



Mitigating Knowledge Discrepancies among Multiple Datasets for Task-agnostic Unified Face Alignment

Jiahao Xia¹ · Min Xu¹ · Wenjian Huang² · Jianguo Zhang² · Haimin Zhang¹ · Chunxia Xiao³

Received: 1 October 2024 / Accepted: 25 June 2025
© The Author(s) 2025

Abstract

Despite the similar structures of human faces, existing face alignment methods cannot learn unified knowledge from multiple datasets with different landmark annotations. The limited training samples in a single dataset commonly result in fragile robustness in this field. To mitigate knowledge discrepancies among different datasets and train a task-agnostic unified face alignment (TUFA) framework, this paper presents a strategy to unify knowledge from multiple datasets. Specifically, we calculate a mean face shape for each dataset. To explicitly align these mean shapes on an interpretable plane based on their semantics, each shape is then incorporated with a group of semantic alignment embeddings. The 2D coordinates of these aligned shapes can be viewed as the anchors of the plane. By encoding them into structure prompts and further regressing the corresponding facial landmarks using image features, a mapping from the plane to the target faces is finally established, which unifies the learning target of different datasets. Consequently, multiple datasets can be utilized to boost the generalization ability of the model. The successful mitigation of discrepancies also enhances the efficiency of knowledge transferring to a novel dataset, significantly boosts the performance of few-shot face alignment. Additionally, the interpretable plane endows TUFA with a task-agnostic characteristic, enabling it to locate landmarks unseen during training in a zero-shot manner. Extensive experiments are carried on seven benchmarks and the results demonstrate an impressive improvement in face alignment brought by knowledge discrepancies mitigation. The code is available at <https://github.com/Jiahao-UTS/TUFA>

Keywords Face alignment · Knowledge discrepancy mitigation · Semantic alignment · Structure prompt · Few-shot learning · Zero-shot-learning

1 Introduction

Facial landmark serves as an important intermediate feature in many downstream tasks, such as face recognition (Xu et al., 2021), facial emotion recognition (Oterboudt, Daoudi, Kacem, Ballihi, and Berretti, 2022; Tanfous, Drira, and Amor, 2020) and facial image synthesis (H. Tang, Shao, Torr, and Sebe, 2023). Despite significant progress in recent years, few existing face alignment methods focus on mitigating the knowledge discrepancies among datasets, resulting in very limited usage of dataset knowledge. Consequently, despite the similar structure of human faces, most existing methods can be only trained on a single dataset. The limited

Communicated by Xin Zhao.

✉ Min Xu
Min.Xu@uts.edu.au
Jiahao Xia
Jiahao.Xia-1@uts.edu.au
Wenjian Huang
huangwj@sustech.edu.cn
Jianguo Zhang
zhangjg@sustech.edu.cn
Haimin Zhang
Haimin.Zhang@uts.edu.au
Chunxia Xiao
cxxiao@whu.edu.cn

¹ Faculty of Engineering and IT, University of Technology Sydney, Ultimo 2007, NSW, Australia

² Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, Guangdong, China

³ School of Computer Science, Wuhan University, Wuhan 430072, Hubei, China

training samples lead to unsatisfied generalization ability. When applying a model to a group of newly defined landmarks, the entire model still needs to be re-trained. As a result, these knowledge discrepancies lead to performance degradation and high costs for downstream tasks.

As shown in Fig. 1, similar to other vision tasks (Pourpanah et al., 2023; G. Huang, Laradji, Vázquez, Lacoste-Julien, and Rodríguez, 2023), there are three types of methods that can train a face alignment model: many-shot, few-shot, and zero-shot learning. Many-shot learning is a widely used method in face alignment. It does not consider sharing knowledge among multiple datasets, directly training a model on a dataset with consistent landmark annotations. Hence, the training can be very costly and time-consuming. To relief the reliance on the number of training samples, few-shot face alignment has been explored in recent works (Browatzki & Wallraven, 2020; He, Bharaj, Ferman, Rhodin, and Garrido, 2023). Similar to few-shot learning in other tasks (G. Huang, Laradji, Vázquez, Lacoste-Julien, and Rodríguez, 2023; Jiang et al., 2023), few-shot face alignment transfers the knowledge learned from a large-scale base dataset to the target landmarks using very few samples. Explicitly mitigating the knowledge discrepancies between the base and target datasets can significantly increase the efficiency of knowledge transferring, while this problem is ignored by existing methods. Further mitigating these discrepancies can extend the learned knowledge from predefined landmarks to the overall face structure, making the model task-agnostic. Consequently, using semantic information as a prompt can locate landmarks that were unseen during training in a zero-shot manner. Unfortunately, no work has been done on this topic yet.

In this paper, we propose a novel strategy to mitigate the knowledge discrepancies among multiple datasets for task-agnostic unified face alignment (TUFA). Instead of treating each facial landmark as an independent regression target, our strategy uses labeled facial landmarks to learn a mapping from an interpretable plane to the target faces. This approach unifies the learning objective and provides a foundation for mitigating knowledge discrepancies. Specifically, we calculate a mean face shape for each dataset. To align the mean shapes from multiple datasets on the same plane according to their semantics and eliminates their semantic ambiguity, we incorporate each mean shape with a group of semantic alignment embeddings. They transform these mean shapes during training and finally find optimal positions on the plane to represents their semantics. Consequently, the semantic ambiguity can be eliminated. Each point from these aligned shapes from different datasets can be view a anchor on the plane. As a result, the learning targets of datasets are unified and their knowledge discrepancy are largely mitigated. By encoding them into structure prompts and regressing the positions of

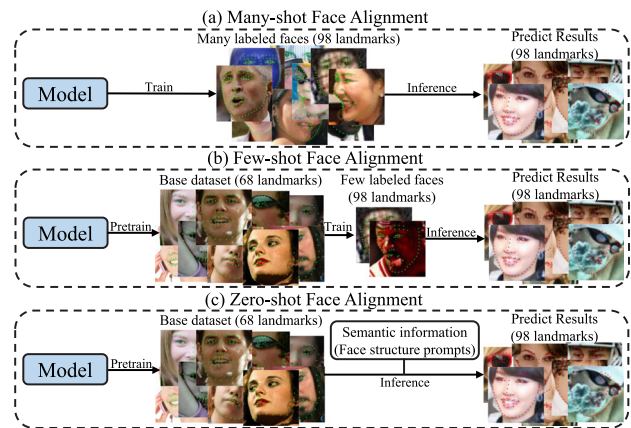


Fig. 1 A schematic diagram of many-shot, few-shot, and zero-shot learning in face alignment. (a) many-shot face alignment directly trains the model on a large-scale dataset labeled with target landmarks. (b) few-shot face alignment pre-trains the model on a base dataset, and then further transfer the model using a few samples labeled with target landmarks. (c) zero-shot face alignment also pre-trains the model on a base dataset, while the model directly locates the target landmarks via semantic information without training

their corresponding landmarks, the mapping to target face can be established using multiple datasets.

Moreover, previous face-alignment methods based on heatmap regression and coordinate regression treat each landmark regression as a separate task, resulting in task-specific knowledge. Although human faces share a similar structure, it remains challenging to extend this knowledge to unseen facial landmarks. Nevertheless, by mitigating the knowledge discrepancies using the proposed strategy, the learned knowledge is not only limited to the labeled landmarks in training but also includes **task-agnostic** knowledge: the entire face structure. The learned task-agnostic knowledge can be efficiently transferred to a group of newly defined landmarks using only a few labeled samples, which significantly boosts the performance of few-shot face alignment. Moreover, by explicitly representing the face on an interpretable plane and encoding its coordinates into structure prompts, this knowledge can also be used to localize unseen facial landmarks without fine-tuning. Therefore, TUFA also achieves zero-shot face alignment for the first time, predicting not only the pre-defined landmarks, but also those landmarks unseen during training, as shown in Fig. 2.

We train TUFA using four widely used datasets and carry out four **within-dataset** validations and three **cross-dataset** validations in a many-shot manner. TUFA achieves state-of-the-art performance on these seven validations using **only one unified model**. This pre-trained model demonstrates strong generalization ability, making it particularly useful for further few-shot face alignment. Thanks to the mitigated discrepancies, the TUFA trained with only 10 labeled samples can even outperform the state-of-the-art few-shot face align-

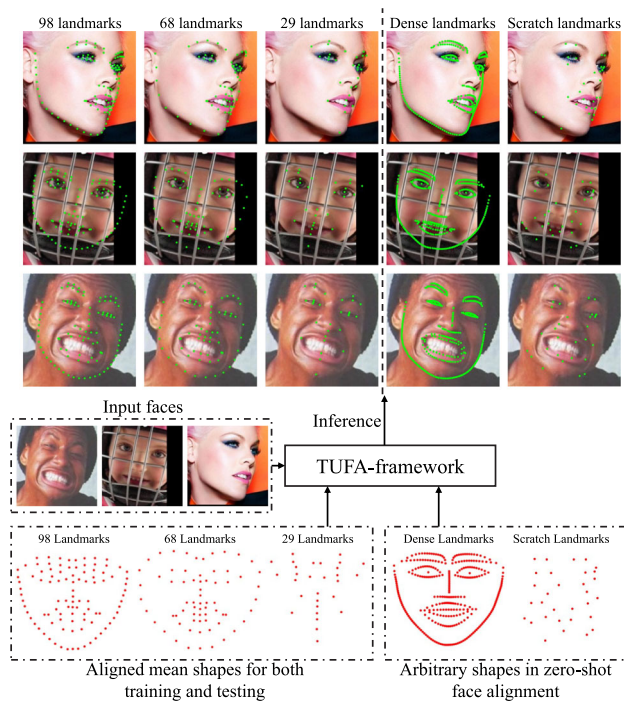


Fig. 2 Teaser of TUFA. TUFA employs the aligned face shapes as anchors to mapping an interpretable plane to target faces. Consequently, it maximizes the knowledge use from multiple datasets to improve the generalization ability. By explicitly representing face structure using this plane, TUFA is able to locate both pre-defined landmarks and those landmarks unseen during training in inference phase

ment methods trained on the entire training set. Moreover, due to its task-agnostic property, TUFA achieves zero-shot face alignment for the first time. The quantitative results show that TUFA even outperforms some fully supervised methods in a **zero-shot** manner.

Overall, the main contributions of this paper can be summarized as follows:

- This paper presents a novel strategy to mitigate knowledge discrepancies among multiple face alignment datasets by unifying their learning objectives and eliminating their semantic ambiguity. Based on this strategy, a task-agnostic unified face alignment model with stronger generalization ability can be trained across multiple datasets.
- To efficiently transfer the knowledge learned across multiple datasets to newly defined landmarks, we propose a few-shot and zero-shot face alignment strategy using the structural information contained in the proposed structure prompts and the learned interpretable plane.
- Extensive experiments and ablation studies are carried out on widely used benchmarks. TUFA achieves state-of-the-art performance in the setting of many-shot, few-shot and zero-shot face alignment, which demonstrates the

strong generalization ability and transferable capability of TUFA.

2 Related Works

It has been noticed that the performance of **many-shot face alignment** has improved significantly in recent years. However, the high cost of dataset annotation has shifted the focus of recent works towards **few-shot & zero-shot face alignment**. Moreover, the existing models still generalize poorly in a cross-dataset validation setting, especially for those landmarks have a big gap compared to the seen landmarks. **Multi-dataset learning in face alignment**, which is a potential strategy to address this problem, has also gained a lot of attention in recent years. Therefore, we present a concise review focusing on these three topics, given that they represent the primary challenges tackled in this paper.

2.1 Many-shot Face Alignment

As a very pivotal technology in computer vision, many-shot face alignment has undergone over two decades of intensive research and exploration. At the very early stage, constrained local model (CLM) (Cristinacce & Cootes, 2006), active shape model (ASM) (Cootes, Taylor, Cooper, and Graham, 1995), active appearance model (AAM) (Cootes, Edwards, and Taylor, 2001; Asthana, Marks, Jones, Tieu, and Rohith, 2011; Liu, 2009) dominated this task. Their core idea is to align a statistical mean shape to a new face and that is why this task is named as face alignment rather than facial landmark detection. To overcome the fragile robustness and generalization, cascaded shape regression (CSR) methods (Burgos-Artiz, Perona, and Dollár, 2013; Xiong & De la Torre, 2013; Z.-H. Feng, Kittler, Christmas, Huber, and Wu, 2017; S. Zhu, Li, Loy, and Tang, 2015; Z. Feng, Kittler, Christmas, Huber, and Wu, 2017) are proposed to achieve face alignment with multi-steps. Lately, the development of CNN and vision transformer (ViT) (Dosovitskiy et al., 2021) significantly boost the performance of face alignment. Recent many-shot face alignment methods can be roughly divided into two categories: heatmap regression methods and coordinate regression methods. They set each landmark as an independent regression target and their learned models are task-specific.

Coordinate regression methods directly project image features into 2D landmark coordinates with fully connected (FC) layers. To extract more fine-grained image feature, cascaded networks (Lv, Shao, Xing, Cheng, and Zhou, 2017; Kowalski, Naruniec, and Trzcinski, 2017; H. Liu, Lu, Guo, Wu, and Zhou, 2020), recurrent networks (Trigeorgis, Snape, Nicolaou, Antonakos, and Zafeiriou, 2016; S. Xiao et al., 2016) and face local patches (H. Liu, Lu, Guo, Wu, and Zhou,

2020; C. Zhu, Wan, Xie, Li, and Gu, 2022) are widely used. The regular shape of human faces serves as an important prompt in face alignment, and recent coordinate regression methods further improve the performance by retaining the prompt in networks. Lin et al. (Lin et al., 2021) represent face structure with an adjacent matrix and Li et al. (Li et al., 2020) make the matrix learnable. Xia et al. (Xia et al., 2022), Xia et al. (Xia et al., 2023) and Prados-Torreblanca et al. (Prados-Torreblanca, Buenaposada, and Baumela, 2022) further encourage the model to learn a case-dependent inherent relation based on the attention mechanism. However, the learned face structure is represented in high-dimensional latent space, making it hard to be understood and edited by human. In this paper, TUFA represents the face structure on an interpretable plane, making the prompt can be edited by human for the first time.

Heatmap regression methods differ significantly from coordinate regression methods, as they generate a high-resolution likelihood heatmap for each landmark using a deep network (J. Wang et al., 2021; Newell, Yang, and Deng, 2016; B. Xiao, Wu, and Wei, 2018; Z. Tang, Peng, Li, and Metaxas, 2020). Although the semantic supervision enables heatmap regression methods to perform better with limited training samples, they still have some drawbacks. For instance, to transfer the predicted heatmaps into 2D coordinates, they consider the pixel with highest intensity as the optimal output. Therefore, the predicted coordinates can only be integers, resulting in a quantitative error due to lower resolution of the output heatmap than the input image. To mitigate the error, Zhang et al. (J. Zhang, Hu, and Feng, 2020), Jin et al. (Jin, Liao, and Shao, 2021) and Lan et al. (Lan, Hu, and Cheng, 2021) predict the decimal part of landmark coordinate on heatmaps with an additional head; Tai et al. (Tai et al., 2019), Chen et al. (Chen, Su, and Ji, 2019), Kumar et al. (Kumar et al., 2020) and Zhou et al. (Zhou et al., 2023) achieve subpixel landmark localization on heatmap by estimating landmark distribution. Furthermore, it is more challenging for heatmap regression methods to model the regular structure of human faces because they do not explicitly learn a representation for each individual landmark. Despite the utilization of facial boundaries as a key indicator in heatmap regression models (Wu et al., 2018; X. Wang, Bo, and Fuxin, 2019; Y. Huang, Yang, Li, Kim, and Wei, 2021) to maintain face structure, these models still overlook the long-term relation that exists among landmarks. Their network structure also makes it impossible for them to predict an arbitrary number of landmarks. Therefore, for any newly defined landmark, these methods must be re-trained.

2.2 Few-shot & Zero-shot Face Alignment

Given that existing methods require both model re-training and dataset re-annotation for predicting a group of newly

defined landmarks, recent research (Qian, Sun, Wu, Qian, and Jia, 2019; Browatzki & Wallraven, 2020; He, Bharaj, Ferman, Rhodin, and Garrido, 2023) has increasingly focused few-shot face alignment. It enables the swift adaptation to new face alignment tasks. Qian et al. (Qian, Sun, Wu, Qian, and Jia, 2019) expand the training set by style transfer so that the model learns a more generalized result with a limited number of training samples. However, the synthetic images have a significant domain gap to the real images, which can lead to performance degradation. And, the style transfer network still requires hundreds of samples for training. Browatzki et al. (Browatzki & Wallraven, 2020) and He et al. (He, Bharaj, Ferman, Rhodin, and Garrido, 2023) set face reconstruction as a pre-text task. With the feature learned from reconstruction, the models can be transferred into face alignment with few-shot examples. Nevertheless, as the experiment results shown in (He, Bharaj, Ferman, Rhodin, and Garrido, 2023), the large discrepancies between face alignment and image reconstruction can result in performance degradation when the samples for fine-tuning are relatively sufficient. Therefore, mitigating knowledge discrepancies in few-shot face alignment is quite crucial, but no study has focused on this problem yet. Self-supervised (Mallis, Sanchez, Bell, and Tzimiropoulos, 2023) and unsupervised landmark detection (Thewlis, Bilen, and Vedaldi, 2017; Y. Zhang et al., 2018; Lorenz, Bereska, Milbich, and Ommer, 2019; Jakab, Gupta, Bilen, and Vedaldi, 2018; He, Wandt, and Rhodin, 2022) have also gained more attention as they eliminate the need for manual annotation. Nevertheless, the unpredictability of semantic definitions for the landmarks discovered during training, which cannot be manually assigned, presents a challenge when applying them to downstream tasks. Zhang et al. (H. Zhang et al., 2024) align textual prompts with visual features to extend landmark detection model to unseen categories in zero-shot manner. However, textual prompts are difficult to describe the structure relations among landmarks, making it hard to extend the model to unseen landmarks. This work implements zero-shot learning in landmark-level for the first time, significantly broadening the application range of face alignment.

3 Multi-datasets Learning in Face Alignment

In this section, we first introduce the overall training pipeline for TUFA and explain its approach to achieving task-agnostic face alignment. Next, we describe the proposed method for mitigating the knowledge discrepancies among different datasets, as well as the loss function of TUFA. Finally, we discuss the strategy for achieving few-shot and zero-shot face alignment using TUFA.

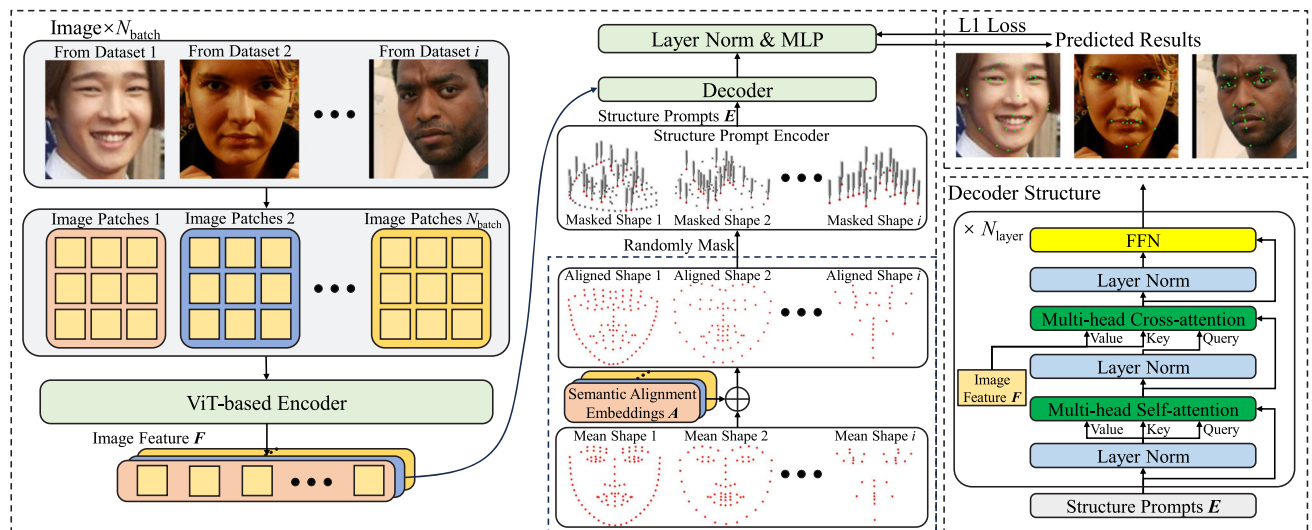


Fig. 3 Overall training pipeline of TUFA. It mainly consists of three parts: a ViT-based encoder, a structure prompt encoder and a decoder. TUFA first splits input images into patches and extracts image features F with a ViT-based encoder. Each image has a corresponding statistical mean shape determined by its dataset. For multi-dataset learning, the learnable semantic alignment embeddings are added to the mean shapes to align their semantics on a shared plane. During each iteration, TUFA

randomly masks these input shapes, keeping only a certain number of landmarks as the input anchors to form a batch. Then, the structure prompt encoder encodes 2D positions of the input anchors into structure prompts E using sine and cosine functions. The structure prompts finally serve as queries, mapping the mean shapes to target landmarks based on the image features. The detailed structure of the decoder is shown at the bottom right

3.1 Task-agnostic Unified Face Alignment

The overall training pipeline of TUFA is demonstrated in Fig. 3. It mainly consists of three parts: a ViT-based encoder for image feature extraction, a structure prompt encoder that generates structure prompts from the 2D coordinates of the input anchors, and a decoder aimed at mapping the input shapes to target faces.

3.1.1 ViT-based Encoder

We employ a ViT-based encoder to extract an image features F that later serve as essential clues in constructing the mapping between the pre-defined plane and various target faces. Given an input image $I \in \mathbb{R}^{H_I \times W_I \times 3}$, the ViT-based encoder first splits it into regular patches with size (H_P, W_P) . Then, each patch is projected into a vector with C dimensions by a CNN layer and is further added with learnable positional embeddings $P \in \mathbb{R}^{L \times C}$ to retain spatial information, where $L = \frac{H_I}{H_P} \times \frac{W_I}{W_P}$. Finally, the encoder learns the image feature $F \in \mathbb{R}^{L \times C}$ via the attention mechanism. F is further fed into the cross-attention block of the decoder.

3.1.2 Structure Prompt Encoder

A human face can be represented on a 2D plane because it is the surface of the human head and has a regular shape. During single dataset learning, to constrain the representation

of the human face to the 2D plane, we calculate a statistical mean shape from the training samples, and the coordinates of the mean shape are normalized in the range of $[-1, 1]$. The mean shape, which explicitly represents the regular structure of faces, later serves as the anchors of this plane. By regressing target landmarks based on the structure prompts of these anchors, we can establish a mapping between the 2D plane and various faces. Moreover, the mean shape also ensures that the face structure representation on the plane is easily understood by humans, allowing the structure prompts to be readily edited. During multi-dataset learning, the additional semantic alignment embeddings A should be incorporated into the mean shapes to mitigate the knowledge discrepancies among multiple datasets. The details of the semantic alignment embeddings A will be discussed in Section 3.2.

Before feeding the mean shapes into the structure prompt encoder, TUFA randomly masks part of these shapes during each iteration, retaining only N_a landmarks as the input anchors (N_a is set to 24, 25% of 98 landmarks). The use of the mean shape masking is threefold: 1) the masked mean shape is much sparser, encouraging TUFA to learn the mapping based on long-term landmark relation rather than the relation between neighbouring landmarks. It can effectively prevent overfitting and facilitate a coherent mapping for improved performance. 2) The computational complexity of the decoder during the training phase is significantly reduced with fewer structure prompts, which enhances training speed and reduces GPU memory consumption. 3) For

multi-dataset learning, it ensures the consistency in the input anchors numbers across different datasets, allowing them to form a training batch.

To encode the 2D coordinates of masked shapes into high-dimensional vectors while retaining their geometric relations, we utilize cosine and sine functions. The high-dimensional vector can be written as:

$$\begin{aligned} E_{(x,2c)}^X &= \sin(x/\tau^{(2c/0.5C)}), \\ E_{(x,2c+1)}^X &= \cos(x/\tau^{(2c/0.5C)}), \end{aligned} \quad (1)$$

$$\begin{aligned} E_{(y,2c)}^Y &= \sin(y/\tau^{(2c/0.5C)}), \\ E_{(y,2c+1)}^Y &= \cos(y/\tau^{(2c/0.5C)}), \end{aligned} \quad (2)$$

where E_x^X and E_y^Y are two vectors with $0.5C$ dimensions, representing the landmark coordinate (x, y) on X-axis and Y-axis respectively. $c \in [0, C/4)$ is the index of dimension and τ is a hyperparameter that determines the wavelengths (τ is set to 10000 in this paper). The final structure prompt E can be formulated as:

$$E_{(x,y)} = \text{Concat}(E_x^X; E_y^Y), \quad (3)$$

where Concat means concatenation process. For any fixed offset $(\Delta x, \Delta y)$, the vectors $E_{x+\Delta x}^X$ and $E_{y+\Delta y}^Y$ can be written as:

$$\begin{bmatrix} E_{(x+\Delta x,2c)}^X \\ E_{(x+\Delta x,2c+1)}^X \end{bmatrix} = \begin{bmatrix} \cos(\frac{\Delta x}{\tau^{2c/0.5C}}) & \sin(\frac{\Delta x}{\tau^{2c/0.5C}}) \\ -\sin(\frac{\Delta x}{\tau^{2c/0.5C}}) & \cos(\frac{\Delta x}{\tau^{2c/0.5C}}) \end{bmatrix} \begin{bmatrix} E_{(x,2c)}^X \\ E_{(x,2c+1)}^X \end{bmatrix}, \quad (4)$$

$$\begin{bmatrix} E_{(y+\Delta y,2c)}^Y \\ E_{(y+\Delta y,2c+1)}^Y \end{bmatrix} = \begin{bmatrix} \cos(\frac{\Delta y}{\tau^{2c/0.5C}}) & \sin(\frac{\Delta y}{\tau^{2c/0.5C}}) \\ -\sin(\frac{\Delta y}{\tau^{2c/0.5C}}) & \cos(\frac{\Delta y}{\tau^{2c/0.5C}}) \end{bmatrix} \begin{bmatrix} E_{(y,2c)}^Y \\ E_{(y,2c+1)}^Y \end{bmatrix}. \quad (5)$$

Thus, $E_{x+\Delta x}^X$ and $E_{y+\Delta y}^Y$ can be viewed as a linear function of E_x^X and E_y^Y respectively, and the transformation matrix is determined by the corresponding offset. This property enables the 2D geometric relationships of the mean shape landmarks to be well retained in the high-dimensional vector for the learning of the mapping.

3.1.3 Decoder

The decoder is the key component for learning the mapping from the 2D plane to target faces. As shown at the bottom left in Fig. 3, the decoder mainly consists of three blocks: multi-head self-attention (MSA) block, multi-head cross-attention (MCA) block, and feed-forward network (FFN). Besides,

there is an extra LayerNorm (Ba, Kiros, and Hinton, 2016) before each block.

The mapping between the 2D plane and target faces should be determined by multiple anchors, not just a single one. Therefore, the MSA block is crucial because it enables the structure prompts to share their geometric information for the mapping determination. The key (K_z), query (Q_z), and value (V_z) of the z -th head in MSA block can be calculated as:

$$K_z = (T_z + E_z)W_z^k, Q_z = (T_z + E_z)W_z^q, V_z = T_z W_z^v, \quad (6)$$

where $E \in \mathbb{R}^{N_a \times C}$ and $T \in \mathbb{R}^{N_a \times C}$ are the face structure prompts and the input to the MSA block respectively. E and T are further divided into N_h (N_h is the number of heads) sequences equally with $C_h = C/N_h$ dimensions. $W_z^k \in \mathbb{R}^{C_h \times C_h}$, $W_z^q \in \mathbb{R}^{C_h \times C_h}$, and $W_z^v \in \mathbb{R}^{C_h \times C_h}$ are three learnable matrices. The output of z -th head H_z can be calculated as:

$$H_z = \text{softmax}\left(\frac{Q_z K_z^T}{\sqrt{C_h}}\right) V_z. \quad (7)$$

The final output of MSA block can be written as:

$$\mathcal{F}_{\text{MSA}}(T) = \text{Concat}(H_1; \dots; H_{N_h})W_{\text{MSA}}, \quad (8)$$

where $W_{\text{MSA}} \in \mathbb{R}^{C \times C}$ is also a learnable matrix for linear projection.

The mapping also depends on the input face. Therefore, the MCA block is essential since it aggregates the image feature F learned by the ViT-based encoder based on the structure prompts, enabling the learned mapping to be case-dependent. The key (K'_z), query (Q'_z) and value (V'_z) of the z -th head in MCA block can be formulated as:

$$K'_z = (F_z + P_z)W_z^{k'}, Q'_z = (T'_z + E_z)W_z^{q'}, V'_z = F_z W_z^{v'}, \quad (9)$$

where $T' \in \mathbb{R}^{N_a \times C}$ and $F \in \mathbb{R}^{L \times C}$ are the input to MCA block and image feature respectively. We reuse the learnable positional embeddings $P \in \mathbb{R}^{L \times C}$ in the ViT-based encoder to retain the spatial information of image. Similar to the MSA block, T' , F and P are further divided into N_h sequences equally with $C_h = C/N_h$ dimensions. $W_z^{k'} \in \mathbb{R}^{C_h \times C_h}$, $W_z^{q'} \in \mathbb{R}^{C_h \times C_h}$, and $W_z^{v'} \in \mathbb{R}^{C_h \times C_h}$ are three learnable matrices. The output of z -th head H'_z can be written as:

$$H'_z = \text{softmax}\left(\frac{Q'_z K_z'^T}{\sqrt{C_h}}\right) V'_z. \quad (10)$$

The final output of MCA block can be formulated as:

$$\mathcal{F}_{\text{MCA}}(T) = \text{Concat}(H'_1; \dots; H'_{N_h})W_{\text{MCA}}, \quad (11)$$




Fig. 4 (a) the mean shapes of 29 & 98 landmarks before/after adding semantic alignment embeddings A . (b) the predicted 29 & 98 landmarks on the test image. The red and blue points indicate the 29 landmarks and 98 landmarks respectively (Color figure online)

where $W_{MCA} \in \mathbb{R}^{C \times C}$ is also a learnable matrix.

Moreover, a FFN block fuses the features aggregated from F in the channel-wise, further enhancing the expressive ability of the network. It allows the decoder to model a more fine-grained mapping between the target face and 2D plane.

Finally, the output features of the decoder are fed into an MLP for regressing the coordinates of the mapped anchors.

3.2 Semantic Alignment Embedding

The calculated mean shape reflects the semantic definitions and statistically geometric relationships between the pre-defined landmarks. Based on the fact that human faces have a regular structure, the semantics of pre-defined landmarks across different datasets can be aligned on the same 2D plane, regardless of the variations in landmark definition and number. This alignment effectively mitigates the knowledge discrepancies among multiple datasets and unifies their learning targets.

Therefore, we introduce extra learnable semantic embeddings $A_i \in \mathbb{R}^{N_D^i \times 2}$ to the mean shape of each dataset, where i is the index of dataset and N_D^i is the number of the pre-defined landmarks in the i -th dataset. A ensures that each landmark of the mean shapes has a corresponding learnable 2D offset. As shown in Eq. 4 and Eq. 5, adding A to the mean shape equals to applying a linear transformation to the initial structure prompts. Because the coordinates of mean shapes are normalized in the range of $[-1, 1]$, the transformation matrices are continuous and unique for any offset. It ensures that the structure prompts after linear transformation are also continuous and unique. As a result, it is possible for TUFA to find an optimal offset for the semantic alignment.

As shown in Fig. 4 (a), after adding the semantic alignment embeddings, the mean shapes of both the 98 and 29 landmarks still retain a clear face structure, which illustrates that these semantic alignment embeddings specifically incorporate the face structure information. The predicted

results shown in Fig. 4 (b) demonstrates that even the landmarks from different datasets, defined similarly, still exhibit semantic variance due to different annotation methods. For instance, the eye corners in the 29 landmarks are always located above the eye corners in the 98 landmarks. This can be also viewed as a kind of knowledge discrepancies. Therefore, directly using unaligned mean shapes for training may lead to semantic ambiguity and degrade performance further. After semantic alignment, the geometric relationship between the predicted 29 & 98 landmarks is consistent with that of the mean shapes of 29 & 98 landmarks, and the knowledge discrepancies are largely mitigated. Therefore, with the semantic alignment embeddings, the semantics of the landmarks from different datasets have been aligned on a plane successfully.

3.3 Loss Function

When the input anchors are dense enough, TUFA can learn a mapping from a 2D plane to a very complex surface by minimizing the distance between the mapped anchors and the corresponding labeled landmarks. We measure the distance using L1 loss as follows:

$$\mathcal{L} = \frac{1}{N_{\text{batch}} N_a} \sum_{j=1}^{N_{\text{batch}}} \sum_{k=1}^{N_a} |(x_{\text{map}}^{jk}, y_{\text{map}}^{jk}) - (x_{\text{gt}}^{jk}, y_{\text{gt}}^{jk})|, \quad (12)$$

where N_{batch} is the batch size, $(x_{\text{map}}^{jk}, y_{\text{map}}^{jk})$ is the position of the k -th mapped anchor in j -th sample, and $(x_{\text{gt}}^{jk}, y_{\text{gt}}^{jk})$ is the corresponding labeled position.

3.4 Zero-shot & Few-shot Face Alignment

The transformer-based decoder enables TUFA to accept an arbitrary number of structure prompts as the input, and the interpretable plane, learned based on the statistical mean shape, allows for easy editing of these structure prompts. Leveraging the structure prompts, the gap between the seen and unseen landmarks can be bridged. To locate an unseen landmark, we can determine its corresponding position on the plane, based on its geometric relationship with the anchors used during training. By encoding this 2D coordinate into a structure prompt, TUFA can locate the coordinate of the unseen landmark across different faces. As shown in Fig. 5, we randomly generate three scratch shapes and locate the corresponding landmarks on the target faces. We can clearly observe that the geometric relationship between the scratch shapes and the training anchors on the plane remains consistent with that geometric relationship on the target faces. Moreover, the semantics of the located landmarks also keep consistent across different faces. This demonstrates the

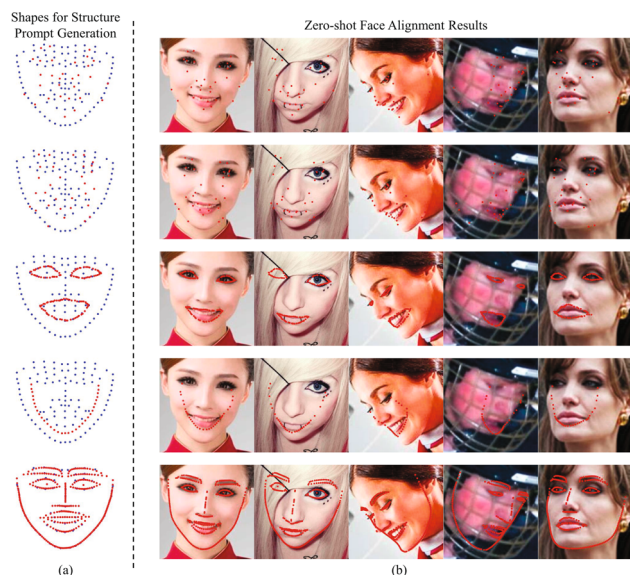


Fig. 5 (a) the **red** points are used to generate face structure prompts for zero-shot face alignment. The **blue** points represent the mean shape of 98 landmarks, which serves as anchors during training. (b) the **red** points represent the zero-shot face alignment results under various conditions (color figure online)

editability of the predicted target in TUFA, as well as the successful implementation of zero-shot face alignment.

The unified learning target also enables TUFA to be easily transferred to a group of newly defined facial landmarks in a few-shot manner. We first calculate a mean shape from the very few training samples. This mean shape is then added with a group of newly defined learnable semantic alignment embeddings. Finally, the pretrained TUFA is fine-tuned end-to-end on the new dataset without any changes of the structure. Despite the semantic difference in the newly defined landmarks, TUFA can still be efficiently transferred to the dataset. This is because TUFA explicitly learns the face structure, unifying the learning target across different datasets. Therefore, TUFA can easily inherit the knowledge learned from pre-trained datasets in few-shot learning, regardless of the differences in landmark semantics. Even with fewer training samples, TUFA significantly outperforms other state-of-the-art few-shot face alignment methods.

4 Experiments

4.1 Datasets

- **WFLW** (Wu et al., 2018): WFLW consists of 10,000 faces (7,500 for training and 2,500 for testing) from WIDER Face (S. Yang, Luo, Loy, and Tang, 2016). Each face is fully manual annotated with 98 landmarks and attribute labels. Compared to other datasets, the WFLW

is more challenging because it contains a large number of samples under extreme conditions, such as heavy occlusion and profile view.

- **300W** (Sagonas, Tzimiropoulos, Zafeiriou, and Pantic, 2013a): 300W consists of 3,837 faces (3,148 for training and 689 for testing) from AFW (X. Zhu & Ramanan, 2012), HELEN (Le, Brandt, Lin, Bourdev, and Huang, 2012), LFPW (Belhumeur, Jacobs, Kriegman, and Kumar, 2011) and IBUG (Sagonas, Tzimiropoulos, Zafeiriou, and Pantic, 2013b). Each face is annotated with 68 landmarks using a semi-automatic methodology (Sagonas, Tzimiropoulos, Zafeiriou, and Pantic, 2013b). The testing set can be further divided into a challenging subset (135 faces) and a common subset (554 faces). Moreover, 300W also provides additional 600 faces named as **300W-private** to test the generalization ability of the trained model.
- **Masked 300W** (C. Zhu, Li, Li, and Dai, 2021): Masked 300W is a variant of 300W (Sagonas, Tzimiropoulos, Zafeiriou, and Pantic, 2013a) used exclusively for testing the performance of face alignment under heavy occlusion. It synthesizes 689 masked faces from the testing set of 300W, and the landmark annotation remains consistent with the original 300W.
- **MERL-RAV** (Kumar et al., 2020): MERL-RAV manually re-annotates 19,314 faces from AFLW (Köstinger, Wohlhart, Roth, and Bischof, 2011) with 68 landmarks, providing 15,449 samples for training and 3,865 samples for testing. MERL-RAV further categorizes facial landmarks into unoccluded, externally occluded and self-occluded landmarks. Only unoccluded and externally occluded landmarks are provided with location information.
- **COFW** (Burgos-Artizzu, Perona, and Dollár, 2013): COFW collects 1,007 challenging faces with an average occlusion of over 23% from a variety of sources. It uses 500 of these faces and 845 samples from LFPW (Belhumeur, Jacobs, Kriegman, and Kumar, 2011) training set for training, and tests on the remaining 507 faces. Each face is annotated with 29 landmarks.
- **COFW68** (Ghiasi & Fowlkes, 2014): COFW68, a variant of COFW (Burgos-Artizzu, Perona, and Dollár, 2013), is used for cross-dataset validation. The testing set of COFW is manually re-annotated with 68 landmarks in this variant.
- **CelebA-aligned** (Z. Liu, Luo, Wang, and Tang, 2015): CelebA-aligned contains 10,000 identities with 200,000 face images. Each face is labeled with 5 landmarks and aligned to center. As suggested by (Thewlis, Bilen, and Vedaldi, 2017), CelebA can be subdivided into three subsets: CelebA training set without MAFL (160,000 images), MAFL training set (19,000 images) and MAFL testing set (1,000 images).

Table 1 Performance comparison against state-of-the-art heatmap regression and coordinate regression methods on WFLW and its subsets. Corresponding parameter size, flops are also reported.
Key: [*Best*, *Second Best*, *=*trained with multiple datasets*]

Method	Backbone	type	Extra face images	Flops ↓	Params ↓	Full Pose	Inter-Ocular NME (%) ↓			Occ.	Blur
							Exp.	Ill.	Mu.		
3FabRec (Browatzki & Wallraven, 2020)*	ResNet18	heatmap	2.1 millions	-	-	5.62	10.23	6.09	5.55	5.68	6.92
ATF (Lan, Hu, and Cheng, 2022)*	HRNetW18C	heatmap	20,000	4.75G	9.66M	4.50	7.54	4.63	4.45	4.20	5.30
PIPNet (Jin, Liao, and Shao, 2021)	ResNet101	heatmap	N/A	10.5G	45.7M	4.31	7.51	4.44	4.19	4.02	5.36
AWing (X. Wang, Bo, and Fuxin, 2019)	4 Hourglass	heatmap	N/A	26.8G	24.15M	4.21	7.21	4.46	4.23	4.02	4.99
HIH (Lan, Hu, and Cheng, 2021)	2 Hourglass	heatmap	N/A	10.38G	14.47M	4.18	7.20	4.19	4.45	3.97	5.00
ADNet (Y. Huang, Yang, Li, Kim, and Wei, 2021)	4 Hourglass	heatmap	N/A	17.04G	13.37M	4.14	6.96	4.38	4.09	4.05	5.06
FarL (Zheng et al., 2022)*	ViT-B/16	heatmap	20 millions	-	-	4.03	6.81	4.32	3.92	3.87	4.70
STAR (Zhou et al., 2023)	4 Hourglass	heatmap	N/A	-	13.37M	4.02	6.79	4.27	3.97	3.83	4.80
AV w. SAN (Qian, Sun, Wu, Qian, and Jia, 2019)*	ResNet152	coordinate	120,000	33.87G	35.02M	4.39	8.42	4.68	4.24	4.37	5.60
SDL (Lin et al., 2021)	HRNetW18C	coordinate	N/A	5.17G	-	4.35	7.42	4.63	4.29	4.22	5.19
SDL (Li et al., 2020)	HRNetW18C	coordinate	N/A	-	-	4.21	7.36	4.49	4.12	4.05	4.98
SLPT (Xia et al., 2022)	HRNetW18C-lite	coordinate	N/A	6.12G	13.19M	4.14	6.96	4.45	4.05	4.00	5.06
SPIGA (Prados-Torreblanca, Buenaposada, and Baumela, 2022)	4 Hourglass	coordinate	N/A	-	-	4.06	7.14	4.46	4.00	3.81	4.95
DSLPT (Xia et al., 2023)	HRNetW18C	coordinate	N/A	7.83G	19.35M	4.01	6.87	4.29	3.99	3.86	4.79
TUFA	ViT-S/16	coordinate	N/A	8.043G	36.00M	4.38	7.51	4.72	4.32	4.32	5.22
TUFA*	ViT-S/16	coordinate	19,942	8.043G	36.00M	4.00	6.57	4.20	3.90	3.86	4.73
TUFA	ViT-S/8	coordinate	N/A	35.17G	36.03M	4.23	7.23	4.53	4.17	4.12	5.05
TUFA*	ViT-S/8	coordinate	19,942	35.17G	36.03M	3.93	6.48	4.11	3.82	3.81	4.68

Table 2 Evaluation results in $FR_{0.1}$ and $AUC_{0.1}$ on WFLW. Key: [Best, Second Best, *=trained with multiple datasets]

Method	$FR_{0.1}(\%) \downarrow$							$AUC_{0.1} \uparrow$						
	Full	Pose	Exp.	Ill.	Mu.	Occ.	Blur	Full	Pose	Exp.	Ill.	Mu.	Occ.	Blur
3FabRec*	8.28	34.35	8.28	6.73	10.19	15.08	9.44	0.484	0.192	0.448	0.496	0.473	0.398	0.434
ATF*	2.52	13.19	2.23	2.44	0.49	5.03	3.88	0.560	0.301	0.546	0.566	0.581	0.487	0.489
PIPNet	-	-	-	-	-	-	-	-	-	-	-	-	-	-
AWing	2.04	9.20	1.27	2.01	0.98	4.21	2.72	0.590	0.334	0.572	0.596	0.602	0.528	0.539
HIH	2.96	15.03	1.59	2.58	1.46	6.11	3.49	0.597	0.342	0.590	0.606	0.604	0.527	0.549
ADNet	2.72	12.72	2.15	2.44	1.94	5.79	3.54	0.602	0.344	0.523	0.580	0.601	0.530	0.548
FaRL*	1.76	-	-	-	-	-	-	0.602	-	-	-	-	-	-
STAR	2.32	11.69	2.24	1.58	0.98	4.76	3.24	0.605	0.362	0.584	0.609	0.622	0.538	0.551
AV w. SAN*	4.08	18.10	4.46	2.72	4.37	7.74	4.40	0.591	0.311	0.549	0.609	0.581	0.516	0.551
SDFL	2.72	12.88	1.59	2.58	2.43	5.71	3.62	0.576	0.315	0.550	0.585	0.583	0.504	0.515
SDL	3.04	15.95	2.86	2.72	1.46	5.29	4.01	0.589	0.315	0.566	0.595	0.604	0.524	0.533
SLPT	2.76	12.27	2.23	1.86	3.40	5.98	3.88	0.595	0.348	0.574	0.601	0.605	0.515	0.535
SPIGA	2.08	11.66	2.23	1.58	1.46	4.48	2.20	0.606	0.353	0.580	0.613	0.622	0.533	0.553
DSLPT	2.52	13.19	2.23	2.44	0.98	4.89	3.49	0.607	0.353	0.586	0.614	0.623	0.535	0.549
TUFA (ViT-S/16)	2.92	14.42	3.50	2.29	2.43	5.84	4.53	0.570	0.296	0.540	0.579	0.571	0.497	0.524
TUFA (ViT-S/16)*	1.72	8.28	1.59	1.29	1.46	3.53	2.59	0.604	0.359	0.584	0.613	0.616	0.534	0.551
TUFA (ViT-S/8)	2.44	12.27	1.91	2.15	1.46	5.30	3.49	0.585	0.321	0.553	0.593	0.591	0.513	0.535
TUFA (ViT-S/8)*	1.52	7.36	1.27	0.86	1.46	2.99	2.33	0.610	0.371	0.592	0.620	0.622	0.540	0.553

4.2 Evaluation Metrics

Following previous works (Wu et al., 2018), (Y. Huang, Yang, Li, Kim, and Wei, 2021), (Xia et al., 2022), (Xia et al., 2023), we use Normalized Mean Error (NME), Failure Rate (FR) and Area Under Curve (AUC) to quantitatively measure the performance of TUFA. NME is the mean of L2 distance between the predicted landmarks and the annotated landmarks. The mean distance is then normalized by a factor, denoted as d_{norm} . For *inter-pupil* NME, d_{norm} represents the distance between pupil centers, while for *inter-ocular* NME, it stands for the distance between outer eye corners. In the case of NME_{box} , d_{norm} is defined as the geometric mean of the labeled box, calculated as $d_{\text{norm}} = \sqrt{W_{\text{box}} \times H_{\text{box}}}$. FR_{α} represents the percentage of the testing cases in which the NME exceeds a certain threshold, denoted as α . AUC indicates the area beneath the Cumulative Errors Distribution (CED) curve from 0 to the threshold of FR_{α} , which can be formulated as $\int_0^{\alpha} f(\epsilon) d\epsilon$.

4.3 Implementation Details

The input image is the face region cropped from the initial image and is then resized to a fixed size (256×256). For training data, we apply augmentation techniques, which include random translation (± 10 pixels), random rotation ($\pm 30^\circ$), random scaling ($\pm 5\%$), random horizontal flipping (50%), random gray (20%), random brightness adjustment

(50%, ± 0.3), random occlusion (50%), random shearing (33%). In TUFA, We employ two types of encoders : ViT-S/16 and ViT-S/8 (Dosovitskiy et al., 2021), both of which are pretrained on ImageNet (Deng et al., 2009) with DINO (Caron et al., 2021). The parameters of the decoder and semantic alignment embeddings are initialized from scratch. By default, the batch size is set to 16 and the layer number of the decoder is set to 6. TUFA is trained with AdamW (Loshchilov and Hutter, 2019), setting the initial learning rate to 1×10^{-4} for ViT-S/8 and 5×10^{-5} for ViT-S/16. For multi-dataset learning, we employ the training sets of WFLW, 300W, MERL-RAV and COFW for training. The model is trained for 100 epochs, with the learning rate decaying by a factor of 0.1 at the 80th and 90th epochs respectively. The learned result is then tested on the WFLW, 300W, 300W-private, masked 300W, MERL-RAV, COFW and COFW68 **simultaneously**. For single-dataset learning, due to the smaller number of training samples, the model is trained for 140 epochs, with the learning rate decaying by a factor of 0.1 at the 120th and 130th epochs respectively.

4.4 Comparisons in Within-dataset Validation

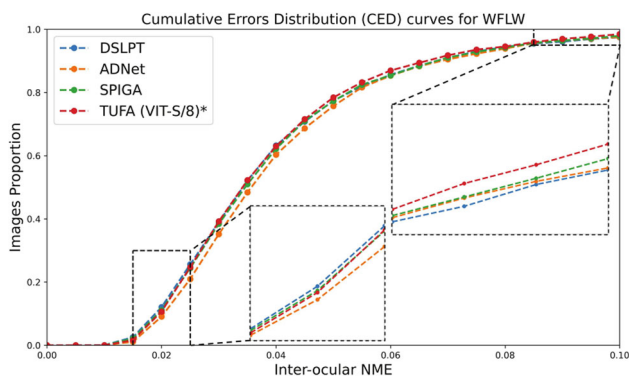
WFLW: we report the inter-ocular NME, parameter size and flops of TUFA and other state-of-the-art methods in Table 1. $FR_{0.1}$ and $AUC_{0.1}$ are tabulated in Table 2. Without increasing any computational complexity, the ability of TUFA to mitigate knowledge discrepancies across multiple

Table 3 Performance comparison against state-of-the-art methods on 300W. Key: [*Best*, *Second Best*, *=trained with multiple datasets]

Method	Inter-Ocular NME (%) ↓		
	Common	Challenging	Fullset
3FabRec (Browatzki & Wallraven, 2020)*	3.36	5.74	3.82
HIH (Lan, Hu, and Cheng, 2021)	2.93	5.00	3.33
PIPNet (Jin, Liao, and Shao, 2021)	2.78	4.89	3.19
ATF (Lan, Hu, and Cheng, 2022)*	2.75	4.86	3.17
AWing (X. Wang, Bo, and Fuxin, 2019)	2.72	4.53	3.07
ADNet (Y. Huang, Yang, Li, Kim, and Wei, 2021)	2.53	4.58	2.93
SLPT (Xia et al., 2022)	2.75	4.90	3.17
FaRL (Zheng et al., 2022)*	2.70	4.64	3.08
SPIGA (Prados-Torreblanca, Buenaposada, and Baumela, 2022)	2.59	4.66	2.99
DSLPT (Xia et al., 2023)	2.57	4.69	2.98
STAR (Zhou et al., 2023)	2.52	4.32	2.87
TUFA (ViT-S/16)	2.92	5.05	3.34
TUFA (ViT-S/16)*	2.68	4.58	3.05
TUFA (ViT-S/8)	2.82	4.90	3.23
TUFA (ViT-S/8)*	2.59	4.45	2.95

Table 4 Performance comparison against state-of-the-art methods on COFW. The $FR_{0.1}$ in Inter-Ocular and Inter-Pupil normalization are reported. Key: [*Best*, *Second Best*, *=trained with multiple datasets]

Method	Inter-Ocular		Inter-Pupil	
	NME ↓	FR ↓	NME ↓	FR ↓
SDFL (Lin et al., 2021)	3.63%	0.00%	-	-
ATF (Lan, Hu, and Cheng, 2022)*	3.32%	-	-	-
AWing (X. Wang, Bo, and Fuxin, 2019)	-	-	4.94%	0.99%
ADNet (Y. Huang, Yang, Li, Kim, and Wei, 2021)	-	-	4.68%	0.59%
SLPT (Xia et al., 2022)	3.32%	0.00%	4.79%	1.18%
DSLPT (Xia et al., 2023)	3.33%	0.20%	4.79%	1.36%
STAR (Zhou et al., 2023)	3.21%	0.00%	4.62%	0.79%
PIP (Jin, Liao, and Shao, 2021)	3.08%	-	-	-
TUFA (ViT-S/16)	3.49%	0.20%	5.03%	0.79%
TUFA (ViT-S/16)*	3.18%	0.20%	4.58%	0.39%
TUFA (ViT-S/8)	3.40%	0.20%	4.91%	1.18%
TUFA (ViT-S/8)*	3.07%	0.20%	4.43%	0.39%

**Fig. 6** The comparisons of CED curves between TUFA and other state-of-the-art methods on WFLW full set.

datasets maximizes the utilization of extra training samples. We observe a significant improvement of 7.09% in NME on the full set compared to TUFA (ViT-S/8) trained on a single dataset, even though TUFA uses fewer extra samples than other methods trained on multiple datasets.

Considering efficiency, we also implemented a TUFA with a lighter backbone, ViT-S/16. Similarly, training on multiple datasets achieves an impressive improvement of 8.68% in NME on the full set compared to single dataset training. Despite having much lower computational complexity (flops) than most heatmap regression methods, it still yields the second best performance in NME and $FR_{0.1}$. It also outperforms DSLPT in NME on the full set and all subsets with comparable flops. Moreover, the entire training process for this

TUFA can be completed within five hours on a single A40 GPU, demonstrating the high efficiency of our method.

Most existing methods, such as DSLPT, SPIGA, ADNet, and STAR, achieve competitive performance based on a multi-stage approach, while TUFA locates facial landmarks using only a single stage. For a comprehensive analysis, we plot the CED curves of several multi-stage methods and TUFA (using multi-dataset learning) in Fig. 6. Note that the official implementation of ADNet performs slightly better than the results reported in their original paper. The CED curves of DSLPT and SPIGA are higher than that of TUFA in the range of [0, 0.025]. It suggests the proportion of the samples with very small NME predicted by DSLPT and SPIGA is larger than that of TUFA, despite that TUFA (using multi-dataset learning) performs better in NME. This indicates that these multi-stage face alignment methods primarily reduce the NME on easy samples for better numerical results. However, their improvements under challenging conditions are relatively insignificant. Despite the fact that these conditions are long-tailed, they are critical as they ultimately determine the quality of face alignment in real-world scenarios. By employing extra data, TUFA significantly improves the performance on these challenging conditions. As a result, TUFA achieves improvements of 26.9%, 36.9% and 33.3% in $FR_{0.1}$ on full set, largepose subset and occlusion subset respectively compared to SPIGA.

300W: as tabulated in Table 3, multi-dataset learning also improves the performance of TUFA on common set, challenging set and full set by 8.1%, 9.18% and 8.67% respectively in the metric of NME. Heatmap regression methods, such as ADNet, Awing and STAR, commonly demonstrate superior performance on 300W, especially the challenging subset, compared to coordinate regression methods. This is because they provide semantic supervision to the network by encoding the annotated landmarks into heatmaps, which delivers better performance with limited training samples. However, the unique properties of TUFA successfully address this limitation of coordinate regression methods, achieving the best performance in NME on the challenging set, as low as 4.45%, even though TUFA is a single-stage face alignment method. This demonstrates that the multi-dataset learning of TUFA promises a very competitive generalization ability. Compared to FaRL and ATF, which are also trained with multiple face datasets, TUFA outperforms them by 4.22% and 6.94% respectively in NME despite fewer extra faces used in training.

COFW: with the very limited number of training samples and an average occlusion of over 23% on the testing faces, COFW presents a significant challenge for all face alignment methods. As shown in Table 4, despite the competitive performance on other datasets, existing state-of-the-art methods, such as DSLPT, ADNet and STAR, often suffer from overfitting on this dataset. Similarly, TUFA when trained with

	Full	Challenging	Largepose	Occlusion	Full	Challenging	Largepose	Occlusion	Full	Challenging	Largepose	Occlusion
DSLPT (Xia et al., 2023)	1.51%	1.51%	1.51%	1.51%	78.33%	78.47%	78.33%	78.33%	1.50%	1.50%	1.50%	1.50%
SPIGA (Torreblanca, Buenaposada, and Baumela, 2022)	1.51%	1.51%	1.51%	1.51%	78.33%	78.47%	78.33%	78.33%	1.50%	1.50%	1.50%	1.50%
DSLPT (Xia et al., 2023)	1.48%	1.48%	1.48%	1.48%	78.87%	78.87%	78.87%	78.87%	1.48%	1.48%	1.48%	1.48%
TUFA (ViT-S/16)	1.63%	1.63%	1.63%	1.63%	76.74%	76.74%	76.74%	76.74%	1.65%	1.65%	1.65%	1.65%
TUFA (ViT-S/16)*	1.61%	1.61%	1.61%	1.61%	77.09%	77.09%	77.09%	77.09%	1.63%	1.63%	1.63%	1.63%
TUFA (ViT-S/8)	1.55%	1.55%	1.55%	1.55%	77.93%	77.93%	77.93%	77.93%	1.56%	1.56%	1.56%	1.56%
TUFA (ViT-S/8)*	1.54%	1.54%	1.54%	1.54%	78.01%	78.01%	78.01%	78.01%	1.56%	1.56%	1.56%	1.56%
DSLPT (Xia et al., 2023)	3.26%	3.26%	3.26%	3.26%	79.37%	79.37%	79.37%	79.37%	3.33%	3.33%	3.33%	3.33%
TUFA (ViT-S/16)	3.37%	3.37%	3.37%	3.37%	76.81%	76.81%	76.81%	76.81%	3.33%	3.33%	3.33%	3.33%
TUFA (ViT-S/16)*	3.31%	3.31%	3.31%	3.31%	77.11%	77.11%	77.11%	77.11%	3.30%	3.30%	3.30%	3.30%
TUFA (ViT-S/8)	3.30%	3.30%	3.30%	3.30%	78.29%	78.29%	78.29%	78.29%	3.21%	3.21%	3.21%	3.21%
TUFA (ViT-S/8)*	3.21%	3.21%	3.21%	3.21%	78.17%	78.17%	78.17%	78.17%	3.21%	3.21%	3.21%	3.21%

single dataset, also tends to overfit the training set, achieving only 3.40% in the metric of NME. Nevertheless, TUFA successfully mitigates the knowledge discrepancies among multiple datasets and utilizes the knowledge learned from other face datasets, thereby significantly improving the NME metric by 9.71% and yielding the best performance. Although ATF also leverages extra datasets, it fails to consider the ambiguity between the landmarks with similar definitions, which inadvertently introduces noise into the training and leads to a less significant improvement.

MERL-RAV: the results on MERL-RAV are shown in Table 5. The proportions of unoccluded, externally occluded and self-occluded landmarks in MERL-RAV full set are 76.56%, 10.81% and 12.63% respectively. Because the position annotations of the self-occluded landmarks are not provided, the numerical results on the full set and its subsets exclude the results of these landmarks. Therefore, these numerical results are largely determined by the performance on the easy samples. As mentioned before, multi-stage methods tend to perform better on the easy samples compared to one-stage methods. Therefore, TUFA only demonstrates comparable performance to these multi-stage methods. Nevertheless, when we investigate the externally occluded landmarks separately, we can find that the multi-dataset learning still improves the performance on these challenging landmarks effectively, despite the relatively fewer extra samples provided by WFLW, 300W and COFW compared to the training set of MEAL-RAV. Moreover, the qualitative results shown in Fig. 7 illustrate TUFA performs well on the self-occluded landmarks based on the knowledge of other datasets, even though MERL-RAV does not provide their position labels for training.

4.5 Comparisons in Cross-dataset Validation

COFW68: the cross-dataset validation aims at evaluating the generalization ability of face alignment methods. In the case of TUFA with single dataset learning, the model is trained on 300W training set. For TUFA trained with multiple datasets, the model used in four within-dataset validations is directly evaluated on COFW68. The evaluation results are shown in Table 6. Despite not re-training the TUFA for COFW68, it still outperforms existing state-of-the-art methods by a significant margin. Compared to single dataset learning, the multi-dataset learning approach of TUFA demonstrates much stronger generalization ability, achieving an impressive improvement of 9.98% in NME.

Masked 300W: the setting of cross-dataset validation on Masked-300W is kept consistent with the setting on COFW68, and the evaluation results are displayed in Table 7. The testing faces are with an average occlusion of over 50%, leading to a significant domain gap between the training and testing faces. To minimize this gap, existing state-of-the-art




Fig. 7 (a) Results in profile view predicted by the TUFA with multi-dataset learning. (b) Results in profile view predicted by the TUFA with single dataset learning. The red points represent the predicted landmarks (Color figure online)

methods have optimized the data augmentation, randomly masking each training sample with several blocks of varying sizes. Although this approach can improve the quantitative results in cases with heavy occlusion, it also results in performance degradation under other conditions. Without any modification in data augmentation techniques, the TUFA learned from multiple datasets still achieves an impressive result, as low as 5.58% in NME. This outcome also demonstrates the remarkable generalization ability and robustness of TUFA.

300W-private: we carry out cross-dataset validation on 300W-private using the same settings as those for COFW and Masked 300W. Similarly, the improvement brought to TUFA by multi-dataset learning is also significant. Even though most samples from 300W-private are under common conditions, which are easily handled by other multi-stage methods, TUFA still yields the best performance across all metrics on the full set. The $FR_{0.08}$ of TUFA is as low as 0.17%, which means only one sample failed to be aligned by TUFA in the entire dataset. This demonstrates that TUFA is highly robust compared to all existing methods.

4.6 Few-shot Face Alignment

We pretrain TUFA (ViT-S/8) on COFW, 300W and MERL-RAV (29 & 68 landmarks) and implement few-shot face alignment on WFLW (98 landmarks) as described in Section 3.4. Note that there are large differences between the semantic definitions of 29 & 68 landmarks and 98 landmarks. We randomly select a certain number of training samples in each experiment. We repeat each experiment five times and report the mean NME with each training set size in Table 9. With only 10 training samples, TUFA still demonstrates superior performance compared to other methods, outperforming 3FabRec and He et al. trained with 7500 annotated faces. When we fine-tune TUFA on WFLW using 100% of the original training set (7500 samples), the NME on the testing set improves to 3.99%. This performance is comparable to the NME of the TUFA when trained with all four datasets together. This result indicates that TUFA can effec-

Table 6 Performance comparison of *cross*-dataset validation on COFW68. The NME and FR_{0.1} under inter-ocular normalization are reported. Key: [**Best**, **Second Best**, *=trained with multiple datasets]

Method	NME(%)↓	FR _{0.1} (%)↓
AV w. SAN (Qian, Sun, Wu, Qian, and Jia, 2019)*	4.43	2.82
LAB (Wu et al., 2018)	4.62	2.17
GlomFace (C. Zhu, Wan, Xie, Li, and Gu, 2022)	4.21	0.79
SDL (Li et al., 2020)	4.22	0.39
SDFL (Lin et al., 2021)	4.18	0.00
SLPT (Xia et al., 2022)	4.10	0.59
DSLPT (Xia et al., 2023)	4.03	0.20
SPIGA (Prados-Torresblanca, Buenaposada, and Baumela, 2022)	3.93	-
TUFA (ViT-S/16)	4.22	0.20
TUFA (ViT-S/16)*	3.72	0.00
TUFA (ViT-S/8)	4.11	0.20
TUFA (ViT-S/8)*	3.70	0.00

Table 7 Performance comparison of *cross*-dataset validation on Masked 300W. Key: [**Best**, **Second Best**, *=trained with multiple datasets, †=data augmentation adjustment]

Method	Inter-Ocular NME (%) ↓		
	Common	Challenging	Fullset
Hourglass (Newell, Yang, and Deng, 2016)	8.17	13.52	9.22
FAN (Bulat & Tzimiropoulos, 2017)	7.36	10.81	8.02
LAB (Wu et al., 2018)	6.07	9.59	6.76
SAAT (C. Zhu, Li, Li, and Dai, 2021)†	5.42	11.36	6.58
GlomFace (C. Zhu, Wan, Xie, Li, and Gu, 2022)†	5.29	8.81	5.98
DSLPT (Xia et al., 2023)	6.01	10.19	6.83
DSLPT (Xia et al., 2023)†	4.78	8.10	5.42
TUFA (ViT-S/16)	6.08	9.52	6.76
TUFA (ViT-S/16)*	5.21	8.25	5.80
TUFA (ViT-S/8)	7.12	9.96	7.67
TUFA (ViT-S/8)*	4.98	8.08	5.58

tively leverage the knowledge learned from pretraining. This is achieved because TUFA successfully unifies the learning objectives of different datasets and mitigates their knowledge discrepancies, rather than treating landmark regression as an independent target. Because of this unique property, TUFA can be transferred to a group of newly defined landmarks efficiently, even with a very limited number of annotated samples. 3FabRec, Autolink and He et al. set image reconstruction as a pre-text task to encourage model to perform better with very limited training images. However, He et al. also found that this pre-text task may also lead to performance degradation when the number of training samples is sufficient. The main reason is that the learning objectives of image reconstruction and face alignment are not aligned. This large knowledge discrepancy significantly degrades performance.

4.7 Zero-shot Face Alignment

CelebA-aligned: to evaluate the semantic consistency of the landmarks predicted by TUFA in a zero-shot manner, we follow the same evaluation protocol as other unsupervised landmark detection methods (Jakab, Gupta, Bilen, and Vedaldi, 2018), (Lorenz, Bereska, Milbich, and Ommer, 2019), (He, Wandt, and Rhodin, 2022) used on CelebA-aligned (Z. Liu, Luo, Wang, and Tang, 2015). We randomly generate a scratch shape with a certain number of points N_{pre} . Then, the TUFA (ViT-S/8), trained with the four datasets, directly predicts the corresponding landmarks on both MAFL training and testing subset. The predicted coordinates on MAFL training subset are used to calculate a matrix that projects the coordinates to the positions of the labeled landmarks. We use this matrix to transfer the predicted landmark coordinates on MAFL testing subset to the coordinates of the labeled landmarks. Finally, we quantitatively measure the performance with inter-ocular NME.

Table 8 Performance comparison of *cross*-dataset validation on 300W-private. The NME, FR_{0.08} and AUC_{0.08} under inter-ocular normalization are reported. Key: [*Best*, *Second Best*, *=trained with multiple datasets]

Method	Indoor subset			Outdoor subset			Full set		
	NME↓	AUC _{0.08} ↑	FR _{0.08} ↓	NME↓	AUC _{0.08} ↑	FR _{0.08} ↓	NME↓	AUC _{0.08} ↑	FR _{0.08} ↓
DAN (Kowalski, Naruniec, and Trzcinski, 2017)	-	-	-	-	-	-	4.30%	47.00%	2.67%
SHN (J. Yang, Liu, and Zhang, 2017)	4.10%	-	-	4.00%	-	-	4.05%	-	-
DCF _E (Valle, Buenapósada, Valdés, and Baumela, 2018)	3.96%	52.28%	2.33%	3.81%	52.56%	1.33%	3.88%	52.42%	1.83%
SPIGA (Prados-Torreblanca, Buenapósada, and Baumela, 2022)	3.43%	57.35%	1.00%	3.43%	57.17%	0.33%	3.43%	57.27%	0.67%
DSLPT (Xia et al., 2023)	3.47%	56.60%	0.33%	3.47%	56.68%	0.00%	3.47%	56.64%	0.17%
TUFA (ViT-S/16)	3.89%	51.39%	0.67%	3.92%	51.11%	1.00%	3.91%	51.25%	0.83%
TUFA (ViT-S/16)*	3.51%	56.09%	0.33%	3.52%	55.96%	0.67%	3.52%	56.03%	0.50%
TUFA (ViT-S/8)	3.86%	51.98%	1.33%	3.86%	51.90%	1.00%	3.86%	51.94%	1.17%
TUFA (ViT-S/8)*	3.40%	57.48%	0.00%	3.42%	57.31%	0.33%	3.41%	57.40%	0.17%

Table 9 Performance comparison in inter-ocular NME (%)↓ with reduced training set on WFLW full set. Key: [*Best*, *Second Best*]

Method	Training set size						
	1	10	20	50	375	1500	7500
AV w. SAN (Qian, Sun, Wu, Qian, and Jia, 2019)	-	-	-	-	-	6.00	4.39
Xiao et al. (B. Xiao, Wu, and Wei, 2018)	43.0	21.9	19.3	17.6	10.6	7.08	5.62
Autolink (He, Wandt, and Rhodin, 2022)	14.9	13.5	13.3	11.2	7.68	7.31	6.35
3FabRec (Browatzki & Wallraven, 2020)	15.8	9.66	-	8.39	7.68	6.51	5.62
He et al. (He, Bharaj, Ferman, Rhodin, and Garrido, 2023)	12.4	9.19	8.62	7.90	6.22	5.61	5.38
TUFA	7.62	5.20	4.90	4.77	4.43	4.28	3.99

Table 10 Performance of TUFA in inter-ocular NME ($N_{pre}=10$) on CelebA under the setting of zero-shot face alignment. The evaluation results of state-of-the-art self-supervised and unsupervised methods also reported. Key: [*Best*, *Second Best*]

Method	type	NME ↓
Thewlis et al. (Thewlis, Bilen, and Vedaldi, 2017)	unsupervised	7.95%
Zhang et al. (Y. Zhang et al., 2018)	unsupervised	3.46%
Lorenz et al. (Lorenz, Bereska, Milbich, and Ommer, 2019)	unsupervised	3.24%
IMM (Jakab, Gupta, Bilen, and Vedaldi, 2018)	unsupervised	3.19%
AutoLink (He, Wandt, and Rhodin, 2022)	unsupervised	3.92±0.69%
Mallis et al. (Mallis, Sanchez, Bell, and Tzimiropoulos, 2023)	self-supervised	3.83%
TUFA	zero-shot	2.65±0.33%

We conduct this experiment 10 times with 10 randomly generated scratch shapes. The mean and variance of the NMEs are reported in Table 10. Although TUFA is not trained with the 160,000 in-domain faces from CelebA training set, unlike other self-supervised or unsupervised methods, it still yields the best performance, significantly outperforming Mallis et al. and IMM, by 30.81% and 16.93% respectively, in terms of NME. This result demonstrates the landmarks predicted by TUFA in the zero-shot manner display better semantic consistency in various cases, as compared to other methods. Moreover, the semantic definitions of the landmarks predicted by TUFA can be assigned by humans via a manually generated shape, whereas those in self-supervised or unsupervised methods are learned randomly. Consequently, TUFA has a much broader application range.

WFLW: to better evaluate the performance of zero-shot face alignment in practical application, we carry out an additional experiment. We first train TUFA (ViT-S/8) on MERL-RAV (68 landmarks). Then, we align the mean shape on WFLW (98 landmarks) to the mean shape on MERL-RAV using affine transformation. Finally, we evaluate TUFA on the WFLW full set directly without re-training. Since TUFA is the first framework for zero-shot face alignment, we can only compare TUFA in the zero-shot setting with some fully-supervised methods to quantitatively demonstrate its performance. The evaluation results of zero-shot face alignment with TUFA, and some fully supervised methods released before 2018, are listed in Table 11. Despite the large disparity in landmark semantics and numbers between WFLW and MERL-RAV, TUFA in the zero-shot setting still outperforms SDM, CFSS and DVLN by 42.66%, 34.95% and 2.96% respectively in the metric of NME. The qualitative results shown in Fig. 8 demonstrate the semantics of the landmarks predicted by the zero-shot face alignment remain highly consistent, regardless of heavy occlusion or profile view. Therefore, TUFA can be easily applied to a group of newly defined landmarks without any annotated training samples and re-training processes, significantly broadening the application range of existing face alignment methods. Moreover, we can also conclude that TUFA learns a more universal knowledge for face alignment, which explicitly represents the regular structure of human face.

4.8 Ablation Studies

Influence of the number of training datasets: the performance of TUFA, when trained with different numbers of datasets, is tabulated in Table 12. The quantitative results on each dataset are significantly improved as the number of training datasets increases, despite the large difference between their annotation forms. This demonstrates TUFA effectively mitigates knowledge discrepancies among multi-



Fig. 8 Zero-shot alignment results predicted by TUFA on WFLW. The green points and red points represent the predicted landmarks and ground truth respectively (Color figure online)

ple datasets and maximizes the utilization of their knowledge. Therefore, even though COFW only provides 1,345 labeled faces, it still improves NME from 3.23% to 3.14% on 300W. This improvement is quite competitive in face alignment task. MERL-RAV provides the largest number of annotated faces, leading to the most significant improvement on 300W, COFW and WFLW. The unique property of TUFA also enables its performance to be further improved with the release of more high-quality face alignment datasets in the future.

Influence of the semantic alignment embeddings: to further investigate the influence of semantic alignment embeddings, we also implement a version of TUFA that does not incorporate these embeddings. Instead, we align the mean shapes from different datasets using affine transformation and train TUFA using these aligned mean shapes, without adding the semantic alignment embeddings. The comparison results between the TUFA with/without these embeddings are shown in Table 13. The aligned mean shape based on affine transformation cannot explicitly mitigate the knowledge discrepancies among multiple datasets and may lead to semantic ambiguity. Consequently, it inevitably introduces a quantitative error into the predicted results. Because 300W, 300W-private and Masked 300W are labeled using semi-supervised methodology (Sagonas, Tzimiropoulos, Zafeiriou, and Pantic, 2013b), their annotations exhibit much smaller variance than manual annotations. Therefore, the quantitative errors result in significant performance degradation on these three datasets. With an average occlusion of over 23%, the manually annotated landmarks on COFW and COFW68 exhibit much larger variance than other datasets. Contrary to the results on 300W, 300W-private and Masked 300W, the evaluation results on these datasets are less sensitive to the quantitative errors. Nevertheless, the absence of semantic alignment still leads to slight performance degradation. Overall, the subtle semantic differences between the landmarks with similar definitions cannot be ignored. The semantic alignment embeddings, while only

Table 11 Performance of TUFA in inter-ocular NME (%)↓ on WFLW under the setting of zero-shot face alignment. The evaluation results of some fully supervised methods before 2018 are also reported. Key: [*Best*, *Second Best*]

Method	Full	Pose	Exp.	Ill.	Mu.	Occ.	Blur
Fully-supervised Face Alignment							
SDM (Xiong & De la Torre, 2013)	10.29	24.10	11.45	9.32	9.38	13.03	11.28
CFSS (S. Zhu, Li, Loy, and Tang, 2015)	9.07	21.36	10.09	8.30	8.74	11.76	9.96
DVLN (Wu & Yang, 2017)	6.08	11.54	6.78	5.73	5.98	7.33	6.88
Zero-shot Face Alignment							
TUFA	5.90	10.03	6.16	5.78	5.50	6.68	6.69

Table 12 The influence of different numbers of datasets used in training. The inter-ocular NME on 300W, COFW and WFLW, and the NME_{box} on MERL-RAV are reported. Key: [*Best*, *Second Best*]

Method	300W	Training Datasets			300W	NME (%)↓		
		COFW	WFLW	MERL		COFW	WFLW	MERL
TUFA 1 (ViT-S/8)	✓	-	-	-	3.23	-	-	-
TUFA 2 (ViT-S/8)	✓	✓	-	-	3.14	3.27	-	-
TUFA 3 (ViT-S/8)	✓	✓	✓	-	3.08	3.18	4.17	-
TUFA 4 (ViT-S/8)	✓	✓	✓	✓	2.95	3.07	3.93	1.54

Table 13 The influence of semantic alignment embeddings. The inter-ocular NME on 300W, COFW, WFLW, COFW68, Masked 300W and 300W-private, and the NME_{box} on MERL-RAV are reported. Key: [*Best*, 300W-M=masked 300W, 300W-P=300W-private]

Method	Semantic Alignment Embeddings	NME (%)↓			MERL	COFW68	300W-M	300W-P
		300W	COFW	WFLW				
TUFA1 (ViT-S/8)	w/o	3.42	3.08	3.97	1.57	3.71	5.87	4.08
TUFA2 (ViT-S/8)	w	2.95	3.07	3.93	1.54	3.70	5.58	3.40

slightly increasing the number of parameters in the training process, effectively eliminate the quantitative error caused by these semantic differences.

Comparison to multi-task learning: To further demonstrate the importance of mitigating knowledge discrepancies, we implement a model using multi-task learning under the same settings and compare it to TUFA. We define a

learnable vector for each landmark to replace the structure prompt for regression. The results are presented in Table 14. Despite the similar semantics among the landmarks from 300W and MERL, they are treated as completely different regression targets in multi-task learning, leading to conflicts during training. Since MERL-RAV does not provide annotations for self-occluded landmarks, the model primarily fits

Table 14 Performance comparison to multi-task learning. The inter-ocular NME on 300W, COFW, WFLW, COFW68, Masked 300W and 300W-private, and the NME_{box} on MERL-RAV are reported. Key: [*Best*, 300W-M=masked 300W, 300W-P=300W-private]

Method	NME (%)↓			MERL	COFW68	300W-M	300W-P
	300W	COFW	WFLW				
Multi-task (ViT-S/8)	3.01	3.26	4.01	1.94	3.73	6.06	3.49
TUFA (ViT-S/8)	2.95	3.07	3.93	1.54	3.70	5.58	3.40

Table 15 The influence of masking ratio on WFLW. The inter-ocular NME, FR_{0.1}, AUC_{0.1} and the computational complexity of Decoder (Flops) are reported. Additionally, the memory usage and training speed on a single A40 are also provided (with the batch size set to 16). Key: [*Best*, *Second Best*]

Metric	masking ration						
	0%	25%	50%	75%	80%	85%	90%
NME(%)↓	4.265	4.256	4.233	4.229	4.229	4.241	4.682
FR _{0.1} (%)↓	2.800	2.360	2.360	2.440	2.320	2.440	3.920
AUC _{0.1} ↑	0.582	0.582	0.584	0.585	0.584	0.582	0.546
Decoder Flops (G) ↓	3.537	3.089	2.663	2.226	2.139	2.053	1.967
Memory usage (MB)↓	14,004	13,792	13,608	13,502	13,490	13,436	13,368
Training Speed (samples/sec) ↑	45.8	46.4	46.8	47.2	47.3	47.3	47.5

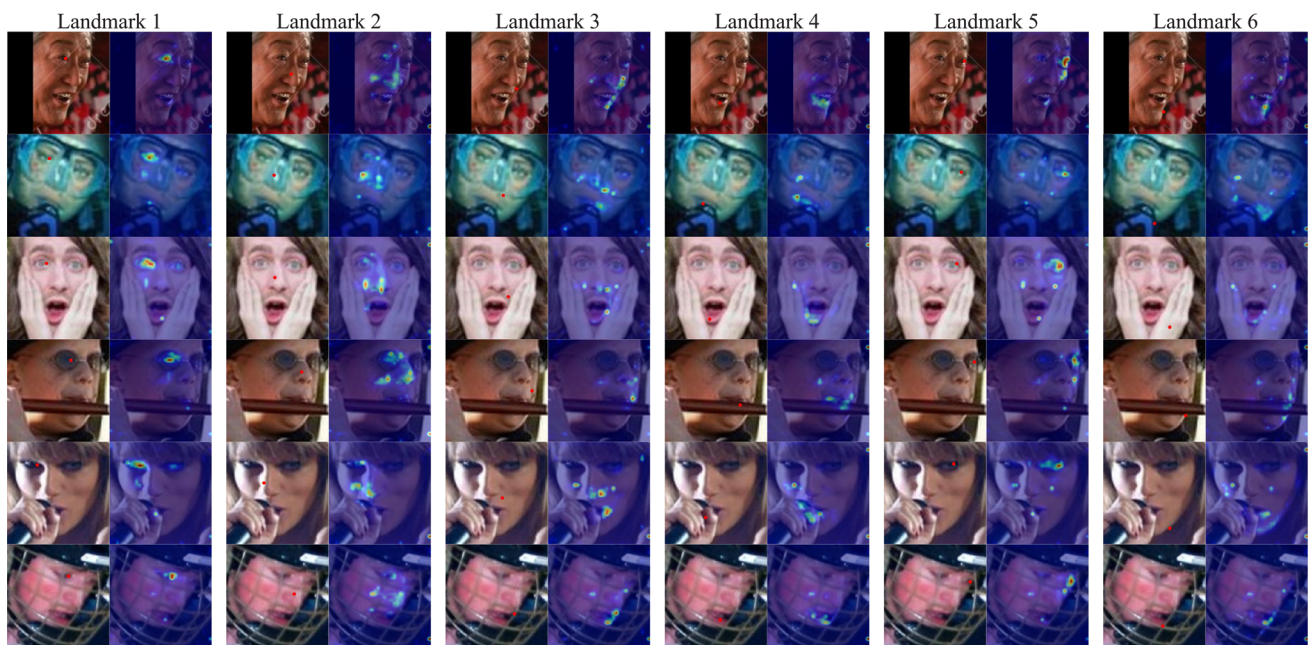


Fig. 9 The zero-shot landmark predicted results of TUFA (column 1, 3, 5, 7, 9 and 11) and the corresponding mean cross attention maps (column 2, 4, 6, 8, 10 and 12). The red points represent the predicted landmarks (color figure online)

300W. Consequently, TUFA improves NME by 20.62% on MERL-RAV, which is significantly greater compared to the improvement on other datasets. Additionally, the comparisons on other datasets illustrate that the proposed strategy can maximize the utilization of knowledge from multiple datasets and boost the generalization ability of model.

Influence of the masking ratio: to further explore the influence of masking ratio, we train TUFA (ViT-S/8) with different masking ratios on WFLW. The results are reported in Table 15. As the masking ratio increases, the performance of TUFA is improved. The primary reason is that randomly masking a certain ration of landmarks encourages TUFA to learn the mapping using the relationship between distant landmarks. This relationship ensures a more coherent mapping between the 2D plane and target faces, while the landmarks that are close together often introduce bias into the learned mapping. Moreover, increasing the masking ratio also significantly reduces the computational complexity of the TUFA decoder. When the ratio is set to 75%, the computational complexity of TUFA decoder is only 2/3 of that at a 0% masking ration. Although TUFA decoder only accounts for just about 10% of the total computational complexity compared to TUFA encoder (31.572G flops), we can still observe a notable reduction in memory usage and an acceleration in training speed brought by the masking training strategy.

The results show TUFA works well within the range of [50%, 85%]. If the masking ratio exceeds 85%, the number of remaining landmarks, which serves as anchors during training, is insufficient to construct a mapping from a 2D plane to a highly complex face surface. As a result, the performance

of TUFA degrades significantly. After comprehensively considering both performance and computational complexity, we set the masking ratio to 0.75 in our model.

Cross attention visualization: we visualize the cross attention maps of a scratch shape predicted by TUFA in Fig. 9. The attention maps indicate the mean weights from the structure prompts to image patches in MCA blocks, and the definition of the scratch is shown in the second row of Fig. 5. Even though the landmarks of the scratch shape are unseen during training, they can still guide TUFA to focus on the corresponding parts of faces regardless of various poses and conditions. It reveals the reliability of the zero-shot face alignment achieved by TUFA. Moreover, if the corresponding parts on faces are occluded or lack significant features (row 3, landmark 3 and 6), TUFA can also look at nearby facial components and utilize their relative positions for landmark locating. That is why TUFA demonstrates much stronger robustness than other methods in very challenging cases.

4.9 Limitations

TUFA can only generalize to unseen landmarks within the face area. If TUFA uses a set of structural prompts generated from points outside the face area on the interpretable plane, it still localize a set of nearest landmarks to these points inside the face area. The main reason is that the training datasets only contain information about facial structure. By using a dataset for dense human pose estimation (Guler, Neverova, and Kokkinos, 2018), it becomes possible to locate unseen landmarks on human bodies in a zero-shot manner.

5 Conclusion

In this paper, we introduce a novel strategy to mitigate knowledge discrepancies among multiple datasets and further achieve task-agnostic, unified face alignment. Compared to existing face alignment methods, the proposed method represents a significant breakthrough in three critical aspects: 1) TUFA successfully mitigates knowledge discrepancies among multiple datasets by aligning their landmark semantics on an interpretable plane and further unifying their learning target to map the plane to target faces. 2) the unified learning target also enables the learned knowledge to be easily transferred to a new set of defined landmarks. Consequently, TUFA significantly boosts the performance of few-shot face alignment. 3) By mitigating the discrepancies, the learned knowledge is extended from target landmarks to face structure. As a result, we can employ structure prompts to further bridge the gap between seen and unseen landmarks, achieving zero-shot face alignment for the first time.

Although the gap between seen and unseen landmarks has been successfully addressed, the gap between different domains, such as varying lighting conditions, has not been explicitly considered in TUFA. In the future, we will explore the domain gap among facial landmarks to develop a more universal face alignment approach.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Data Availability The datasets used in our paper are available in: WFLW: <https://wywu.github.io/projects/LAB/WFLW.html>, 300W: [https://wywu.github.io/projects/LAB/WFLW.html](https://ibug.doc.ic.ac.uk/resources/300-W/#::text=Automatic%20facial%20landmark%20detection%20is%20a,Masked%20300W:https://github.com/zhuocly/saat,MERL-RAV:https://github.com/abhil1kumar/MERL-RAV_dataset,COFW:https://data.caltech.edu/records/bc0bf-nc666#:text=A%20novel%20face%20dataset%20focused%20on,COFW68:https://github.com/golnazghiasi/cofw68-benchmark,CelebA:https://mmlab.ie.cuhk.edu.hk/projects/CelebA.html), 300W: https://ibug.doc.ic.ac.uk/resources/300-W/#::text=Automatic%20facial%20landmark%20detection%20is%20a,Masked%20300W:https://github.com/zhuocly/saat,MERL-RAV:https://github.com/abhil1kumar/MERL-RAV_dataset,COFW:https://data.caltech.edu/records/bc0bf-nc666#:text=A%20novel%20face%20dataset%20focused%20on,COFW68:https://github.com/golnazghiasi/cofw68-benchmark,CelebA:https://mmlab.ie.cuhk.edu.hk/projects/CelebA.html

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Asthana, A., Marks, T.K., Jones, M.J., Tieu, K.H., & Rohith, M. (2011). Fully automatic pose-invariant face recognition via 3d pose normalization. *Proc. IEEE Int. Conf. Comput. Vis.* (p.937-944).
- Ba, J.L., Kiros, J.R., Hinton, G.E. (2016). Layer normalization. [arXiv:1607.06450](https://arxiv.org/abs/1607.06450)
- Belhumeur, P.N., Jacobs, D.W., Kriegman, D.J., & Kumar, N. (2011). Localizing parts of faces using a consensus of exemplars. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.545-552).
- Browatzki, B., & Wallraven, C. (2020). 3fabrec: Fast few-shot face alignment by reconstruction. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.6109-6119).
- Bulat, A., & Tzimiropoulos, G. (2017). How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). *Proc. IEEE Int. Conf. Comput. Vis.* (p.1021-1030).
- Burgos-Artizzu, X.P., Perona, P., & Dollár, P. (2013). Robust face landmark estimation under occlusion. *Proc. IEEE Int. Conf. Comput. Vis.* (p.1513-1520).
- Caron, M., Touvron, H., Misra, I., Jegou, H., Mairal, J., Bojanowski, P., & Joulin, A. (2021). Emerging properties in self-supervised vision transformers. *Proc. IEEE Int. Conf. Comput. Vis.* (p.9630-9640).
- Chen, L., Su, H., & Ji, Q. (2019). Face alignment with kernel density deep neural network. *Proc. IEEE Int. Conf. Comput. Vis.* (p.6991-7001).
- Cootes, T., Edwards, G., & Taylor, C. (2001). Active appearance models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(6), 681-685. <https://doi.org/10.1109/34.927467>
- Cootes, T., Taylor, C., Cooper, D., & Graham, J. (1995). Active shape models-their training and application. *Comput. Vis. Image Understand.*, 61(1), 38-59. <https://doi.org/10.1006/cviu.1995.1004>
- Cristinacce, D., & Cootes, T.F. (2006). Feature detection and tracking with constrained local models. *Proc. British Mach. Vis. Conf.* (p.95.1-95.10). (<https://doi.org/10.5244/C.20.95>)
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.248-255).
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., & Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. *Proc. Int. Conf. Learn. Representations*.
- Feng, Z., Kittler, J., Christmas, W., Huber, P., Wu, X. (2017). Dynamic attention-controlled cascaded shape regression exploiting training data augmentation and fuzzy-set sample weighting. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.3681-3690).
- Feng, Z.-H., Kittler, J., Christmas, W., Huber, P., Wu, X.-J. (2017). Dynamic attention-controlled cascaded shape regression exploiting training data augmentation and fuzzy-set sample weighting. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.3681-3690).
- Ghiasi, G., & Fowlkes, C.C. (2014). Occlusion coherence: Localizing occluded faces with a hierarchical deformable part model. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.1899-1906).
- Guler, R.A., Neverova, N., & Kokkinos, I. (2018). DensePose: Dense Human Pose Estimation in the Wild. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.7297-7306).
- He, X., Bharaj, G., Ferman, D., Rhodin, H., & Garrido, P. (2023). Few-shot geometry-aware keypoint localization. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (p.21337-21348).
- He, X., Wandt, B., & Rhodin, H. (2022). Autolink: Self-supervised learning of human skeletons and object outlines by linking keypoints. S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), *Proc. advances neural inf. process. syst.* (Vol. 35, pp. 36123-36141). Curran Associates, Inc.
- Huang, G., Laradji, I., Vázquez, D., Lacoste-Julien, S., & Rodríguez, P. (2023). A survey of self-supervised and few-shot object detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(4), 4071-4089. <https://doi.org/10.1109/TPAMI.2022.3199617>
- Huang, Y., Yang, H., Li, C., Kim, J., & Wei, F. (2021). Adnet: Leveraging error-bias towards normal direction in face alignment. *Proc. IEEE Int. Conf. Comput. Vis.* (p.3060-3070).

- Jakab, T., Gupta, A., Bilen, H., & Vedaldi, A. (2018). Unsupervised learning of object landmarks through conditional image generation. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), *Proc. advances neural inf. process. syst.* (Vol. 31). Curran Associates, Inc.
- Jiang, S., Zhu, Y., Liu, C., Song, X., Li, X., & Min, W. (2023). Dataset bias in few-shot image recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(1), 229–246. <https://doi.org/10.1109/TPAMI.2022.3153611>
- Jin, H., Liao, S., & Shao, L. (2021). Pixel-in-pixel net: Towards efficient facial landmark detection in the wild. *Int. J. Comput. Vis.* <https://doi.org/10.1007/s11263-021-01521-4>
- Kowalski, M., Naruniec, J., & Trzcinski, T. (2017). Deep alignment network: A convolutional neural network for robust face alignment. *Proc. IEEE conf. comput. vis. pattern recognit. workshops* (p.2034-2043).
- Kumar, A., Marks, T.K., Mou, W., Wang, Y., Jones, M., Cherian, A., & Feng, C. (2020). Luvli face alignment: Estimating landmarks' location, uncertainty, and visibility likelihood. *Proc. IEEE conf. comput. vis. pattern recognit.* (p.8233-8243).
- Köstinger, M., Wohlhart, P., Roth, P.M., & Bischof, H. (2011). Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. *Proc. IEEE int. conf. comput. vis. workshops* (p.2144-2151).
- Lan, X., Hu, Q., & Cheng, J. (2021). Revisiting quantization error in face alignment. *Proc. IEEE int. conf. comput. vis. workshops* (p.1521-1530).
- Lan, X., Hu, Q., & Cheng, J. (2022). Atf: An alternating training framework for weakly supervised face alignment. *IEEE Trans. Multimed., Early Access*, 1-1. <https://doi.org/10.1109/TMM.2022.3164798>
- Le, V., Brandt, J., Lin, Z., Bourdev, L., & Huang, T.S. (2012). Interactive facial feature localization. *Proc. eur. conf. comput. vis.* (p.679-692).
- Li, W., Lu, Y., Zheng, K., Liao, H., Lin, C., Luo, J., & Miao, S. (2020). Structured landmark detection via topology-adapting deep graph learning. *Proc. eur. conf. comput. vis.* (pp. 266–283). Cham: Springer International Publishing.
- Lin, C., Zhu, B., Wang, Q., Liao, R., Qian, C., Lu, J., & Zhou, J. (2021). Structure-coherent deep feature learning for robust face alignment. *IEEE Trans. Image Process.*, 30, 5313–5326. <https://doi.org/10.1109/TIP.2021.3082319>
- Liu, H., Lu, J., Guo, M., Wu, S., & Zhou, J. (2020). Learning reasoning-decision networks for robust face alignment. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(3), 679–693. <https://doi.org/10.1109/TPAMI.2018.2885298>
- Liu, X. (2009). Discriminative face alignment. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31(11), 1941–1954. <https://doi.org/10.1109/TPAMI.2008.238>
- Liu, Z., Luo, P., Wang, X., & Tang, X. (2015). Deep learning face attributes in the wild. *Proc. IEEE int. conf. comput. vis.* (p.3730-3738).
- Lorenz, D., Bereska, L., Milbich, T., & Ommer, B. (2019). Unsupervised part-based disentangling of object shape and appearance. *Proc. IEEE conf. comput. vis. pattern recognit.* (p.10947-10956).
- Loshchilov, I., & Hutter, F. (2019). Decoupled weight decay regularization. *Proc. int. conf. learn. representations*.
- Lv, J., Shao, X., Xing, J., Cheng, C., & Zhou, X. (2017). A deep regression architecture with two-stage re-initialization for high performance facial landmark detection. *Proc. IEEE conf. comput. vis. pattern recognit.* (p.3691-3700).
- Mallis, D., Sanchez, E., Bell, M., & Tzimiropoulos, G. (2023). From keypoints to object landmarks via self-training correspondence: A novel approach to unsupervised landmark discovery. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(7), 8390–8404. <https://doi.org/10.1109/TPAMI.2023.3234212>
- Newell, A., Yang, K., & Deng, J. (2016). Stacked hourglass networks for human pose estimation. *Proc. eur. conf. comput. vis.* (pp. 483–499).
- Otherdout, N., Daoudi, M., Kacem, A., Ballihi, L., & Berretti, S. (2022). Dynamic facial expression generation on hilbert hypersphere with conditional wasserstein generative adversarial nets. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(2), 848–863. <https://doi.org/10.1109/TPAMI.2020.3002500>
- Pourpanah, F., Abdar, M., Luo, Y., Zhou, X., Wang, R., Lim, C. P., & Wu, Q. M. J. (2023). A review of generalized zero-shot learning methods. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(4), 4051–4070. <https://doi.org/10.1109/TPAMI.2022.3191696>
- Prados-Torrealblanca, A., Buenaposada, J.M., & Baumela, L. (2022). Shape preserving facial landmarks with graph attention networks. *Proc. british mach. vis. conf.*
- Qian, S., Sun, K., Wu, W., Qian, C., & Jia, J. (2019). Aggregation via separation: Boosting facial landmark detector with semi-supervised style translation. *Proc. IEEE int. conf. comput. vis.* (p.10152-10162).
- Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., & Pantic, M. (2013a). 300 faces in-the-wild challenge: The first facial landmark localization challenge. *Proc. IEEE int. conf. comput. vis. workshops* (p.397-403).
- Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., & Pantic, M. (2013b). A semi-automatic methodology for facial landmark annotation. *Proc. IEEE conf. comput. vis. pattern recognit. workshops* (p.896-903).
- Tai, Y., Liang, Y., Liu, X., Duan, L., Li, J., Wang, C., & Chen, Y. (2019). Towards highly accurate and stable face alignment for high-resolution videos. *Proc. AAAI conf. artif. intell.* (Vol. 33, pp. 8893–8900).
- Tanfous, A. B., Drira, H., & Amor, B. B. (2020). Sparse coding of shape trajectories for facial expression and action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(10), 2594–2607. <https://doi.org/10.1109/TPAMI.2019.2932979>
- Tang, H., Shao, L., Torr, P. H., & Sebe, N. (2023). Bipartite graph reasoning gans for person pose and facial image synthesis. *Int. J. Comput. Vis.*, 131(3), 644–658. <https://doi.org/10.1007/s11263-022-01722-5>
- Tang, Z., Peng, X., Li, K., & Metaxas, D. N. (2020). Towards efficient unets: A coupled and quantized approach. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(8), 2038–2050. <https://doi.org/10.1109/TPAMI.2019.2907634>
- Thewlis, J., Bilen, H., & Vedaldi, A. (2017). Unsupervised learning of object landmarks by factorized spatial embeddings. *Proc. IEEE int. conf. comput. vis.* (p.3229-3238).
- Trigeorgis, G., Snape, P., Nicolaou, M.A., Antonakos, E., & Zafeiriou, S. (2016). Mnemonic descent method: A recurrent process applied for end-to-end face alignment. *Proc. IEEE conf. comput. vis. pattern recognit.* (p.4177-4187).
- Valle, R., Buenaposada, J.M., Valdés, A., & Baumela, L. (2018). A deeply-initialized coarse-to-fine ensemble of regression trees for face alignment. *Proc. eur. conf. comput. vis.* (pp. 609–624). Cham.
- Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., & Xiao, B. (2021). Deep high-resolution representation learning for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(10), 3349–3364. <https://doi.org/10.1109/TPAMI.2020.2983686>
- Wang, X., Bo, L., & Fuxin, L. (2019). Adaptive wing loss for robust face alignment via heatmap regression. *Proc. IEEE int. conf. comput. vis.* (p.6970-6980).
- Wu, W., Qian, C., Yang, S., Wang, Q., Cai, Y., & Zhou, Q. (2018). Look at boundary: A boundary-aware face alignment algorithm. *Proc. IEEE conf. comput. vis. pattern recognit.* (p.2129-2138).
- Wu, W., & Yang, S. (2017). Leveraging intra and inter-dataset variations for robust face alignment. *Proc. IEEE conf. comput. vis. pattern recognit. workshops* (p.2096-2105).
- Xia, J., Qu, W., Huang, W., Zhang, J., Wang, X., & Xu, M. (2022). Sparse local patch transformer for robust face alignment and landmarks

- inherent relation learning. *Proc. ieee conf. comput. vis. pattern recognit.* (p.4052-4061).
- Xia, J., Xu, M., Zhang, H., Zhang, J., Huang, W., Cao, H., & Wen, S. (2023). Robust face alignment via inherent relation learning and uncertainty estimation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(8), 10358–10375. <https://doi.org/10.1109/TPAMI.2023.3260926>
- Xiao, B., Wu, H., & Wei, Y. (2018). Simple baselines for human pose estimation and tracking. *Proc. eur. conf. comput. vis.* (pp. 472–487).
- Xiao, S., Feng, J., Xing, J., Lai, H., Yan, S., & Kassim, A. (2016). Robust facial landmark detection via recurrent attentive-refinement networks. *Proc. eur. conf. comput. vis.* (pp. 57–72). Cham.
- Xiong, X., & De la Torre, F. (2013). Supervised descent method and its applications to face alignment. *Proc. ieee conf. comput. vis. pattern recognit.* (p.532-539).
- Xu, X., Meng, Q., Qin, Y., Guo, J., Zhao, C., Zhou, F., & Lei, Z. (2021). Searching for alignment in face recognition. *Proc. aaai conf. artif. intell.* (Vol. 35, pp. 3065–3073).
- Yang, J., Liu, Q., & Zhang, K. (2017). Stacked hourglass network for robust facial landmark localisation. *Proc. ieee conf. comput. vis. pattern recognit. workshops* (p.2025-2033).
- Yang, S., Luo, P., Loy, C.-C., & Tang, X. (2016). Wider face: A face detection benchmark. *Proc. ieee conf. comput. vis. pattern recognit.*
- Zhang, H., Xu, L., Lai, S., Shao, W., Zheng, N., Luo, P., & Zhang, K. (2024). Open-vocabulary animal keypoint detection with semantic-feature matching. *Int. J. Comput. Vis.*, 1–18. DOI: <https://doi.org/10.1007/s11263-024-02126-3>
- Zhang, J., Hu, H., & Feng, S. (2020). Robust facial landmark detection via heatmap-offset regression. *IEEE Trans. Image Process.*, 29, 5050–5064. <https://doi.org/10.1109/TIP.2020.2976765>
- Zhang, Y., Guo, Y., Jin, Y., Luo, Y., He, Z., & Lee, H. (2018). Unsupervised discovery of object landmarks as structural representations. *Proc. ieee conf. comput. vis. pattern recognit.* (p.2694-2703).
- Zheng, Y., Yang, H., Zhang, T., Bao, J., Chen, D., Huang, Y., & Wen, F. (2022). General facial representation learning in a visual-linguistic manner. *Proc. ieee conf. comput. vis. pattern recognit.* (p.18697-18709).
- Zhou, Z., Li, H., Liu, H., Wang, N., Yu, G., & Ji, R. (2023). Star loss: Reducing semantic ambiguity in facial landmark detection. *Proc. ieee conf. comput. vis. pattern recognit.* (p.15475-15484).
- Zhu, C., Li, X., Li, J., Dai, S. (2021). Improving robustness of facial landmark detection by defending against adversarial attacks. *Proc. ieee int. conf. comput. vis.* (p.11731-11740).
- Zhu, C., Wan, X., Xie, S., Li, X., Gu, Y. (2022). Occlusion-robust face alignment using a viewpoint-invariant hierarchical network architecture. *Proc. ieee conf. comput. vis. pattern recognit.* (p.11102-11111).
- Zhu, S., Li, C., Loy, C.C., Tang, X. (2015). Face alignment by coarse-to-fine shape searching. *Proc. ieee conf. comput. vis. pattern recognit.* (p.4998-5006).
- Zhu, X., & Ramanan, D. (2012). Face detection, pose estimation, and landmark localization in the wild. *Proc. ieee conf. comput. vis. pattern recognit.* (p.2879-2886).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.