

Towards Omni-Supervised Face Alignment for Large Scale Unlabeled Videos

Congcong Zhu,^{1,2} Hao Liu,^{2,3*} Zhenhua Yu,^{2,3} Xuehong Sun^{2,3}

¹School of Computer Engineering and Science, Shanghai University, Shanghai, 200444, China

²School of Information Engineering, Ningxia University, Yinchuan, 750021, China

³Collaborative Innovation Center for Ningxia Big Data and Artificial Intelligence

Co-founded by Ningxia Municipality and Ministry of Education, Yinchuan, 750021, China

congcong.zhu_nxu@outlook.com, {liuhao, zhyu, sunxh}@nxu.edu.cn

Abstract

In this paper, we propose a spatial-temporal relational reasoning networks (STRRN) approach to investigate the problem of omni-supervised face alignment in videos. Unlike existing fully supervised methods which rely on numerous annotations by hand, our learner exploits large scale unlabeled videos plus available labeled data to generate auxiliary plausible training annotations. Motivated by the fact that neighbouring facial landmarks are usually correlated and coherent across consecutive frames, our approach automatically reasons about discriminative spatial-temporal relationships among landmarks for stable face tracking. Specifically, we carefully develop an interpretable and efficient network module, which disentangles facial geometry relationship for every static frame and simultaneously enforces the bi-directional cycle-consistency across adjacent frames, thus allowing the modeling of intrinsic spatial-temporal relations from raw face sequences. Extensive experimental results demonstrate that our approach surpasses the performance of most fully supervised state-of-the-arts.

Introduction

Face alignment (*a.k.a.*, facial landmark detection) aims to detect multiple facial landmarks for the given facial image or video sequence, which has dominated a crucial step in many facial analysis tasks such as face identification and recognition (Hu, Lu, and Tan 2014; Liu et al. 2019a) and face animation (Roth, Tong, and Liu 2015; Liu et al. 2016). While significant works have been devoted to face alignment recently (Liu et al. 2019b; Wu et al. 2018; Kumar and Chellappa 2018), the practical performance hardly meets the precise requirements in the scenario of performing large scale unlabeled videos. The major reasons are two-fold: 1) Existing methods intensely rely on the sheer volume of training annotations. 2) It is tedious to annotate the spatial-temporal structures on the massive set of frames. Hence, we are supposed to propose an accurate and robust face alignment algorithm to automatically annotate the large amounts of unlabeled face videos.

*Corresponding Author is Hao Liu.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

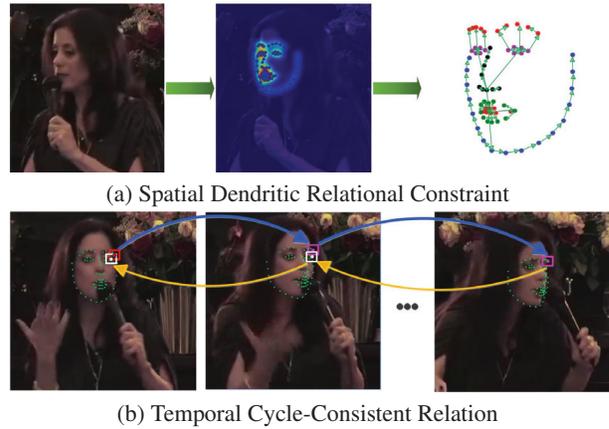


Figure 1: Insight of our proposed STRRN. Our method processes the large scale unlabeled video by automatically distilling the spatial-temporal knowledge, *i.e.*, a) exploiting the facial geometry relationship by holding a dendritic structure, b) enforcing the bi-directional cycle-consistency temporally on consecutive frames. We leveraged these self-supervised spatial-temporal relations to generate extra reliable training annotations for model update.

Throughout recent literatures, face alignment is dominated by the regression-based approaches (Trigeorgis et al. 2016; Wu et al. 2018; Kumar and Chellappa 2018; Guo, Lu, and Zhou 2018; Liu et al. 2019b; 2018), which typically seek discriminative feature-to-shape mappings with preserving the shape constraint. Methods such as LAB (Wu et al. 2018) and PCD-CNN (Kumar and Chellappa 2018) impose the relative geometry among multiple landmarks, making the plug-in module pose-invariant for accurate performance. However, these fully-supervised methods require large amounts of precise hand-crafted annotations for training and cannot be directly adopted for unlabeled data. Besides the spatial relationship, video-based methods such as CCR (Sánchez-Lozano et al. 2016) and TSTN (Liu et al. 2018) enforces the consistency relation temporally across adjacent frames, which improves the stableness versus the jitter problem during tracking. One major issue in this type of method is

that the training efficiency still relies on the tremendous volume of per-frame annotations, where it is challenging to manually annotate numerous frames. To address this issue, we investigate into the omni-supervised learning towards face alignment inspired by (Radosavovic et al. 2018; Dong et al. 2018). Our basic idea aims at reasoning disentangling spatial-temporal relationship by exploiting large scale unlabeled videos plus all available training annotations. The self-reasoned spatial-temporal relation allows us to trust their predictions on the unseen video data. The intuitive idea is presented in Fig. 1.

To tackle the omni-supervised face alignment, we propose a spatial-temporal relational reasoning networks (STRRN) approach, which automatically distills interpretable knowledge of large scale unlabeled video data. To produce plausible training annotations for model update, our approach reasons about meaningful relations of unlabeled videos in both the spatial and temporal dimensions accordingly. For modeling of the spatial relation, we first divide the whole face into different semantic components (*e.g.*, eyes, nose, mouth, facial cheek and etc.), where the nose is assumed as the root. Then our STRRN disentangles the component-based appearance and geometric information with preserving a dendritic structure. For modeling of the temporal relation, we ensure that our model tracks forward until the target and it should arrive the starting position in the backward order. This principally enforces the cycle-consistent temporal relation on consecutive frames for reliable tracking. To learn the network parameters, we adopt a cooperative and competitive strategy to exploit the complementary information from both the tracking module and the backbone detector. Fig. 2 shows the detailed network architecture of the proposed STRRN. To evaluate the effectiveness of our proposed approach, we carry out extensive experiments on folds of large scale unlabeled video datasets and experimental results indicate compelling performance versus state-of-the-art methods.

Related Work

We briefly review some related literatures of existing face alignment methods and knowledge distillation models.

Conventional Face Alignment: Existing face alignment methods are roughly classified into image-based (Cao et al. 2012; Xiong and la Torre 2013; Zhu et al. 2015; Ren et al. 2014) and video-based (Sánchez-Lozano et al. 2016; Tzimiropoulos 2015). Earlier shallow models intend to seek a sequence of discriminative linear feature-to-shape mappings, so that the initialized shape is adjusted to the target one in a coarse-to-fine manner. Representative cascade regression-based methods include explicit shape regression (ESR) (Cao et al. 2012), supervised descent method (SDM) (Xiong and la Torre 2013) and coarse-to-fine shape searching (CFSS) (Zhu et al. 2015). To make the image-based methods suitable for video data, one common solution is to regard the outcomes of previous frames as initializations for the following frames via a tracking-by-detection method (Wang et al. 2015). However, this method could only extract 2D spatial appearances from static images and cannot explicitly exploit the temporal information

on consecutive frames. To circumvent this problem, video-based methods (Sánchez-Lozano et al. 2016; Tzimiropoulos 2015; Haris Khan, McDonagh, and Tzimiropoulos 2017) learn to memorize and flow the temporal consistency information across frames, which improves the robustness to the jitter problem in visual tracking. However, one major issue in these methods is that these linear mappings are too weak to exploit the complex and nonlinear relationship between image pixels and shape variations in unconstrained environments.

Deep Neural Networks for Face Alignment: These years have witnessed that deep learning (Krizhevsky, Sutskever, and Hinton 2012; Lecun, Bengio, and Hinton 2015) contributes breakthroughs in face alignment (Guo, Lu, and Zhou 2018; Trigeorgis et al. 2016; Wu et al. 2018; Kumar and Chellappa 2018; Sun, Wang, and Tang 2013; Zhang et al. 2014; Jourabloo et al. 2017; Liu et al. 2019b; 2018; Peng et al. 2016), which achieves improvements by taking advantages of nonlinear regression and end-to-end feature learning. Representative deep learning-based methods include mnemonic descent method (MDM) (Trigeorgis et al. 2016), two-stream transformer networks (Liu et al. 2018), pose conditioned dendritic convolutional neural networks (PCD-CNN) (Kumar and Chellappa 2018), Hour-Glass Network (HGN) (Newell, Yang, and Deng 2016) and a look-at-boundary (LAB) method (Wu et al. 2018). However, most methods are performed by fully supervised learning and likely give rise to the upper-bounded performance with existing labeled data (Radosavovic et al. 2018). Moreover, it is difficult to manually annotate the massive set of adjacent frames in practice. More recently, reinforcement learning (Williams 1992) is introduced to face alignment recently (Liu et al. 2019b; Guo, Lu, and Zhou 2018), which teaches a policy to select actions sequentially by obtaining a greater feedback value each time. The same sense has been discussed in (Xiong and la Torre 2015) that cascade regression can be considered as a simple and nature case of imitation learning. In contrast to prior efforts, our approach aims to reason about meaningful knowledge from amounts of unlabeled data. Meanwhile, both the spatial geometry relationship and temporal dependency information are simultaneously exploited in a unified deep learning architecture. As a result, our architecture serves as a scalable solution for processing large scale unlabeled data source.

Knowledge Distillation Models: Knowledge distillation targets on making predictions on unlabeled data and further using them to update the model, which has been adopted in various scenarios such as visual detection (Radosavovic et al. 2018; Chou, Chien, and Chen 2018; Chu et al. 2016) and model compression (Romero et al. 2015; Gupta, Hoffman, and Malik 2016). For example, Romero *et al.* (Gupta, Hoffman, and Malik 2016) adopts a shallow and wider teacher model to learn a thin and deep student model. Without training a large set of models, Radosavovic *et al.* (Radosavovic et al. 2018) presented a data distillation method to inferring predictions from unlabeled data by multiple transformations, resulting new training annotations. Our work is mostly related to a supervised-by-registration (SBR) (Dong et al. 2018) method where the detector’s predictions on unlabeled

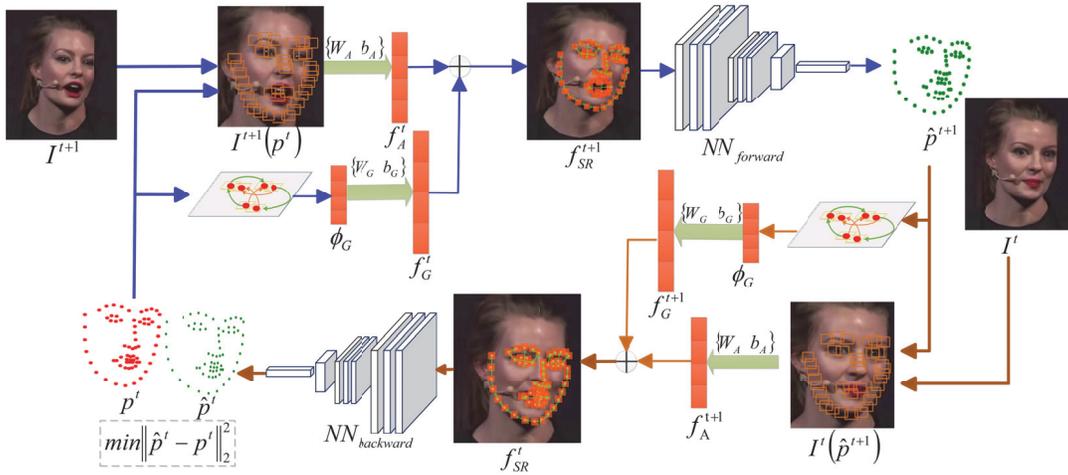


Figure 2: Network design of our STRRN. Overall, our method processes in a bi-directional tracking routing, which is visualized in the blue arrow (forward) and brown arrow (backward). For each routing, our architecture incorporates with the spatial and temporal modules. Accordingly, the spatial module is fed with the current frame and initialized landmarks (tracked results of previous frame). Then the module reasons about both the spatial appearance and geometry relationship, resulting discriminative features with preserving a dendritic-structure. Having obtained these features, we directly feed them to a neural networks (NN) structure to track landmarks in the temporal module. The network parameters of both directions are shared, thus allowing the temporal modeling of cycle-consistent relation for stabiling tracking. During learning phase, the STRRN is self-supervised based on the outcomes of the backbone detector. This figure is best viewed in color pdf file.

beled data are leveraged to train itself. However, SBR (Dong et al. 2018) ignores the spatial geometric relationship of neighbouring landmarks, thus likely making bias predictions during tracking. On contrast, in our approach, we automatically reason the meaningful spatial-temporal knowledge from large scale unlabeled videos by following the success of the relation networks framework (Hu et al. 2018). As a result, the dendritic geometry in the spatial dimension and cycle-consistency in the temporal dimension are simultaneously exploited to generate reliable new annotations.

Learning Spatial-Temporal Relational Reasoning Networks

In our approach, the main goal aims to address the omniscient supervised face alignment for large scale unlabeled videos. The basic idea is to distill spatial-temporal knowledge from numerous amounts of unlabeled frames, thus generating extra reasonable training annotations. To achieve this, we propose a spatial-temporal relational reasoning networks (STRRN) architecture which consists of the spatial and temporal modules. Specifically, our architecture typically reasons about tree-structural relationship spatially based on the appearance and geometry features. Temporally, we enforce the bi-directional cycle-consistency relation across consecutive frames for robust landmark tracking.

Fig. 2 demonstrates the network design which incorporates with the bi-directional executions of the STRRN. Suppose we have an unlabeled video set $\mathcal{D} = \{\mathbf{I}_i^t\}_{1:N}^{1:T}$ with N video clips, where $\{\mathbf{I}_i^t\}_{1:T} = \{\mathbf{I}_i^1, \mathbf{I}_i^2, \dots, \mathbf{I}_i^t, \dots, \mathbf{I}_i^T\}$ denotes each face sequence that consists of T frames. It is valuable to stress that we assume

that each face frame to be tracked is already cropped by the learned face detector (Zhang et al. 2016a). Let $\mathbf{p}^t = [(x_1, y_1), (x_2, y_2), \dots, (x_L, y_L)]^t \in \mathbb{R}^{2L \times 1}$ denote the coordinates of the facial shape vector for the t th frame, where L specifies the landmark quality within the whole face (e.g., 68 landmarks in 300-W (Sagonas et al. 2016)). For notation brevity, we assume there is only one unlabeled video with T frames and thus ignore the video index i . The main objective of face tracking in videos is to transform face sequence \mathbf{I}^t (a set of local patches cropped based on the previous outcomes) to the shape coordinates \mathbf{p}^t at time t . One alternative to model the shape coordinates is to regress a set of heat maps by feeding the entire face image like (Kumar and Chellappa 2018; Wu et al. 2018), which indicates the superiority on promoting the alignment performance. However, face tracking in videos is always sensitive to the detection. That is, the low-quality detections probably cause the drifting issue for tracking results. How to achieve the synergy between detection and tracking is investigated in (Haris Khan, McDonagh, and Tzimiropoulos 2017; Guo, Lu, and Zhou 2018), which appears beyond the range of this work. Next, we describe the network specification for the spatial and temporal modules accordingly in our proposed STRRN. Moreover, we present the detailed training procedure.

Spatial Relational Reasoning Module: It has been studied that conditioning structural information of neighbouring landmarks provides much potential to promote the performance (Belhumeur et al. 2011; Liu et al. 2017; Kumar and Chellappa 2018). Inspired by this, the motivation of our spatial module aims at disentangling a dendritic geometry relationship based on the tracked results of precious frame and

then tracking on the current face frame. To exploit the geometry relation, we first divide the whole landmarks \mathbf{p} into C components using the tracked results of previous frame, *e.g.*, facial cheek, left eye, right eye, nose, mouth, by strictly following the standard of the 300-W annotations¹. Based on the initialized landmarks, we extract a set of raw patches and divide them to C groups. Let $\mathbf{I}^t(\mathbf{p}_c^t, d)$ represent raw patches cropped from the t -th frame via the initialized shape vector \mathbf{p} , where c denotes the group identity up to C and the patch size is assigned to d (ignored for simplicity). Within each group, we compute the coupled *appearance feature* and *geometry feature* via deep neural networks. Specifically, we embed the raw patches to the immediate features in each group as (ignoring the patch index):

$$\mathbf{f}_A^t = \sigma(\mathbf{W}_A \cdot \mathbf{I}^t(\mathbf{p}^{t-1}) + \mathbf{b}_A), \quad (1)$$

where the parameters $\{\mathbf{W}_A, \mathbf{b}_A\}$ specify the spatial appearance weights and $\sigma(\cdot)$ is the nonlinear function, *e.g.*, ReLU (Krizhevsky, Sutskever, and Hinton 2012), respectively.

For the description of the geometry relation, we compute the relation feature by the comparisons of the paired (m, n) patches within each group. To make it invariant to affine transformations due to different poses, we assume the nose as the base location for robust feature learning. More specifically, we compute the geometry features (ϕ_G^m, ϕ_G^n) (simply ignoring the time-stamp index) of the paired patches by concatenating the relative distance vector $[\log(|x^m - x^n|), \log(|y^m - y^n|)]$ and the pose-invariant vector $[\log(|x^m - x^*|), \log(|y^m - y^*|)]$ (x^* and y^* specify the horizontal and vertical coordinates of the nose, respectively). To exploit the complexly nonlinear spatial relationship, we embed these geometric features into the high-dimensional feature representations as

$$\mathbf{f}_G^t(m) = \sigma(\mathbf{W}_G \cdot \text{CONCAT}[\phi_G^m, \phi_G^n] + \mathbf{b}_G), \quad (2)$$

where $\{\mathbf{W}_G, \mathbf{b}_G\}$ are the spatial geometric weights and $\text{CONCAT}(\cdot)$ denotes the concatenation operator. For each group, we perform the spatial relation features via addition,

$$\mathbf{f}_{\text{SR}}^t = \text{CONCAT}[\mathbf{f}_G^t, \mathbf{f}_A^t(1), \dots, \mathbf{f}_A^t(m)], \quad (3)$$

where \mathbf{f}_G^t specifies the holistic geometric feature concatenation across all groups. Note that $\mathbf{f}_A^t(m)$ specifies the appearance feature for m -th group where m is the landmark scalability for each group.

Lastly, we concatenate all the embedded features into one long feature vector \mathbf{f}_{SR}^t across m groups for the t th frame. Furthermore, we feed the spatial relation feature \mathbf{f}_{SR} into a tiny convolutional network to estimate landmark positions of following frames.

In summary, the spatial module learns the correlation of neighbouring landmarks by grouping them separately. By exploiting local appearance for each facial group, it guides landmarks moving towards more plausible prediction. Moreover, our module reasons with the nose-centered geometric relations, which enforces view-consistency for robust landmark tracking versus various view changes (Zavan et al. 2016).

¹<https://ibug.doc.ic.ac.uk/resources/300-W/>

Temporal Relational Reasoning Module: In the term of the temporal module, our goal is to reason about the bi-directional cycle-consistent relation between adjacent frames. In other words, the module ensures that our tracker proceeds forward along the video and it should arrive at the starting position in the backward order. Thus, this allows our model to reason with the cycle-consistency constraint of unlabeled videos for stabilizing temporal modeling (Meister, Hur, and Roth 2018). To achieve this, we learn two trackers parameterized by the designed deep neural networks, playing the video forward and backward accordingly. In order to minimize the discrepancy of tracked results in both forward and backward orders, we employ a cycle-consistency check function to evaluate the tracking reliability.

As demonstrated in Fig. 2, there are two steps for each tracker. First, we regard the tracked results \mathbf{p}^t to perform the spatial relation feature $\mathbf{f}_{\text{SR}}^{t+1}$ on \mathbf{I}^{t+1} via the proposed spatial module. Sequentially, we feed the feature $\mathbf{f}_{\text{SR}}^{t+1}$ to a forward-NN to estimate the landmarks $\hat{\mathbf{p}}^{t+1}$ on the $(t+1)$ th frame. In the similar manner on the t th frame \mathbf{I}^t , we employ the tracked landmarks \mathbf{p}^{t+1} as the initialization to perform the spatial relation feature \mathbf{f}_{SR}^t . Having obtained \mathbf{f}_{SR}^t , we pass it to the backward-NN and achieve the feedback results $\hat{\mathbf{p}}^t$. Both the forward and backward tracking executions are formulated as follows:

$$\hat{\mathbf{p}}^{t+1} = \text{NN}_{\text{backward}}(\mathbf{f}_{\text{SR}}^{t+1}), \quad \hat{\mathbf{p}}^t = \text{NN}_{\text{forward}}(\mathbf{f}_{\text{SR}}^t), \quad (4)$$

where $\text{NN}_{\text{forward}}(\cdot)$ and $\text{NN}_{\text{backward}}(\cdot)$ denote the neural networks with multiple layers, aiming to estimate the coordinates of the landmarks. Moreover, the parameters both forward and backward networks are shared. To make the tracking process efficient, we leveraged a squeezed deep architecture which is equipped with only a few layers of convolutions, max pooling, nonlinear ReLU activation and fully connection.

To further evaluate the reliability of the cyclic execution, we compute the cycle-consistency check function, which ensures that the tracked landmarks should be re-traced back again to the start. By doing this, our premise is that if the tracked results of the forward-tracker are plausible, the landmarks should return to the same positions after tracking in the backward order. In this way, our method achieves reliable tracking results with the bi-directional cycle-consistency check temporally across adjacent frames. The check loss is computed as follows:

$$\mathcal{L} = \|\hat{\mathbf{p}}^t - \mathbf{p}^t\|_2^2, \quad (5)$$

where $\|\cdot\|$ denotes the ℓ^2 norm to measure the discrepancy between the tracked and feedback results.

Training Procedure: In our approach, we investigate into the omni-supervised setting of face alignment particularly for large scale unlabeled video data. Targeting on omni-supervised learning, our method learns to trust the predictions across the massive amounts of frames and moreover includes these proxy predictions as training annotations. Since our tracker module may not always succeed, the predictions are not directly considered as the reliable training supervision signals. Hence, to make the prediction robust for the model update, we leverage a backbone detector, *e.g.*,

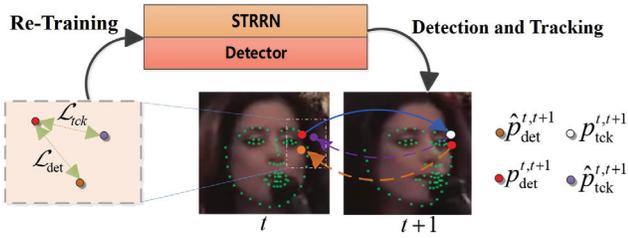


Figure 3: Training procedure of our STRRN. Basically, there are two complementary steps: 1) We pass frames onto our STRRN by distilling the spatial-temporal relation and perform the tracked results. 2) We utilize the backbone detector to evaluate the tracked results. Consequently, our tracker is self-supervised by justifying the reliability of both the cycle-consistency and the detected results. This really allows our model to trust these reasoned predictions for robust and efficient training.

MDM (Trigeorgis et al. 2016) and HGN (Newell, Yang, and Deng 2016), learned by amounts of annotated images, to evaluate the tracking reliability. With the detector, our tracker is iteratively refined by the reliable reasoned training data itself, offering much potential to surpass the performance of existing fully-supervised approaches.

As illustrated in Fig. 3, our training procedure justifies two-fold scenarios: 1) The higher reliability will be given when the tracking results $\mathbf{p}_{\text{tck}}^t$ go back again to the start $\hat{\mathbf{p}}_{\text{tck}}^{t-1}$, where the discrepancy \mathcal{L}_{tck} enforces the temporal continuity in our STRRN. 2) A negative feedback is provided when the discrepancy \mathcal{L}_{det} between backward-tracked results $\hat{\mathbf{p}}_{\text{tck}}^{t-1}$ and those $\mathbf{p}_{\text{det}}^{t-1}$ generated by the backbone image-based detector. This means the performance degrades to the pre-trained backbone detector when we undergo imprecise and unreliable tracking results.

To exploit the model update, we select the reliable outputs of both tracking and detection by minimize the complete loss functions of \mathcal{L}_{det} and \mathcal{L}_{tck} as follows:

$$\mathcal{L}_{\text{ensemble}} = \sum_{t=1}^T \mathcal{L}_{\text{tck}}^t + \lambda \sum_{t=1}^T \mathcal{L}_{\text{det}}^t, \quad (6)$$

where λ is the weight of the ensemble for both detected and tracked results. With (6), our method achieves to be self-supervised by reasoning with the reliable tracking. This principally benefits from the competitive and cooperative executions of our tracker and the backbone detector. Moreover, the complementary feedback of both parts leads to reliable new training annotations \mathcal{D}_{det} and \mathcal{D}_{tck} , where \mathcal{D}_{det} and \mathcal{D}_{tck} denote the detected and tracked for module update. *Algorithm 1* shows the detailed pseudocode for training procedure.

Experiments

To justify the effectiveness of the proposed STRRN, we represent folds of experimental results and analysis based on three downloaded large scale video datasets.

Algorithm 1: Training Procedure of Our STRRN

Input: Training set: $\mathcal{D} = \{\mathbf{I}^t\}_{t=1:T}$ (ignoring sample i for simplicity) and \mathcal{S} (samples from 300-W).
Output: Network parameters: $\{\mathbf{W}_A, \mathbf{b}_A\}$, $\{\mathbf{W}_G, \mathbf{b}_G\}$, $\text{NN}_{\text{forward}}$ and $\text{NN}_{\text{backward}}$.

- 1 //Training the backbone Detector (Both MDM and HGN were employed in the experiments);
- 2 **for** $t \in \text{video} < T$ **do**
- 3 /*Detecting on the t th frame:*/
- 4 $\mathbf{p}_{\text{det}}^t = \text{Detector}(\mathbf{I}^t)$;
- 5 /*Forward tracking on the t th frame:*/
- 6 $\mathbf{p}_{\text{tck}}^t = \text{Forward-STRRN}(\mathbf{I}^t, \mathbf{p}_{\text{det}}^{t-1})$;
- 7 /*Backward tracking based on detected landmarks:*/
- 8 $\hat{\mathbf{p}}_{\text{det}}^{t-1} = \text{Backward-STRRN}(\mathbf{I}^{t-1}, \mathbf{p}_{\text{det}}^t)$;
- 9 /*Backward tracking based on tracked landmarks:*/
- 10 $\hat{\mathbf{p}}_{\text{tck}}^{t-1} = \text{Backward-STRRN}(\mathbf{I}^{t-1}, \mathbf{p}_{\text{tck}}^t)$;
- 11 /*Criticizing the tracking reliability by the detection according to (5):*/
- 12 $\mathcal{L}_{\text{det}} = \|\hat{\mathbf{p}}_{\text{det}}^{t-1} - \mathbf{p}_{\text{det}}^{t-1}\|_2^2$, $\mathcal{L}_{\text{tck}} = \|\hat{\mathbf{p}}_{\text{tck}}^{t-1} - \mathbf{p}_{\text{det}}^{t-1}\|_2^2$;
- 13 /*Generating new training annotations by thresholding \mathcal{T} :*/
- 14 **if** $\mathcal{L}_{\text{det}}, \mathcal{L}_{\text{tck}} < \mathcal{T}$ **then**
- 15 **if** $\mathcal{L}_{\text{tck}} > \mathcal{L}_{\text{det}}$ **then**
- 16 $\mathcal{D}_{\text{tck}} \leftarrow \mathcal{D}_{\text{tck}} \cup (\mathbf{I}^t, \mathbf{p}_{\text{det}}^t)$;
- 17 **end**
- 18 $\mathcal{D}_{\text{det}} \leftarrow \mathcal{D}_{\text{det}} \cup (\mathbf{I}^t, \mathbf{p}_{\text{tck}}^t)$;
- 19 **end**
- 20 **end**
- 21 /*Data distilling and model re-training:*/
- 22 Train(Detector, $\mathcal{D}_{\text{det}} \cup \mathcal{S}$); Train(STRRN, \mathcal{D}_{tck});
- 23 **Return:** $\{\mathbf{W}_A, \mathbf{b}_A\}$, $\{\mathbf{W}_G, \mathbf{b}_G\}$.

Evaluation Datasets: *300VW* (Shen et al. 2015): The 300 Videos in the Wild (300VW) Dataset was collected specific for video-based face alignment. The dataset contains 114 videos captured in various conditions and each video has around 25-30 images every second. By following (Shen et al. 2015), we utilized 50 sequences for training and the remaining 64 sequences were used for testing. The whole testing set was divided into three categories: well-lit, mild unconstrained and challenging. In particular, the Category Three exploits many difficult cases of face sequences, which highlights the superiority of the proposed approach. It is valuable to note that we leveraged the 300-W (Sagonas et al. 2016) training set to initialize our backbone detector. In our approach, we mainly improved the tracker module by exploiting large scale unlabeled videos, which is quite different from SBR (Dong et al. 2018) principally to reinforce the landmark detector. Hence, we do not justified the superiority of our method on the image-based datasets such as 300-W (Sagonas et al. 2016).

YouTube-Face (Wolf, Hassner, and Maoz 2011) and *YouTube-Celebrities* (Kim et al. 2008): We also leveraged two large scale unlabeled video datasets including YouTube-Face (Wolf, Hassner, and Maoz 2011) and YouTube-

Table 1: Comparisons of averaged errors of our proposed STRRN with the state-of-the-arts (68-lms, in chronological ranking). *It can be seen from the results that our unsupervised STRRN outperforms both the unsupervised SBR and other fully supervised approaches. Moreover, the performance of both the backbone MDM and HGN is promoted by integrating with our STRRN.*

Method	Cate-1	Cate-2	Cate-3
CFSS (2015)	7.68	6.42	13.67
TCDCN (2015)	7.66	6.77	14.98
CCR (2016)	7.26	5.89	15.74
iCCR (2016)	6.71	4.00	12.75
MDM (2016)	9.88	8.14	9.81
HGN (2016)	6.12	5.33	7.04
TSTN (2018)	5.36	4.51	12.84
SBR (2018)	5.77	4.90	12.66
STRRN (Backbone-MDM)	5.03	4.74	6.63
STRRN (Backbone-HGN)	5.31	4.77	8.06
† Semi-STRRN ⁻⁵⁰	4.21	4.18	5.16
† Semi-STRRN ⁻²⁵	4.49	4.05	5.88
† Semi-STRRN ⁻¹⁰	4.67	4.00	5.93

† -50 , -25 and -10 denote the percentages of 300VW annotations were employed for our STRRN training, respectively.

Celebrities(Kim et al. 2008) for evaluation. Specifically, YouTube-Face consists of 3425 video clips with 1595 person identities, while YouTube-Celebrities consists of videos with 35 celebrities which undergo large variations due to different poses, illumination and partial occlusions. By following the setting in SBR (Dong et al. 2018) which is also designed suitable to unlabeled videos, we filtered these videos with low resolution and trained our STRRN unsupervisedly.

Experimental Setup: Following the evaluation protocol employed in (Liu et al. 2018), we computed the root mean squared error (RMSE) normalized by the inter-pupil distance (Xiong and la Torre 2013; Sagonas et al. 2016) and cumulative error distribution (CED) curves for comparisons in our experiments. To be specific, we performed the RMSEs of all frames within each category in 300VW and then averaged them as final performance. Moreover, we leveraged the CED curves of RMSE errors to quantitatively evaluate the performance. Besides, as discussed in MDM (Trigeorgis et al. 2016), there is no consistent way of the normalized factors for alignment, *e.g.*, inter-ocular distance (Xiong and la Torre 2013; Sagonas et al. 2016) or inter-pupil distance (Zhu et al. 2015; Cao et al. 2012). To clarify the results, we also computed the area-under-the-curve (AUC) by respectively thresholding the RMSEs at 0.05 and 0.08.

For each video to be aligned, we rescaled each face frame in the size of $[315 \times 315]$. Our spatial relational reasoning module accepts L raw patches, where L is assigned 68 according to the 300-W annotation and each patch is rescaled in size of 30×30 . For hyper-parameters in our STRRN, we empirically set the discounted factor λ to 0.4 and the thresholding \mathcal{T} to the normalized RMSE 0.02 during generating extra training annotations. More details will be made publicly in our release model and source code.

Comparisons with State-of-the-art Methods: We com-

pared our proposed STRRN with image-based methods (Xiong and la Torre 2013; Zhu et al. 2015; Zhang et al. 2016b; Trigeorgis et al. 2016) and video-based methods (Sánchez-Lozano et al. 2016; Liu et al. 2018), where these methods were fully supervised by large amounts of annotations. For fair comparisons, we re-produced the results with their released codes or cropped the reported results from their original papers. Note that only our STRRN and SBR (Dong et al. 2018) only leveraged 300-W images for model training.

To investigate the effectiveness of our proposed approach, we compared our approach with the state-of-the-art face alignment methods, which were designed for both static images and tracking in videos via a fully-supervised manner. For fair comparisons, we first discarded all annotations for model training by following the unsupervised setting. Hence, we did not compare our STRRN with the newly-proposed supervised method dubbed FHR (Tai et al. 2019). To further highlight the advantage of our approach, we trained our model by only a subset of annotations, which is termed as *Semi-STRRN* in Table 1, which tabulates the comparisons of our method with the state-of-the-arts in both unsupervised and semi-supervised on the 300VW datasets. From these results, we see that our proposed STRRN significantly outperforms other face alignment methods by a large margin especially on the challenging 300VW Category Three. This is mainly because our cycle-consistent execution exploits more interpretable cues to learn discriminative spatial-temporal features for robust face tracking. Moreover, seeing from Table 1 which tabulates the comparisons of STRRN versus semi-STRRN, we found that even with weakly-labeled supervision signals, we still perform very compelling results on the challenging cases including large poses, diverse expressions and severe occlusions. This also demonstrates the effectiveness of the proposed spatial-temporal relation modeling of the bidirectional orders, where these learned temporal relations are helpful to reinforce the tracking stabilization.

Lastly, we investigated the effectiveness of the temporal relational reasoning module in our STRRN. We used SBR (Dong et al. 2018) as the baseline method. Both our method and SBR (Dong et al. 2018) were learned on 300VW without any supervision, and tested on the selected 517-th face sequence of the most challenging category. Fig. 4(a) shows the sequential errors of our STRRN and SBR (Dong et al. 2018) across a series of time stamps. From these curves, we achieve a clear and stable results by a large margin compared with SBR (Dong et al. 2018), which shows the stabilization of ours versus various temporal motions. This is mainly because the testing sequences undergo continuous occlusion over time steps. Clearly, our STRRN achieves to reason about the robust spatial-temporal feature representation across frames for stabilizing landmark tracking, while SBR (Dong et al. 2018) always degrades especially on the facial cheek. The similar sense is visualized on the qualitative comparisons in Fig. 4(b), where our STRRN achieves the coherency cues in contrast to SBR (Dong et al. 2018).

Investigation on Unlabeled Video Datasets: We further compared our proposed STRRN on different datasets in-

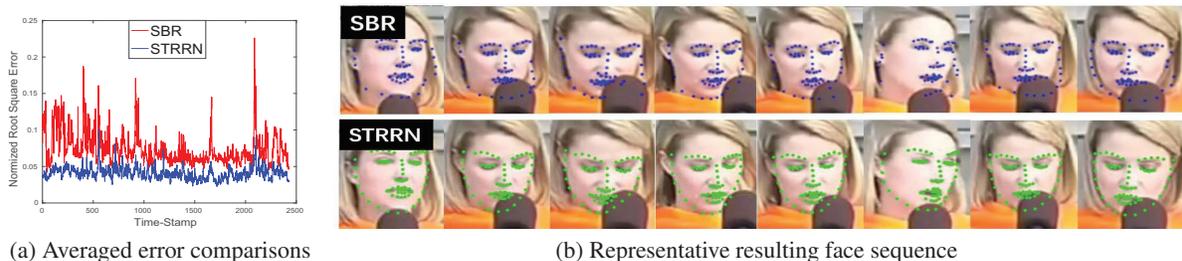


Figure 4: Comparisons of our STRRN and SBR on the challenging 517-th sequence of the 300VW Category Three. Both-fold results demonstrate that our model achieves low errors especially on occluded facial cheek thanks to the spatial constraint.

Table 2: Comparisons of STRRN on different training datasets (AUC (%) thresholding RMSEs at 0.05 and 0.08).

Video Dataset	AUC@0.05	AUC@0.08
300VW	26.67	78.07
300VW+YouTubeCelebrities	28.31	77.24
300VW+YouTubeFace	36.33	80.43
All Combined	39.48	85.06

cluding 300VW (Shen et al. 2015), YouTube-Face (Wolf, Hassner, and Maoz 2011) and YouTube-Celebrities (Kim et al. 2008). By carefully following the employed setting in SBR (Dong et al. 2018), we regarded our STRRN trained on the 300VW dataset without any annotations as the baseline method, where 68 landmarks were employed for evaluation. Even without any every-frame annotations, our method trained with either of large scale unlabeled video data outperforms that trained with only 300VW. Moreover, Table 2 tabulates the AUC comparisons at the thresholding RMSEs at 0.05 and 0.08 on different datasets. From Table 2, the outperformed results prove that we succeed in reasoning the meaningful spatial-temporal knowledge from unlabeled videos, thus making reliable extra annotations for model update.

Ablation Study: To further investigate the ablation experiments of the proposed components in our STRRN, we computed the comparisons of CED curves on various training strategies with respect to our STRRN. The partition method for each facial component is consistent with that was defined in our spatial module. For the baseline method, we selected the detector HGN (Newell, Yang, and Deng 2016) that was fine-tuned by our STRRN. Table 3 demonstrates the CED comparisons and two-fold conclusions can be made from Table 1) $\lambda=0.4$ achieves higher performance than training by \mathcal{L}_{tck} or \mathcal{L}_{det} independently, and 2) \mathcal{L}_{tck} is relatively important than \mathcal{L}_{det} for robust spatial-temporal relation reasoning. Moreover, Table 3 also tabulates the comparisons of our STRRN versus independent spatial and temporal modules, which indicates the contributions of the complementary information of both modules.

Computational Efficiency: Our STRRN requires the computation complexity of $\mathcal{O}(N_R^2)$ for the geometry construction, where N_R denotes the landmark scalability within each facial component. According to the 300-W annotation, N_R maximally reaches approximate $20 \times (20 + 1)$ for

Table 3: Ablation experiments of our STRRN, where the 300VW dataset was employed for evaluation.

Method	Cate-1	Cate-2	Cate-3
STRRN (w/o \mathcal{L}_{tck})	7.92	5.78	8.78
STRRN (w/o \mathcal{L}_{det} , $\lambda=0$)	7.92	7.08	8.26
STRRN (\mathcal{L}_{tck} and \mathcal{L}_{det} , $\lambda=1.6$)	7.02	6.91	8.21
STRRN (w/o Temporal Module)	8.76	7.98	8.93
STRRN (w/o Spatial Module)	8.45	7.84	8.76
STRRN (S & T Module, $\lambda=0.4$)	5.03	4.74	6.63

the mouth region. For the efficient training, our architecture was implemented under the parallel-computing deep learning Tensorflow. The whole training procedure processes at about 60ms each frame with a GPU of single NVIDIA GTX 1080 Ti graphic computation card (11G memory). Excluding the time of the face detection part, our model runs at 30 frames per second on one CPU with the Intel(R) Core(TM) i5-6500 CPU@3.20GHz and requires around 2G memory usage for runtime data loading.

Conclusion

In this work, we have investigated into the omni-supervised face alignment particular for large scale unlabeled video data. We have proposed a spatial-temporal relational reasoning networks method to reason about the spatial-temporal relation from the unlabeled videos. Extensive experimental results have showed that our approach have surpassed the performance of most state-of-the-art methods. How to apply reinforcement learning (Büchler, Brattoli, and Ommer 2018; Liu et al. 2019b; Guo, Lu, and Zhou 2018) to address the 3D face alignment (Jourabloo et al. 2017) with omni-supervision is desirable in our future works.

Acknowledgment

This work was supported in part by the National Science Foundation of China under Grant 61806104, in part by the Natural Science Foundation of Ningxia under Grant 2018AAC03035, in part by the Scientific Research Projects of Colleges and Universities of Ningxia under Grant NGY2018050, and in part by the Third Batch of Ningxia Youth Talents Supporting Program under Grant TJGC2018028

References

- Belhumeur, P. N.; Jacobs, D. W.; Kriegman, D. J.; and Kumar, N. 2011. Localizing parts of faces using a consensus of exemplars. In *CVPR*, 545–552.
- Büchler, U.; Brattoli, B.; and Ommer, B. 2018. Improving spatiotemporal self-supervision by deep reinforcement learning. In *ECCV*, 797–814.
- Cao, X.; Wei, Y.; Wen, F.; and Sun, J. 2012. Face alignment by explicit shape regression. In *CVPR*, 2887–2894.
- Chou, C.; Chien, J.; and Chen, H. 2018. Self adversarial training for human pose estimation. In *APSIPA*, 17–30.
- Chu, X.; Ouyang, W.; Li, H.; and Wang, X. 2016. Structured feature learning for pose estimation. In *CVPR*, 4715–4723.
- Dong, X.; Yu, S.; Weng, X.; Wei, S.; Yang, Y.; and Sheikh, Y. 2018. Supervision-by-registration: An unsupervised approach to improve the precision of facial landmark detectors. In *CVPR*, 360–368.
- Guo, M.; Lu, J.; and Zhou, J. 2018. Dual-agent deep reinforcement learning for deformable face tracking. In *ECCV*, 783–799.
- Gupta, S.; Hoffman, J.; and Malik, J. 2016. Cross modal distillation for supervision transfer. In *CVPR*, 2827–2836.
- Haris Khan, M.; McDonagh, J.; and Tzimiropoulos, G. 2017. Synergy between face alignment and tracking via discriminative global consensus optimization. In *ICCV*.
- Hu, H.; Gu, J.; Zhang, Z.; Dai, J.; and Wei, Y. 2018. Relation networks for object detection. In *CVPR*, 3588–3597.
- Hu, J.; Lu, J.; and Tan, Y. 2014. Discriminative deep metric learning for face verification in the wild. In *CVPR*, 1875–1882.
- Jourabloo, A.; Ye, M.; Liu, X.; and Ren, L. 2017. Pose-invariant face alignment with a single cnn. In *ICCV*.
- Kim, M.; Kumar, S.; Pavlovic, V.; and Rowley, H. A. 2008. Face tracking and recognition with visual constraints in real-world videos. In *CVPR*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*, 1106–1114.
- Kumar, A., and Chellappa, R. 2018. Disentangling 3d pose in a dendritic cnn for unconstrained 2d face alignment. *CVPR*.
- Lecun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature* 521(7553):436–444.
- Liu, F.; Zeng, D.; Zhao, Q.; and Liu, X. 2016. Joint face alignment and 3d face reconstruction. In *ECCV*, 545–560.
- Liu, H.; Lu, J.; Feng, J.; and Zhou, J. 2017. Learning deep sharable and structural detectors for face alignment. *TIP* 26(4):1666–1678.
- Liu, H.; Lu, J.; Feng, J.; and Zhou, J. 2018. Two-stream transformer networks for video-based face alignment. *TPAMI* 40(11):2546–2554.
- Liu, H.; Lu, J.; Feng, J.; and Zhou, J. 2019a. Ordinal deep learning for facial age estimation. *TCSVT* 29(2):486–501.
- Liu, H.; Lu, J.; Guo, M.; Wu, S.; and Zhou, J. 2019b. Learning reasoning-decision networks for robust face alignment. *TPAMI*, in press.
- Meister, S.; Hur, J.; and Roth, S. 2018. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *AAAI*, 7251–7259.
- Newell, A.; Yang, K.; and Deng, J. 2016. Stacked hourglass networks for human pose estimation. In *ECCV*, 483–499.
- Peng, X.; Feris, R. S.; Wang, X.; and Metaxas, D. N. 2016. A recurrent encoder-decoder network for sequential face alignment. In *ECCV*, 38–56.
- Radosavovic, I.; Dollár, P.; Girshick, R. B.; Gkioxari, G.; and He, K. 2018. Data distillation: Towards omni-supervised learning. In *CVPR*, 4119–4128.
- Ren, S.; Cao, X.; Wei, Y.; and Sun, J. 2014. Face alignment at 3000 FPS via regressing local binary features. In *CVPR*, 1685–1692.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2015. Fitnets: Hints for thin deep nets.
- Roth, J.; Tong, Y.; and Liu, X. 2015. Unconstrained 3d face reconstruction. In *CVPR*, 2606–2615.
- Sagonas, C.; Antonakos, E.; Tzimiropoulos, G.; Zafeiriou, S.; and Pantic, M. 2016. 300 faces in-the-wild challenge: database and results. *IVC* 47:3–18.
- Sánchez-Lozano, E.; Martínez, B.; Tzimiropoulos, G.; and Valstar, M. F. 2016. Cascaded continuous regression for real-time incremental face tracking. In *ECCV*, 645–661.
- Shen, J.; Zafeiriou, S.; Chrysos, G. G.; Kossaiji, J.; Tzimiropoulos, G.; and Pantic, M. 2015. The first facial landmark tracking in-the-wild challenge: Benchmark and results. In *ICCVW*, 1003–1011.
- Sun, Y.; Wang, X.; and Tang, X. 2013. Deep convolutional network cascade for facial point detection. In *CVPR*, 3476–3483.
- Tai, Y.; Liang, Y.; Liu, X.; Duan, L.; Li, J.; Wang, C.; Huang, F.; and Chen, Y. 2019. Towards highly accurate and stable face alignment for high-resolution videos. In *AAAI*.
- Trigeorgis, G.; Snape, P.; Nicolaou, M. A.; Antonakos, E.; and Zafeiriou, S. 2016. Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In *CVPR*, 4177–4187.
- Tzimiropoulos, G. 2015. Project-out cascaded regression with an application to face alignment. In *CVPR*.
- Wang, X.; Yang, M.; Zhu, S.; and Lin, Y. 2015. Regionlets for generic object detection. *TPAMI* 37(10):2071–2084.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8:229–256.
- Wolf, L.; Hassner, T.; and Maoz, I. 2011. Face recognition in unconstrained videos with matched background similarity. In *CVPR*, 529–534.
- Wu, W.; Qian, C.; Yang, S.; Wang, Q.; Cai, Y.; and Zhou, Q. 2018. Look at boundary: A boundary-aware face alignment algorithm. In *CVPR*.
- Xiong, X., and la Torre, F. D. 2013. Supervised descent method and its applications to face alignment. In *CVPR*, 532–539.
- Xiong, X., and la Torre, F. D. 2015. Global supervised descent method. In *CVPR*, 2664–2673.
- Zavan, F. H. B.; Nascimento, A. C. P.; e Silva, L. P.; Bellon, O. R. P.; and Silva, L. 2016. 3d face alignment in the wild: A landmark-free, nose-based approach. In *ECCV Workshop*, 581–589.
- Zhang, J.; Shan, S.; Kan, M.; and Chen, X. 2014. Coarse-to-fine auto-encoder networks (CFAN) for real-time face alignment. In *ECCV*, 1–16.
- Zhang, K.; Zhang, Z.; Li, Z.; and Qiao, Y. 2016a. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Process. Lett.* 23(10):1499–1503.
- Zhang, Z.; Luo, P.; Loy, C. C.; and Tang, X. 2016b. Learning deep representation for face alignment with auxiliary attributes. *TPAMI* 38(5):918–930.
- Zhu, S.; Li, C.; Loy, C. C.; and Tang, X. 2015. Face alignment by coarse-to-fine shape searching. In *CVPR*, 4998–5006.