

Multi-spectral Facial Landmark Detection

Jin Keong, Xingbo Dong and Zhe Jin

School of Information Technology

Monash University Malaysia

47500 Subang Jaya, Selangor, Malaysia

Email: {keong.jin, xingbo.dong, jin.zhe}@monash.edu

Khawla Mallat and Jean-Luc Dugelay

Department of Digital Security

EURECOM

Sophia-Antipolis, France

Email: {mallat,dugelay}@eurecom.fr

Abstract—Thermal face image analysis is favorable for certain circumstances. For example, illumination-sensitive applications, like nighttime surveillance; and privacy-preserving demanded access control. However, the inadequate study on thermal face image analysis calls for attention in responding to the industry requirements. Detecting facial landmark points are important for many face analysis tasks, such as face recognition, 3D face reconstruction, and face expression recognition. In this paper, we propose a robust neural network enabled facial landmark detection, namely Deep Multi-Spectral Learning (DMSL). Briefly, DMSL consists of two sub-models, i.e. face boundary detection, and landmark coordinates detection. Such an architecture demonstrates the capability of detecting the facial landmarks on both visible and thermal images. Particularly, the proposed DMSL model is robust in facial landmark detection where the face is partially occluded, or facing different directions. The experiment conducted on Eurecom’s visible and thermal paired database shows the superior performance of DMSL over the state-of-the-art for thermal facial landmark detection. In addition to that, we have annotated a thermal face dataset with their respective facial landmark for the purpose of experimentation.

I. INTRODUCTION

In recent years, infrared spectrum (IR) imagery has been utilized to facilitate the development of computer vision (CV) applications mainly due to its light-insensitive nature. For instance, IR camera can capture images under unfavorable illumination situations, hence it can be used at night or even under completely dark conditions. Most conventional IR imagery today uses an active approach which senses reflected electromagnetic radiation (EMR) within the Near infrared (NIR) or Short-Wave infrared (SWIR) spectrum. This approach relies on the external IR illumination [1]. As a result, it can be deteriorated by different illumination conditions. Thermal, a.k.a Long-Wave infrared ($7\mu\text{m} - 14\mu\text{m}$ wavelength), is regarded as a promising direction to extend imaging technology [2]. Thermal imagery is a passive IR technology that primarily captures the emitted EMR specifically the heat energy from the object, hence thermal spectrum is insensitive to illumination variations. However, it can be affected by ambient temperature.

Bourlai and Cukik [3] presented some scenarios which require images captured from multiple electromagnetic spectra to perform facial recognition. This suggests that multi-spectral

imagery technology can be employed to tackle some CV problems that involve working under a variety of illumination conditions. Chang et al. [4] showed that the images from multiple spectra can be fused together to perform face recognition under poor illumination conditions. They also presented a way to fuse visible and thermal face images. There are some scenarios where the visible and thermal imagery can work in conjunction with each other for a more robust performance than a single spectral system. For instance, in real time facial image analysis, we can have a dual camera setup equipped with visible and thermal cameras such as (FLIR Duo Pro R 640), this camera has a thermal sensor as well as visible light sensor designed for camera drone. Depending on the illumination conditions, we can switch in between both configurations, or combine the input to perform a seamless analysis.

One of the main prerequisites of facial image analysis is facial landmark detection as it plays a key role in many tasks, such as automatic face recognition [5] [6], expression recognition [7], 3D face reconstruction [8]. Many facial landmark detectors have been designed for face images in the visible spectrum and they are proved to be very reliable. However, these detectors fail to achieve good performance for thermal images [9]. To achieve multi-spectral analysis, there is a need for an efficient thermal facial landmark detector as well. In recent years, deep learning methods have been successfully applied to many areas such as image recognition due to their automatic learning capabilities, and have brought significant improvements in these areas. Therefore, this paper focuses on the thermal facial landmark detection based on deep learning methods.

Due to the expensive capture devices and limited dataset, there are only few works on thermal facial landmark detection. Active Appearance Models [10] is a relatively traditional method initially used to model faces given facial landmark positions [11]. This method is a handcrafted model designed to compute the mean shape and appearance metric of a given dataset and use to it to predict the landmark of an unseen image. Like most handcrafted techniques, this method does not work well with high data variability.

A current trend towards thermal facial landmark detection is to apply deep learning models that are designed for detecting visible facial landmarks and adapted to the thermal image domain. Poster et al.[12] noted that the state-of-the-art visible

facial landmark detection models work poorly with thermal images due to relatively low information contained in thermal images. This work shows that thermal images aligned with modern landmark detection algorithms often fail to achieve thermal-to-visible face verification results compared to manually aligned imagery. The models being tested in this study include Deep Alignment Network (DAN) [13], Multi-task Convolutional Neural Network (MTCNN) [14], and Multi-class Patch-based fully convolutional neural network (PBC) [12].

Another solution is to transform thermal images into visible images, and then to detect the landmarks from the transformed images using existing landmark detectors designed for visible images [9]. The transformation is achieved by developing a CycleGAN model [15]. This method is shown to be unreliable, for example, CycleGAN mode collapse. This solution also offers room for inconsistency as it involves transforming a low information image into a higher information domain. To tackle this problem, an end-to-end landmark detection model based on two-stage training mechanism is proposed in [9]. The first stage of this solution is to train a U-Net [16] to outline the face of a given image. A fully connected network is attached to the U-Net at the second stage to estimate facial landmark coordinates from the outline. An auxiliary output layer is used to detect the facial expression to enhance the training outcome. This solution however, are only able to detect front facing faces with no occlusion. Even self occlusion from slight face tilting to different direction results in landmark distortion.

Based on the aforementioned discussion, thermal facial landmark detection is still a challenging task. In this paper, we proposed a Deep Multi-Spectral Learning (DMSL) model for facial landmark detection. The proposed model is a stack of two independent sub-models, an auxiliary model for detecting face boundary, and a main model for detecting facial landmarks. DMSL is a unified model that can detect facial landmarks from both thermal and visible images, this leads to efficient and convenient solution for multi-spectral facial image analysis. Another core strength is that this model is robust for thermal facial landmark detection when a face is not front facing or when there is occlusion. Additionally, we annotated a thermal image dataset by a 68 points landmark configuration. This dataset is used for the experimentation of our proposed deep learning model based on U-Net to detect facial landmarks on thermal images.

The rest of this paper is organized as follows. In Section II, we present the proposed model architecture and pipeline. In section III, we describe the dataset used for experimentation. The steps for pre-processing, and dataset annotation, are outlined in section IV. In section V, we describe the metrics used to evaluate our model. In section VI, we explain the experiment process in detail as well as illustrate the performance and effectiveness of the model. Finally we conclude with a discussion in section VII.

II. MULTI-SPECTRAL FACIAL LANDMARK DETECTOR

A. Model intuition

Compared to visible images, thermal images contain relatively low information and has a narrower pixel intensity range. When a network is being trained directly to detect facial landmarks, it is easy for the network to overfit and recognise area such as neck or large beard as facial landmarks.

To address this problem we use an auxiliary model to highlight the face region within a given image and blackout other pixels first. This allows the main landmark detector to focus on the targeted region for more precise facial landmark detection.

B. Model task formulation

Let the auxiliary model of the DMSL that outputs face boundary be M_b . When given a face image F , the auxiliary outputs an array B that contains four values $B = [\beta_1, \beta_2, \beta_3, \beta_4]$. These values represent the (x, y) coordinate of the leftmost highest point, width w , and height h of the boundary respectively. An example is depicted in figure 1.

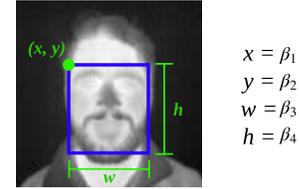


Fig. 1: Array B value depiction.

Boundary loss L_b represented by equation 1 is used as the loss function for M_b :

$$L_b = \frac{1}{V} \sum_{i=1}^V (\beta_i - \hat{\beta}_i)^2 \quad (1)$$

where V denotes the number of output values. In the case of face boundary detection, $V = 4$. β denotes the ground truth value and $\hat{\beta}$ denotes predicted value. This is the mean squared error between ground truth and predicted value.

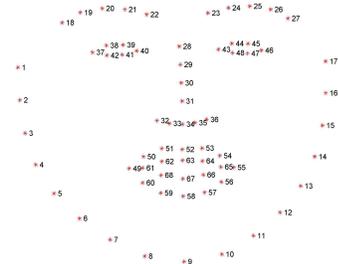


Fig. 2: Illustration of 68 point facial landmark configuration. Figure taken from [17].

The main model for facial landmark detection referred to as M_l is trained to output an array Λ that contains P facial

landmark coordinates when given a face images F . We used a well established 68 point facial landmark configuration from Multi-PIE [18] shown in figure 2. Each point is represented by an (x, y) coordinate. Λ is an array of x and y coordinates arranged in the form of $\Lambda = [x_1, y_1, x_2, y_2, x_3, y_3, \dots, x_P, y_P]$, where $P = 68$.

M_l is trained using landmark loss L_l implemented by [9]. L_l is defined as:

$$L_l = \frac{1}{P} \sum_{i=1}^P ((x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2) \quad (2)$$

where P denotes the number of landmark points. This calculates with the squared difference between ground truth coordinates Λ , and predicted coordinates $\hat{\Lambda}$.

C. Model architecture

We use a modified version of U-Net developed by [16]. U-Net is originally used for biomedical image segmentation to tackle the issue of limited availability of data within the domain. This solves our problem of having only small datasets. Another reason to use U-Net is because our task is similar to image segmentation. In biomedical image segmentation, the goal is for the model to outline the target area required by user. This fits our purpose because we need to outline the prominent landmark features on the face and then to find the corresponding coordinates.

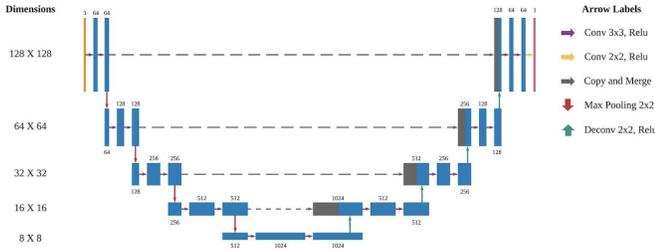


Fig. 3: The configured U-Net structure.

U-Net is a neural network that is only composed of convolutional layers. There are two major paths in this architecture, down-sampling and up-sampling. Figure 3 shows our configuration of U-Net. The left side of the U shape structure is the down-sampling path, which is similar to typical convolutional layers in convolutional neural network (CNN). One down-sampling step involves two 3×3 , stride 1 convolution and a 2×2 max pooling with stride 2. An up-sampling step starts with a 2×2 deconvolution with stride 2, and followed by two 3×3 , stride 1 convolution. At every up-sampling step, features from the down-sampling path with the same dimension are concatenated to the up-sampled feature. This is represented by the gray arrow in figure 3. Finally, we have a 1×1 , stride 1 convolution operation to create an output with the same dimension as input but with depth set to one. With this configuration, the output dimension should be the same as the

input, which in our case is 128×128 . The activation function used throughout is ReLU.

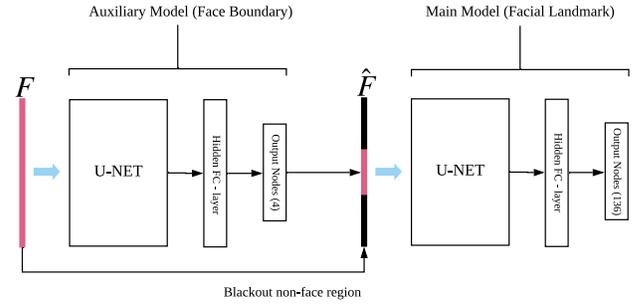


Fig. 4: Illustration of our proposed pipeline.

Figure 4 shows the pipeline of DMSL. F represents an input image and \hat{F} represents the image with non-face region blacked out using the output from the auxiliary model. From figure 4, Auxiliary model illustrates M_b . The output dimension of U-Net in M_b is flattened to 16384×1 . It is then connected to a fully connected network with a hidden layer with 1024 nodes and an output layer with 4 nodes for predicting array B . M_l is represented by the main model in figure 4. The difference between the two models is that the output layer in M_l has 136 nodes for the prediction of array Λ . The sigmoid function is used for all fully connected layers.

III. DATASET

Eurecom's VIS-TH visible and thermal paired face database (VIS-TH) is used in this study [19]. The dataset is composed of 50 subjects with varying ethnicity, sex, and age. The camera is set to capture two images simultaneously, one capturing EMR from the thermal spectrum and another capturing visible EMR.

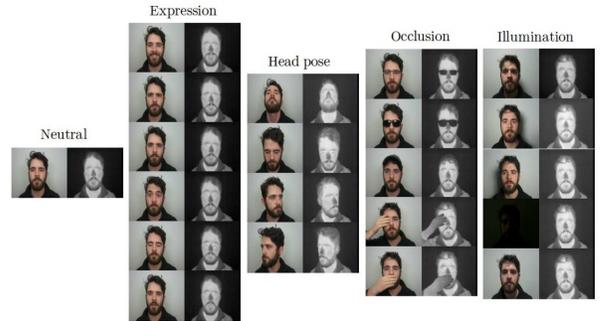


Fig. 5: Example of visible and thermal image pair with all variations. Figure taken from [19]

The image pair for all variations is shown in figure 5. Each subject has 21 visible (VIS) images and 21 thermal (TH) images. This results in a total of 2100 images included in the dataset. It is worthy to highlight that VIS-TH dataset fits our settings due to the pixel-to-pixel registered images for both VIS and TH face. This makes it convenient for annotation purposes.

IV. PRE-PROCESSING, AUGMENTATION, AND ANNOTATION

All images are cropped into a 1 : 1 ratio image without cutting part of the subject face to match the input dimension of the U-Net convolution layer (120×12). To mitigate the issue of having a small dataset, we generate a vertically mirrored version for each image within the dataset. This increases the final number of images to 4200. There is a particular variation where an image is taken under all lights off condition. The resulting black VIS image are omitted from the experiments, making the total usable images count 4100.

Dlib-ml [20] is used to annotate facial landmark on all VIS images. Taking advantage of this dataset having both TH and VIS image being captured simultaneously, the annotated data is subsequently superimposed on their TH image counterpart. This gives us ground truth array Λ . Figure 6 shows the landmark plot for multiple images. The first row shows the VIS images being used for detection, while the second row represents plots of the superimposed coordinates on TH image. This is used as ground truth array B .



Fig. 6: Example of annotated images.

In many cases, there are some misalignment of landmark points by Dlib. To counter this, we manually calibrate inaccurate points to fit the face within an image. From the adjusted landmark coordinates, we find the leftmost highest point, rightmost point and, the lowest point from it. From this, we can calculate the width and height of face. These value along with the two coordinate value from the leftmost highest point is used to represent the boundary of face.

V. EVALUATION METRICS

The performance of DMSL is evaluated by measuring normalized mean error (NME) used by [17]. It normalises the distance between ground truth and predicted value with the distance of both eyes. This is represented by equation 3

$$NME = \frac{1}{N} \sum_{i=1}^N \frac{\|\Lambda_i - \hat{\Lambda}_i\|}{P \times D_i} \quad (3)$$

In equation 3, $\|\Lambda - \hat{\Lambda}\|$ represents the L2 norm between ground truth, Λ and predicted value, $\hat{\Lambda}$. D_i represents the interocular distance for that particular face. This is calculated by the euclidean distance between the outer corner of two eyes. For our configuration, the right eye is represented by point 37 while the left eye is marked with point 46 as illustrated

in figure 2. P denotes the number of facial landmark points. Finally, N represents the number of sample images used.

VI. EXPERIMENT AND RESULTS

A. Training and testing protocol

We train and test the model using a 10-fold cross validation. The 50 subjects are divided into 10 groups by their ID number. Each iteration we use 1 group for testing, 2 for validating, and the rest for training.

The experiments are conducted on a Linux server equipped with one NVIDIA Titan Xp graphics processing unit (GPU). The deep learning models are constructed and tested with TensorFlow 2.0.0-rc0 using Python 3.7.

B. Training

To produce an optimised DMSL model, Both M_b and M_l are trained individually using a two stage training procedure similar to the training process in [9]. This is to maximise the advantage of using a U-Net as feature extractor. In essence, we have two individual models and both require a two stage training procedure.

1) *Face boundary detection model M_b* : The first stage is to train a U-Net to outline the face boundary in the original image. A boundary mask shown in figure 7(b) is generated from our face boundary ground truth annotation.

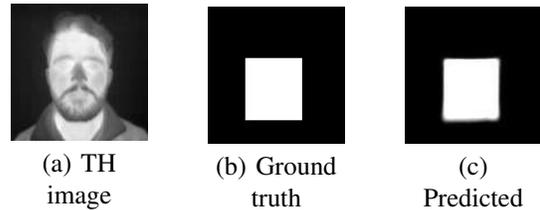


Fig. 7: Example of face boundary mask.

The goal is to train the U-Net to get an output similar to the boundary mask shown in figure 7(b). Figure 7(c) shows the output by this U-Net. This is achieved by minimising U-Net loss defined in equation 4.

$$L_u = -\frac{1}{N} \sum_{i=1}^N m_i \log \hat{m}_i + (1 - m_i) \log(1 - \hat{m}_i) \quad (4)$$

where, m , and \hat{m} represents ground truth mask and predicted mask respectively with N represents the number of used sample images. This computes the softmax cross-entropy loss between the two masks.

The second stage is for training the attached fully connected network to output the B array. The parameters of the convolutional U-Net are frozen and the fully connected network is trained by minimising boundary loss represented by equation 1.

Figure 8 shows a comparison between images cropped using values predicted by the auxiliary model against values from ground truth. These cropped images will be the input of the main model that detects facial landmark.

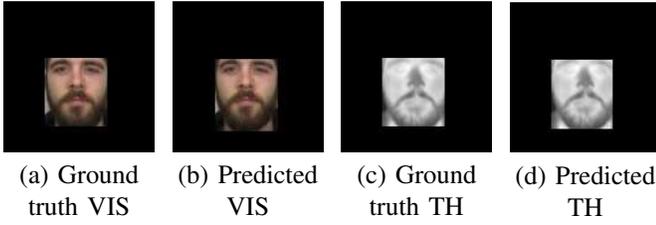


Fig. 8: Example of non-face region blacked out using ground truth value and M_b generated value.

2) *Facial landmark detection model M_L* : The main model uses the same two stage training procedure. The labels for first stage are mask generated from the landmark coordinates with algorithm 1. Figure 9(b) is an example of the mask generated.

Algorithm 1: Generate Landmark Mask

input : Coordinates, $Dimension_x$, $Dimension_y$
output: Mask

```

Initialize 2D array (Mask) according to  $Dimension_x$ ,
and  $Dimension_y$ ;
while  $i \leq \text{length}(\text{Coordinates})$  do
   $x \leftarrow \text{Coordinates}[i] \times Dimension_x$ ;
   $y \leftarrow \text{Coordinates}[i + 1] \times Dimension_y$ ;
  for  $j = 1, j \leq X - Dimension_x$  do
    for  $k = 1, k \leq Y - Dimension_y$  do
      if  $\text{Mask}[j][k] < 255$  then
         $\text{Mask}[j][k] +=$ 
           $0.5^{\max(|x-j|, |y-k|)} \times 255$ ;
      end
    end
  end
end

```

Algorithm 1 is based on the method used by [9]. The idea behind this mask is to highlight all the 68 facial landmark points. The surrounding area are set to gradually dim creating a glowing effect. This helps the model to outline the facial landmarks.

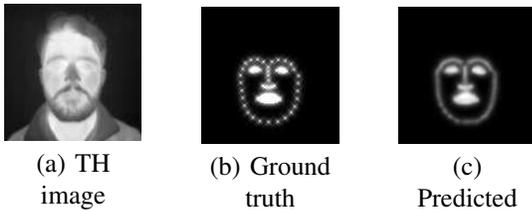


Fig. 9: Example of facial landmark mask.

The goal in stage one is to train a U-Net to output an image similar to Figure 9(b). The training is done by minimising U-Net loss. Figure 9(c) shows an example of the U-Net output.

Stage two is for training the attached fully connected network of the main model to output the Λ array. The convolu-

tional U-Net parameters are frozen, while the fully connected network are trained by minimising the loss represented by equation 2.

C. Performance of Landmark Detection

Figure 10 shows the plotted facial landmark output by our model. The model is able to detect facial landmarks with samples facing different directions.



Fig. 10: Sample plots of output from our model with different face pose.

The model also works fairly well when certain areas such as eyes or mouth are occluded. This is shown in figure 11.



Fig. 11: Sample plots of output from our model with different face occlusions.

We evaluate the performance of our model against Dlibml [20], the optimal AAM method examined in [10], and a two stage fully connected multi-task U-Net [9]. The output of both models is evaluated using our manually annotated ground truth.

| | TH | VIS |
|----------------|--------|--------|
| AAM [10] | 0.1311 | 0.1434 |
| Dlib [20] | 0.0293 | 0.0581 |
| Chu et al. [9] | 0.0222 | 0.0556 |
| Ours | 0.0210 | 0.0544 |

TABLE I: Model performance comparison in terms of NME

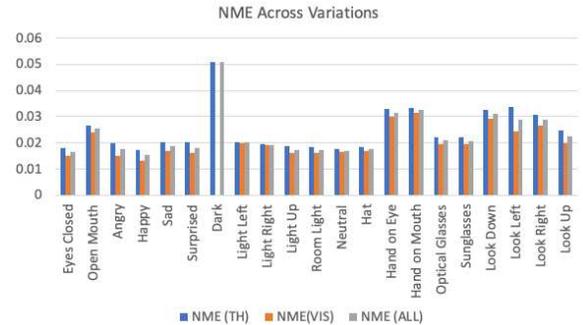


Fig. 12: Normalised mean error comparison between variations.

Comparison between each method is shown in table I. Our model is able to achieve a better result for TH images and a competitive result for VIS images. The samples that Dlib-ml and AAM models are unable to detect are excluded. By doing so, we ensure a fair comparison of flaws. These samples consist of 711 TH images and 139 VIS images including the mirrored version. VIS images show a higher NME because the portion of occluded images present are higher compared to TH images. Figure 12 illustrates the performance of our DMSL model in detail. This shows the average error value for each variation with all available images included.

VII. DISCUSSION AND CONCLUSION

An interesting observation from the experiments is that by training a model with images from both spectrum yields better results for thermal landmark detection. This is illustrated in figure 12 where the class *Dark* that is trained with a lack of a usable VIS image has relatively poor performance when compared to other variations.

Experimental results have shown that our DMSL model can outperform existing landmark detection methods. However, we still have a small amount of cases where the model outputs ill-shaped or compressed landmark coordinates. An example of this output is illustrated in figure 13. This could be potentially mitigated in the future by experimenting with more loss functions and adjusting the fully connected network.

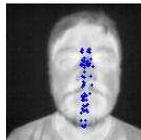


Fig. 13: Ill-shaped facial landmark output.

The main contribution of this study is that we have developed a DMSL model for facial landmark detection. This model can detect facial landmarks in TH and VIS images regardless of which direction the target is facing. In addition, the model can still detect landmarks on faces that are partially occluded by objects such as sunglasses, hand, and hat. Both of these capabilities are not present in existing state-of-the-art thermal facial landmark systems. Another contribution of this study is to provide a facial landmark annotation for a TH image dataset [19].

This study can be further applied for an array of future research involving multi-spectral or thermal imagery. This includes but is not limited to facial recognition, thermal-to-visible face generation, and real-time surveillance. Since we are using information collected from heat generated by a human body, this research can be potentially integrated with facial biometrics tasks with body temperature detection. This could help with public health investigation for disease detection.

ACKNOWLEDGMENT

This research was supported by Fundamental Research Grant Scheme (FRGS/1/2018/ICT02/MUSM/03/3) and NVIDIA Corporation donation of the Titan Xp GPU.

REFERENCES

- [1] X. P. Maldague, "Introduction to ndt by active infrared thermography," *Materials Evaluation*, vol. 60, no. 9, pp. 1060–1073, 2002.
- [2] R. S. Ghiass, O. Arandjelović, A. Bendada, and X. Maldague, "Infrared face recognition: A comprehensive review of methodologies and databases," *Pattern Recognition*, vol. 47, no. 9, pp. 2807–2824, 2014.
- [3] T. Bourlai and B. Cukic, "Multi-spectral face recognition: Identification of people in difficult environments," in *2012 IEEE International Conference on Intelligence and Security Informatics*. IEEE, 2012, pp. 196–201.
- [4] H. Chang, A. Koschan, M. Abidi, S. G. Kong, and C.-H. Won, "Multispectral visible and infrared imaging for face recognition," in *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2008, pp. 1–6.
- [5] L. Wiskott, N. Krüger, N. Kuiger, and C. Von Der Malsburg, "Face recognition by elastic bunch graph matching," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 775–779, 1997.
- [6] M. Husken, M. Brauckmann, S. Gehlen, and C. Von der Malsburg, "Strategies and benefits of fusion of 2d and 3d face recognition," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)-Workshops*. IEEE, 2005, pp. 174–174.
- [7] H. Jung, S. Lee, J. Yim, S. Park, and J. Kim, "Joint fine-tuning in deep neural networks for facial expression recognition," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2983–2991.
- [8] J. Choi, G. Medioni, Y. Lin, L. Silva, O. Regina, M. Pamplona, and T. C. Faltemier, "3d face reconstruction using a single or multiple views," in *2010 20th International Conference on Pattern Recognition*. IEEE, 2010, pp. 3959–3962.
- [9] W.-T. Chu and Y.-H. Liu, "Thermal Facial Landmark Detection by Deep Multi-Task Learning," in *2019 IEEE 21st International Workshop on Multimedia Signal Processing (MMSP)*, 2019, pp. 1–6.
- [10] M. Kopaczka, K. Acar, and D. Merhof, "Robust facial landmark detection and face tracking in thermal infrared images using active appearance models," in *VISIGRAPP (4: VISAPP)*, 2016, pp. 150–158.
- [11] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," in *European conference on computer vision*. Springer, 1998, pp. 484–498.
- [12] D. Poster, S. Hu, N. Nasrabadi, and B. Riggan, "An examination of deep-learning based landmark detection methods on thermal face imagery," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.
- [13] M. Kowalski, J. Naruniec, and T. Trzcinski, "Deep alignment network: A convolutional neural network for robust face alignment," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 88–97.
- [14] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.
- [15] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [16] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 9351. Springer Verlag, 2015, pp. 234–241.
- [17] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, "300 faces in-the-wild challenge: The first facial landmark localization challenge," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2013, pp. 397–403.
- [18] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-pie," *Image and Vision Computing*, vol. 28, no. 5, pp. 807–813, 2010.
- [19] K. Mallat and J.-L. Dugelay, "A benchmark database of visible and thermal paired face images across multiple variations," in *International Conference of the Biometrics Special Interest Group, {BIOSIG} 2018, Darmstadt, Germany, September*, ser. LNI. {GI} / {IEEE}, pp. 199–206.
- [20] D. E. King, "Dlib-ml: A Machine Learning Toolkit," *Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.