

Facial anatomical landmark detection using regularized transfer learning with application to Fetal Alcohol Syndrome recognition

Fu, Zeyu; Jiao, Jianbo; Suttie, Michael; Noble, J. Alison

DOI:

[10.1109/JBHI.2021.3110680](https://doi.org/10.1109/JBHI.2021.3110680)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Fu, Z, Jiao, J, Suttie, M & Noble, JA 2022, 'Facial anatomical landmark detection using regularized transfer learning with application to Fetal Alcohol Syndrome recognition', *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 4, 9531054, pp. 1591-1601. <https://doi.org/10.1109/JBHI.2021.3110680>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Facial Anatomical Landmark Detection using Regularized Transfer Learning with Application to Fetal Alcohol Syndrome Recognition

Zeyu Fu, *Member, IEEE*, Jianbo Jiao, *Member, IEEE*, Michael Suttie, and J. Alison Noble.

Abstract—Fetal alcohol syndrome (FAS) caused by prenatal alcohol exposure can result in a series of cranio-facial anomalies, and behavioral and neurocognitive problems. Current diagnosis of FAS is typically done by identifying a set of facial characteristics, which are often obtained by manual examination. Anatomical landmark detection, which provides rich geometric information, is important to detect the presence of FAS associated facial anomalies. This imaging application is characterized by large variations in data appearance and limited availability of labeled data. Current deep learning-based heatmap regression methods designed for facial landmark detection in natural images assume availability of large datasets and are therefore not well-suited for this application. To address this restriction, we develop a new regularized transfer learning approach that exploits the knowledge of a network learned on large facial recognition datasets. In contrast to standard transfer learning which focuses on adjusting the pre-trained weights, the proposed learning approach regularizes the model behavior. It explicitly reuses the rich visual semantics of a domain-similar source model on the target task data as an additional supervisory signal for regularizing landmark detection optimization. Specifically, we develop four regularization constraints for the proposed transfer learning, including constraining the feature outputs from classification and intermediate layers, as well as matching activation attention maps in both spatial and channel levels. Experimental evaluation on a collected clinical imaging dataset demonstrate that the proposed approach can effectively improve model generalizability under limited training samples, and is advantageous to other approaches in the literature.

Index Terms—Landmark detection, knowledge transfer, regularization, fetal alcohol syndrome

I. INTRODUCTION

Accurate localization of anatomical landmarks is a fundamental step for numerous medical imaging applications such as image registration and shape analysis [1]–[5]. It also has the potential to advance the early recognition of those affected by prenatal alcohol exposure [6], [7]. Prenatal alcohol exposure (PAE) damages the developing central nervous system,

This work was done in conjunction with the Collaborative Initiative on Fetal Alcohol Spectrum Disorders (CIFASD), which is funded by grants from the National Institute on Alcohol Abuse and Alcoholism (NIAAA). This work was supported by NIH grant U01AA014809 and EPSRC grant EP/M013774/1.

Ethical approval for this study was obtained from Oxford Tropical Research Ethics Committee, University of Oxford (Date of approval: 31/03/17 and Protocol number: 519-17).

Z. Fu, J. Jiao and J.A. Noble are with the Department of Engineering Science, University of Oxford, Oxford, UK, Emails: {zeyu.fu, jianbo.jiao, alison.noble}@eng.ox.ac.uk.

M. Suttie is with the Nuffield Department of Women's and Reproductive Health, Big Data Institute, University of Oxford, Oxford, UK, E-mail: michael.suttie@wrh.ox.ac.uk.




Fig. 1. Conventional clinical approaches for extracting diagnostic facial features of fetal alcohol syndrome (FAS). (Image credit: [10], [12])




Fig. 2. Illustration of anatomical landmarks used in this study.

resulting in a continuum of effects [8]. Conditions that result from PAE are collectively known as fetal alcohol spectrum disorders (FASDs). Fetal alcohol syndrome (FAS) is the most severe outcome in this spectrum, characterized by restricted growth, neurobehavioral deficits, and facial dysmorphism [9]. The diagnosis of FAS requires at least 2 of 3 cardinal facial features; a smooth philtrum, a thin upper lip, and a reduced palpebral fissure length (PFL) [8], [9].

Conventional clinical approaches [10], [11] for facial assessment rely on manual examination, using a ruler for PFL measurement, and a Likert-scale comparison chart for the lip and philtrum, as depicted in Fig. 1. Clinical examination is necessarily subjective, with accuracy heavily reliant on the skill and experience of the clinician. Consequently, there is poor inter-rater reliability, and underdiagnosis and misdiagnosis of FASDs are a common occurrence [11]. Therefore, automatically localizing anatomical landmarks, as shown in Fig. 2, is seen as a crucial part of the clinical workflow to provide objective measurements for recognition of FASD.

With the rapid development of deep learning, recent years have witnessed important advances in automation of various medical imaging tasks [13]. Included in this, recent approaches for anatomical landmark detection [2]–[5], [7], have often leveraged the power of deep neural networks, due to the ability to capture hierarchical representations from raw data, and



Fig. 3. Examples of facial phenotypes in our challenging dataset.

have achieved promising performance. However, anatomical landmark detection in our imaging application remains a challenge due to the large appearance variation of faces, different age groups and ethnic background, as exemplified in Fig. 3. One feasible solution to mitigate the challenge is to utilize backbone CNN architectures (e.g. VGG16 [14] or ResNet50 [15]) trained on a large-scale and diverse image dataset, such as VGGFace2 [16], which can be fine-tuned or specialized with additional task-specific layers to promote the optimization of facial anatomical landmark detection [17], [18]. Finetuning, as a common paradigm in transfer learning [19], aims to benefit the target task by providing a good initialization, but it can require exhaustive tuning or a set of ad-hoc hyper-parameters to achieve good performance [20], [21]. Furthermore, limited availability of labeled data, which often occurs in the medical domain, can result in overfitting in transfer learning, leading to poor accuracy for the desired task [4].

To address this, L^2 regularization, also known as weight decay, has been frequently suggested to reduce the effect of overfitting by regularizing the parameter search space [20]. However, such a regularizer may not perform well in transfer learning, especially, and as in our case, when transferring knowledge across different tasks (classification \rightarrow regression). We illustrate this in Fig. 4 (b) which shows that standard fine-tuning with weight decay ignores useful facial semantics or patterns (see Fig. 4 (a)) which are related to the target task of landmark detection, but instead can pay attention to irrelevant regions. To this end, we argue that the target (fine-tuned) model should not only focus on adjusting the pre-trained weights, but also learn how to “behave” as a source model does [21], so that facial semantics can be explicitly reused to improve learning of the landmark detector, as shown in Fig. 4 (c). Therefore, we propose to retain feature responses of a source classification model given the target task data to regularize landmark detection optimization, as shown in Section III-B. Inspired by knowledge distillation for model compression [22]–[24], we develop four regularization strategies during fine tuning, including constraining the feature outputs from classification and intermediate layers, as well as matching activation attention maps in both spatial and channel levels which are shown to provide more comprehensive representations, as depicted in Fig. 4 (d). Extensive experimental evaluations reported in Section IV, demonstrate




Fig. 4. Illustration of top-level activation attention maps for different approaches given an image in the target dataset, where (a) represents the original source model, (b) is the fine-tuned model with weight decay, (c) denotes the proposed regularization constraint on intermediate layer outputs, and (d) demonstrates the proposed regularization constraint on spatial attention maps.

that the proposed framework generalizes well in anatomical landmark detection with limited training samples and favorably outperforms other solutions from the literature.

This paper considerably extends our preliminary work in [25], developing a unified regularized transfer learning framework for localizing anatomical landmarks in fetal alcohol syndrome. The new contributions of this paper relative to [25] are, 1) we present a more detailed justification of the proposed framework as well as its importance with additional illustrations. 2) we propose two additional regularization constraints based on spatial and channel attention maps which improve localization accuracy relative to results reported in [25]. The technical details of each solution in the proposed framework are explained. 3) Experimentally, we provide an in-depth analysis of different proposed regularization strategies, analyze the effectiveness of different pre-trained source models for the proposed framework, study the impact of different parameter selection and include detailed comparisons with further related methods both quantitatively and qualitatively. 4) we provide a detailed discussion to discuss the strengths and the limitations of our approaches and suggest their translational potentials.

The rest of the paper is structured as following: Section II discusses the most closely related work to the proposed approach. Section III formulates the problem considered and details the proposed solution via regularized transfer learning. In Section IV, we perform extensive experimental evaluations. Finally, following the Discussion in Section V, Section VI draws conclusions and suggests some future directions of research.

II. RELATED WORK

A. Anatomical Landmark Detection

In the literature, there has been significant progress on automatic landmark (keypoint) detection of natural images, which has in turn advanced human pose estimation, and facial alignment. Existing methods for this task include random forest [1], [26] and cascaded shape regression [27], [28]. Convolutional neural networks (CNN) based regression has become a recently popular approach for keypoint localization [18], [29]–[31]. Two CNN-based architecture designs have emerged: direct coordinate regression [30], [32] and heatmap regression [29], [33], where the latter usually outperforms the former, due to the advantage of preserving higher spatial resolution for accurate localization. Methods based on U-Net

[13], a stacked hourglass [29], and simple-baseline [33] CNN architectures have achieved good performance on the challenging keypoint detection benchmarks (e.g. MPII benchmark [34]).

Different from generic keypoint localization in natural images, anatomical landmark detection in medical imaging is often constrained by the scarcity of labeled data, as annotating abundant landmarks in the medical domain is expensive. To accommodate this, some existing methods have considered cascaded CNNs following a coarse-to-fine strategy [5], [35], [36]. Although cascaded designs can give good results, they require additional computational efforts, where repeatedly extracting similar low-level features may not be necessary [37]. On the other hand, attention mechanisms which were originally proposed for machine translations [38], have been exploited to enhance the CNN-based anatomical landmark detection [3], [39], [40]. Specifically, Zhong et al. [39] developed a two-stage U-Net with attention guidance to localize anatomical landmarks. Chen et al. [3] and Liu et al. [40] exploited self-attention modules [38], [41] for pelvis and cephalometric landmark detection respectively. More recently, Huang et al. [7] added a differentiable spatial to numerical transform layer to an encoder-decoder network. Vlontzos et al. [42] introduced multi-agent reinforcement learning for multiple landmark detection by designing a collaborative Deep Q-Network. Unlike the aforementioned methods, the proposed learning approach focuses on internally enriching the feature representations via regularized transfer learning for anatomical landmark detection, requiring neither attention mechanisms nor cascaded processing.

B. Knowledge Transfer

Fine-tuning [33] is a standard practice in transfer learning for refining a pre-trained model to be suitable for a new target task. Typically a small learning rate is applied and some model parameters may need to be frozen to reduce overfitting. However, empirically modifying existing model parameters is not a direct method to reuse the representational features from a source task model, which may limit good generalization to a small new task training dataset [43].

In contrast, knowledge distillation, originally proposed for model compression [22], can be considered as a more direct way to retain the learning experiences from a source task model by encouraging predictions of a student model to be close to those of a teacher model. This technique has been extended and applied to various applications, including hint learning [23], incremental learning [43]–[45], privileged learning [46], domain adaptation [47] and human expert knowledge distillation [48]. Distillation methods focus on producing a compact model by operating the knowledge transfer across common tasks [22], [23], [48]. Different from the aforementioned methods, the proposed learning approach in this paper regularizes the transfer learning across different tasks. In other words, the proposed regularized knowledge transfer can be interpreted as cross-task representation learning, where “cross-task” means that the learning process is made between the source and target tasks corresponding to a different label space.

To boost the performance of the target task, joint training of the source and the target tasks, termed multi-task learning [32] has been proposed, where model parameters for both tasks are jointly optimized to learn a shared representation to benefit the target task. However that method requires data and labels of both tasks to be available during training, which can be difficult to satisfy for modelling tasks from clinical data. For our proposed design, we do not require access to the data from the source domain.

III. METHOD

In this section, we first introduce the task of anatomical landmark detection, then describe the proposed regularized transfer learning solution for this task.

A. Problem Formulation

Consider a supervised anatomical landmark detection task, where we are given a training dataset with N pairs of training samples, denoted as $\mathcal{D} = \{\mathbf{I}_i, \mathbf{p}_i\}_{i=1}^N$, where $\mathbf{I}_i \in \mathbb{R}^{H \times W \times 3}$ represents a 2D facial image with height H and width W , and $\mathbf{p}_i = [(x_1, y_1), (x_2, y_2), \dots, (x_K, y_K)] \in \mathbb{R}^{2 \times K}$ denotes K numbered labeled landmark coordinates in each of the N images ($i = 1 \rightarrow N$).



Inspired by its recent success in keypoint localization [29], [33], we use a heatmap to regress the keypoints. To do this, we need to transform these landmark coordinates into heatmaps. Specifically, labeled coordinates are downsampled to $1/4$ of the input size ($\mathbf{p}_i \leftarrow \mathbf{p}_i/4$), and then are transformed to a set of heatmaps $\mathbf{G}_i \in \mathbb{R}^{(H/4) \times (W/4) \times K}$. Each heatmap $\mathbf{g}_{i,k} \in \mathbb{R}^{(H/4) \times (W/4)}$, $k \in \{1, \dots, K\}$ is defined as a 2D Gaussian kernel centered on the k -th landmark coordinate $(x_{i,k}, y_{i,k})$, which is given as,

$$\mathbf{g}_{i,k}(a, b) = \exp \left(-\frac{(a - x_{i,k})^2 + (b - y_{i,k})^2}{2\sigma^2} \right) \quad (1)$$

where a and b denotes the entry of matrix $\mathbf{g}_{i,k}$ in the a -th row and b -th column, and σ is Gaussian kernel scale parameter.

Given the dataset $\mathcal{D} = \{\mathbf{I}_i, \mathbf{G}_i\}_{i=1}^N$, the goal is then to learn a non-linear mapping function between an input image and a set of heatmaps. We follow [18], [33], and design a CNN-based encoder-decoder network $\mathcal{T}_{\theta_e, \theta_d} : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^{(H/4) \times (W/4) \times K}$ with learnable parameters θ_e and θ_d , to achieve the goal. To mitigate against facial variability in our dataset, we consider a ResNet-50 model [15] pre-trained on a large-scale and diverse VGGFace2 [16] dataset as a backbone for feature extraction (encoder) $f_{\theta_e} : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^m$, with the classification branch $g_{\theta_c} : \mathbb{R}^m \rightarrow \mathbb{R}^C$ being disabled, where m denotes the dimensionality of the encoder output, and C represents the number of classes in the source classification task. For the decoder, three deconvolutional layers are employed to recover the spatial resolution, each of which is with the dimension of 256, a 4×4 kernel and a stride of 2, and is followed by a batch normalization and a rectified linear unit (ReLU) [49] activation function. Then, a 1×1 convolutional layer with linear activation is applied to obtain the predicted heatmaps $h_{\theta_d} : \mathbb{R}^m \rightarrow \mathbb{R}^{(H/4) \times (W/4) \times K}$. Consequently, the network $\mathcal{T}_{\theta_e, \theta_d} := h_{\theta_d}(f_{\theta_e}(\mathbf{I}))$ can be learned by minimizing

Source Task: Face Recognition



Target Task: Anatomical Landmark Detection

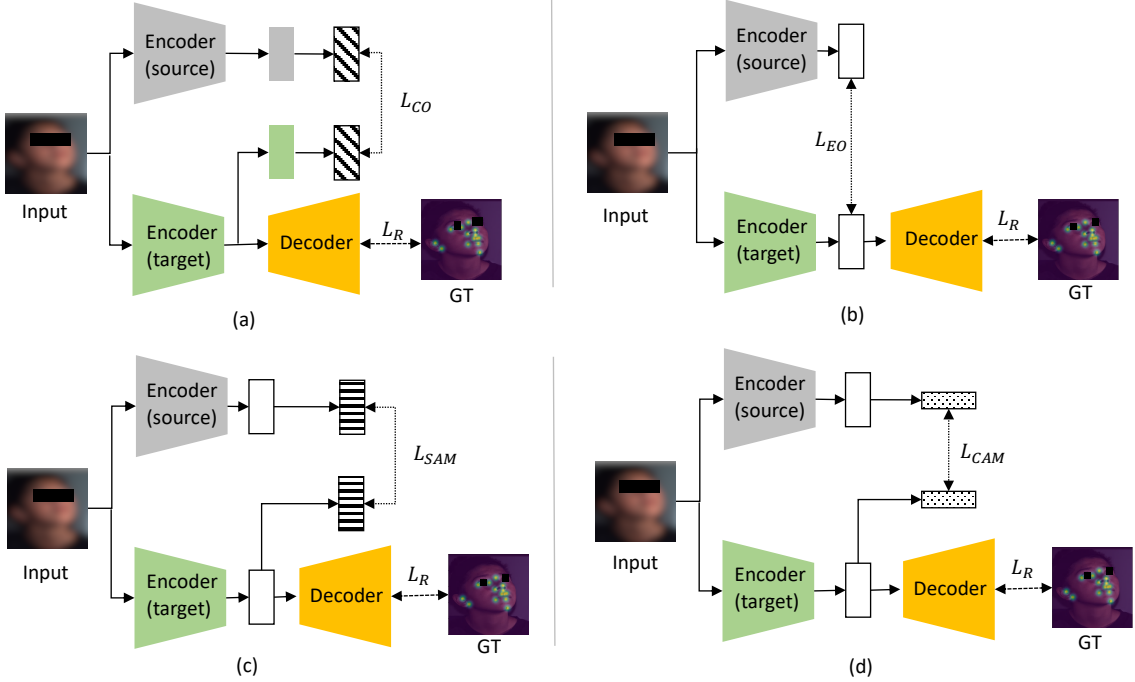


Fig. 5. Illustration of proposed regularized transfer learning for the anatomical landmark detection. Face recognition is considered as the source task to provide rich facial semantics. For the target task, (a)-(d) illustrate four regularization constraints on the classification output (L_{CO}), the encoder output (L_{EO}), the spatial attention map (L_{SAM}) and the channel attention map (L_{CAM}), respectively.

the following loss between the predicted heatmaps and the corresponding ground truth,

$$L_R = \frac{1}{N} \sum_{i=1}^N \|\mathbf{G}_i - h_{\theta_d}(f_{\theta_e}(\mathbf{I}_i))\|_F^2 \quad (2)$$

where F denotes the Frobenius norm.

The above underpins our primary learning objective for the task of anatomical landmark detection. Standard training strategies for such a transfer learning approach usually rely on fine-tuning, where the entire network or only the decoder is trained on the target data. Considering the limited available labelled data in our case, an L2 penalty (weight decay) is usually utilized in the fine-tuning to reduce parameter overfitting. However, merely penalizing the model weights may not capture the original knowledge embedded in the source network. In our case this may lead to a loss of useful facial descriptions and the distraction from other irrelevant regions, as illustrated in Fig. 4. We present proposed solutions to address this problem in the next subsection.

B. Regularized Transfer Learning for Anatomical Landmark Detection

In order to reuse knowledge from the source network in the landmark detector model, we introduce a new regularized transfer learning (RTL) approach which aligns the network features (behaviors) of the source and the target models during target model training. To this end, we compute predictions of the source network on the target dataset \mathcal{D}^t , which are matched with those from the target network through a regularization loss L_γ . Hence, the total loss is updated as,

$$L = L_R + \lambda L_\gamma \quad (3)$$

where $\lambda > 0$ is a weighting parameter. If $\lambda = 0$, Eq. (3) reduces to standard fine-tuning. Fig. 5 shows the overall architecture of proposed regularized transfer learning for anatomical landmark detection. Face recognition is considered as the source task to provide rich facial semantics. For the target task, as shown in Fig. 5 (a)-(d), there are four different regularization constraints that are developed and explored for the proposed learning approach.

1) *Constraint on Classifier Output (RTL-CO)*: Motivated by [22] [43], this variant constrains the distance between the

classification outputs for the design of L_γ , as shown in Fig. 5 (a). To do this, the previously abandoned classification branch $g_{\theta_c} : \mathbb{R}^m \rightarrow \mathbb{R}^C$ in Section III-A is re-enabled for training. Similar to the distillation loss [22], we use a temperature parameter μ with *softmax* function denoted as σ to smooth the predictions. While the original cross-entropy function is replaced by the following term,

$$L_{CO} = \frac{1}{N} \sum_{i=1}^N \left\| \sigma \left(\frac{g_{\theta_c}^s(f_{\theta_e}^s(\mathbf{I}_i))}{\mu} \right) - \sigma \left(\frac{g_{\theta_c}^t(f_{\theta_e}^t(\mathbf{I}_i))}{\mu} \right) \right\|_2^2 \quad (4)$$

The main reason for this replacement is that the task (label space) between the source and target domains is different, and applying cross-entropy function as a strong regularizer could make the optimization lean towards the source task, thereby affecting the target task performance. Moreover, our empirical results showed that using the L2 norm improves the cross-entropy function by 0.45 pixels regarding the mean square errors. Note that the superscripts s and t refer to the source and target (fine-tuned) networks respectively.

2) *Constraint on Encoder Output (RTL-EO)*: this variant considers the role of the regularization constraint, as shown in Fig. 5 (b). The goal is to align the features from encoder outputs, similar to hint learning [23]. Compared to L_{CO} , aligning the intermediate feature maps which preserve spatial configuration can be more effective and helpful for the landmark localization [47]. In this case, we adopt the cosine similarity for feature alignment as described in Eq. (5) [50], as we conjecture that penalizing higher-order angular differences would benefit matching the relational information between the source and target tasks.

$$L_{EO} = 1 - \frac{f_{\theta_e}^s(\mathbf{I}_i) \cdot f_{\theta_e}^t(\mathbf{I}_i)}{\|f_{\theta_e}^s(\mathbf{I}_i)\| \times \|f_{\theta_e}^t(\mathbf{I}_i)\|} \quad (5)$$

3) *Constraint on Spatial Attention Map (RTL-SAM)*: In addition to purely matching the predicted scores or feature maps between source and target tasks, it is important to understand how retainment of source task facial knowledge is beneficial to the target task. In this regard, we propose to constrain spatial attention maps which indicate where the meaningful features are emphasized given the input data. This is inspired by [24], [51] which demonstrated attention maps can be more interpretative than estimated feature maps. Fig. 4 (c) and Fig. 4 (d) demonstrate the advantage of *RTL-SAM* over that of *RTL-EO* on an example image. We specifically employ activation-based spatial attention maps [24] for the proposed *RTL-SAM*, as shown in Fig. 5 (c). Given an activation map $\mathbf{E}_i \in \mathbb{R}^{H \times W \times B} = f_{\theta_e}(\mathbf{I}_i)$ for the i -th sample of the target dataset, the corresponding spatial attention map $\mathbf{A}_i \in \mathbb{R}^{H \times W}$ is computed as [24],

$$\mathbf{A}_i = \sum_{b=1}^B |\mathbf{E}_i(b)|^2 \quad (6)$$

where B denotes the number of feature channels, $|\cdot|$ and $(\cdot)^2$ are the element-wise absolute value and power operators, respectively. We then reshape the $\mathbf{A}_i \in \mathbb{R}^{H \times W}$ into $\hat{\mathbf{A}}_i \in$

\mathbb{R}^{HW} and use it to formulate the following loss function [24],

$$L_{SAM} = \frac{1}{N} \sum_{i=1}^N \left\| \frac{\hat{\mathbf{A}}_i^s}{\|\hat{\mathbf{A}}_i^s\|_2} - \frac{\hat{\mathbf{A}}_i^t}{\|\hat{\mathbf{A}}_i^t\|_2} \right\|_2^2 \quad (7)$$

4) *Constraint on Channel Attention Map (RTL-CAM)*: The final variant considers channel-wise attention maps [51], [52] as a regularization constraint in our proposed framework. See Fig. 5 (d). Here we squeeze the spatial dimension of an activation map $\mathbf{E}_i \in \mathbb{R}^{H \times W \times B}$ to achieve a context vector $\mathbf{Q}_i \in \mathbb{R}^B$ [52], as given below,

$$\mathbf{Q}_i = \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W |\mathbf{E}_i(h, w)| \quad (8)$$

Similar to L_{SAM} , we formulate the regularization loss L_{CAM} as below,

$$L_{CAM} = \frac{1}{N} \sum_{i=1}^N \left\| \frac{\mathbf{Q}_i^s}{\|\mathbf{Q}_i^s\|_2} - \frac{\mathbf{Q}_i^t}{\|\mathbf{Q}_i^t\|_2} \right\|_2^2 \quad (9)$$

In practice, the four regularization terms L_{CO} , L_{EO} , L_{SAM} and L_{CAM} can be combined in the total loss function as described in 3. The effect on performance of doing this will be assessed in the following section. For the model inference, only the trained target encoder and decoder are utilized to predict heatmaps $\mathbf{H}_i = h_{\theta_d}^t(f_{\theta_e}^t(\mathbf{I}_i))$, and they are further processed via an *argmax* function to obtain final landmark locations $\hat{\mathbf{p}}_i$. Note that other methods such as fitting another radial basis function to obtain subpixel localization can be also utilized in here for keypoint extraction.

IV. EXPERIMENTS

A. Dataset

The data used in the experiments was collected by the Collaborative Initiative on Fetal Alcohol Spectrum Disorders (CIFASD)¹, a global multi-disciplinary consortium focused on furthering the understanding of FASD. It contains subjects from 4 sites across the USA, aged between 4 and 18 years. Each subject was imaged using a commercially available static 3D photogrammetry system from 3dMD². Since directly processing a high spatial resolution 3D mesh is computationally intensive, we utilize the 2D images captured during 3D acquisition, which are used as UV mapped textures for 3D surface reconstruction. We used in total 1549 facial scans for our experiments, and for each scan, there are 14 FAS associated landmarks which had been manually annotated by an expert (see Fig. 2 for details). These scans were randomly split into a training/validation set (70%/10%), and a test set (20%) for the experiments.

B. Implementation Details

All the images were cropped and resized to 256×256 pixels, and intensity normalized with mean and standard deviation (std) values provided by the VGGFace2 [16] for the network

¹<https://cifasd.org/>

²<http://www.3dmd.com/>

TABLE I
QUANTITATIVE COMPARISON BETWEEN PROPOSED AND OTHER RELATED EXISTING APPROACHES USING $MSE \pm STD$ AND FAILURE RATE ON THE TEST SET. BEST RESULTS ARE MARKED IN BOLD.

Method	$MSE \pm Std$ (↓)	Failure Rate (↓)			AUC (↑)		
		1.0 pixel	1.2 pixel	1.4 pixel	1.0 pixel	1.2 pixel	1.4 pixel
FE [33]	1.82 ± 0.50	99.03%	94.52%	80.00%	0.01	0.01	0.02
DSNT [7]	1.41 ± 0.26	96.13%	74.19%	50.65%	0.01	0.02	0.07
FTP [33]	1.16 ± 0.26	73.87%	40.32%	16.77%	0.04	0.10	0.19
U-Net [13]	0.93 ± 0.39	30.00%	10.00%	4.84%	0.15	0.26	0.36
HG [29]	0.88 ± 0.39	24.84%	12.58%	6.77%	0.20	0.30	0.39
FT [33]	0.86 ± 0.24	24.19%	10.65%	3.23%	0.18	0.29	0.38
<i>RTL-CO</i>	0.84 ± 0.25	24.84%	5.81%	1.61%	0.20	0.31	0.40
<i>RTL-EO</i>	0.83 ± 0.25	21.29%	7.74%	1.61%	0.21	0.32	0.41
<i>RTL-SAM</i>	0.78 ± 0.22	13.55%	4.84%	1.61%	0.24	0.35	0.44
<i>RTL-CAM</i>	0.82 ± 0.24	20.00%	5.81%	1.61%	0.21	0.32	0.42

training, validation and testing. Data augmentation was performed with randomly horizontal flip (50%) and scaling (0.8). During training, the Adam optimizer [53] was used for the optimization ($\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 1e^{-8}$) with a mini-batch size of 2 for 150 epochs. A polynomial decay learning rate was used with an initial value of 0.001. Models saved with the best accuracy on the validation set were used for evaluation. For more details of model training, Fig. 12 in Section VII-A depicts a learning curve of model performance on the train/validation sets. The temperature parameter μ in Eq. (4) was set to 2. We experimentally set up the weighting parameter $\lambda = 0.002$ used in Eq. (3) and the kernel scale parameter $\sigma = 1.5$ used in Eq. (1), and we used this parameter setting throughout the remaining experiments. The proposed learning algorithm was implemented using TensorFlow 1.6.0 and the model trained using an Nvidia GTX 1060 GPU.

C. Evaluation Metrics

We firstly employ the Mean Squared Error (MSE) for evaluation, which is a commonly-used metric to assess landmark detection. It is defined as:

$$MSE = \frac{1}{N_e} \sum_{i=1}^{N_e} \frac{1}{K} \|\mathbf{p}_i - \hat{\mathbf{p}}_i\|_2 \quad (10)$$

where N_e is the number of images in the test set, and \mathbf{p}_i and $\hat{\mathbf{p}}_i$ denote the ground truth and prediction landmark, respectively. Note that the original normalization factor estimated by interocular distance (Euclidean distance between outer eye corners) [54] is not included in this evaluation, as annotations for the other eye are not available, as shown in Fig. 7. In addition, we utilize the Cumulative Errors Distribution (CED) curve with two additional metrics of Area-Under-the-Curve (AUC) and Failure Rate (FR). When the Euclidean distance between the estimated landmark and its ground truth location (point-to-point Euclidean error) is greater than a pre-defined threshold, it is considered as a mis-detection. We consider three thresholds of 1.0 pixels, 1.2 pixels and 1.4 pixels.

D. Comparison with Existing Approaches

In this section, we evaluate the effect of the proposed regularized transfer learning (RTL) methods (*RTL-CO*, *RTL-EO*, *RTL-SAM*, and *RTL-CAM*) and compare these variants




Fig. 6. Illustration of cumulative error distribution (CED) curve on the test set. Best viewed in colored version.

with several related approaches from the literature. The related approaches are: a two-stage Hourglass network (HG) [29], a modified U-Net for landmark localization [13], and a fully-convolutional network with a differentiable spatial-to-numerical transform (DSNT) [7]. In addition, three variants of transfer learning [33] without regularization ($\lambda = 0$) are considered: Feature Extraction (FE) [33] with freezing the encoder, Fine Tuning Parts (FTP) [33] without freezing the final convolutional layer of the encoder, and Fine Tuning (FT) [33] without freezing any layer. For a fair comparison, all the evaluated methods use the same weight decay rate of 10^{-4} .

Table I summarizes the quantitative evaluation on the test set by reporting the MSE, failure rate, and AUC for each compared method. Performance metrics with (↓) or (↑) indicate that lower is better or higher is better, respectively. From the results, we observe that all the proposed new approaches favorably surpass the previous methods from the literature for most metrics. It is noteworthy that *RTL-SAM* reports the best performance among all the compared approaches. Compared to FT [33], which is closely related to our work, the proposed *RTL-SAM* improves the MSE by 0.08 pixels (9.3%), and the failure rate by 11.29%, 5.81%, and 1.62% for the three thresholds. This demonstrates that the proposed regularization constraint is beneficial for transfer learning across domain-similar source/target tasks.

In Fig. 6, we show the cumulative distributions of landmark




Fig. 7. Qualitative evaluation of landmark detection (first and third rows) and heatmap regression (second and fourth rows) on the test set. Best viewed in color version.

TABLE II
ABLATION STUDIES ON DIFFERENT PROPOSED REGULARIZATION CONSTRAINTS. THE UPPER PART SHOWS THE RESULTS FOR THE INDIVIDUAL LANDMARK LOCALIZATION USING $MSE \pm STD$ (PIXEL), WHILE THE LOWER PART PRESENTS THE FAILURE RATE WITH THE ERROR THRESHOLD OF 1.2 PIXEL. BEST RESULTS ARE SHOWN IN BOLD.

Landmark	Configurations						
	<i>RTL-CO</i>	<i>RTL-EO</i>	<i>RTL-SAM</i>	<i>RTL-CAM</i>	<i>RTL-CO-EO</i>	<i>RTL-CO-SAM</i>	<i>RTL-CO-CAM</i>
L1	0.64±0.58	0.63±0.59	0.59±0.56	0.65±0.58	0.64±0.58	0.61±0.60	0.62±0.56
L2	0.68±0.54	0.74±0.54	0.67±0.56	0.78±0.55	0.79±0.55	0.74±0.58	0.77±0.64
L3	0.82±0.63	0.89±0.60	0.74±0.59	0.77±0.84	0.82±1.12	0.75±0.60	0.73±0.72
L4	0.65±0.60	0.61±0.56	0.66±0.93	0.68±0.55	0.62±0.56	0.63±0.57	0.64±0.58
L5	0.71±0.57	0.64±0.56	0.67±0.55	0.69±0.56	0.67±0.58	0.66±0.57	0.73±0.56
L6	0.75±0.58	0.74±0.57	0.68±0.56	0.69±0.56	0.73±0.56	0.78±0.56	0.79±0.58
L7	1.48±1.17	1.46±1.13	1.20±1.02	1.46±1.22	1.44±1.10	1.58±1.29	1.28±1.05
L8	0.74±1.02	0.72±0.56	0.73±0.56	0.78±0.56	0.74±0.55	0.74±0.55	0.68±0.56
L9	1.40±1.05	1.39±1.15	1.17±1.19	1.28±1.17	1.30±1.19	1.50±1.33	1.32±1.13
L10	0.64±0.57	0.60±0.55	0.63±0.59	0.60±0.58	0.64±0.58	0.63±0.57	0.71±0.57
L11	0.89±1.06	0.75±0.60	0.80±0.58	0.84±0.60	0.83±0.61	0.83±0.57	0.91±1.07
L12	0.89±0.72	0.92±1.35	0.85±0.62	0.84±0.66	0.82±0.64	0.95±0.69	0.94±0.63
L13	0.72±0.60	0.70±0.61	0.74±0.60	0.63±0.59	0.71±0.62	0.69±0.60	0.76±0.60
L14	0.78±0.59	0.86±0.72	0.78±0.58	0.80±0.63	0.85±0.62	0.76±0.57	0.86±0.60
Average	0.84±0.25	0.83±0.24	0.78±0.22	0.82±0.24	0.83±0.25	0.85±0.25	0.86±0.33

detection errors for the methods considered. This reveals the advantage of the proposed approaches over the compared previous ones, which is consistent with the findings in Table I. From the CED curve, we observe that more than 95% of the testing images are successfully predicted within the error range of 1.3 pixels, and over 98% of the test set is well estimated within an error range of 1.4 pixels.

Fig. 7 presents a qualitative comparison between different approaches. From the figure, we observe that the predicted landmarks and heatmaps from the proposed methods generally align very well with the ground truth, and are less prone to ambiguities when landmarks are in close proximity (for instance, the upper lip).

E. Ablation Studies

1) *Impact of Different Regularization Constraints:* In this section, we present an ablation study to analyze the effectiveness of each proposed regularization constraint, in which

different combined regularization constraints are also added for evaluation. In addition to the original four regularization constraints, we study the effectiveness of combinations of both the classifier and encoder outputs (*RTL-CO-EO*), the classifier output with spatial attention map (*RTL-CO-SAM*), and the classifier output with channel attention map (*RTL-CO-CAM*).

Table II presents a detailed comparison of the different proposed approaches by listing the performance scores per individual landmark. From the results, we can see that *RTL-EO* and *RTL-SAM* achieve the highest number of best scores 4 and 5 landmarks respectively. This indicates that preserving and highlighting the feature outputs with spatial information from intermediate layers can be potentially more beneficial for the localization task. It can be also seen that for all the proposed variants the MSE ranged from 0.59 pixels to 1.58 pixels for different landmarks, where landmarks L1 and L2 corresponding to the exocanthion and endocanthion are well estimated with relatively lower scores than the others.




Fig. 8. AUC scores of cumulative error distribution (CED) curve versus point-to-point error threshold for different proposed regularization constraints. Best viewed in color version.

TABLE III

ABLATION STUDIES ON DIFFERENT PRE-TRAINED SOURCE TASKS. A FAILURE CASE HERE IS CONSIDERED IF THE POINT-TO-POINT EUCLIDEAN ERROR IS GREATER THAN 1.2 PIXEL. BEST RESULTS ARE MARKED IN BOLD.

Method	MSE \pm Std (\downarrow)	FR (\downarrow)	AUC (\uparrow)
FT-ImageNet [33]	0.85 \pm 0.37	9.68%	0.31
FT-YoutubeFaces [33]	0.86 \pm 0.30	8.06%	0.30
FT-VGGFace2 [33]	0.86 \pm 0.24	10.65%	0.29
<i>RTL-SAM-ImageNet</i>	0.82 \pm 0.32	8.39%	0.33
<i>RTL-SAM-YoutubeFaces</i>	0.84 \pm 0.25	8.06%	0.31
<i>RTL-SAM-VGGFace2</i>	0.78\pm0.22	4.84%	0.35

However, similar to the visual findings in Fig. 7, landmarks L7 and L9 marked as the crista philtri on both sides are seen as the most challenging cases.

Fig. 8 compares the effect of modelling with the different proposed regularization constraints using the AUC score calculated from the CED curve. Interestingly, the results indicate that combined approaches which introduce additional computational cost performed similarly or even worse than those of individual ones.

2) *Impact of Different Source Tasks*: To investigate the impact of different source tasks on the proposed framework, we re-implement methods FT [33] and our proposed *RTL-SAM* using the network ResNet-50 pre-trained on ImageNet [55] and a sub-dataset of YoutubeFaces [56] (facial alignment with 3462 images). We denote the evaluated methods as FT-ImageNet, FT-YoutubeFace, *RTL-SAM-ImageNet* and *RTL-SAM-YoutubeFaces*. We then compare them with those previously obtained using VGGFace2 [16] pre-trained weights. We observe from Table III, that by utilizing ImageNet and YoutubeFaces pre-trained weights, our proposed regularization strategy *RTL-SAM* outperforms FT in most evaluated metrics, which further confirms the effectiveness of our proposed method. Moreover, it is noteworthy that the proposed method with VGGFace2 pre-trained weights (*RTL-SAM-VGGFace2*) achieves the best performance among all compared methods. Interestingly, *RTL-SAM-YoutubeFaces* using the source task which is closely relevant to the target task, did not show advantage over *RTL-SAM-ImageNet* and *RTL-SAM-VGGFace2*.



Fig. 9. Illustration of top activation attention maps for models with different pre-trained source tasks: the source model (ResNet-50) pre-trained on ImageNet (ResNet-50-ImageNet), the source model (ResNet-50) pre-trained on VGGFace2 (ResNet-50-VGGFace2), the proposed regularization constraint on spatial attention maps transferred from ImageNet pre-trained weights (*RTL-SAM-ImageNet*), and the proposed regularization constraint on spatial attention maps transferred from VGGFace2 pre-trained weights (*RTL-SAM-VGGFace2*).




Fig. 10. Illustration of CED curve on validation set for different values of weighting parameter λ , where corresponding results using MSE \pm Std are shown in the parenthesis.

To provide further insight on this finding, we visualize the top activation attention maps for the models derived with the source tasks of ImageNet and VGGFace2, as shown in Fig. 9. We can see that the attention map generated from the source model (ResNet-50 [15]) pre-trained on VGGFace2 [16] (ResNet-50-VGGFace2) concentrates more on the desired facial region compared to ResNet-50-ImageNet. After refining the pre-trained models to the target task, the attention map produced from *RTL-SAM-VGGFace2* seems to be more informative than that of *RTL-SAM-ImageNet* for inferring the anatomical landmarks.

3) *Impact of λ* : To study the influence of weighting parameter choice on transfer learning regularization, we test a set of values for the parameter λ ranging from 0 to 2×10^{-5} on the validation set using the proposed *RTL-EO*, which has the simplest implementation among proposed methods. Note that when $\lambda = 0$, the designed model is evaluated without regularization, which is equivalent to standard fine tuning. Fig. 10 presents the results for varying parameter λ . This shows that the proposed approach reduces the localization error when adjusting the value of λ from 2×10^{-5} to 2×10^{-3} . We also observe that the setting of $\lambda = 2 \times 10^{-3}$, denoted as the red solid line, achieves the best performance considering both the CED curve and MSE. However, a larger value of $\lambda = 2 \times 10^{-2}$, regarded as a strong regularization during training, could affect model convergence for the landmark detection.

4) *Impact of σ* : Here we conduct a sensitivity analysis to study the impact of Gaussian kernel scaling parameter σ on the

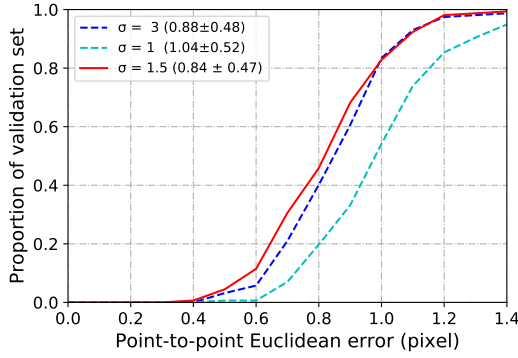


Fig. 11. Illustration of CED curve on validation set for different values of kernel scaling parameter σ , where corresponding results using $\text{MSE} \pm \text{Std}$ are shown in the parenthesis.

keypoint localization accuracy. Similar to the above analysis, we also utilize the proposed *RTL-EO* with $\lambda = 2 \times 10^{-3}$ to test a range of values for the parameter σ on the validation set. As seen from Fig. 11, we observe that the best performance evaluated in both metrics of CED and MSE is obtained with $\sigma = 1.5$, and the results in general are varied within a reasonable range.

V. DISCUSSION

A. Strengths and Limitations

In this work, we have shown the importance of transferring knowledge from a pre-trained source classification model for learning the anatomical landmark detection with a limited availability of training data. This observation was validated by the experimental study in Section IV-D, where transfer learning methods [33], achieve noticeably better performance than other compared solutions. More importantly, the strength of this study is that the proposed methods are shown to better preserve facial semantics embedded in the source model for learning the anatomical landmark detection compared to finetuning. This can be best seen from the evaluations both quantitatively and qualitatively in Section IV-D, where our proposed methods outperform the method of FT [33]. This property of the proposed framework is particularly useful to reduce ambiguity in landmark detection, where landmarks which are in close proximity can be semantically distinguished by reusing pre-computed facial semantics.

However, the prerequisite for the proposed framework is the utilization of a pre-trained source task model, which may not be always available for other clinical applications. In Section IV-E2, we have demonstrated that the proposed approach performs better when incorporating the knowledge from VGGFace2 which is more relevant to the target task compared to that of ImageNet. Another important finding is that the proposed knowledge transfer with larger source datasets (ImageNet 14 million images and VGGFace2 3.31 million images) can generalize better on the target task regardless of task relevance.

Moreover, evaluated results in Section IV-E3 have shown that the proposed framework with different settings of the weighting parameter has a generally positive impact on the

transfer learning process for anatomical landmark detection, but these empirical configurations may be only specific to the addressed clinical application. Hence, how to adaptively determine the required amount of domain knowledge for a given target task remains an open question.

B. Translational Potential

The motivating clinical need for accurate, anatomical landmark placement in fetal alcohol syndrome (FAS) analysis is 1) to identify linear and angular facial anthropometric measurements such as palpebral fissure length (PFL) and 2) to provide a surface-based analysis of facial form to compare an individual to unaffected controls. We link this need to our research findings. The proposed framework delivered particularly good performance for landmarks L1 and L2 (exocanthion and endocanthion), as presented in Section IV-E1, which is important for accurate estimation of PFL. Previous analysis using a dense surface model (DSM) for the surface-based study of facial shape has shown the ability to identify subtle facial dysmorphism across FAS [6], [11], [57], [58]. However, to date, their construction has been reliant upon the manual placement of a sparse set of anatomical landmarks. Based on the evaluations in Fig. 6, more than 95% of the testing cases have been accurately estimated by the proposed framework, where *RTL-SAM* and *RTL-EO* perform the best. We believe that the proposed framework can provide an automated alternative to the manual anthropometric placement, and can be readily embedded in a clinical prediction tool owing to its real-time speed about 44 frames per second (FPS). Specifically the proposed methods of *RTL-SAM* and *RTL-EO* are recommended for the integration of a DSM into the clinical workflow for better identification of subjects with FAS.

VI. CONCLUSIONS

In this paper, we presented a new regularized transfer learning framework to address the problem of anatomical landmark detection in FAS. The proposed learning approach reused the rich visual semantics from a domain-similar source model as guidance for learning the landmark detector. Four regularization constraints were developed and investigated in the proposed framework, including constraining the feature outputs from classification and intermediate layers, as well as matching activation attention maps in both the spatial and channel levels. Experimental results show that the proposed learning framework performs well with limited training samples and outperforms other compared solutions. The generalization of the proposed learning framework to other clinical applications with the availability of domain or medical expert knowledge would be an interesting research avenue. Future work will integrate the automated anatomical landmarks with a dense surface model (DSM) into the clinical workflow for better early identification of those with FAS.

REFERENCES

- [1] D. Han, Y. Gao, G. Wu, P.-T. Yap, and D. Shen, "Robust anatomical landmark detection for MR brain image registration," in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2014, pp. 186–193.

- [2] K. Oh, I. Oh, T. V. n. Le, and D. LEE, "Deep anatomical context feature learning for cephalometric landmark detection," *IEEE Journal of Biomedical and Health Informatics*, pp. 1–1, 2020.
- [3] R. Chen, Y. Ma, N. Chen, D. Lee, and W. Wang, "Cephalometric landmark detection by attentive feature pyramid fusion and regression-voting," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2019, pp. 873–881.
- [4] J. Zhang, M. Liu, and D. Shen, "Detecting anatomical landmarks from limited medical imaging data using two-stage task-oriented deep neural networks," *IEEE Transactions on Image Processing*, vol. 26, no. 10, pp. 4753–4764, 2017.
- [5] J. M. H. Noothout, B. D. De Vos, J. M. Wolterink, E. M. Postma, P. A. M. Smeets, R. A. P. Takx, T. Leiner, M. A. Viergever, and I. Išgum, "Deep learning-based regression and classification for automatic landmark localization in medical images," *IEEE Transactions on Medical Imaging*, pp. 1–1, 2020.
- [6] M. Suttie *et al.*, "Facial dysmorphism across the fetal alcohol spectrum," *Pediatrics*, vol. 131, no. 3, pp. e779–e788, 2013.
- [7] R. Huang, M. Suttie, and J. A. Noble, "An automated CNN-based 3D anatomical landmark detection method to facilitate surface-based 3D facial shape analysis," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI) Workshops*, 2019, pp. 163–171.
- [8] K. L. Jones, L. K. Robinson, L. N. Bakhireva *et al.*, "Accuracy of the diagnosis of physical features of fetal alcohol syndrome by pediatricians after specialized training," *Pediatrics*, vol. 131, no. 3, pp. e1734–e1738, 2006.
- [9] H. E. Hoyme *et al.*, "Updated guidelines for fetal alcohol spectrum disorders," *Pediatrics*, vol. 138, no. 2, 2016.
- [10] S. J. Astley, "Palpebral fissure length measurement: accuracy of the FAS facial photographic analysis software and inaccuracy of the ruler," *Journal of Population Therapeutics and Clinical Pharmacology*, vol. 22, no. 1, pp. e9–e26, 2015.
- [11] P. Hammond *et al.*, "Discriminating power of localized three-dimensional facial morphology," *American journal of human genetics*, vol. 77, no. 6, pp. 999–1010, 2005.
- [12] D. Wargowski. Screening, Diagnosis and Assessment of FAS. [Online]. Available: https://www.youtube.com/watch?v=044Zxy3_Ou8
- [13] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 2015, pp. 234–241.
- [14] A. Z. K. Simonyan, "Very deep convolutional networks for large-scale image recognition," in *Proc. of International Conference on Learning Representations (ICLR)*, 2015, pp. 1–14.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [16] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vggface2: A dataset for recognising faces across pose and age," in *IEEE International Conference on Automatic Face Gesture Recognition*, 2018, pp. 67–74.
- [17] O. Wiles, A. Koepke, and A. Zisserman, "Self-supervised learning of a facial attribute embedding from video," in *British Machine Vision Conference (BMVC)*, 2018.
- [18] Y. Zhao, Y. Liu, C. Shen, Y. Gao, and S. Xiong, "MobileFAN: Transferring deep hidden representation for face alignment," *Pattern Recognition*, vol. 100, pp. 107–114, 2020.
- [19] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [20] X. Li, Y. Grandvalet, and F. Davoine, "Explicit inductive bias for transfer learning with convolutional networks," in *International Conference on Machine Learning (ICML)*, vol. 80, 2018, pp. 2830–2839.
- [21] X. Li, H. Xiong, H. Wang, Y. Rao, L. Liu, Z. Chen, and J. Huan, "DELTA: Deep learning transfer using feature map with attention for convolutional networks," in *Proc. of International Conference on Learning Representations (ICLR)*, 2019, pp. 1–13.
- [22] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *Conference on Neural Information Processing Systems (NeurIPS) Workshops*, 2015.
- [23] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "Fitnets: Hints for thin deep nets," in *Proc. of International Conference on Learning Representations (ICLR)*, 2015, pp. 1–13.
- [24] S. Zagoruyko and N. Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," in *Proc. of International Conference on Learning Representations (ICLR)*, 2017, pp. 1–13.
- [25] Z. Fu, J. Jiao, M. Suttie, and J. A. Noble, "Cross-task representation learning for anatomical landmark detection," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI) Workshops*, 2020, pp. 1–9.
- [26] M. Urschler, T. Ebner, and D. Štern, "Integrating spatial configuration into heatmap regression based CNNs for landmark localization," *Medical Image Analysis*, vol. 43, pp. 23–36, 2018.
- [27] X. Cao, Y. Wei, F. Wen, and J. Sun, "Face alignment by explicit shape regression," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 2887–2894.
- [28] P. Dollár, P. Welinder, and P. Perona, "Cascaded pose regression," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 1078–1085.
- [29] A. Newell, K. Yang, and J. Deng, "Stacked hourglass networks for human pose estimation," in *European Conference on Computer Vision (ECCV)*, 2016, pp. 483–499.
- [30] Z.-H. Feng, J. Kittler, M. Awais, P. Huber, and X.-J. Wu, "Wing loss for robust facial landmark localisation with convolutional neural networks," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [31] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [32] Z. Zhang, P. Luo, C. C. Loy, and X. Tang, "Learning deep representation for face alignment with auxiliary attributes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 5, pp. 918–930, 2016.
- [33] B. Xiao, H. Wu, and Y. Wei, "Simple baselines for human pose estimation and tracking," in *European Conference on Computer Vision (ECCV)*, 2018, pp. 472–487.
- [34] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele, "2d human pose estimation: New benchmark and state of the art analysis," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 3686–3693.
- [35] J. Zhang, M. Liu, and D. Shen, "Detecting anatomical landmarks from limited medical imaging data using two-stage task-oriented deep neural networks," *IEEE Transactions on Image Processing*, vol. 26, no. 10, pp. 4753–4764, 2017.
- [36] C. Payer, D. Štern, H. Bischof, and M. Urschler, "Integrating spatial configuration into heatmap regression based CNNs for landmark localization," *Medical Image Analysis*, vol. 54, pp. 207–219, 2019.
- [37] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker, and D. Rueckert, "Attention U-Net: Learning where to look for the pancreas," in *International Conference on Medical Imaging with Deep Learning (MIDL)*, 2018, pp. 1–10.
- [38] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. omez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [39] Z. Zhong, J. Li, Z. Zhang, Z. Jiao, and X. Gao, "An attention-guided deep regression model for landmark detection in cephalograms," in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2019, pp. 540–548.
- [40] C. Liu, H. Xie, S. Zhang, Z. Mao, J. Sun, and Y. Zhang, "Misshapen pelvis landmark detection with local-global feature learning for diagnosing developmental dysplasia of the hip," *IEEE Transactions on Medical Imaging*, pp. 1–1, 2020.
- [41] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 7794–7803.
- [42] A. Vlontzos, A. Alansary, K. Kamnitsas, D. Rueckert, and B. Kainz, "Multiple landmark detection using multi-agent reinforcement learning," in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2019, pp. 262–270.
- [43] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 12, pp. 2935–2947, 2018.
- [44] P. Dhar, R. V. Singh, K.-C. Peng, Z. Wu, and R. Chellappa, "Learning without memorizing," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [45] A. Patra and J. A. Noble, "Hierarchical class incremental learning of anatomical structures in fetal echocardiography videos," *IEEE Journal of Biomedical and Health Informatics*, vol. 24, no. 4, pp. 1046–1058, 2020.
- [46] B. S. V. V. David Lopez-Paz, Léon Bottou, "Unifying distillation and privileged information," pp. 1–10, 2016.

- [47] S. Gupta, J. Hoffman, and J. Malik, “Cross modal distillation for supervision transfer,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [48] A. Patra, Y. Cai, P. Chatelain, H. Sharma, L. Drukker, A. T. Papa-georgiou, and J. A. Noble, “Efficient ultrasound image analysis models with sonographer gaze assisted distillation,” in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2019, pp. 394–402.
- [49] X. Glorot, A. Bordes, and Y. Bengio, “Deep sparse rectifier neural networks,” in *Proc. of Artificial Intelligence and Statistics (AISTATS)*, 2011, pp. 315–323.
- [50] W. Park, D. Kim, Y. Lu, and M. Cho, “Relational knowledge distillation,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [51] S. Woo, J. Park, J.-Y. Lee, and I. So Kweon, “CBAM: Convolutional block attention module,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [52] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 7132–7141.
- [53] D. P. Kingma and J. Ba, “Adam: a method for stochastic optimization,” in *Proc. of International Conference on Learning Representations (ICLR)*, 2015, pp. 1–15.
- [54] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, “300 faces in-the-wild challenge: The first facial landmark localization challenge,” in *IEEE International Conference on Computer Vision (ICCV) Workshops*, 2013, pp. 397–403.
- [55] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255.
- [56] L. Wolf, T. Hassner, and I. Maoz, “Face recognition in unconstrained videos with matched background similarity,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [57] M. Suttie *et al.*, “Facial curvature detects and explicates ethnic differences in effects of prenatal alcohol exposure,” *Alcohol Clin Exp Res*, vol. 41, no. 8, pp. 1471–1483, 2017.
- [58] —, “Combined face-brain morphology and associated neurocognitive correlates in fetal alcohol spectrum disorders,” *Alcohol Clin Exp Res*, vol. 42, no. 9, pp. 1769–1782, 2018.

VII. APPENDIX

A. Visualization of Training Details




Fig. 12. Illustration of training and validation loss curve.

B. Analysis of Adding Skip Connections

The baseline network architecture [33] used in the proposed learning framework does not include skip connections like U-Net [13]. Here, we run a pilot test to study whether adding the skip connections between the encoder and decoder can improve the keypoint localization accuracy. For simplicity, we modify one of the baseline methods FT-ImageNet for this analysis. The results in Table IV show that both MSE \pm Std and AUC score of FT-ImageNet are further improved by adding the skip connections. This validates the effectiveness of inclusion

skip connections and suggests that such network design should be preferred for high-resolution keypoint localization in the future.

TABLE IV

ANALYSIS OF ADDING SKIP CONNECTIONS. A FAILURE CASE HERE IS CONSIDERED IF THE POINT-TO-POINT EUCLIDEAN ERROR IS GREATER THAN 1.2 PIXEL. BEST RESULTS ARE MARKED IN BOLD.

Method	MSE \pm Std (\downarrow)	FR (\downarrow)	AUC (\uparrow)
Without skip connections	0.85 \pm 0.37	9.68%	0.31
With skip connections	0.83\pm0.27	11.9%	0.33