

Recurrence without Recurrence: Stable Video Landmark Detection with Deep Equilibrium Models

Paul Micaelli
University of Edinburgh
paul.micaelli@ed.ac.uk

Arash Vahdat
NVIDIA
avahdat@nvidia.com

Hongxu Yin
NVIDIA
dannyy@nvidia.com

Jan Kautz
NVIDIA
jkautz@nvidia.com

Pavlo Molchanov
NVIDIA
pmolchanov@nvidia.com

Abstract

Cascaded computation, whereby predictions are recurrently refined over several stages, has been a persistent theme throughout the development of landmark detection models. In this work, we show that the recently proposed Deep Equilibrium Model (DEQ) can be naturally adapted to this form of computation. Our Landmark DEQ (LDEQ) achieves state-of-the-art performance on the challenging WFLW facial landmark dataset, reaching 3.92 NME with fewer parameters and a training memory cost of $\mathcal{O}(1)$ in the number of recurrent modules. Furthermore, we show that DEQs are particularly suited for landmark detection in videos. In this setting, it is typical to train on still images due to the lack of labelled videos. This can lead to a “flickering” effect at inference time on video, whereby a model can rapidly oscillate between different plausible solutions across consecutive frames. By rephrasing DEQs as a constrained optimization, we emulate recurrence at inference time, despite not having access to temporal data at training time. This Recurrence without Recurrence (RwR) paradigm helps in reducing landmark flicker, which we demonstrate by introducing a new metric, normalized mean flicker (NMF), and contributing a new facial landmark video dataset (WFLW-V) targeting landmark uncertainty. On the WFLW-V hard subset made up of 500 videos, our LDEQ with RwR improves the NME and NMF by 10 and 13% respectively, compared to the strongest previously published model using a hand-tuned conventional filter.

1. Introduction

The field of facial landmark detection has been fueled by important applications such as face recognition [49, 71], facial expression recognition [33, 37, 39, 43, 79], and face

alignment [55, 87, 89]. Early approaches to landmark detection relied on a statistical model of the global face appearance and shape [19, 20, 25], but this was then superseded by deep learning regression models [21, 26, 34, 40, 44, 46, 54, 63, 68, 69, 73, 75, 77]. Both traditional and modern approaches have relied upon cascaded computation, an approach which starts with an initial guess of the landmarks and iteratively produces corrected landmarks which match the input face more finely. These iterations typically increase the training memory cost linearly, and do not have an obvious stopping criteria. To solve these issues, we adapt the recently proposed Deep Equilibrium Model [9–11] to the setting of landmark detection. Our Landmark DEQ (LDEQ) achieves state-of-the-art performance on the WFLW dataset, while enjoying a natural stopping criteria and a memory cost that is constant in the number of cascaded iterations.

Furthermore, we explore the benefits of DEQs in landmark detection from facial videos. Since obtaining landmark annotation for videos is notoriously expensive, models are virtually always trained on still images and applied frame-wise on videos at inference time. When a sequence of frames have ambiguous landmarks (*e.g.*, occluded faces or motion blur), this leads to flickering landmarks, which rapidly oscillate between different possible configurations across consecutive frames. This poor temporal coherence is particularly problematic in applications where high precision is required, which is typically the case for facial landmarks. These applications include face transforms [1, 2], face reenactment [84], video emotion recognition [33, 37, 39, 79], movie dubbing [28] or tiredness monitoring [35]. We propose to modify the DEQ objective at inference time to include a new recurrent loss term that encourages temporal coherence. We measure this improvement on our new WFLW-Video dataset (WFLW-V), demonstrating superiority over traditional filters, which typically reduce flickering at the cost of reducing landmark accuracy.

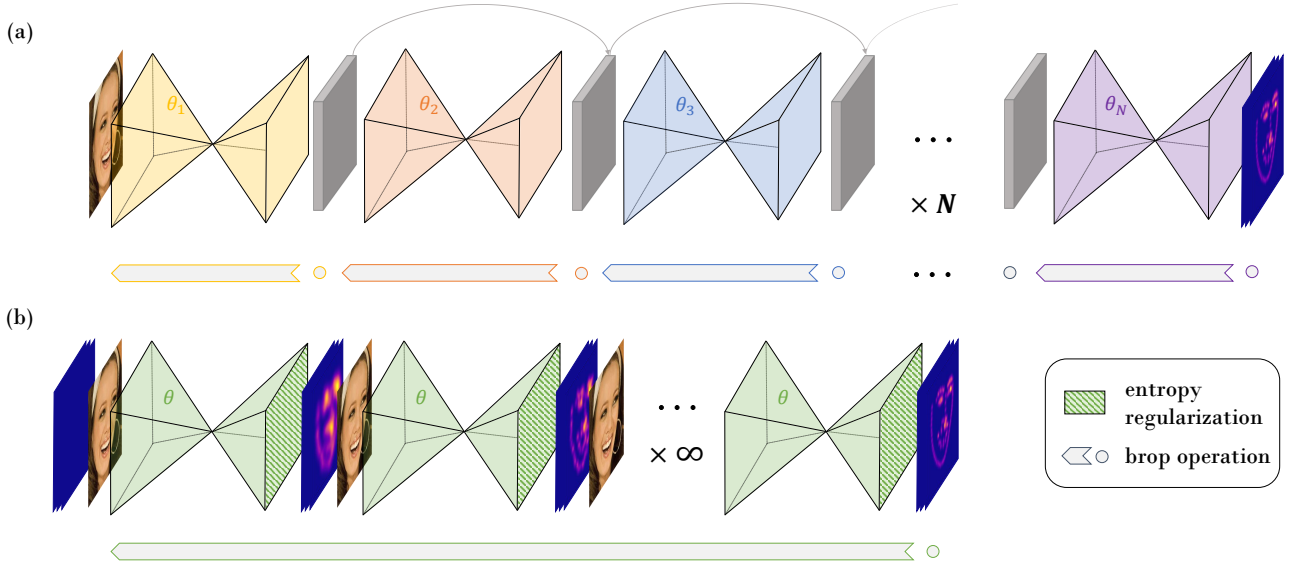


Figure 1. (a) Common Stacked-Hourglass architecture [51], whereby each cascaded stage increases the memory cost and the number of backpropagation operations. (b) Our LDEQ model, which adapts an equilibrium model [9] to the landmark detection setting, enjoying a constant memory cost with the number of refining stages. At each stage, we compute landmark probability heatmaps, and encourage convergence to an equilibrium by lowering their entropy. We release our code here: https://github.com/polo5/LDEQ_RwR.

2. Related work

Deep Equilibrium Models [9] are part of the family of implicit models, which learn implicit functions such as the solution to an ODE [17, 24], or the solution to an optimization problem [3, 22, 72]. These functions are called implicit in the sense that the output cannot be written as a function of the input explicitly. In the case of DEQs, the output is the root of a function, and the model learned is agnostic to the root solver used. Early DEQ models were too slow to be competitive, and much work since has focused on better architecture design [10] and faster root solving [11, 27, 29]. Since the vanilla formulation of DEQs does not guarantee convergence or uniqueness of an equilibrium, another branch of research has focused on providing convergence guarantees [52, 56, 74], which usually comes at the cost of a performance drop. Most similar to our work is the recent use of DEQs for videos in the context of optical flow estimation [8], where slightly better performance than LSTM-based models was observed.

A common theme throughout the development of landmark detection models has been the idea of cascaded compute, which has repeatedly enjoyed a better performance compared to single stage models [71]. This is true in traditional models like Cascaded Pose Regression (CPR) [7, 15, 23, 65, 78], but also in most modern landmark detection models [34, 40, 42, 69, 73, 75, 80], which usually rely on the Stacked-Hourglass backbone [51], or an RNN structure [41, 66]. In contrast to these methods, our DEQ-based model can directly solve for an infinite number of refine-

ment stages at a constant memory cost and without gradient degradation. This is done by relying on the implicit function theorem, as opposed to tracking each forward operation in autograd. Furthermore, our model naturally supports adaptive compute, in the sense that the number of cascaded stages will be determined by how hard finding an equilibrium is for a specific input, while the number of stages in, say, the Stacked-Hourglass backbone, must be constant at all times.

Our recurrence without recurrence approach is most closely related to test time adaption methods. These have been most commonly developed in the context of domain adaptation [64, 70], reinforcement learning [32], meta-learning [86], generative models [36, 50], pose estimation [45] or super resolution [6]. Typically, the methods above finetune a trained model at test time using a form of self-supervision. In contrast, our model doesn't need to be finetuned at test time: since DEQs solve for an objective function in the forward pass, this objective is simply modified at test time.

3. DEQs for landmark detection

Consider learning a landmark detection model F parameterized by θ , which maps an input image x to landmarks z . Instead of directly having z as 2D landmarks, it is common for z to represent L heatmaps of size $D \times D$, where D is a hyperparameter (usually 64) and L is the number of landmarks to learn. While typical machine learning models can explicitly write down the function $z = F(x; \theta)$, in

the DEQ approach this function is expressed implicitly by requiring its output to be a fixed point of another function $f(\mathbf{z}, \mathbf{x}; \boldsymbol{\theta})$:

$$F : \mathbf{x} \rightarrow \mathbf{z}^* \quad \text{s.t.} \quad \mathbf{z}^* = f(\mathbf{z}^*, \mathbf{x}; \boldsymbol{\theta}) \quad (1)$$

where \mathbf{z}^* denotes the fixed point, or equivalently the root of $g(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta}) = f(\mathbf{z}, \mathbf{x}; \boldsymbol{\theta}) - \mathbf{z}$. The function f must have inputs and outputs of similar shape, but beyond this restriction there is still limited understanding of its desired properties in machine learning. For simplicity, we build f from the ubiquitous hourglass module h (similar to a Unet):

$$f(\mathbf{z}, \mathbf{x}; \boldsymbol{\theta}) = \sigma(h([\mathbf{x}, \mathbf{z}]; \boldsymbol{\theta})) \quad (2)$$

where h inputs the concatenation of \mathbf{x} and \mathbf{z} , and σ is a normalization function. For clarity, we omit from our notation that image x is downsampled with a few convolutions to match the shape of \mathbf{z} .

To evaluate f in the forward pass, we must solve for its fixed point \mathbf{z}^* . When f is a contraction mapping, $f \circ f \circ f \circ \dots \circ f \circ \mathbf{z}^{(0)}$ converges to a unique \mathbf{z}^* for any initial heatmap $\mathbf{z}^{(0)}$. In practice, it is neither tractable or helpful to directly take an infinite number of fixed point iteration steps. Instead, it is common to achieve the same result by leveraging quasi Newtonian solvers like Broyden’s method [13] or Anderson acceleration [4], which find \mathbf{z}^* in fewer iterations. Similarly to the original DEQ model, we use $\mathbf{z}^{(0)} = \mathbf{0}$ when training our LDEQ on still images.

Guaranteeing the existence of a unique fixed point by enforcing contraction restrictions on f is cumbersome, and better performance can often be obtained by relying on regularization heuristics that are conducive to convergence, such as weight normalization and variational dropout [10]. In our landmark model, we did not find these tricks helpful, and instead used a simple normalization of the heatmaps to $[0, 1]$ at each refining stage:

$$\sigma(\mathbf{z}) = \exp\left(\frac{\mathbf{z} - \max(\mathbf{z})}{T}\right) \quad (3)$$

where T is a temperature hyperparameter. This layer also acts as an entropy regularizer, since it induces low-entropy (“peaked”) heatmaps, which we found to improve convergence of root solvers.

We contrast our model in Fig. 1 to the popular Stacked-Hourglass backbone [51]. Contrary to this model, our DEQ-based model uses a single hourglass module which updates \mathbf{z} until an equilibrium is found. The last predicted heatmaps, \mathbf{z}^* , are converted into 2D landmark points $\hat{\mathbf{p}} = \Phi(\mathbf{z}^*)$ using the softargmax function $\Phi(\mathbf{z})$ proposed in [48]. These points are trained to match the ground truth \mathbf{p} , and so the DEQ landmark training problem can be seen as a constrained optimization:

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg \min_{\boldsymbol{\theta}} \mathcal{L}_{\text{MSE}}(\Phi(\mathbf{z}^*), \mathbf{p}) \\ \text{s.t.} \quad \mathbf{z}^* &= f(\mathbf{z}^*, \mathbf{x}; \boldsymbol{\theta}) \end{aligned} \quad (4)$$

To differentiate our loss function \mathcal{L}_{MSE} through this root solving process, the implicit function theorem is used [9] to derive

$$\frac{\partial \mathcal{L}_{\text{MSE}}}{\partial \boldsymbol{\theta}} = -\frac{\partial \mathcal{L}_{\text{MSE}}}{\partial \mathbf{z}^*} \mathcal{J}_{g\mathbf{z}^*}^{-1} \frac{\partial f(\mathbf{z}^*, \mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \quad (5)$$

where the first two terms on the RHS can be expressed as the solution to a fixed point problem as well. Solving for this root in the backward pass means that we do not need to compute or store the expensive inverse Jacobian term $\mathcal{J}_{g\mathbf{z}^*}^{-1}$ directly [9, 90]. Importantly, this backward-pass computation only depends on \mathbf{z}^* and doesn’t depend on the operations done in the forward pass to reach an equilibrium. This means that these operations do not need to be tracked by autograd, and therefore that training requires a memory cost of $\mathcal{O}(1)$, despite differentiating through a potentially infinite recurrence.

4. Recurrence without recurrence

Low temporal coherence (i.e. a large amount of flicker) is illustrated in Fig. 2. This can be a nuisance for many applications that require consistently precise landmarks for video. In this section, we describe how our LDEQ model can address this challenge by enabling recurrence at test time without recurrence at training time (RwR).

Recall that in the DEQ formulation of Sec. 3, there is no guarantee that a unique fixed point solution exists. This can be a limitation for some applications, and DEQ variants have been proposed to allow provably unique solutions at the cost of additional model complexity [53, 74]. In this work, we instead propose a new paradigm: we leverage the potentially large solution space of DEQs after training to allow for some additional objective at inference time. This new objective is used to disambiguate which fixed point of $f(\mathbf{z}, \mathbf{x}; \boldsymbol{\theta}^*)$ is found, in light of extra information present at test time. We demonstrate this approach for the specific application of training a landmark model on face images and evaluating it on videos. In this case, DEQs allow us to include a recurrent loss term at inference time, which isn’t achievable with conventional architectures.

Let $f(\mathbf{z}, \mathbf{x}; \boldsymbol{\theta}^*)$ be our DEQ model trained as per the formulation in Eq. (4). We would now like to do inference on a video of N frames $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$. Consider that a given frame \mathbf{x}_n has a corresponding set of fixed points $\mathcal{Z}_n^* = \{\mathbf{z} \text{ s.t. } f(\mathbf{z}, \mathbf{x}_n; \boldsymbol{\theta}^*) = \mathbf{z}\}$, representing plausible landmark heatmaps. If we select some $\mathbf{z}_n^* \in \mathcal{Z}_n^*$ at random for each frame n , the corresponding heatmaps $\mathbf{z}_1^*, \mathbf{z}_2^*, \dots, \mathbf{z}_N^*$ often exhibit some flickering artefacts, whereby landmarks rapidly change across contiguous frames (see Fig. 2). We propose to address this issue by choosing the fixed point at frame n that is closest to the fixed point at frame $n - 1$. This can be expressed by the following constrained optimization:

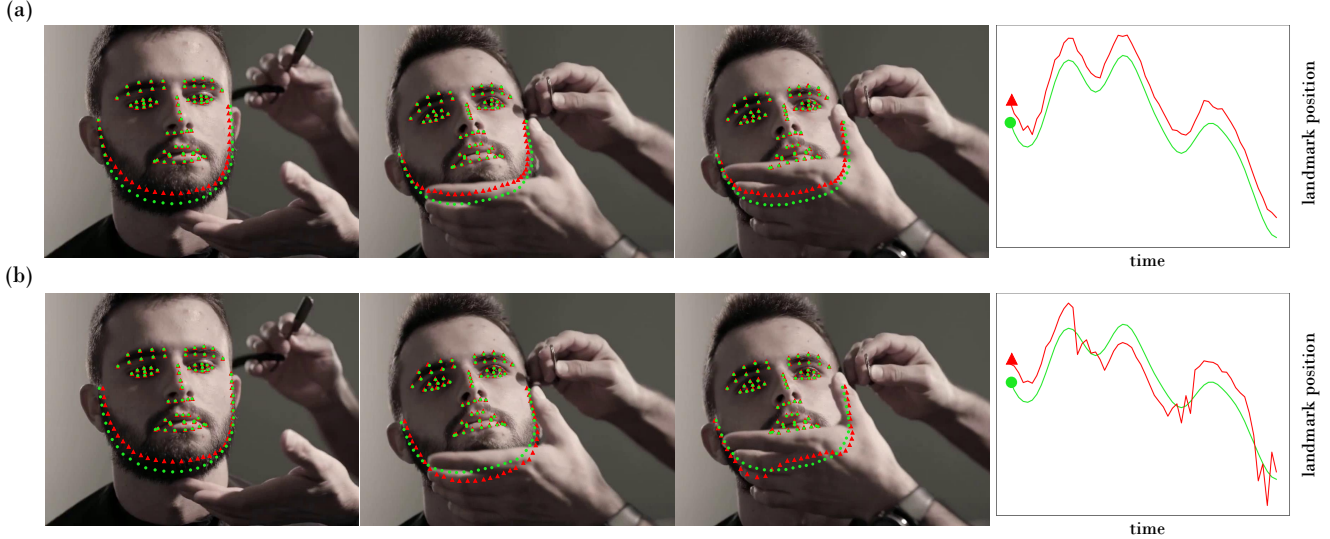


Figure 2. Model predictions (\blacktriangle) and ground truth landmarks (\bullet) for (a) a model with high temporal coherence and (b) a model with the same accuracy but exhibiting a worse temporal coherence, due to ambiguity at the chin. This causes flickering around the ground truth, as illustrated in the landmark trajectory (right). This flickering is common when video frames are evaluated individually, as opposed to recurrently.

$$z_n^* = \arg \min_z \|z - z_{n-1}^*\|_2^2 \quad (6a)$$

$$\text{s.t. } f(z, x_n; \theta^*) = z \quad (6b)$$

The problem above is equivalent to solving for the saddle point of a Lagrangian as follows:

$$\min_z \max_\lambda \|z - z_{n-1}^*\|_2^2 + \lambda^T (f(z, x_n; \theta^*) - z) \quad (7)$$

where λ are Lagrange multipliers. Effectively, we are using the set of fixed points Z_n^* in Eq. (6b) as the trust region for the objective in Eq. (6a). In practice, adversarial optimization is notoriously unstable, as is typically observed in the context of GANs [5, 30, 62]. Furthermore, this objective breaks down if $Z_n^* = \emptyset$ for any x_n . We can remedy both of these problems by relaxing the inference time optimization problem to:

$$\min_z \|f(z, x_n; \theta^*) - z\|_2^2 + \frac{\alpha}{2} \|z - z_{n-1}^*\|_2^2 \quad (8)$$

where α is a hyperparameter that trades off fixed-point solver error vs. the shift in heatmaps across two consecutive frames. This objective can be more readily tackled with Newtonian optimizers like L-BFGS [47]. When doing so, our DEQ at inference time can be described as a form of OptNet [3], albeit without any of the practical limitations (e.g., quadratic programs) related to making gradient calculations cheap.

Converting our root solving problem into an optimization problem during the forward pass of each frame can significantly increase inference time. Thankfully, the objective in Eq. (8) can also be solved by using root solvers. First, note that it is equivalent to finding the MAP estimate given a log likelihood and prior:

$$\log p(x_n | z; \theta^*) \propto -\|f(z, x_n; \theta^*) - z\|_2^2 \quad (9a)$$

$$p(z) = \mathcal{N}(z; z_{n-1}, \alpha^{-1} \mathbb{1}) \quad (9b)$$

It has been demonstrated in various settings that this prior, when centered on the initialization to a search algorithm, can be implemented by early stopping this algorithm [12, 31, 58, 61]. As such, we can approximate the solution to Eq. (6a-6b) by simply initializing the root solver with $z_n^{(0)} = z_{n-1}^*$ (“reuse”) and imposing a hard limit on the number of steps that it can take (“early stopping”). We call this approach Recurrence without Recurrence (RwR), and illustrate it in Fig. 3.

5. Video landmark coherence

In this section we describe the metric (NMF) and the dataset (WFLW-V) that we contribute to measure the amount of temporal coherence in landmark videos. These are later used to benchmark the performance of our RwR paradigm against alternatives in Sec. 6.2.

5.1. NMF: a metric to track temporal coherence

The performance of a landmark detection model is typically measured with a single metric called the Normal-

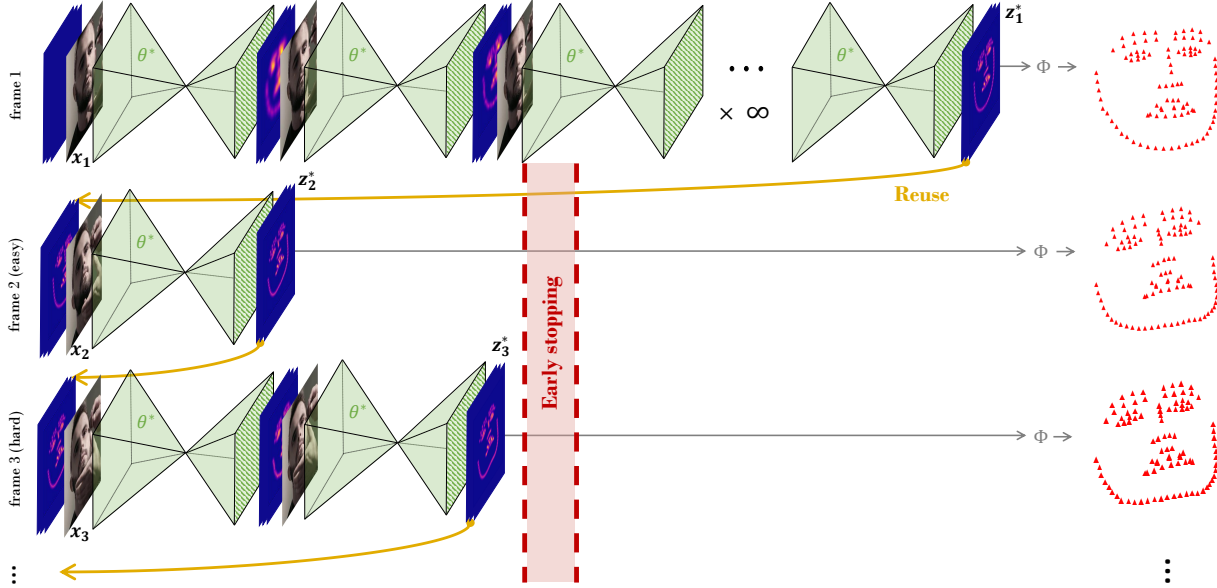


Figure 3. Recurrence without recurrence (RwR) on video data, from an LDEQ that was trained on still images. We use the initialization $z_1^{(0)} = \mathbf{0}$ for the first frame, and then reuse $z_n^{(0)} = z_{n-1}^*$ for $n > 1$. Combined with early stopping, this is equivalent to regularizing the fixed point z_n so that it is more temporally coherent with all its predecessors.

ized Mean Error (NME). Consider a video sequence of N frames, each containing L ground truth landmarks. A single landmark point is a 2D vector denoted $\mathbf{p}_{n,l} \in \mathcal{R}^2$, where n and l are the frame and landmark index respectively. Let $\mathbf{r}_{n,l} = \mathbf{p}_{n,l} - \hat{\mathbf{p}}_{n,l}$ be the residual vector between ground truth landmarks and predicted landmarks $\hat{\mathbf{p}}_{n,l}$. The NME simply averages the ℓ_2 norm of this residual across all landmarks and all frames:

$$\text{NME}_n = \frac{1}{L} \sum_{l=1}^L \frac{\|\mathbf{r}_{n,l}\|}{d_0} \quad (10a)$$

$$\text{NME} = \frac{1}{N} \sum_{n=1}^N \text{NME}_n \quad (10b)$$

Here d_0 is usually the inter-ocular distance, and aims to make the NME better correlate with the human perception of landmark error. We argue that this metric alone is insufficient to measure the performance of a landmark detector in videos. In Fig. 2 we show two models of equal NME but different coherence in time, with one model exhibiting *flickering* between plausible hypotheses when uncertain. This flickering is a nuisance for many applications, and yet is not captured by the NME. This is in contrast to random noise (*jitter*) which is unstructured and already reflected in the NME metric.

To measure temporal coherence, we propose a new metric called the Normalized Mean Flicker (NMF). We design this metric to be orthogonal to the NME, by making it agnostic to the magnitude of $\mathbf{r}_{n,l}$, and only focusing on the

change of $\mathbf{r}_{n,l}$ across consecutive frames:

$$\text{NMF}_n = \sqrt{\frac{1}{L} \sum_{l=1}^L \frac{\|\mathbf{r}_{n,l} - \mathbf{r}_{n-1,l}\|^2}{d_1^2}} \quad (11a)$$

$$\text{NMF} = \sqrt{\frac{1}{N} \sum_{n=2}^N \text{NMF}_n^2} \quad (11b)$$

We replace the means in the NME with a root mean square to better represent the human perception of flicker. Indeed, this penalizes a short sudden changes in $\mathbf{r}_{n,l}$ compared to the same change smoothed out in time and space. The value d_1^2 is chosen to be the face area. This prevents a long term issue with the NME score, namely the fact that large poses can have an artificially large NME due to having a small d_0 .

5.2. A new landmark video dataset: WFLW-V

Due to the tedious nature of producing landmark annotations for video data, there are few existing datasets for face landmark detection in videos. Shen *et al.* have proposed 300-VW [60], a dataset made up of ~ 100 videos using the 68-point landmark scheme from the 300-W dataset [57]. Unfortunately, two major issues make this dataset unpopular: 1) it was labelled using fairly weak models from [18] and [67] which results in many labelling errors and high flicker (see Appendix B), and 2) it only contains 100 videos of little diversity, many of which being from the same speaker, or from different speakers in the same environment. Taken together, these two issues mean that the



Figure 4. Example of ground truth labels (●) obtained semi-automatically from an ensemble of 45 models (▲). Ensembling these diverse models provides ground truth labels that do not flicker and can thus be used to measure flickering against. Note how landmark points with the most uncertainty are the most prone to having flickering predictions across consecutive frames.

performance of modern models on 300-WV barely correlates with performance on real-world face videos.

We propose a new video dataset for facial landmarks: WFLW-Video (WFLW-V). It consists of 1000 Youtube creative-commons videos (*i.e.*, an order of magnitude more than its predecessor) and covers a wide range of people, expressions, poses, activities and background environments. Each video is 5s in length. These videos were collected by targeting challenging faces, where ground truth landmarks are subject to uncertainty. The dataset contains two subsets, *hard* and *easy*, made up of 500 videos each. This allows debugging temporal filters and smoothing techniques, whose optimal hyperparameters are often different for videos with little or a lot of flicker. This split was obtained by scraping 2000 videos, and selecting the top and bottom 500 videos based on the variance of predictions in the ensemble. To evaluate a landmark model on the WFLW-V dataset, we train it on the WFLW training set, and evaluate it on all WFLW-V videos. This pipeline best reflects the way large landmark models are trained in the industry, where labelled video data is scarce.

Contrary to the 68-landmark scheme of the 300-VW dataset, we label videos semi-automatically using the more challenging 98-landmark scheme from the WFLW dataset [75], as it is considered the most relevant dataset for future research in face landmark detection [38]. To produce ground truth labels, we train an ensemble of 45 state-of-the-art models using a wide range of data augmentations

of both the test and train set of WFLW (amounting to 10,000 images). We use a mix of large Unets, HRNets [63] and HRFormers [83] to promote landmark diversity. The heatmaps of these models are averaged to produce the ground truth heatmap, allowing uncertain models to weight less in the aggregated output. Ensembling models provides temporal coherence without the need for using hand-tuned filtering or smoothing, which are susceptible to misinterpreting signal for noise (*e.g.* closing eye landmarks mimic high frequency noise).

We provide examples of annotated videos in Fig. 4. Note that regions of ambiguity, such as occluded parts of the face, correspond to a higher variance in landmark predictions. While the ground truth for some frames may be subjective (*e.g.* occluded mouth), having a temporally stable “best guess” is sufficient to measure flicker of individual models. We found that 45 models in the ensemble was enough to provide a low error on the mean for all videos in WFLW-V. While we manually checked that our Oracle was highly accurate and coherent in time, note that it is completely impractical to use it for real-world applications due to its computational cost ($\sim 2B$ parameters).

We checked each frame manually for errors, which were rare. When we found an error in annotation, we corrected it by removing inaccurate models from the ensemble for the frames affected, rather than re-labeling the frame manually. We found this approach to be faster and less prone to human subjectivity when dealing with ambiguous faces, such

Method		Params (M)	Full	Large poses	Expressions	Illumination	Makeup	Occlusion	Blur
LAB [75]	CVPR 2018	12.3	5.27	10.24	5.51	5.23	5.15	6.79	6.32
Wing [26]	CVPR 2018	25	4.99	8.43	5.21	4.88	5.26	6.21	5.81
MHHN [69]	TIP 2020	-	4.77	9.31	4.79	4.72	4.59	6.17	5.82
DecaFA [21]	ICCV 2019	10	4.62	8.11	4.65	4.41	4.63	5.74	5.38
HRNet [63]	TPAMI 2020	9.7	4.60	7.94	4.85	4.55	4.29	5.44	5.42
AS [54]	ICCV 2019	35	4.39	8.42	4.68	4.24	4.37	5.60	4.86
LUVLI [40]	CVPR 2020	-	4.37	7.56	4.77	4.30	4.33	5.29	4.94
AWing [73]	ICCV 2019	24.2	4.36	7.38	4.58	4.32	4.27	5.19	4.96
SDFL [46]	TIP 2021	-	4.35	7.42	4.63	4.29	4.22	5.19	5.08
SDL [44]	ECCV 2020	-	4.21	7.36	4.49	4.12	4.05	4.98	4.82
ADNet [34]	ICCV 2021	13.4	4.14	6.96	4.38	4.09	4.05	5.06	4.79
SLPT [77]	CVPR 2022	19.5	4.12	6.99	4.37	4.02	4.03	5.01	4.79
HIH [42]	-	22.7	4.08	6.87	4.06	4.34	3.85	4.85	4.66
LDEQ (ours)	-	21.8	3.92	6.86	3.94	4.17	3.75	4.77	4.59

Table 1. Performance of our model and previous state-of-the-art on the various WFLW subsets, using the NME metric (\downarrow) for comparison. Models using pre-training on other datasets have been excluded for fair comparison [14, 82, 88].

as occluded ones. Occlusion has been characterized as one of the main remaining challenges in modern landmark detection [76], being most difficult in video data [60]. Finally, the stability of our oracle also depends on the stability of the face bounding box detector. We found the most popular detector, MTCNN [85] to be too jittery, and instead obtained stable detection by bootstrapping our oracle landmarks into the detector. More details about scraping and curating our dataset can be found in Appendix A.

6. Experiments

The aim of these experiments is to demonstrate that: 1) LDEQ is a state-of-the-art landmark detection model for high precision settings like faces, and 2) the LDEQ objective can be modified at inference time on videos to include a recurrence loss. This increases temporal coherence without decreasing accuracy, a common pitfall of popular filters.

6.1. Landmark accuracy

We compare the performance of LDEQ to state-of-the-art models on the WFLW dataset [75], which is based on the WIDER Face dataset [81] and is made up of 7500 train images and 2500 test images. Each image is annotated with a face bounding box and 98 2D landmarks. Compared to previous datasets, WFLW uses denser facial landmarks, and introduces much more diversity in poses, expressions, occlusions and image quality (the test set is further divided into subsets reflecting these factors). This makes it uniquely appropriate for training landmark detectors meant to be deployed on real-world data, like the videos of WFLW-V.

The common evaluation metrics for the WFLW test set are the Normalized Mean Error (see Eq. (10b)), the Area Under the Curve (AUC), and the Failure Rate (FR). The AUC is computed on the cumulative error distribution curve (CED), which plots the fraction of images with NME less or equal to some cutoff, vs. increasing cutoff values. We report AUC_{10} , referring to a maximum NME cutoff of 10 for the CED curve. Higher AUC is better. The FR_X metric is equal to the percentage of test images whose NME is larger than X . We report FR_{10} ; lower is better.

We train our LDEQ for 60 epochs using the pre-cropped WFLW dataset as per [42], and the common data augmentations for face landmarks: rotations, flips, translations, blurs

Metric	Method	Full	Pose	Exp.	Illum.	Mu.	Occ.	Blur
FR_{10} (\downarrow)	LAB	7.56	28.83	6.37	6.73	7.77	13.72	10.74
	HRNet	4.64	23.01	3.50	4.72	2.43	8.29	6.34
	AS	4.08	18.10	4.46	2.72	4.37	7.74	4.40
	LUVLi	3.12	15.95	3.18	2.15	3.40	6.39	3.23
	AWing	2.84	13.50	2.23	2.58	2.91	5.98	3.75
	SDFL	2.72	12.88	1.59	2.58	2.43	5.71	3.62
	SDL	3.04	15.95	2.86	2.72	1.45	5.29	4.01
	ADNet	2.72	12.72	2.15	2.44	1.94	5.79	3.54
	SLPT	2.72	11.96	1.59	2.15	1.94	5.70	3.88
	HIH	2.60	12.88	1.27	2.43	1.45	5.16	3.10
	LDEQ	2.48	12.58	1.59	2.29	1.94	5.36	2.84
AUC_{10} (\uparrow)	LAB	0.532	0.235	0.495	0.543	0.539	0.449	0.463
	HRNet	0.524	0.251	0.510	0.533	0.545	0.459	0.452
	AS	0.591	0.311	0.549	0.609	0.581	0.516	0.551
	LUVLi	0.557	0.310	0.549	0.584	0.588	0.505	0.525
	AWing	0.572	0.312	0.515	0.578	0.572	0.502	0.512
	SDFL	0.576	0.315	0.550	0.585	0.583	0.504	0.515
	SDL	0.589	0.315	0.566	0.595	0.604	0.524	0.533
	ADNet	0.602	0.344	0.523	0.580	0.601	0.530	0.548
	SLPT	0.596	0.349	0.573	0.603	0.608	0.520	0.537
	HIH	0.605	0.358	0.601	0.613	0.618	0.539	0.561
	LDEQ	0.624	0.373	0.614	0.631	0.631	0.552	0.574

Table 2. AUC_{10} and FR_{10} on the WFLW test set.

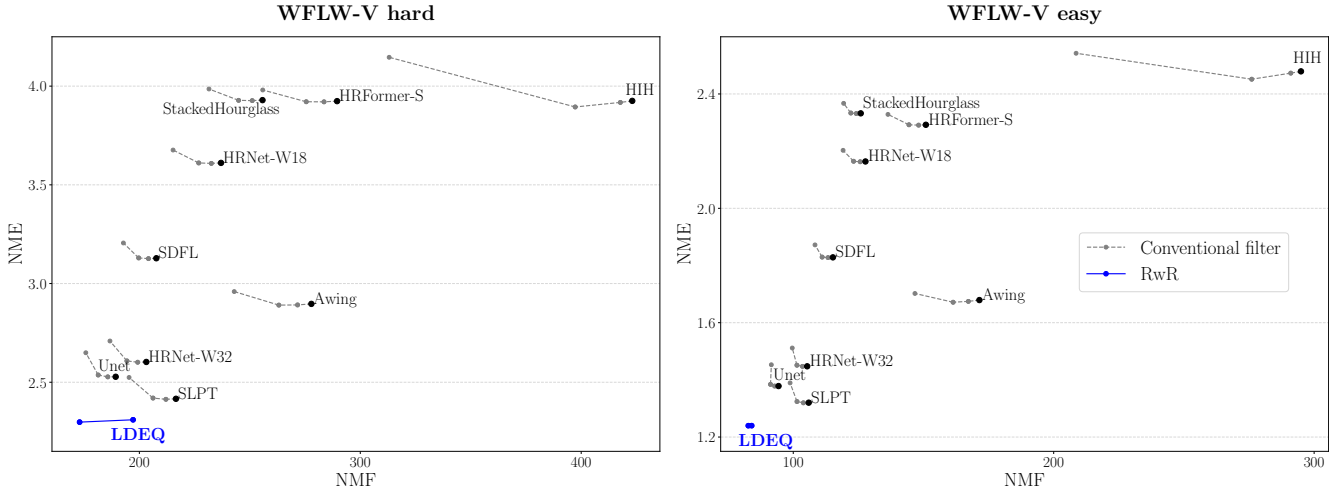


Figure 5. NME (accuracy) vs NMF (temporal coherence) for various models, on the hard (left) and easy (right) WFLW-V subsets. We compare our RwR scheme (blue) to a conventional filter (exponential moving average) using three different smoothing coefficients (gray). For hard videos susceptible to flicker, RwR on LDEQ decreases NMF by 12% without increasing NME, contrary to the conventional filter alternative. For easy videos that contain little to no flicker, conventional filters can increase both NME and NMF, while our model converges to the same fixed point with or without RwR. These results are given in tabular form in Appendix C.

and occlusions. We found the Anderson and fixed-point-iteration solvers to work best over the Broyden solver used in the original DEQ model [9, 10]. By setting a normalization temperature of $T = 2.5$, convergence to fixed points only takes around 5 solver iterations. The NME, AUC and FR performance of LDEQ can be found in tables 1 and 2. We outperform all existing models on all three evaluation metrics, usually dominating individual subsets as well, which is important when applying LDEQ to video data.

6.2. Landmark temporal coherence

We evaluate the performance of LDEQ on the WFLW-V dataset, for both landmark accuracy (NME) and temporal coherence (NMF). We use RwR with early stopping after 2 solver iterations, allowing some adaptive compute (1 solver iteration) in cases where two consecutive frames are almost identical. Our baselines include previous state-of-the-art models that have publicly available weights [42, 46, 73, 77], as well as common architectures of comparable parameter count, which we trained with our own augmentation pipeline. We apply a conventional filtering algorithm, the exponential moving average, to our baselines. This was found to be more robust than more sophisticated filters like Savitzky-Golay filter [59] and One Euro filter [16].

The NME vs. NMF results are shown in Fig. 5 for all models, for the hard and easy WFLW-V subsets. For videos that contain little uncertainty in landmarks (WFLW-V easy), there is little flickering and conventional filtering methods can mistakenly smooth out high frequency signal (e.g. eyes and mouth moving fast). For videos subject to more flickering (WFLW-hard), these same filtering tech-

niques do indeed improve the NMF metric, but beyond a certain smoothing factor this comes at the cost of increasing the NME. In contrast, LDEQ + RwR correctly smooths out flicker for WFLW-W hard without compromising performance on WFLW-V easy. This improvement comes from the fact that the RwR loss in Eq. (8) contains both a smoothing loss plus a log likelihood loss that constrains output landmarks to be plausible solutions, while conventional filters only optimize for the former.

7. Conclusion

We adapt Deep Equilibrium Models to landmark detection, and demonstrate that our LDEQ model can reach state-of-the-art accuracy on the challenging WFLW facial dataset. We then bring the attention of the landmark community to a common problem in video applications, whereby landmarks flicker across consecutive frames. We contribute a new dataset and metric to effectively benchmark solutions to that problem. Since DEQs solve for an objective in the forward pass, we propose to change this objective at test time to take into account new information. This new paradigm can be used to tackle the flickering problem, by adding a recurrent loss term at inference that wasn't present at training time (RwR). We show how to solve for this objective cheaply in a way that leads to state-of-the-art video temporal coherence. We hope that our work brings attention to the potential of deep equilibrium models for computer vision applications, and in particular, the ability to add loss terms to the forward pass process, to leverage new information at inference time.

References

- [1] Mediapipe 3d face transform. <https://developers.googleblog.com/2020/09/mediapipe-3d-face-transform.html>. Accessed: 2022-11-09. **1**
- [2] Mediapipe face mesh. https://google.github.io/mediapipe/solutions/face_mesh. Accessed: 2022-11-09. **1**
- [3] Brandon Amos and J. Zico Kolter. OptNet: Differentiable optimization as a layer in neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 136–145. PMLR, 2017. **2, 4**
- [4] Donald G. Anderson. Iterative procedures for nonlinear integral equations. *J. ACM*, 12(4):547–560, oct 1965. **3**
- [5] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223. PMLR, 06–11 Aug 2017. **4**
- [6] Michal Irani Assaf Shocher, Nadav Cohen. “zero-shot” super-resolution using deep internal learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. **2**
- [7] Akshay Asthana, Stefanos Zafeiriou, Shiyang Cheng, and Maja Pantic. Incremental face alignment in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1859–1866, 2014. **2**
- [8] Shaojie Bai, Zhengyang Geng, Yash Savani, and J. Zico Kolter. Deep equilibrium optical flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. **2**
- [9] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. **1, 2, 3, 8**
- [10] Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Multiscale deep equilibrium models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. **1, 2, 3, 8**
- [11] Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Stabilizing equilibrium models by jacobian regularization. In *International Conference on Machine Learning (ICML)*, 2021. **1, 2**
- [12] Christopher Bishop. Regularization and complexity control in feed-forward networks. In *Proceedings International Conference on Artificial Neural Networks ICANN’95*, volume 1, pages 141–148. EC2 et Cie, January 1995. **4**
- [13] C. G. Broyden. A class of methods for solving nonlinear simultaneous equations. *Mathematics of Computation*, 19:577–593, 1965. **3**
- [14] Adrian Bulat, Enrique Sanchez, and Georgios Tzimiropoulos. Subpixel heatmap regression for facial landmark localization. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2021. **7**
- [15] Xudong Cao, Yichen Wei, Fang Wen, and Jian Sun. Face alignment by explicit shape regression. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2887–2894, 2012. **2**
- [16] Géry Casiez, Nicolas Roussel, and Daniel Vogel. 1 € filter: A simple speed-based low-pass filter for noisy input in interactive systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’12, page 2527–2530, New York, NY, USA, 2012. Association for Computing Machinery. **8**
- [17] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. **2**
- [18] Grigoris G. Chrysos, Epameinondas Antonakos, Stefanos Zafeiriou, and Patrick Snape. Offline deformable face tracking in arbitrary videos. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*, pages 954–962, 2015. **5**
- [19] T.F. Cootes, G.J. Edwards, and C.J. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):681–685, 2001. **1**
- [20] D. Cristinacce and T. F. Cootes. Feature detection and tracking with constrained local models. In *Proceedings of the British Machine Vision Conference*, pages 95.1–95.10. BMVA Press, 2006. doi:10.5244/C.20.95. **1**
- [21] Arnaud Dapogny, Matthieu Cord, and Kevin Bailly. Decafa: Deep convolutional cascade for face alignment in the wild. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6892–6900, 2019. **1, 7**
- [22] Josip Djolonga and Andreas Krause. Differentiable learning of submodular models. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. **2**
- [23] Piotr Dollár, Peter Welinder, and Pietro Perona. Cascaded pose regression. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1078–1085, 2010. **2**
- [24] Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. **2**
- [25] G.J. Edwards, C.J. Taylor, and T.F. Cootes. Interpreting face images using active appearance models. In *Proceedings Third IEEE International Conference on Automatic Face and Gesture Recognition*, pages 300–305, 1998. **1**
- [26] Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, and Xiao-Jun Wu. Wing loss for robust facial landmark localisation with convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*, pages 2235–2245. IEEE, 2018. **1, 7**
- [27] Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel McKenzie, Stanley J. Osher, and Wotao Yin. Fixed point networks: Implicit depth models with jacobian-free backprop. *CoRR*, abs/2103.12803, 2021. **2**
- [28] P. Garrido, L. Valgaerts, H. Sarmadi, I. Steiner, K. Varanasi, P. Pérez, and C. Theobalt. Vdub: Modifying face video of

- actors for plausible visual alignment to a dubbed audio track. *Comput. Graph. Forum*, 34(2):193–204, may 2015. 1
- [29] Zhengyang Geng, Xin-Yu Zhang, Shaojie Bai, Yisen Wang, and Zhouchen Lin. On training implicit models. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 24247–24260. Curran Associates, Inc., 2021. 2
- [30] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 4
- [31] Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. In *International Conference on Learning Representations*, 2018. 4
- [32] Nicklas Hansen, Rishabh Jangir, Yu Sun, Guillem Alenyà, Pieter Abbeel, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. Self-supervised policy adaptation during deployment. In *International Conference on Learning Representations*, 2021. 2
- [33] Behzad Hasani and Mohammad H. Mahoor. Facial expression recognition using enhanced deep 3d convolutional neural networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2278–2288, 2017. 1
- [34] Yangyu Huang, Hao Yang, Chong Li, Jongyoo Kim, and Fangyun Wei. Adnet: Leveraging error-bias towards normal direction in face alignment. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3060–3070, 2021. 1, 2, 7
- [35] Rateb Jabbar, Mohammed Shinoy, Mohamed Kharbeche, Khalifa Al-Khalifa, Moez Krichen, and Kamel Barkaoui. Driver drowsiness detection model using convolutional neural networks techniques for android application. In *2020 IEEE International Conference on Informatics, IoT, and Enabling Technologies (ICIOT)*, pages 237–242, 2020. 1
- [36] Hanwen Jiang, Shaowei Liu, Jiashun Wang, and Xiaolong Wang. Hand-object contact consistency reasoning for human grasps generation. In *Proceedings of the International Conference on Computer Vision*, 2021. 2
- [37] Heechul Jung, Sihaeng Lee, Junho Yim, Sunjeong Park, and Junmo Kim. Joint fine-tuning in deep neural networks for facial expression recognition. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2983–2991, 2015. 1
- [38] Kostiantyn Khabarlak and Larysa Koriashkina. Fast facial landmark detection and applications: A survey. *Journal of Computer Science and Technology*, 22, 2022. 6
- [39] Dae Ha Kim, Min Kyu Lee, Dong Yoon Choi, and Byung Cheol Song. Multi-modal emotion recognition using semi-supervised learning and multiple neural networks in the wild. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction, ICMI '17*, page 529–535, New York, NY, USA, 2017. Association for Computing Machinery. 1
- [40] Abhinav Kumar, Tim K. Marks, Wenxuan Mou, Ye Wang, Michael Jones, Anoop Cherian, Toshiaki Koike-Akino, Xiaoming Liu, and Chen Feng. Luvli face alignment: Estimating landmarks’ location, uncertainty, and visibility likelihood. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 7
- [41] Hanjiang Lai, Shengtao Xiao, Yan Pan, Zhen Cui, Jiashi Feng, Chunyan Xu, Jian Yin, and Shuicheng Yan. Deep recurrent regression for facial landmark detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(5):1144–1157, 2018. 2
- [42] Xing Lan, Qinghao Hu, and Jian Cheng. HIH: towards more accurate face alignment via heatmap in heatmap. *CoRR*, abs/2104.03100, 2021. 2, 7, 8
- [43] Shan Li and Weihong Deng. Deep facial expression recognition: A survey. *IEEE Transactions on Affective Computing*, pages 1–1, 2020. 1
- [44] Weijian Li, Yuhang Lu, Kang Zheng, Haofu Liao, Chihung Lin, Jiebo Luo, Chi-Tung Cheng, Jing Xiao, Le Lu, Chang-Fu Kuo, and Shun Miao. Structured landmark detection via topology-adapting deep graph learning. In *ECCV*, 2020. 1, 7
- [45] Yizhuo Li, Miao Hao, Zonglin Di, Nitesh Bharadwaj Gundavarapu, and Xiaolong Wang. Test-time personalization with a transformer for human pose estimation. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 2
- [46] Chunze Lin, Beier Zhu, Quan Wang, Renjie Liao, Chen Qian, Jiwen Lu, and Jie Zhou. Structure-coherent deep feature learning for robust face alignment. *IEEE Transactions on Image Processing*, 30:5313–5326, 2021. 1, 7, 8
- [47] Dong C. Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *MATHEMATICAL PROGRAMMING*, 45:503–528, 1989. 4
- [48] Diogo C. Luvizon, Hedi Tabia, and David Picard. Human pose regression by combining indirect part detection and contextual information. *CoRR*, abs/1710.02322, 2017. 3
- [49] Iacopo Masi, Yue Wu, Tal Hassner, and Prem Natarajan. Deep face recognition: A survey. In *2018 31st SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 471–478, 2018. 1
- [50] Jiteng Mu, Weichao Qiu, Adam Kortylewski, Alan L. Yuille, Nuno Vasconcelos, and Xiaolong Wang. A-SDF: learning disentangled signed distance functions for articulated shape representation. In *ICCV*, pages 12981–12991, 2021. 2
- [51] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. *CoRR*, abs/1603.06937, 2016. 2, 3
- [52] Chirag Pabbaraju, Ezra Winston, and J Zico Kolter. Estimating lipschitz constants of monotone deep equilibrium models. In *International Conference on Learning Representations*, 2021. 2
- [53] Chirag Pabbaraju, Ezra Winston, and J Zico Kolter. Estimating lipschitz constants of monotone deep equilibrium models. In *International Conference on Learning Representations*, 2021. 3

- [54] Shengju Qian, Keqiang Sun, Wayne Wu, Chen Qian, and Ji-aya Jia. Aggregation via separation: Boosting facial landmark detector with semi-supervised style translation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10152–10162, 2019. [1](#), [7](#)
- [55] Mallikarjun B R, Visesh Chari, and C.V. Jawahar. Efficient face frontalization in unconstrained images. In *2015 Fifth National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics (NCVPRIPG)*, pages 1–4, 2015. [1](#)
- [56] Max Revay, Ruigang Wang, and Ian R. Manchester. Lipschitz bounded equilibrium networks. *CoRR*, abs/2010.01732, 2020. [2](#)
- [57] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. *2013 IEEE International Conference on Computer Vision Workshops*, pages 397–403, 2013. [5](#)
- [58] Reginaldo J. Santos. Equivalence of regularization and truncated iteration for general ill-posed problems. *Linear Algebra and its Applications*, 236:25–33, 1996. [4](#)
- [59] Abraham. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964. [8](#)
- [60] Jie Shen, Stefanos Zafeiriou, Grigoris G. Chrysos, Jean Kossai, Georgios Tzimiropoulos, and Maja Pantic. The first facial landmark tracking in-the-wild challenge: Benchmark and results. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*, pages 1003–1011, 2015. [5](#), [7](#)
- [61] J. Sjöberg and L. Ljung. Overtraining, regularization, and searching for minimum in neural networks. *IFAC Proceedings Volumes*, 25(14):73–78, 1992. 4th IFAC Symposium on Adaptive Systems in Control and Signal Processing 1992, Grenoble, France, 1-3 July. [4](#)
- [62] Akash Srivastava, Lazar Valkov, Chris Russell, Michael U. Gutmann, and Charles Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [4](#)
- [63] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5686–5696, 2019. [1](#), [6](#), [7](#)
- [64] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9229–9248. PMLR, 13–18 Jul 2020. [2](#)
- [65] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep convolutional network cascade for facial point detection. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3476–3483, 2013. [2](#)
- [66] George Trigeorgis, Patrick Snape, Mihalios A. Nicolaou, Epameinondas Antonakos, and Stefanos Zafeiriou. Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4177–4187, 2016. [2](#)
- [67] Georgios Tzimiropoulos. Project-out cascaded regression with an application to face alignment. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3659–3667, 2015. [5](#)
- [68] Roberto Valle, José Miguel Buenaposada, Antonio Valdés, and Luis Baumela. A deeply-initialized coarse-to-fine ensemble of regression trees for face alignment. In *ECCV*, 2018. [1](#)
- [69] Jun Wan, Zhihui Lai, Jun Liu, Jie Zhou, and Can Gao. Robust face alignment by multi-order high-precision hourglass network. *IEEE Transactions on Image Processing*, 30:121–133, 2021. [1](#), [2](#), [7](#)
- [70] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2021. [2](#)
- [71] Mei Wang and Weihong Deng. Deep face recognition: A survey. *CoRR*, abs/1804.06655, 2018. [1](#), [2](#)
- [72] Po-Wei Wang, Priya L. Donti, Bryan Wilder, and J. Zico Kolter. Satnet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6545–6554. PMLR, 2019. [2](#)
- [73] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive wing loss for robust face alignment via heatmap regression. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6970–6980, 2019. [1](#), [2](#), [7](#), [8](#)
- [74] Ezra Winston and J. Zico Kolter. Monotone operator equilibrium networks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10718–10728. Curran Associates, Inc., 2020. [2](#), [3](#)
- [75] Wayne Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, and Qiang Zhou. Look at boundary: A boundary-aware face alignment algorithm. In *CVPR*, 2018. [1](#), [2](#), [6](#), [7](#)
- [76] Yue Wu and Qiang Ji. Facial landmark detection: A literature survey. *Int. J. Comput. Vision*, 127(2):115–142, feb 2019. [7](#)
- [77] Jiahao Xia, Weiwei Qu, Wenjian Huang, Jianguo Zhang, Xi Wang, and Min Xu. Sparse local patch transformer for robust face alignment and landmarks inherent relation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4052–4061, June 2022. [1](#), [7](#), [8](#)
- [78] Xuehan Xiong and Fernando De la Torre. Supervised descent method and its applications to face alignment. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 532–539, 2013. [2](#)
- [79] Jingwei Yan, Wenming Zheng, Zhen Cui, Chuangao Tang, Tong Zhang, Yuan Zong, and Ning Sun. Multi-clue fusion

- for emotion recognition in the wild. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction, ICMI '16*, page 458–463, New York, NY, USA, 2016. Association for Computing Machinery. 1
- [80] Jing Yang, Qingshan Liu, and Kaihua Zhang. Stacked hourglass network for robust facial landmark localisation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2025–2033, 2017. 2
- [81] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Wider face: A face detection benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 7
- [82] Baosheng Yu and Dacheng Tao. Heatmap regression via randomized rounding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 7
- [83] Yuhui Yuan, Rao Fu, Lang Huang, Weihong Lin, Chao Zhang, Xilin Chen, and Jingdong Wang. Hrformer: High-resolution transformer for dense prediction. In *NeurIPS*, 2021. 6
- [84] Jiangning Zhang, Xianfang Zeng, Yusu Pan, Yong Liu, Yu Ding, and Changjie Fan. Freenet: Multi-identity face reenactment. *CoRR*, abs/1905.11805, 2019. 1
- [85] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, Oct 2016. 7
- [86] Marvin Mengxin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: Learning to adapt to domain shift. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 2
- [87] Yizhe Zhang, Ming Shao, Edward K. Wong, and Yun Fu. Random faces guided sparse many-to-one encoder for pose-invariant face recognition. In *2013 IEEE International Conference on Computer Vision*, pages 2416–2423, 2013. 1
- [88] Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. *CoRR*, abs/2112.03109, 2021. 7
- [89] Xiangyu Zhu, Zhen Lei, Junjie Yan, Dong Yi, and Stan Z. Li. High-fidelity pose and expression normalization for face recognition in the wild. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 787–796, 2015. 1
- [90] Matt Johnson Zico Kolter, David Duvenaud. http://implicit-layers-tutorial.org/deep_equilibrium_models/. 3