

Automatic Assessment of Infant Face and Upper-Body Symmetry as Early Signs of Torticollis

Michael Wan^{1,2}, Xiaofei Huang², Bethany Tunik³, Sarah Ostadabbas^{2,4}

¹Roux Institute, Northeastern University, Portland, ME, USA

²Augmented Cognition Lab (ACLab), Northeastern University, Boston, MA, USA

³Board-Certified Clinical Specialist in Pediatric Physical Therapy, Boston, MA, USA

⁴Corresponding author: ostadabbas@ece.neu.edu

Abstract—We apply computer vision pose estimation techniques developed expressly for the data-scarce infant domain to the study of torticollis, a common condition in infants for which early identification and treatment is critical. Specifically, we use a combination of facial landmark and body joint estimation techniques designed for infants to estimate a range of geometric measures pertaining to face and upper body symmetry, drawn from an array of sources in the physical therapy and ophthalmology research literature in torticollis. We gauge performance with a range of metrics and show that the estimates of most these geometric measures are successful, yielding strong to very strong Spearman’s ρ correlation with ground truth values. Furthermore, we show that these estimates, derived from pose estimation neural networks designed for the infant domain, cleanly outperform estimates derived from more widely known networks designed for the adult domain¹.

I. INTRODUCTION

Torticollis is a common condition in infants and children, characterized by a persistent neck tilt or twist to one side. Its most common form, congenital muscular torticollis (CMT), has an estimated incidence of 3.9% to 16% [11]. Early treatment is critical: outcomes are best when CMT is diagnosed and physical therapy treatment started before the infant is three months old, and conversely, if untreated or treated later, CMT can lead to face, skull, or spine deformities, pain and limited motion, and the need for invasive interventions and surgery [17]. Screening, diagnosis, and monitoring during treatment for CMT require laborious professional assessments, so with the recent onset of computer vision technology specifically studying infant face and body poses for health and developmental applications [4], [8], [10], [13], [20], it is natural to wonder whether algorithmic techniques can help enable remote monitoring or automated screening and diagnosis to augment clinical expertise. In this paper, as a first step towards such applications, we explore viability of using computer vision techniques to assess a set of geometric measures of symmetry in the face and upper body, previously identified in the medical literature as being relevant to CMT or the similarly presenting (non-congenital) ocular torticollis condition, purely from casual photographs of infants in their natural environments.

The geometric symmetry measures we consider are illustrated in Fig. 1. We carefully researched and selected




Fig. 1. We study the effectiveness of using deep learning computer vision techniques to assess a range of geometric facial and upper body measures of symmetry—illustrated schematically here—which are drawn from medical research literature on torticollis in infants and children. The assessments employ recent advances in infant-domain estimation of facial and upper body landmarks—also illustrated faintly—and we demonstrate that this yields better results than landmark estimation methods largely trained on adult data.

these from among measures studied by physical therapy and ophthalmology researchers to enable qualitative assessment of changes over time and in response to treatments [1], [6], [14], [22], including specifically to quantify outcomes after surgery [3]. In [15], the authors even study the reliability of the procedure of extracting such measurements from still photographs itself. The general contention is not, of course, that measurements from individual still photographs are fully determinative of torticollis conditions, but rather that aggregated over time they can be employed alongside other tools as part of its detection and treatment. We hope that by establishing the viability of algorithmic assessments of these geometric measurements of symmetry from still photographs, we will open the door to future applications in automated monitoring and screening, including from videos.

The technical tools that we employ for this proof-of-concept experiment are based in computer vision pose estimation from still images, and in particular, face and body landmark estimation. Mature solutions to these tasks exist but are generally based in deep learning from primarily adult faces and bodies. As alluded earlier, specialized methods tailored for the unique faces and bodies of infants have

¹Code and data available at <https://github.com/ostadabbas/Infant-Upper-Body-Postural-Symmetry>. 979-8-3503-4544-5/23/\$31.00 ©2023 IEEE

only begun to crop up in recent years, in recognition of the significant domain gap between infant and adults from the point of view of computer vision representations. In this paper, we employ an infant face landmark estimation model from [20] together with an infant body landmark estimation model from [8], both of which employ domain adaptation techniques to tune existing adult-focused models to the infant domain. We work with a subset of the InfAnFace dataset from [20], and compare the values of our six geometric symmetry measures derived from both predicted and ground truth landmarks. We propose modifications of these measures from their original definitions in the literature to enable compatibility with the landmarks used in pose estimation techniques.

Our findings show that predictions derived from the recent infant-domain pose models exhibit “strong” or “very strong” Spearman’s ρ ranked correlation with the ground truth values on a precisely labeled test set of infant faces in the wild, with best performance on the gaze angle (ga) between the line connecting the outer corners of the eyes and the midsternal plumb line, and on the habitual head deviation (hhd), the angle between the eye line and the acromion process (shoulder) line. Predictions of three other measurements (including non-angles) were strong, but we found only moderate success in the predictions of the orbit slopes angle (osa), the angle between the lines connecting the outer and inner corners of the eyes—this being arguably the most subtle metric. Based on this and our further analysis involving more performance metrics in Section V, we conclude that computer vision infant pose estimation techniques can successfully measure a range of quantities pertaining to torticollis.

II. BACKGROUND: TORTICOLLIS AND INFANT DEVELOPMENT

A. Quantifying torticollis

While there is a large corpus of research and established methodology in the diagnosis and treatment of torticollis, it is largely based on in-person physical assessments by experts and follow-ups with imaging or other deeper techniques [11]. By contrast, our work is inspired by a smaller cluster of papers dealing explicitly with measuring signs and symptoms of torticollis geometrically from still images.

We start with congenital muscular torticollis (CMT). The author in [14] studied the effectiveness of a specific therapeutic intervention for CMT by comparing changes in an infant’s “habitual head deviation” (also “head tilt”) as measured by hand from still photographs. The same author later studied the reliability of this photograph-based method itself, in [15]. Separately, [3] measured the “gaze angle” and “transformational deformity” of child subjects, again from still photographs, to gauge improvement in response to surgical intervention. Measurements from photographs offer researchers a repeatable, objective way to quantify the change in severity of torticollis after an intervention.

Sometimes torticollis is not congenital (present from birth) but rather acquired, as is the case with ocular torticollis, where the abnormal head posture is adopted to compensate

for a defect in vision. In such cases, diagnoses often occur only in adulthood, and can be informed by examination of head pose in childhood photographs. Correspondingly, ophthalmologists have also studied the quantification of facial asymmetry via geometric measurements from still images. An overview of such methods and quantities is given in [1], who in turn cite [6] for definitions of facial measurements such as the “orbit slopes angle,” “relative face size,” and “facial bulk mass,” and [22] for definitions of the “facial angle” and the “nasal tip deviation.” These measurements are studied not as a means to quantify the effect of interventions, but as part of a more comprehensive study on the differential diagnosis of ocular torticollis and other conditions related to facial asymmetries, especially superior oblique palsy [1].

B. Computer vision for infant health and development

We are not aware of prior computer vision research intended to detect torticollis or gauge head asymmetry in infants. The closest in spirit might be a pair of related papers, [19], [24], in which researchers employ computer vision to analyze head posture and tremor with a view towards algorithmic understanding of cervical dystonia (also known as spasmodic torticollis), with incidence largely in the adult population. Accordingly, those studies can take advantage of far more mature adult-focused head pose estimation techniques like OpenFace 2.0 [2], whereas our efforts are highly constrained by data scarcity in the infant domain. In the infant domain, there is prior work on bodies but not faces: [4] develops an infant-specific body pose estimation deep network to extract body motion information from infant videos, in an attempt to assess infant neuromotor risk; and [10] uses 3D infant pose estimation techniques to assess infant body symmetry, with a view towards applications in infant development and torticollis, but without a specific implementation of such. As mentioned, all of this work comes in the context of recent attention in computer vision to the small data domain problem of infant pose comprehension, for both faces [20] and bodies [7], [8], [9], [13].

III. CONCEPTS: MEASUREMENTS OF SYMMETRY

From the sources cited in Section II-A, we identified all clearly defined geometric symmetry measures used by researchers in the study of torticollis and facial asymmetries—six in all. We altered the definitions to base them explicitly on the 68 facial landmark coordinates and two body joint (shoulder) coordinates used by our pose estimators, as illustrated and enumerated in Fig. 2, which also enables more consistent comparisons. The final measures are defined in Table I and illustrated in Fig. 1. In the rest of this section, we clarify and discuss these definitions.

A. Assumptions and context

Since we work with both ground truth and estimated landmark coordinates in two dimensions, we generally assume that every infant is facing forward into the camera, so that the infant’s face plane is roughly parallel with the camera image plane. In principle, all of the measurements we consider are

TABLE I
GEOMETRIC MEASURES OF FACE AND UPPER BODY SYMMETRY PERTAINING TO TORTICOLLIS (DETAILS IN SECTION III)

Measure	Code	Description	Definition	Source
Orbit slopes angle	osa	Angle between outer cathnus (eye corner) line and inner cathnus line	$\angle (\overrightarrow{P_{36}P_{45}}, \overrightarrow{P_{39}P_{42}})$	[1], [6]
Relative face size	rfs	Left divided by right outer-canthus-to-mouth-corner lengths	$\frac{\ \overrightarrow{P_{45}P_{54}}\ }{\ \overrightarrow{P_{36}P_{48}}\ }$	[1], [6]
Facial angle	fa	Angle between the eye line and mouth corners line	$\angle (\overrightarrow{P_{36}^{\mu}P_{39}^{\mu}}, \overrightarrow{P_{42}^{\mu}P_{45}^{\mu}})$	[1], [22]
Gaze angle	ga	Angle between outer cathnus line and midsternal plumb line	$\angle (\overrightarrow{P_{36}P_{45}}, \perp \overrightarrow{P_{68}P_{69}})$	[3]
Translational deformity	td	Distance between chin and midsternal plumb line, normalized by the distance between the outer cathnuses	$\frac{\ \overrightarrow{P_8, \perp \overrightarrow{P_{68}P_{69}}}\ }{\ \overrightarrow{P_{36}P_{45}}\ }$	[3]
Habitual head deviation	hhd	Angle between eye line and acromion process (shoulder) line	$\angle (\overrightarrow{P_{36}^{\mu}P_{39}^{\mu}}, \overrightarrow{P_{42}^{\mu}P_{45}^{\mu}})$	[14], [15]




Fig. 2. Face and body (shoulder) landmarks available from ground truth annotations, with index numbers for each landmark point in correspondence with the definitions in Table I.

well-defined for three dimensional face and bodies, but we do not have access to three dimensional face landmarks in the infant domain, and thus we must assume this alignment to ensure that these measurements are well-defined from the two dimensional landmarks.

B. Symbols and functions

In Table I, $\angle(u, v)$ represents the signed angle in degrees between vectors u and v , relative to a fixed orientation (say, clockwise)². Furthermore, P_i denotes the i th landmark (with $i \in \{0, \dots, 67\}$ corresponding to facial landmarks and $i \in \{68, 69\}$ the shoulders); $\overrightarrow{PQ}$ denotes the vector between two points P and Q , $P_{i,j}^{\mu}$ is the midpoint between P_i and P_j ; $\perp u$ is the perpendicular vector to u (say, taken clockwise); $\|\cdot\|$ is the L^2 norm; and $\|P, u\|$ is the L^2 distance between a point P and the line spanned by a vector u . (Distances can be computed in pixels, but the final geometric measures

²We want to measure the *signed* angle because, for instance, we would like a clockwise and a counterclockwise angles of equal magnitude to be considered different, for the purposes of quantifying predictions. Note that we have $\angle(u, v) = -\angle(v, u)$ for all u and v .

based on distance are unitless because they are normalized, as described next.)

C. Definitional choices

The descriptions in Table I are lightly adapted from the source texts for clarity and uniformity. These descriptions were then formalized into the geometric definitions in Table I by assigning face and body features to landmarks used in our models (as in Fig. 2), with choices in specific cases as follows:

- For facial angle (fa) and habitual head deviation (hhd), the eye line is interpreted as being the line between the midpoints of the eyes, as defined for each eye by the midpoint of its corners, to ensure consistency regardless of whether the eye is open or closed.
- For relative face size (rfs), [1] defines this to be the greater divided by the lesser of the two outer-canthus-to-mouth-corner lengths, but we choose to remove the extra “logical” qualifier and always divide the left length by the right length to enable more precise evaluation of predictions.
- For gaze angle (ga) and translational deformity (td), we take the perpendicular bisector to the two shoulder joints as the midsternal plumb line.
- For translational deformity (td), since we do not have the luxury of standardized photograph sizes as in [3], we normalize by the distance between the outer cathnuses (eye corners), a relatively stable quantity.

D. Omissions

We do not include the following two measurements found in our literature search. The “facial bulk mass” of one side compared to the other from [6] is alluded to in [1] but not precisely defined, and is perhaps not intended to be inferred from photographs. The “nasal tip deviation” from [22] and described in [1] is also not made fully precise, and was deemed too difficult to model reliably. Note that we do *not* exclude quantities simply because they cannot be accurately measured algorithmically with our methods (and indeed, we find that the orbit slopes angle (osa) cannot be).

TABLE II
CONFIGURATION OF OUR LANDMARK ESTIMATION MODELS

Model	Task	Backbone	Training Data
HRNet [21]	Adult Face	HRNet	300-W Train [16]
DarkPose [23]	Adult Body	HRNet	COCO [12]
HRNet-R90JT [20]	Infant Face	HRNet	InfAnFace Train [20]
FiDIP [8]	Infant Body	DarkPose	SyRIP [8]

IV. EXPERIMENTAL SETUP: INFANT DATA AND POSE ESTIMATION

A. Selecting and annotating infant faces

As the testbed for our study, we work with the infant annotated faces (InfAnFace) dataset [20], a comprehensive dataset of 410 infant faces labeled with 68 facial landmark coordinates and four binary pose attributes, designed specifically to alleviate the shortage of annotated data in the infant face domain. The images from the InfAnFace are captured from internet image and video sources, and represent a diverse range of infants in natural environments “in the wild.” We manually select a subset of images from InfAnFace Test satisfying the requirements from Section III-A that the infant be fully front facing, leaving us with 36 images to work with. We manually augment the ground truth facial landmark labels with two additional labels for shoulder joints, to enable us to apply the definitions in Table I.

B. Face and body landmark estimation

To study the effectiveness of algorithmic assessment of our geometric measures, we perform experiments with two sets of pose estimation models, the first established models trained largely on adult data, and the second recent models designed specifically for the infant domain. The models and training data are summarized in Table II.

In the adult domain, for facial landmark estimation, we employ the high-resolution network (HRNet) [18], [21], an influential multi-resolution convolutional neural network designed to maintain high-resolution representations throughout inference, which is handy for landmark estimation for high-resolution images. For body joint estimation, we use DarkPose [23], which modifies the standard convolutional heatmap regression approach to landmark estimation to alleviate bias in coordinate representation arising from arbitrary pixel quantization during encoding and decoding. Specifically, we adopt a DarkPose model with an HRNet backbone.

In the infant domain, for facial landmark estimation, we use the HRNet-R90JT model from [20], also built on HRNet and adapted to the infant domain via fine-tuning and data augmentation directed at the unique features of infant faces. Finally, for body joint estimation, we use the fine-tuned domain-adapted infant pose (FiDIP) model from [8], which uses synthetic data and domain adversarial methods to adapt body landmark estimation (in this case, the HRNet-backed DarkPose model) to the infant domain.

V. RESULTS: PERFORMANCE AND ANALYSIS

A. Performance metrics

Often in machine learning, metrics are chosen to enable robust comparison of ever-improving models over time,

TABLE III
ESTIMATION PERFORMANCE OF GEOMETRIC MEASURES OF SYMMETRY FROM INFANT AND ADULT-BASED MODELS

Metric		Models	osa	rfs	fa	ga	td	hhd
Spearman's ρ	$\uparrow$	Infant	0.36	0.61	0.60	0.79	0.53	0.80
			*	***	***	***	***	***
		Adult	-0.10	0.58	0.53	0.78	0.36	0.78
			***	***	***	*	***	
BCA (%)	$\uparrow$	Infant	50.0	77.8	75.0	80.6	88.9	77.8
		Adult	41.7	75.0	66.7	77.8	69.4	77.8
MAE	$\downarrow$	Infant	1.98	0.04	2.18	2.29	0.06	2.63
		Adult	4.73	0.07	4.51	5.25	0.08	4.71
RMSE	$\downarrow$	Infant	2.70	0.05	3.38	3.28	0.09	3.51
		Adult	12.88	0.17	11.27	15.61	0.12	11.71
Measure type			°	ratio	°	°	ratio	°
Ground truth mean			0.80	1.00	-0.34	87.64	0.11	-2.62
* $p \leq 0.05$, ** $p \leq 0.01$, *** $p \leq 0.001$; best performance in bold								

on large datasets serving as public benchmarks. In our case, because of the relative novelty of the techniques, estimation task, and test set, and because of our interest in immediate medical applications, we rely on a range of more interpretable metrics to give us a sense of absolute performance without needing comparisons. In particular, we will evaluate performance using Spearman’s ρ correlation coefficient, the binary classification accuracy of predicting whether a quantity is above or below the mean, and the mean absolute error. We do also report the root mean squared error, for purposes of robust future comparison.

Spearman’s ρ rank correlation coefficient between two variables is defined as the Pearson’s r correlation between the internal rankings for each variable. In the context of deep learning tools trained and tested on imprecise labels and small data domains, high Pearson’s r correlations between predictions and ground truth is too much to ask for, but adopting Spearman’s ρ allows us to relax the requirement for pinpoint accuracy and focus on the relative ranked correctness of the predictions, which is fully compatible with our goal of enabling screening or diagnosis of torticollis. In a similar vein, we assess the *binary classification accuracy* (BCA) of a set of predictions $(\tilde{y}_i)_{i \in I}$ relative to the ground truth $(y_i)_{i \in I}$ of a given variable, defined simply as the proportion of i in the sample index I for which $\tilde{y}_i > \mu \iff y_i > \mu$, where $\mu = \mu((y_i)_{i \in I})$ is the ground truth sample mean. This coarse measure gauges how often predictions accurately determine whether a quantity is greater or less than the ground truth mean (so unlike Spearman’s ρ , it is not purely relative), and its simplicity allows for ease of interpretation. We also measure the mean absolute error (MAE) and the root mean squared error (RMSE)—the former being more interpretable, and the latter being more robust for subsequent comparisons.

B. Results and discussion

Table III tabulates estimation performance of our geometric measures of symmetry, as gauged by the aforementioned metrics, for both the infant- and adult-based models. See Table IV for a guideline for interpreting Spearman’s ρ scores, from the often cited [5].

On the whole, our results show that the infant pose estima-

TABLE IV
INTERPRETATION OF SPEARMAN’S ρ , FROM [5]

Spearman’s ρ	Correlation
≥ 0.70	Very strong relationship
0.40–0.69	Strong relationship
0.30–0.39	Moderate relationship
0.20–0.29	Weak relationship
0.01–0.19	No or negligible relationship
Descriptions apply to both positive and negative values	

tion models yield predictions of five of the six measures—all except the orbit slopes angle (osa)—with a high degree of fidelity. Spearman’s ρ correlations are generally strong or very strong and binary classification accuracies high (75.0–88.9%). Mean absolute errors are within 3° for angles and 0.1 for ratios, and root mean squared errors are low as well. The predictions of the remaining quantity, the orbit slopes angle (osa), only moderately correlate with the ground truth, with $\rho = 0.36$. For the orbit slopes angle, the BCA is no better than random guessing at 50%, but the low MAE and RMSE values suggest the poor classification accuracy may be due to miscalibration. Geometrically, the orbit slopes angle is quite subtle, measuring slight differences between the line connecting the outer cathnuses (eye corners) and the inner cathnuses, and it appears that this distinction is too fine to be reliably measured with the infant landmark estimation model.

Turning to the adult pose estimation models, Table III shows that *for every geometric symmetry measure and every prediction performance metric, the predictions from the adult models fare worse than those from the infant models*, with the exception of one tie. Note that this superiority is not maintained when cross-comparing performances of *different* measures—for instance, the infant model prediction of face angle (fa) has lower Spearman’s ρ and BCA scores than the adult prediction of gaze angle (ga). Indeed, the adult model predictions of half of the measurements—the relative face size (rfs), gaze angle (ga), and habitual head deviation (hhd)—can be considered to be successful under Spearman’s ρ and BCA, although their absolute errors (MAE and RMSE) are still high.

To complement the summary statistics from Table III, we include full scatter plots of predicted vs ground truth values for all six measures under both sets of models, in Fig. 3. The scale of these scatter plots is chosen to highlight the presence of major outlier mis-predictions from the adult models, and the subtler performance differences between measurements revealed by Table III are harder to perceive. Finally, a visualization of ground truth and predictions the geometric elements involved in the computation of our six measures of interest can be found Fig. 4. The predictions from the infant model appear to be fairly accurate, but since it is hard to visually determine angles and ratios at a glance, it is difficult to make definitive conclusions about performance on this qualitative basis alone. What such figures and the performance scatter plots do make it clear, though, is that a major cause of poor performance from the adult pose estimation models stems from a handful of instances of failed facial landmark estimation.

VI. CONCLUSION

Motivated by applications to the detection and treatment of torticollis in infants, we have studied the extent to which recent deep learning based computer vision algorithms designed specifically for the infant domain can measure a set of geometric facial and upper-body symmetry measures previously defined by torticollis researchers, from still images alone. After carefully crafting a test set of infant faces and honing the definitions to line up with standard facial landmarks, we found that most of these measurements can be successfully predicted by pose estimation models designed for infants. We also showed that these models outperform corresponding models designed for and trained on adult data, demonstrating the importance of using tailored neural network models in the data-scarce infant domain.

One possible next step for computer vision research into torticollis would be to use these predicted geometric measurements as features in machine learning algorithms for screening or diagnosis of torticollis, on clinical subjects or subjects “in the wild,” supervised by assessments from medical experts. The underlying data could take the form of either still images or videos. Another potential direction would be to develop these estimation techniques for infants in wider range of poses, not restricted to those clearly facing forward. This would likely require advances in infant pose estimation, especially three-dimensional infant facial landmark estimation, which is currently out of reach.

REFERENCES

- [1] M. R. Akbari, M. Khorrami Nejad, F. Askarizadeh, F. F. Pour, M. Ranjbar Pazooki, and M. R. Moeinitabar. Facial asymmetry in ocular torticollis. *Journal of Current Ophthalmology*, 27(1-2):4–11, Mar. 2015.
- [2] T. Baltrušaitis, A. Zadeh, Y. C. Lim, and L.-P. Morency. Openface 2.0: Facial behavior analysis toolkit. *13th IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 59–66, 2018.
- [3] A. Bhaskar, H. U., and H. Desai. Congenital muscular torticollis: Use of gaze angle and translational deformity in assessment of facial asymmetry. *Indian Journal of Orthopaedics*, 51(2):123–130, Apr. 2017.
- [4] C. Chambers, N. Seethapathi, R. Saluja, H. Loeb, S. R. Pierce, D. K. Bogen, L. Prosser, M. J. Johnson, and K. P. Kording. Computer Vision to Automatically Assess Infant Neuromotor Risk. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 28(11):2431–2442, Nov. 2020.
- [5] C. P. Dancey and J. Reidy. *Statistics without maths for psychology*. Pearson education, 2007.
- [6] E. M. Helveston, D. Krach, D. A. Plager, and F. D. Ellis. A new classification of superior oblique palsy based on congenital variations in the tendon. *Ophthalmology*, 99(10):1609–1615, 1992.
- [7] N. Hesse, S. Pujades, M. J. Black, M. Arens, U. G. Hofmann, and A. S. Schroeder. Learning and tracking the 3d body shape of freely moving infants from rgb-d sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:2540–2551, 2020.
- [8] X. Huang, N. Fu, S. Liu, and S. Ostadabbas. Invariant Representation Learning for Infant Pose Estimation with Small Data. *16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, Dec. 2021. arXiv: 2010.06100.
- [9] X. Huang, S. Liu, M. Wan, N. Fu, D. L. Pino, B. Modayur, and S. Ostadabbas. Appearance-independent pose-based posture classification in infants. *26th International Conference on Pattern Recognition Workshop, Towards a Complete Analysis of People: From Face and Body to Clothes (ICPR-W T-CAP)*, Aug. 2022.
- [10] X. Huang, M. Wan, L. Luan, B. Tunik, and S. Ostadabbas. Computer Vision to the Rescue: Infant Postural Symmetry Estimation from Incongruent Annotations. *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, Jan. 2023.




Fig. 3. Scatter plots of predictions vs. ground truth of our six geometric measures of symmetry (definitions given in Table I), with the scale chosen to emphasize the strong effect of outliers on the predictions derived from adult pose estimation models (bottom) compared to those derived from the infant pose estimation models (top). See Table III for a more quantitative characterization of performance for each quantity. Identity lines drawn for reference.



Fig. 4. Ground truth and predicted positions of some geometric elements related to our measures of symmetry, superimposed on the underlying face and shoulder landmark estimations. Samples chosen to illustrate outlier cases of bad facial landmark predictions, which disproportionately contribute to the poor quantitative performance from the pose estimation based models, especially for the models trained on adult data. Outside of these outlier cases, performance is difficult to adjudicate by visual inspection, and thus our results analysis in Section V is guided by quantitative performance metrics.

- [11] S. L. Kaplan, C. Coulter, and B. Sargent. Physical Therapy Management of Congenital Muscular Torticollis: A 2018 Evidence-Based Clinical Practice Guideline From the APTA Academy of Pediatric Physical Therapy. *Pediatric Physical Therapy*, 30(4):240–290, Oct. 2018.
- [12] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. pages 740–755, 2014.
- [13] S. Liu, X. Huang, N. Fu, C. Li, Z. Su, and S. Ostadabbas. Simultaneously-collected multimodal lying pose dataset: Enabling in-bed human pose monitoring. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2022.
- [14] M. Rahlin. TAMO Therapy as a Major Component of Physical Therapy Intervention for an Infant With Congenital Muscular Torticollis: A Case Report. *Pediatric Physical Therapy*, 17(3):209–218, 2005.
- [15] M. Rahlin and B. Sarmiento. Reliability of Still Photography Measuring Habitual Head Deviation From Midline in Infants With Congenital Muscular Torticollis. *Pediatric Physical Therapy*, 22(4):399–406, 2010.
- [16] C. Sagonas, E. Antonakos, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: database and results. *Image and Vision Computing*, 47(3):18, 2016. 300-W, the First Automatic Facial Landmark Detection in-the-Wild Challenge.
- [17] B. Sargent, S. L. Kaplan, C. Coulter, and C. Baker. Congenital Muscular Torticollis: Bridging the Gap Between Research and Clinical Practice. *Pediatrics*, 144(2):e20190582, Aug. 2019.
- [18] K. Sun, B. Xiao, D. Liu, and J. Wang. Deep high-resolution representation learning for human pose estimation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [19] J. P. Vu, E. Cisneros, H. Y. Lee, L. Le, Q. Chen, X. A. Guo, R. Rouzbehani, J. Jankovic, S. Factor, C. G. Goetz, R. L. Barbano, J. S. Perlmutter, H. A. Jinnah, S. Piro Richardson, G. T. Stebbins, R. Elble, C. L. Comella, and D. A. Peterson. Head tremor in cervical dystonia: Quantifying severity with computer vision. *Journal of the Neurological Sciences*, 434:120154, 2022.
- [20] M. Wan, S. Zhu, L. Luan, P. Gulati, X. Huang, R. Schwartz-Mette, M. Hayes, E. Zimmerman, and S. Ostadabbas. InfAnFace: Bridging the Infant–Adult Domain Gap in Facial Landmark Estimation in the Wild. *26th International Conference on Pattern Recognition (ICPR)*, 2022.
- [21] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang, W. Liu, and B. Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2019.
- [22] M. E. Wilson and J. Hoxie. Facial asymmetry in superior oblique muscle palsy. *Journal of Pediatric Ophthalmology & Strabismus*, 30(5):315–318, 1993.
- [23] F. Zhang, X. Zhu, H. Dai, M. Ye, and C. Zhu. Distribution-aware coordinate representation for human pose estimation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [24] Z. Zhang, E. Cisneros, H. Y. Lee, J. P. Vu, Q. Chen, C. N. Benadof, J. Whitehill, R. Rouzbehani, D. T. Sy, J. S. Huang, T. J. Sejnowski, J. Jankovic, S. Factor, C. G. Goetz, R. L. Barbano, J. S. Perlmutter, H. A. Jinnah, B. D. Berman, S. P. Richardson, G. T. Stebbins, C. L. Comella, and D. A. Peterson. Hold that pose: capturing cervical dystonia’s head deviation severity from video. *Annals of Clinical and Translational Neurology*, 9(5):684–694, May 2022.