disease in Childhood*, 98:707–713, 2013. 1
- [64] Qingqiang Wu, Guanghua Xu, Fan Wei, Jiachen Kuang, Penglin Qin, Zejiang Li, and Sicong Zhang. Supine infant pose estimation via single depth image. *IEEE Transactions on Instrumentation and Measurement*, 71, 2022. 2
- [65] Qian Zhang, Yuan Xue, and Xiaolei Huang. Online training for body part segmentation in infant movement videos. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, 2019. 3
- [66] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal Relational Reasoning in Videos. In *European Conference on Computer Vision (ECCV)*, 2018. 3