A Functional Regression approach to Facial Landmark Tracking

Enrique Sánchez-Lozano, Georgios Tzimiropoulos*, Brais Martinez*,
Fernando De la Torre and Michel Valstar

Abstract—Linear regression is a fundamental building block in many face detection and tracking algorithms, typically used to predict shape displacements from image features through a linear mapping. This paper presents a Functional Regression solution to the least squares problem, which we coin Continuous Regression, resulting in the first real-time incremental face tracker. Contrary to prior work in Functional Regression, in which B-splines or Fourier series were used, we propose to approximate the input space by its first-order Taylor expansion, yielding a closed-form solution for the continuous domain of displacements. We then extend the continuous least squares problem to correlated variables, and demonstrate the generalisation of our approach. We incorporate Continuous Regression into the cascaded regression framework, and show its computational benefits for both training and testing. We then present a fast approach for incremental learning within Cascaded Continuous Regression, coined iCCR, and show that its complexity allows real-time face tracking, being 20 times faster than the state of the art. To the best of our knowledge, this is the first incremental face tracker that is shown to operate in real-time. We show that iCCR achieves state-of-the-art performance on the 300-VW dataset, the most recent, large-scale benchmark for face tracking.

Index Terms—Continuous Regression, Face Tracking, Functional Regression, Functional Data Analysis.

✦

1 INTRODUCTION

LINEAR REGRESSION is a standard tool in many Computer Vision problems, such as pose estimation [16], and object tracking [48], and it is the foremost approach for facial landmark detection, especially after the introduction of the so-called Cascaded Regression method [11], [43], [46], [52], [54]. The goal of facial landmark detection is to locate a sparse set of facial landmarks in still images or videos. It is a problem of wide interest to the Computer Vision community because many higher level face analysis tasks, such as facial expression recognition [15], and face recognition [59], are critically affected by the performance of facial landmark detection systems.

Arguably, the most popular method for facial landmark detection is the Supervised Descent Method (SDM) [52], in which a set of Linear Regressors are used in a cascade, each approximating the average descent direction for a specific input space. SDM generates samples by randomly perturbing the training data, i.e., each regressor is learnt by minimising the least squares error on a training set of annotated images and known displacements. Many extensions to SDM have been proposed [2], [22], [38], [46], yet little attention has been paid to the limitations of the *sampling-based approach to linear regression*, inherent to all aforementioned methods. More specifically, it has been shown that the sampling-based approach is directly related to the following three limitations:

- The least squares formulation theoretically needs an exponential number of samples with respect to the dimensionality of the input vector in order to avoid biased models.
- Training an SDM model is computationally expensive as it requires the data to be sampled per cascade level, with both memory and time needed increasing dramatically with the number of perturbations.
- Incremental learning for SDM has been reported to be extremely slow [2], [53]. This limits the capacity of SDM and its extensions for learning on-line and in real-time, which for an application such as face tracking is impractical and hence highly unsatisfactory.

In order to overcome the aforementioned limitations, in this paper we propose Continuous Regression, a Functional Regression solution to the Least-Squares problem, and show its application to real-time face tracking. Contrary to SDM, the proposed Continuous Regression solution only requires the data to be sampled at the ground-truth landmarks, i.e., no sampling at perturbations is required. This way, one can sample and store the ground-truth data only once, and then train each cascade level, or even a new model under a different configuration extremely quickly. While the SDM can take hours to train, our proposed Continuous Regression can train different models in seconds, once the data is extracted. Moreover, and contrary to existing cascaded regression methods, our continuous formulation allows for real-time incremental learning, which as we demonstrate is crucial for attaining state-of-the-art tracking performance on the most challenging subset of the 300-VW dataset [41].

Continuous Regression approaches the standard least squares formulation from a continuous perspective where

- E. Sánchez-Lozano, G. Tzimiropoulos, B.Martinez and M. Valstar are with the School of Computer Science. University of Nottingham. E-mail: Enrique.SanchezLozano@nottingham.ac.uk
- F. De la Torre is with the Robotics Institute. Carnegie Mellon University.
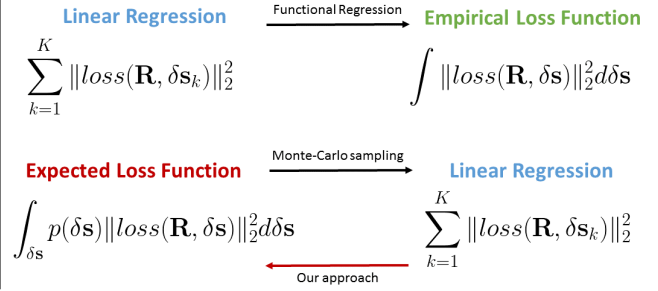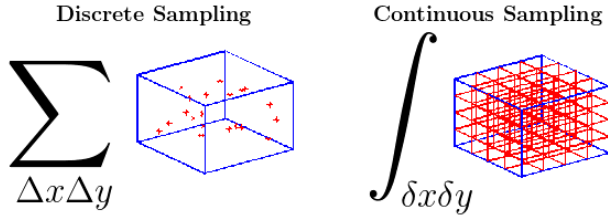- *Both authors contributed equally

Fig. 1. **Left:** Difference between sampling-based regression and continuous regression. The continuous regression accounts for all the samples within a neighbourhood, whereas sampling-based needs to sample the data from a given distribution. **Right:** Our new approach to Continuous Regression can be seen as the inverse of a Monte Carlo sampling estimation. We will see that the probability density function defined by $p(\delta \mathbf{s})$ defines the volume within which samples are taken.

the target variable is treated as a continuum. Unlike existing approaches to Functional Regression, which approximate the input space by means of B-splines [26], Fourier Series [33], or Radial Basis Functions [31], in this paper, we propose a first-order Taylor expansion of the feature space. This way, the input space becomes linear with respect to shape displacements yielding a closed-form solution. We go beyond prior work in Functional Regression, and study the case of correlated variables which naturally arises within landmark localisation and tracking. To this end, we introduce a "data term", tasked with correlating the different dimensions over which the problem is solved.

We then extend Continuous Regression within the cascaded regression framework, and demonstrate that training our method has notable computational benefits compared to that of standard SDM, without sacrificing accuracy. Finally, we devise an approach for incremental learning in the continuous domain, and show that its complexity allows for real-time tracking. To the best of our knowledge, our Incremental Cascaded Continuous Regression method (iCCR), is the first cascaded regression method that incorporates real-time incremental learning capabilities. Crucially, we show the importance of incremental learning in achieving state-of-the-art results on the 300-VW dataset.

This manuscript re-formulates the Continuous Regression over uncorrelated variables, which has been presented in [37], and further extends it to the domain of correlated variables, presented in [39]. Building on [37], [39], we present a complete formulation for the Continuous Regression problem, and then provide a geometric interpretation and link to previous works. An overview of our work is depicted in Fig. 1.

1.1 Contributions

Our main contributions are as follows:

- We propose a complete formulation for Continuous Regression including a closed-form solution based on a first-order Taylor expansion of the input space. Notably, prior work in Functional Regression has only considered approximating the input space with pre-defined basis functions.
- We then go one step further in the classical Functional Regression formulation and propose a novel optimisation problem for Continuous Regression

which allows the target variable to be correlated. We also analytically derive the solution to this problem.
- We incorporate Continuous Regression into the Cascaded Regression framework, demonstrating its computational benefits over SDM. We coin our method Cascaded Continuous Regression (CCR).
- We derive the incremental learning updates for CCR, coined iCCR. We show that iCCR is capable of real-time incremental face tracking, being also an order of magnitude less complex than previous works based on incremental SDM.
- We evaluate iCCR on the 300VW data set [41] and show the importance of incremental learning in achieving state-of-the-art performance.

2 RELATED WORK

Our work applies and extends concepts from Functional Regression to the problem of face tracking. This section provides a brief description of Functional Regression methods and how they are related to our work. We follow with a review of related face alignment and tracking methods.

2.1 Prior Work in Functional Regression Analysis

Functional Data Analysis. Functional Data Analysis (FDA) [32] is a branch of statistics that aims to model realisations of stochastic processes as continuous functions [32]. FDA assumes that observations, and/or responses, are outcomes of continuous processes [28], and attempts to parameterise functions by means of basis functions, rather than as a set of samples. Several basis functions have been proposed to date for different FDA domains, such as Radial Basis Functions [31], Fourier Series [33], or B-splines [26].

A typical use of Functional Data Analysis is in the field of Linear Regression, called *Functional Regression* [57], in which either the responses or the observed data, or both, are modelled as continuous functions. Functional Regression on longitudinal data can be seen as an extension of Multivariate Polynomial Regression to complex time series. Typically, basis functions are used to approximate the observed data, although some elegant solutions have also been proposed to model responses as well [28]. However, in practice it is not computationally feasible to approximate image features by means of Radial Basis Functions or Fourier Series. We

propose to approximate the input features by their Taylor expansion with respect to the variable of interest (variation on shapes) to address this. To the best of our knowledge, this is the first time that this approximation is applied in the context of Functional Regression.

FDA in Computer Vision. To the best of our knowledge, there are very few works that apply ideas from FDA to Computer Vision problems, and none address the problem of facial landmark localisation. The work that is perhaps closest to ours (in some of its preliminary assumptions) is PCA applied over continuous subspaces presented by Levin and Shashua [21]. The main idea of [21] is to extend the covariance matrix of the input data, necessary for the computation of PCA, to all possible linear combinations of the training samples. That is, given a set of training examples (images) $\mathbf{a}_i$, $i = 1...d$, the covariance is defined as $Cov(W) = \frac{1}{V} \int_{\mathbf{a} \in W} \mathbf{a}\mathbf{a}^T d\mathbf{a}$, where W is the set of convex combinations of training examples, and V is the volume of the polytope defined by W. The inverse of the volume outside the integral indicates that a uniform density function when sampling the points $\mathbf{a} \in W$ is assumed. In this work, we will set out how one can formulate the Least-Squares problem for Continuous Regression using a covariance matrix, providing a link between our work and [21]. Our solution, however, does not necessarily need to be formulated using a uniform distribution.

2.2 Prior Work on Face Tracking

Facial landmark tracking methods have often been adaptations of facial landmark detection methods. For example, Active Shape Models [9], Active Appearance Models (AAM) [11], [27], Constrained Local Models (CLM) [40], or the Supervised Descent Method (SDM) [52] were all presented as detection algorithms. It is thus natural to group facial landmark tracking algorithms in the same way as detection algorithms, i.e. into discriminative and generative methods.

On the generative side, AAMs have often been used for tracking. Because model fitting relies on gradient descent, it suffices to start the fitting from the previously tracked shape[1]. AAMs have been particularly regarded as very reliable for person specific tracking, but not for generic tracking (i.e., tracking faces unseen during training) [17]. Recently [44], [47] showed however that an improved optimisation procedure and the use of in-the-wild images for training can lead to satisfactorily person independent AAM. Eliminating the piecewise-affine representation and adopting a part-based model led to the Gauss-Newton Deformable Part Model (GN-DPM) [45] which is the state-of-the-art AAM.

Discriminative methods directly learn a mapping from image features to shape displacements, effectively eliminating the need to minimise a reconstruction error. A certain family of discriminative methods rely on training local classifier-based models of appearance, with the local responses being then constrained by a shape model [9], [13], [40]. These algorithms are typically cast into the Constrained Local Model (CLM) formulation [13], [25], [40].

One of the simplest tools for discriminative methods is Linear Regression, which was firstly used to bypass the need

---

1. "Implementation tricks" can be found in [43], which provides a very detailed account of how to optimise an AAM tracker

of an appearance model when fitting AAMs to images [11]. Herein, we refer to linear regression as a sampling-based method, in which the training is performed by generating random perturbations of the target variable (the shape displacements). Despite its simplicity and its fast computation, a single regressor has been shown to be a poor choice to account for all input variance. In order to overcome this limitation, boosted regression methods were proposed in [12], [42] where a set of weak regressors are added to a master model, resulting in stronger regressors [12], [24]. Furthermore, in order to reduce the input variance of each of the weak regressors, [16] proposed the use of weakly invariant features. The use of such features requires boosted regression be split into different stages, in which a different regressor has to be learnt for each feature extraction step. This novel approach was firstly exploited for the task of pose estimation [16] and then was successfully applied to face alignment in [6], [7]. However, the most successful form of Cascaded Regression is the Supervised Descent Method (SDM) [52], which applies a sampling-based linear regression for each cascade level, and in which each level can be seen as an average descent direction for the full shape displacement estimation. The use of SIFT [23] features, and the direct estimation of the full shape variation, made the SDM the state-of-the-art method for face alignment. Successive works have further shown impressive efficiency [19], [34] and reliable performance [46], [55].

However, how to best exploit discriminative cascaded regression for tracking and, in particular, how to efficiently integrate incremental learning, is still an open problem. More importantly, SDM has not yet overcome one of the main intrinsic limitations of sampling-based linear regression, to wit, the generation of biased models. This is especially important when it comes to tracking, since the initialisation is given by the points of the previous frame, and therefore there is no guarantee that the output of the model will serve as a good initialisation point for the subsequent frame. Eventually, this might cause the tracker to drift. Also, the SDM is prone to fail for large out-of-plane head poses. In order to overcome this limitation, recent methods have proposed to pay special attention to the way the fitting is initialised. In [54], a Global Supervised Descent Method applies a partition to the shape space, in order to generate a different cascaded regression model for each possible domain of homogeneous descent directions.

Regarding evaluation and benchmarking, the only large-scale face tracking benchmark that exists to date is the 300 Videos in the Wild (300-VW) challenge [41]. Two methods were shown to clearly outperform all others. In [56], a multi-view cascaded regression is used in which a different model is trained for each different partition of the space of poses. This approach, along with an accurate re-initialisation system achieved the best results in the challenge. In [50], a multi-stage regression approach is used, which guides the initialisation process based on the localisation of a subset of landmarks with a strong semantic meaning. The resulting method achieved results very close to those of [56].

Finally, it is worth mentioning the advent of facial landmark localisation using Deep Learning [5], [29], [51]. However, despite achieving impressive accuracy, these techniques are far from being able to operate at real-time speed.

3 LINEAR REGRESSION FOR FACE ALIGNMENT

This section reviews Cascaded Regression training using Linear Regression, which is the learning problem Continuous Regression builds upon. In the following, we will represent a face image by $\mathbf{I}$. A face shape $\mathbf{s} \in \mathbb{R}^{2n}$ is a vector describing the location of the n landmarks considered. We will represent vector shapes as $\mathbf{s} = (x_1, \ldots, x_n, y_1, \ldots y_n)^T$. We also define $\mathbf{x} = f(\mathbf{I}, \mathbf{s}) \in \mathbb{R}^d$ as the feature vector representing shape $\mathbf{s}$. An asterisk represents the ground truth, e.g., $\mathbf{s}_j^*$ is the ground truth shape for image j.

Cascaded Regression [6], [16], [52] is an iterative regression method in which the output of regression at level l is used as input for level $l + 1$, and each level uses image features depending on the current shape estimate. The most widely used form of Cascaded Regression is the Supervised Descent Method (SDM, [52]), which learns a cascade of linear regressors on SIFT features.

Training SDM for face alignment requires a set of M images $\{\mathbf{I}_j\}_{j=1\ldots M}$ and their corresponding K perturbed shapes $\{\mathbf{s}_{j,k}^{(0)}\}_{j=1\ldots M, k=1\ldots K}$. When the fitting is carried out from the output of a face detector, the different perturbations, $k = 1 \ldots K$, are generated by applying a random perturbation to the bounding boxes, capturing the variance of the face detector. When the fitting is initialised from the fitting of the previous frame, as is the case of tracking, the training perturbations are generated by applying a random perturbation to the *ground-truth* points, ideally capturing the variance of facial shapes between consecutive frames. For each cascade level l (starting with $l = 0$), a regressor is trained and applied to the training data, resulting in an update to the input shapes. Therefore, a new set of training shapes $\{\mathbf{s}_{j,k}^{(l+1)}\}_{j=1\ldots M, k=1\ldots K}$ is generated for level $l + 1$, in which the updated perturbed shapes are on average closer to their corresponding ground-truth. This way, the input variance and the training error are decreased as one descends through the cascade.

Mathematically speaking, if we are given a set of perturbations $\delta \mathbf{s}_{j,k}^{(l)} = \mathbf{s}_{j,k}^{(l)} - \mathbf{s}_j^*$ between the initial shapes for level l, and their corresponding ground-truth, a regressor $\mathbf{R}^{(l)}$ is learnt that aims to minimise the least-squares error:

$$\underset{\mathbf{R}}{\arg\min} \sum_{j=1}^{M} \sum_{k=1}^{K} \| \delta \mathbf{s}_{j,k}^{(l)} - \mathbf{R} f(\mathbf{I}_j, \mathbf{s}_j^* + \delta \mathbf{s}_{j,k}^{(l)}) \|_2^2, \quad (1)$$

where the bias term is implicitly included by appending a 1 to the feature vector. If we store the features $\mathbf{x}_{j,k} = f(\mathbf{I}_j, \mathbf{s}_{j,k}^{(l)})$ into matrix $\mathbf{X}^{(l)}$, and the corresponding displacements $\delta \mathbf{s}_{j,k}^{(l)} = \mathbf{s}_{j,k}^{(l)} - \mathbf{s}_j^*$ into matrix $\mathbf{Y}^{(l)}$, the solution to Eq. (1) can be written in closed form:

$$\mathbf{R}^{(l)} = \mathbf{Y}^{(l)} (\mathbf{X}^{(l)})^T \left(\mathbf{X}^{(l)} (\mathbf{X}^{(l)})^T \right)^{-1}. \quad (2)$$

The regressor $\mathbf{R}^{(l)}$ is then applied to each of the extracted features $\mathbf{x}_{j,k}$, and a new training set is generated, in which the training shapes are now given as:

$$\mathbf{s}_{j,k}^{(l+1)} = \mathbf{s}_{j,k}^{(l)} - \mathbf{R}^{(l)} f(\mathbf{I}_j, \mathbf{s}_{j,k}^{(l)}). \quad (3)$$

The process is repeated until the average of the differences $\delta \mathbf{s}_{j,k}^{(l)} = \mathbf{s}_{j,k}^{(l)} - \mathbf{s}_j^*$ no longer decreases. Typically, the training error is minimised after 4 or 5 cascade levels.

During testing, the initial shape $\mathbf{s}^{(0)}$ is forwarded to the first regressor $\mathbf{R}^{(0)}$, to generate a new shape $\mathbf{s}^{(1)}$, which is subsequently forwarded to $\mathbf{R}^{(1)}$. For each level l, a new shape $\mathbf{s}^{(l+1)}$ is generated by applying regressor $\mathbf{R}^{(l)}$ to the features extracted at $\mathbf{s}^{(l)}$:

$$\mathbf{s}^{(l+1)} = \mathbf{s}^{(l)} - \mathbf{R}^{(l)} f(\mathbf{I}, \mathbf{s}^{(l)}). \quad (4)$$

Provided the prediction statistics of $\mathbf{R}^{(l)}$ are known for all cascade levels, the training of Cascaded Regression can be done independently for each of them (par-SDM, [2]). As pointed out in [2], for a sufficiently large and general training set, we can first train SDM, following the sequential algorithm shown above. After SDM training, we can compute the mean $\boldsymbol{\mu}^{(l)}$ and covariance $\boldsymbol{\Sigma}^{(l)}$ of the differences that result from updating the training shapes: $\delta \mathbf{s}_{j,k}^{(l)} = \mathbf{s}_{j,k}^{(l)} - \mathbf{s}_j^*$. Then, for each cascade level, we can generate a set of random perturbations $\delta \mathbf{s}_{j,k}^{(l)}$, drawn from a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}^{(l)}, \boldsymbol{\Sigma}^{(l)})$, and then compute each $\mathbf{R}^{(l)}$ using Eq. (2). This way, computing each regressor is independent from the other levels, meaning that the training can be performed in parallel.

As empirically shown by [2], such a model achieves similar accuracy as a model trained in a sequential manner. In addition, [2] showed that one can use these statistics to train a par-SDM using a larger training set, without compromising accuracy with respect to a sequentially trained SDM. This implies that, for a sufficiently large original training set, the computed statistics properly model the displacements for each level. We will see that our Continuous Regression can be straightforwardly introduced within the Cascaded Regression framework, with the advantage that our proposed method only needs the data to be sampled once, whereas both the SDM and the par-SDM require a sampling step for each Cascade Level.

4 CONTINUOUS REGRESSION

In this section we present a complete formulation and solution for Continuous Regression, in which the shape displacements are treated as continuous variables. We propose the use of a first-order Taylor expansion approximation of the feature space, which yields a closed-form solution in the continuous domain for the Least-Squares problem. We will see that our Taylor-based closed-form solution depends only on the features extracted from the ground-truth shapes and their derivatives. Contrary to existing approaches, **our Continuous Regression solution does not need to sample over perturbations**.

Let us consider the Linear Regression problem shown in Eq. (1). The extension of the Least-Squares problem to the continuous domain of shape displacements is defined as:

$$\underset{\mathbf{R}}{\arg\min} \sum_{j=1}^{M} \int_{-a_1}^{a_1} \ldots \int_{-a_{2n}}^{a_{2n}} \| \delta \mathbf{s} - \mathbf{R} f(\mathbf{I}_j, \mathbf{s}_j^* + \delta \mathbf{s}) \|_2^2 d\delta \mathbf{s}, \quad (5)$$

where[2] the line element is defined as $\delta \mathbf{s} = \delta x_1 \ldots \delta x_n, \delta y_1 \ldots \delta y_n$, the limits are bounded, and defined independently for each dimension. Without loss of generality, we will assume that the limits are symmetric,

2. For the sake of clarity, we drop the dependence on the level

and the displacement assumed for each of the points is therefore unbiased.

Unlike previous works on Functional Regression, which approximate the input feature space by B-splines or Fourier series, we propose to approximate the input feature space by its first-order Taylor expansion:

$$f(\mathbf{I}_j, \mathbf{s}_j^* + \delta\mathbf{s}) \approx f(\mathbf{I}_j, \mathbf{s}_j^*) + \mathbf{J}_j^*\delta\mathbf{s}, \qquad (6)$$

where $\mathbf{J}_j^* = \frac{\partial f(\mathbf{I}_j, \mathbf{s})}{\partial \mathbf{s}}|_{(\mathbf{s}=\mathbf{s}_j^*)} \in \mathbb{R}^{d \times 2n}$, evaluated at $\mathbf{s} = \mathbf{s}_j^*$, is the Jacobian of the feature representation of image $\mathbf{I}_j$, with respect to shape coordinates $\mathbf{s}$, at $\mathbf{s}_j^*$. To compute the Jacobian of the image features (pixels) with respect to the x-positions, we have to use some approximation, e.g. the Sobel filter. We can also compute the empirical derivative of other complex features, such as SIFT or HOG. Indeed, gradient-based features, such as SIFT or HOG, are smoother than pixels, and are therefore more feasible to be approximated by a first-order Taylor expansion. Experiments will show that the use of HOG features suffices to attain state of the art performance, reinforcing the idea that a first-order Taylor expansion is sufficient. Herein, we will define the empirical derivative of the image features with respect to the x (or y) coordinate, as follows:

$$\nabla f_x = \frac{\partial f(\mathbf{I}, x)}{\partial x} \approx \frac{f(\mathbf{I}, x + \Delta x) - f(\mathbf{I}, x - \Delta x)}{2\Delta x} \qquad (7)$$

Then, the i-th column of the Jacobian $\mathbf{J}^*$, for $i = 1, \dots, n$, is given by ∇f_{x_i}, whereas the $(i+n)$-th column of the Jacobian is given by ∇f_{y_i}. In practice, Δx is set to the minimum displacement, i.e., $\Delta x = 1$. Given the approximation of Eq. (6), the original problem can be expressed as:

$$\int \|\delta\mathbf{s} - \mathbf{R}\left(f(\mathbf{I}_j, \mathbf{s}_j^*) + \mathbf{J}_j^*\delta\mathbf{s}\right)\|_2^2 d\delta\mathbf{s} \approx \qquad (8)$$

$$\int [\delta\mathbf{s}^T \mathbf{A}_j \delta\mathbf{s} + 2\delta\mathbf{s}^T \mathbf{b}_j + \mathbf{x}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^*]d\delta\mathbf{s}, \qquad (9)$$

where $\mathbf{A}_j = (\mathbb{I} - \mathbf{R}\mathbf{J}_j^*)^T(\mathbb{I} - \mathbf{R}\mathbf{J}_j^*)$ and $\mathbf{b}_j = \mathbf{J}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^* - \mathbf{R}\mathbf{x}_j^*$. Given the independence between each dimension, the linear term equals to zero for symmetric limits. For the quadratic term, the solution stems from the fact that

$$\int \delta\mathbf{s}^T \mathbf{A}_j \delta\mathbf{s} d\delta\mathbf{s} = \sum_{u,v} \int A_j^{uv} \delta s_u \delta s_v d\delta\mathbf{s}, \qquad (10)$$

where A_j^{uv} is the $\{u,v\}$-th entry of $\mathbf{A}_j$, and s_u, s_v refer indistinctly to any pair of elements of $\mathbf{s}$. We can see that

$$\int A_j^{uv} \delta s_u \delta s_v d\delta\mathbf{s} = \begin{cases} A_j^{uu}\frac{a_u^2}{3} \prod_k 2a_k & \text{if } u = v \\ 0 & \text{otherwise} \end{cases} \qquad (11)$$

Let $V = \prod_k 2a_k$, and let $\mathbf{\Sigma} \in \mathbb{R}^{2n}$ be a diagonal matrix the entries of which are defined by $\frac{a_u^2}{3}$. We can see that $\int \delta\mathbf{s}^T \mathbf{A}_j \delta\mathbf{s} d\delta\mathbf{s} = V \sum_u A_j^{uu}\frac{a_u^2}{3}$. We can further observe that $\sum_u A_j^{uu}\frac{a_u^2}{3} = Tr(\mathbf{A}_j\mathbf{\Sigma})$. Similarly, we can see that $\int \mathbf{x}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^* d\delta\mathbf{s} = \mathbf{x}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^* V$. Then, the solution to Eq. (9) is given by:

$$Tr(\mathbf{A}_j\mathbf{\Sigma})V + \mathbf{x}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^* V. \qquad (12)$$

The minimisation of Eq. (13) w.r.t. $\mathbf{R}$ has a closed-form as follows:

$$\mathbf{R} = \mathbf{\Sigma}\left(\sum_{j=1}^M \mathbf{J}_j^{*T}\right)\left(\sum_{j=1}^M \mathbf{x}_j^*\mathbf{x}_j^{*T} + \mathbf{J}_j^*\mathbf{\Sigma}\mathbf{J}_j^{*T}\right)^{-1}. \qquad (13)$$

The solution shown in Eq. (13) only accounts for the ground-truth features and their corresponding Jacobians. We can see that, once the data has been sampled, training a new regressor for different chosen limits is straightforward, as it does not require recalculation of perturbations. This implies that, for the Cascaded Regression framework, we only need to sample the images once. This is a huge time saving compared to the sampling-based SDM.

5 CONTINUOUS REGRESSION ON CORRELATED VARIABLES

This section extends Continuous Regression to the space of correlated variables. We will show that it is possible to find a closed-form solution for a general case, by assuming a different measure (Section 5.1). We then study the validity of the Taylor expansion to approximate the input features (Section 5.2). We later introduce our solution into the Cascaded Regression framework, and show how our Cascaded Continuous Regression (CCR) implies a huge training time saving (Section 5.3). We also study the importance of the new solution (Section 5.4). Then, we provide a theoretical link to previous work (Section 5.5) and a geometrical interpretation (Section 5.6).

5.1 Reformulating Continuous Regression

The main problem of existing Functional Regression approaches is that the integral is defined with respect to the Lebesgue measure, which implicitly assumes that an (unnormalised) uniform distribution, in which samples are uncorrelated, is used. Thus, it is not possible to solve it for correlated dimensions. To overcome this limitation, we need to solve the integral with respect to a different measure μ. We can readily see that the proper measure is the *probability measure*, i.e., $\mu = \text{Pr}$. Therefore, the problem becomes

$$\sum_{j=1}^M \int_{\delta\mathbf{s}} \|\delta\mathbf{s} - \mathbf{R}f(\mathbf{I}_j, \mathbf{s}_j^* + \delta\mathbf{s})\|_2^2 d\text{Pr}(\delta\mathbf{s}). \qquad (14)$$

Applying the Riemannian form, we can write Eq. (14) as

$$\sum_{j=1}^M \int_{\delta\mathbf{s}} \|\delta\mathbf{s} - \mathbf{R}f(\mathbf{I}_j, \mathbf{s}_j^* + \delta\mathbf{s})\|_2^2 p(\delta\mathbf{s})d\delta\mathbf{s}, \qquad (15)$$

where now $p(\delta\mathbf{s})$ accounts for the pdf of the sampling distribution. Interestingly, we can see that by formulating the integral with respect to a probability measure, we are actually formulating Continuous Regression by means of the average *expected loss function*, rather than from the classical Functional Regression perspective, which minimises the *empirical loss function*. The expected loss is actually the function from which the empirical loss is derived, for example, in [52], the expected loss function is reduced to the empirical loss by applying a Monte Carlo sampling approximation.

Again, we will approximate the feature space by its first-order Taylor expansion. The integrals are now generally unbounded (they become bounded for e.g. a uniform distribution). Following the steps carried out in Section 4, we expand the features, and group linear, quadratic, and independent terms, with respect to $\delta\mathbf{s}$.

$$\int_{\delta\mathbf{s}} p(\delta\mathbf{s})\|\delta\mathbf{s} - \mathbf{R}f(\mathbf{I}_j, \mathbf{s}_j^* + \delta\mathbf{s})\|_2^2 d\delta\mathbf{s} \approx$$
$$\approx \int_{\delta\mathbf{s}} p(\delta\mathbf{s})\left[\delta\mathbf{s}^T \mathbf{A}_j \delta\mathbf{s} + 2\delta\mathbf{s}^T \mathbf{b}_j + \mathbf{x}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^*\right]d\delta\mathbf{s}, \quad (16)$$

where, as previously, $\mathbf{A}_j = (\mathbb{I} - \mathbf{R}\mathbf{J}_j^*)^T(\mathbb{I} - \mathbf{R}\mathbf{J}_j^*)$ and $\mathbf{b}_j = \mathbf{J}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^* - \mathbf{R}\mathbf{x}_j^*$. We only require the pdf $p(\delta\mathbf{s})$ to be parameterised by its mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$, so that, for any symmetric matrix $\mathbf{A}$, the following properties hold [4]:

$$\int_{\delta\mathbf{s}} p(\delta\mathbf{s})d\delta\mathbf{s} = 1, \quad \int_{\delta\mathbf{s}} \delta\mathbf{s}p(\delta\mathbf{s})d\delta\mathbf{s} = \boldsymbol{\mu},$$
$$\int_{\delta\mathbf{s}} p(\delta\mathbf{s})\delta\mathbf{s}^T\mathbf{A}\delta\mathbf{s}d\delta\mathbf{s} = Tr(\mathbf{A}\boldsymbol{\Sigma}) + \boldsymbol{\mu}^T\mathbf{A}\boldsymbol{\mu}. \quad (17)$$

It is then straightforward to compute the error, for the j-th training example:

$$\int_{\delta\mathbf{s}} p(\delta\mathbf{s})\|\delta\mathbf{s} - \mathbf{R}f(\mathbf{I}_j, \mathbf{s}_j^* + \delta\mathbf{s})\|_2^2 d\delta\mathbf{s} \approx$$
$$Tr(\mathbf{A}_j\boldsymbol{\Sigma}) + \boldsymbol{\mu}^T\mathbf{A}_j\boldsymbol{\mu} + 2\boldsymbol{\mu}^T\mathbf{b}_j + \mathbf{x}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^*. \quad (18)$$

Again, $\mathbf{R}$ is obtained after minimising Eq. (18) w.r.t. $\mathbf{R}$, in which derivatives are obtained as follows:

$$\frac{\partial}{\partial\mathbf{R}}Tr(\mathbf{A}_j\boldsymbol{\Sigma}) = 2\mathbf{R}\mathbf{J}_j^*\boldsymbol{\Sigma}\mathbf{J}_j^{*T} - 2\boldsymbol{\Sigma}\mathbf{J}_j^{*T}$$
$$\frac{\partial}{\partial\mathbf{R}}\boldsymbol{\mu}^T\mathbf{A}_j\boldsymbol{\mu} = 2\mathbf{R}\mathbf{J}_j^*\boldsymbol{\mu}\boldsymbol{\mu}^T\mathbf{J}_j^{*T} - 2\boldsymbol{\mu}\boldsymbol{\mu}^T\mathbf{J}_j^{*T}$$
$$\frac{\partial}{\partial\mathbf{R}}2\boldsymbol{\mu}^T\mathbf{b}_j = 4\mathbf{R}\mathbf{x}_j^*\boldsymbol{\mu}^T\mathbf{J}_j^{*T} - 2\boldsymbol{\mu}\mathbf{x}_j^{*T}$$
$$\frac{\partial}{\partial\mathbf{R}}\mathbf{x}_j^{*T}\mathbf{R}^T\mathbf{R}\mathbf{x}_j^* = 2\mathbf{R}\mathbf{x}_j^*\mathbf{x}_j^{*T}. \quad (19)$$

This leads to the closed-form solution:

$$\mathbf{R} = \left(\sum_{j=1}^{M} \boldsymbol{\mu}\mathbf{x}_j^{*T} + (\boldsymbol{\Sigma} + \boldsymbol{\mu}\boldsymbol{\mu}^T)\mathbf{J}_j^{*T}\right) \cdot$$
$$\left(\sum_{j=1}^{M} \mathbf{x}_j^*\mathbf{x}_j^{*T} + 2\mathbf{x}_j^*\boldsymbol{\mu}^T\mathbf{J}_j^{*T} + \mathbf{J}_j^*(\boldsymbol{\Sigma} + \boldsymbol{\mu}\boldsymbol{\mu}^T)\mathbf{J}_j^{*T}\right)^{-1}. \quad (20)$$

The similarities between the closed form error in Eq. (18) and that of Eq. (12) are clear. If $p(\delta\mathbf{s})$ is defined as a zero-mean uniform distribution with diagonal covariance matrix with entries defined as $\frac{a^2}{3}$ (i.e., defined for the limits $a_1, \ldots, a_{2n}$), then Eq. (20) and Eq. (13) would be the same. It is worth highlighting that the solution does not depend on the actual sampling pdf, but rather on its first and second order moments, i.e. Continuous Regression bypasses the question of which distribution should be used when sampling the data.

The proposed formulation has important computational advantages besides the theoretical differences between Continuous Regression and sampling-based methods like SDM. In particular, once the training images have been sampled,

training a new regressor under a different configuration requires very little computation as opposed to the sampling-based approach. The reason for this is that Continuous Regression does not require the sampling process to be repeated, instead it only changes the sampling statistics. For instance, in the context of Cascaded Regression, given the statistics for each level of the cascade with respect to the ground truth, there is no need to sample the images more than once, which yields a huge time saving. Furthermore, we can see that there is a matrix form that can help compute Eq. (20). Let us introduce the following shorthand notation: $\mathbf{M} = [\boldsymbol{\mu}, \boldsymbol{\Sigma} + \boldsymbol{\mu}\boldsymbol{\mu}^T]$, $\mathbf{B} = \left(\begin{smallmatrix} 1 & \boldsymbol{\mu}^T \\ \boldsymbol{\mu} & \boldsymbol{\Sigma}+\boldsymbol{\mu}\boldsymbol{\mu}^T \end{smallmatrix}\right)$, $\mathbf{D}_j^* = [\mathbf{x}_j^*, \mathbf{J}_j^*]$ and $\bar{\mathbf{D}}^* = [\mathbf{D}_1^*, \ldots, \mathbf{D}_M^*]$. Then:

$$\mathbf{R} = \mathbf{M}\left(\sum_{j=1}^{M} \mathbf{D}_j^*\right)^T \left(\bar{\mathbf{D}}^*\hat{\mathbf{B}}\bar{\mathbf{D}}^{*T}\right)^{-1} \quad (21)$$

where $\hat{\mathbf{B}} = \mathbf{B} \otimes \mathbb{I}_M$. Since $\hat{\mathbf{B}}$ is sparse, computing Eq. (21), once the data is given, is done in a matter of seconds.

5.2 Validity of Taylor expansion

In order to validate the Taylor approximation to represent the image features, we have designed the following experiment. For a given set of initial statistics $\boldsymbol{\Sigma}^l$, a set of random perturbations $f(\mathbf{I}, \mathbf{s}^* + \delta\mathbf{s})$ was created, where $\delta\mathbf{s}$ was drawn from $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}^l)$. Then, we measured the distance between the Taylor approximation to the ground-truth features, as well as to the features extracted from the same image at different locations. Comparing these two distances allows to distinguish whether the Taylor expansion gives a reasonable estimation or a random estimation. The distance between the Taylor approximation at $\delta\mathbf{s}_i$ and the features collected at $\delta\mathbf{s}_j$ (including $i = j$) is given as:

$$dist_{i,j} = \frac{\|f(\mathbf{I}, \mathbf{s}^* + \delta\mathbf{s}_j) - (f(\mathbf{I}, \mathbf{s}^*) + f'(\mathbf{I}, \mathbf{s}^*)\delta\mathbf{s}_i)\|}{\|f(\mathbf{I}, \mathbf{s}^* + \delta\mathbf{s}_j) + (f(\mathbf{I}, \mathbf{s}^*) + f'(\mathbf{I}, \mathbf{s}^*)\delta\mathbf{s}_i)\|}, \quad (22)$$

This distance is expected to increase with $\delta\mathbf{s}$. However, despite this increase, we expect samples to be distinguishable from samples taken from other images.

We have used a random subset of 1000 images from the training set (see Section 7 for further details). We have evaluated the distances at the locations defined by a set of 5 different $\boldsymbol{\Sigma}^l$. More specifically, for $l = 3 \ldots 7$, $\boldsymbol{\Sigma}^l$ is defined as a diagonal matrix the elements of which are set to 2^{l-1} (which represents the variance of the displacement per landmark, in pixels). The results are shown in Fig. 2. It can be seen that even when the variance is fairly high, the approximated samples are still distinguishable from samples collected at other locations (i.e. red lines have lower error than corresponding blue lines). This validates the Taylor approximation for a far long margin.

In order to also check the capabilities of a regressor trained using a Taylor approximation, we have trained two regressors for each of the $\boldsymbol{\Sigma}^l$ described above (one using the sampling-based approach, and one using Continuous Regression). For each of the images, a new set of random locations was generated, and the regressors were used to predict the shape displacements. Results are shown in Fig. 3. It can be seen that both regressors have similar accuracy,
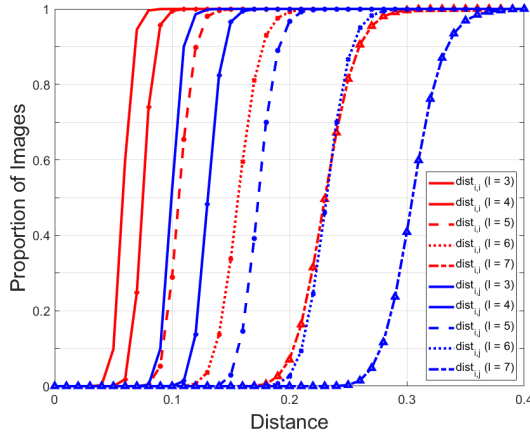
Fig. 2. Results obtained for the Taylor approximation. Red curves correspond to the distances described in Eq. (22), when $i = j$, for a varying set of initial perturbations ranging from $l = 3$ to $l = 7$. Blue curves correspond to the distances described in Eq. (22) when $i \neq j$.

thus validating the assumptions made through the paper. It has to be noted that these validations hold for such a large range thanks to the smoothness and the constrained environment that face tracking generally offers. When the displacement is far from the ground-truth, the distance between the approximation and the sampled data increases, but so it does when collecting the data somewhere else. Therefore, there is still the possibility of training a regressor capable of reducing the error, as shown in Fig. 3.
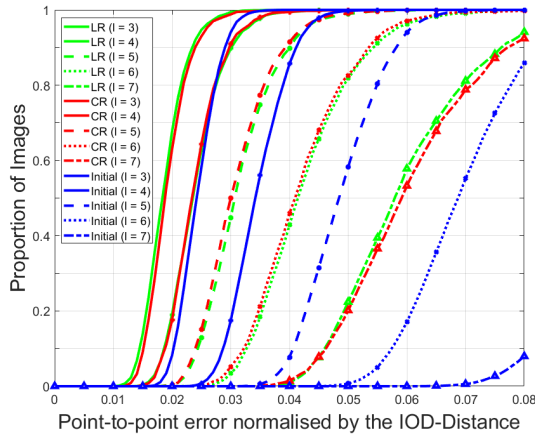

Fig. 3. Results attained by a sampling-based linear regression (green curves) and a regressor trained using Continuous Regression (red curves) for the given initial perturbations (blue curves)

5.3 Cascaded Continuous Regression

We can readily extend Continuous Regression to the Cascaded Regression framework. To do so, we point out that the statistics defined in Eq. (20) correspond to those used in the parallel SDM training proposed in [2], described in Section 3. In this context, we can train each cascade level using the corresponding $\boldsymbol{\mu}^{(l)}$ and $\boldsymbol{\Sigma}^{(l)}$. We coin our method Cascaded Continuous Regression (CCR). The main

advantage of our CCR with respect to SDM is that we only need to sample the training images once, given that Eq. (20) only needs to account for the ground-truth features and their corresponding Jacobians. That is to say, the **sampling process in the CCR training is done only once**.

Denoting the cost of sampling an image as $\mathcal{O}(q)$, for a set of M images, K perturbations, and L cascade levels, the total sampling cost of SDM is $\mathcal{O}(LKMq)$. In our proposed CCR, the sampling is done only once, which is $\mathcal{O}(5qM)$, given that extracting the ground-truth features and Jacobians is $\mathcal{O}(5q)$. Therefore, our CCR training presents a computational advantage, with respect to SDM, of $LK/5$. In our setting, $L = 4$ and $K = 10$, meaning that the Continuous Regression sampling is 8 times faster (in FLOPS). However, we have to note that, if we were to train a different model **for a different set of statistics**, we would not need to sample any data again. This means that, under a different configuration, we only need Eq. (21) to be computed, which can be done in seconds. In this case, we can see that the total sampling cost of SDM training would remain $\mathcal{O}(LKMq)$, while retraining a CCR does not entail any sampling, thus having a null cost. In our configuration, $M \approx 7000$ images, meaning that **the computational saving of retraining a CCR, with respect to SDM, is** $\mathcal{O}(10^6)$.

In the case that no prior SDM has been trained, Algorithm 1 summarises the training of CCR. Given the training images and ground-truth points, CCR needs the data to be pre-computed, by computing all $\mathbf{D}_j^*$. Then, the process of CCR training basically consists of updating the initial statistics $\boldsymbol{\mu}^0$ and $\boldsymbol{\Sigma}^0$ (computed for a training set of annotated videos), to generate a new regressor for each cascade level. The statistics for each level are generated by computing the difference of the outputs of the previous level with respect to the corresponding ground-truth points. For the first level, a set of initial shapes $\mathbf{s}_j^0$ are generated by randomly perturbing the ground-truth with $\boldsymbol{\mu}^0$ and $\boldsymbol{\Sigma}^0$.

Algorithm 1: CCR training

Data: $\{\mathbf{I}_j, \mathbf{s}_j^*, \mathbf{s}_j^0\}_{j=1:M}$, $\boldsymbol{\mu}^0$ and $\boldsymbol{\Sigma}^0$; L levels

1   Pre-compute: Extract $\bar{\mathbf{D}} = \{\mathbf{D}_j^*\}_{j=1:M}$

2   Pre-compute: $\hat{\mathbf{D}} = \left(\sum_{j=1}^{M} \mathbf{D}_j^*\right)^T$

3   **for** $l = 0 : L - 1$ **do**

4   Compute $\mathbf{M}_l$ and $\mathbf{B}_l$

5      $\mathbf{R}_l = \mathbf{M}_l \hat{\mathbf{D}} \left(\bar{\mathbf{D}} \hat{\mathbf{B}}_l \bar{\mathbf{D}}^T\right)^{-1}$

6   Apply $\mathbf{R}_l$ to $\mathbf{s}_j^{(l)}$ to generate $\mathbf{s}_j^{(l+1)}$

7   Compute distances $\delta\mathbf{s}_j = \mathbf{s}_j^{(l+1)} - \mathbf{s}_j^*$

8   Update $\boldsymbol{\mu}^{l+1} = mean(\{\delta\mathbf{s}\})$ and
     $\boldsymbol{\Sigma}^{l+1} = cov(\{\delta\mathbf{s}\})$

9   **end**

5.4 The impact of the data term

To empirically demonstrate the theoretical influence of the data term which allows for correlated variables, we compared the performance of a CCR model trained using the "correlated" solution in Eq. (20) (cor-CCR), with a CCR model trained using the "uncorrelated" version described
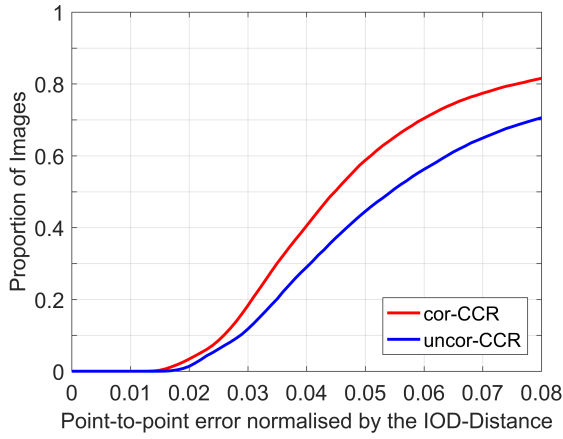
Fig. 4. Cumulative Error Distribution (CED) curve for both the uncor-CR (red) and cor-CR (blue). Results are shown for the 49 points configuration. The contribution of a full covariance matrix is clear

in Section 4 (uncor-CCR). Both models were trained on the Helen database [20], using Algorithm 1. In both cases, the initial training set is generated by perturbing the ground-truth shapes according to the statistics measured at the training partition of 300-VW dataset [41] (see Section 7 for further details). In uncor-CCR, the statistics for each level were forced to have zero-mean and diagonal covariance. Otherwise the two versions are the same. We evaluated both models in the most challenging subset of the 300-VW test partition. Results in Fig. 4 show the Cumulative Error Distribution (CED). It is immediately clear that the contribution of cor-CCR is significant, and thus that shape dimensions are significantly correlated. The result of this experiment clearly demonstrates the importance of solving Continuous Regression in spaces of correlated variables.

5.5   Connection to related work

In this section, we link our work to that of [21]. More specifically, we show that one can first compute the functional covariance in a similar fashion to that of [21], and then link it to the Least-Squares solution. To do so, we first observe that the normal equations shown in Eq. (13) can actually be written by means of a covariance matrix as [49][3]:

$$Cov(X, Y) = Cov(X, X)\mathbf{R}, \tag{23}$$

where we used the term Cov to refer to the *data* covariance, as opposite to $\mathbf{\Sigma}$ referring to the covariance of the sampling pdf. While [21] applies a polytope approximation to estimate the covariance matrix, we rely on the first-order Taylor expansion of the feature space. We have to note that [21] approximates the covariance matrix by considering combinations of training samples, while in our case we want to generate perturbations on the training data, and therefore

3. Without loss of generality, we can assume that the input data and the shape displacements are zero-mean

the polytope approximation would not fit our problem. We can compute the functional covariance matrix as follows:

$$Cov(X, X) = \sum_j \int_{\delta\mathbf{s}} p(\delta\mathbf{s})f(\mathbf{I}, \mathbf{s}_j^* + \delta\mathbf{s})f(\mathbf{I}, \mathbf{s}_j^* + \delta\mathbf{s})^T d\delta\mathbf{s}. \tag{24}$$

If we approximate the input features by its first-order Taylor expansion, we can see that:

$$f(\mathbf{I}, \mathbf{s}_j^* + \delta\mathbf{s})f(\mathbf{I}, \mathbf{s}_j^* + \delta\mathbf{s})^T \approx$$
$$\mathbf{x}_j^*\mathbf{x}_j^{*T} + \mathbf{x}_j^*\delta\mathbf{s}^T\mathbf{J}_j^{*T} + \mathbf{J}_j^*\delta\mathbf{s}\mathbf{x}_j^{*T} + \mathbf{J}_j^*\delta\mathbf{s}\delta\mathbf{s}^T\mathbf{J}_j^{*T}. \tag{25}$$

We can use Eq. (17), and the fact that $\int p(\delta\mathbf{s})\delta\mathbf{s}\delta\mathbf{s}^T = \mathbf{\Sigma} + \boldsymbol{\mu}\boldsymbol{\mu}^T$ to further expand the covariance as:

$$Cov(X, X) = \sum_j \mathbf{x}_j^*\mathbf{x}_j^{*T} + 2\mathbf{x}_j^*\boldsymbol{\mu}^T\mathbf{J}_j^{*T} + \mathbf{J}_j^*(\mathbf{\Sigma} + \boldsymbol{\mu}\boldsymbol{\mu}^T)\mathbf{J}_j^{*T}, \tag{26}$$

which is, exactly, the invertible part of Eq. (18). We can readily see that the covariance can be expressed as:

$$Cov(X, X) = \bar{\mathbf{D}}^*\hat{\mathbf{B}}(\bar{\mathbf{D}}^*)^T. \tag{27}$$

Similarly, we can see that:

$$Cov(X, Y) = \sum_j \int_{\delta\mathbf{s}} \delta\mathbf{s}f(\mathbf{I}, \mathbf{s}_j^* + \delta\mathbf{s})^T d\delta\mathbf{s} \approx$$
$$\approx \sum_j \boldsymbol{\mu}\mathbf{x}_j^{*T} + (\mathbf{\Sigma} + \boldsymbol{\mu}\boldsymbol{\mu}^T)\mathbf{J}_j^{*T}. \tag{28}$$

Obtaining Eq. (20) from Eq. (23) is straightforward given Eq. (28) and Eq. (27). Interestingly, we can take advantage of this approach to generate our PCA models for each cascade level. These models are used to reduce the dimensionality of the feature space, and therefore need to be computed in the perturbed image space. This way, we can easily generate each of the models just by applying an eigendecomposition of the functional covariance of Eq. (27).

5.6   Geometric interpretation

Finally, we want to elaborate on the meaning of the "data term" in the Continuous Regression framework, i.e., the role of a pdf given that an *infinite* set of samples is taken. To do so, we can analyse the solution presented in this paper from the Information Geometry perspective [1], [30], in which a geometrical interpretation can be applied to the probability measure. In the uncorrelated case, i.e. when using a diagonal covariance matrix, the manifold of shape displacements is a $2n$-dimensional parallelepiped, spanned by the Cartesian basis, defined for each of the shape dimensions. When we extend the Continuous Regression to full covariance matrices, we are rotating the manifold of shape displacements according to the eigenvectors of $\mathbf{\Sigma}$, and the origin of the coordinates axes is displaced to $\boldsymbol{\mu}$. Fig. 5 illustrates this geometrical interpretation graphically.

6   INCREMENTAL CASCADED CONTINUOUS REGRESSION

This section introduces the process of updating a CCR model using a new set of images and estimated shapes. We will show that our incremental CCR (which we coin iCCR) has a complexity that yields real-time performance.
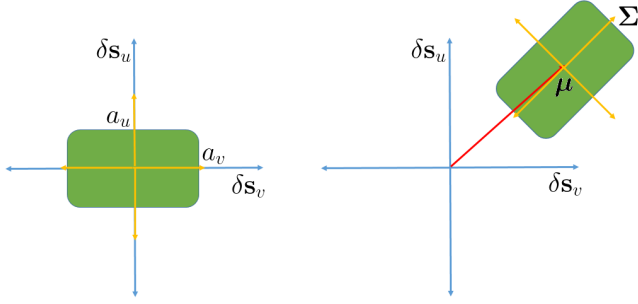
Fig. 5. Difference between classical Functional Regression (**Left**) and CR in correlated variables (**Right**). The green area represents the volume within which samples are taken. Left: a diagonal covariance matrix, with entries defined as $\frac{a^2}{3}$. Right: a full covariance matrix and a non-zero mean vector. The sampling space is translated to the centre, defined by $\boldsymbol{\mu}$, and rotated according to the eigenvectors of $\boldsymbol{\Sigma}$.
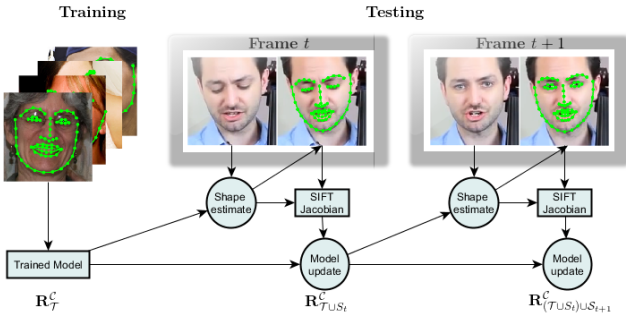

Fig. 6. Overview of our incremental cascaded continuous regression algorithm (iCCR). The originally model $\mathbf{R}_{\mathcal{T}}$ learnt offline is updated with each new frame, thus sequentially adapting to the target face.

To the best of our knowledge, **the iCCR update is the first Cascaded Regression method achieving real-time capabilities**. Results shown in Section 7 demonstrate the importance of incremental learning to achieve state-of-the-art results, thus illustrating the advantages of a real-time incremental learning algorithm.

Once a regressor has been trained, we might want to incorporate new images to the model, without the need of recomputing Eq. (20). This is very important for the task of face tracking, since it has been reported that generic models perform worse than person-specific models. Since training a person-specific model a priori is not an option in most use cases, adding samples to a generic model as the face is being tracked may help the model to track a specific person better in future frames. An overview of the incremental learning procedure is depicted in Fig. 6.

However, we can not retrain the full models by incorporating new images to the original training set, as it would be far too slow, and therefore is not computationally tractable. Some previous works [2], [53] have attempted to incorporate the online learning to their current models, by applying the recursive least squares to a trained model. However, both methods are very slow, and thus are impractical for real-time tracking. It should be noted that the good results reported by [2], [53] rely on the fact that videos were tracked offline, meaning that the update methods could take their time to learn while following frames politely waited for their turn to be tracked. Such affordances are not present when tracking from a live video stream, e.g. a webcam.

Let us devise the incremental learning update rule for the Continuous Regression. To do so, let us assume that we have a regressor $\mathbf{R}_{\mathcal{T}}$, trained using Eq. (21), on a training set $\mathcal{T}$. Also, we denote the covariance matrix for the training data as $\mathbf{V}_{\mathcal{T}} := Cov(X, X) = \bar{\mathbf{D}}_{\mathcal{T}}^* \hat{\mathbf{B}} (\bar{\mathbf{D}}_{\mathcal{T}}^*)^T$. The incremental (also known as online) learning aims to update $\mathbf{R}_{\mathcal{T}}$ with a set of $\mathcal{S}$ new images, for which the ground-truth image features and Jacobians are extracted. Let $\mathbf{D}_{\mathcal{S}}^*$ be the data corresponding to the updating samples. We can define a forgetting factor λ, which sets the influence of a given frame (the lower the λ, the more influence a frame has on the update), and define $\mathbf{W}_\lambda = \lambda \mathbb{I}_{2n+1}$, with $\mathbb{I}_{2n+1}$ the $(2n+1)$-dimensional matrix. Then, the new regressor would be computed as:

$$\mathbf{R}_{\mathcal{T} \cup \mathcal{S}} = \mathbf{M} \left(\sum_{j=1}^{M} \mathbf{D}_j^* + \mathbf{W}_\lambda^{-1} \mathbf{D}_{\mathcal{S}}^* \right)^T (\mathbf{V}_{\mathcal{T} \cup \mathcal{S}})^{-1}, \quad (29)$$

where

$$(\mathbf{V}_{\mathcal{T} \cup \mathcal{S}})^{-1} = \left(\mathbf{V}_{\mathcal{T}} + \mathbf{D}_{\mathcal{S}}^* \mathbf{B} \mathbf{W}_\lambda^{-1} \mathbf{D}_{\mathcal{S}}^{*T} \right)^{-1}. \quad (30)$$

We can apply the Woodbury identity [4] to Eq. (30):

$$\mathbf{V}_{\mathcal{T} \cup \mathcal{S}}^{-1} = \mathbf{V}_{\mathcal{T}}^{-1} -$$

$$\mathbf{V}_{\mathcal{T}}^{-1} \mathbf{D}_{\mathcal{S}}^* \left(\mathbf{W}_\lambda \mathbf{B}^{-1} + \mathbf{D}_{\mathcal{S}}^{*T} \mathbf{V}_{\mathcal{T}}^{-1} \mathbf{D}_{\mathcal{S}}^* \right)^{-1} \mathbf{D}_{\mathcal{S}}^{*T} \mathbf{V}_{\mathcal{T}}^{-1}. \quad (31)$$

This way, obtaining the inverse of the covariance matrix is computationally feasible. During tracking, the set $\mathcal{S}$ consists of a tracked image with its estimated landmarks, assuming these to have been correctly fitted. We can readily see that the sampling cost of the iCCR update is fixed to $\mathcal{O}(5q)$, and does not depend on the number of cascade levels, whereas the incremental SDM (iSDM), requires the sampling process to be carried out for each cascade level. We can see that Eq. (31) needs to compute the inverse of the matrix $\mathbf{W}_\lambda \mathbf{B}^{-1} + \mathbf{D}_{\mathcal{S}}^{*T} \mathbf{V}_{\mathcal{T}}^{-1} \mathbf{D}_{\mathcal{S}}^*$, which is $(2n+1)$-dimensional. This operation has a complexity of $\mathcal{O}((2n+1)^3)$. That is, the computational complexity of inverting that matrix is cubic with respect to the number of points. However, we can alleviate this computation by working with a PDM [10], instead of predicting over the points in a direct way. In a PDM, a shape $\mathbf{s}$ is parameterised in terms of $\mathbf{p} = [\mathbf{q}, \mathbf{c}] \in \mathbb{R}^m$, where $\mathbf{q} \in \mathbb{R}^4$ represents the rigid parameters and $\mathbf{c}$ represents the flexible shape parameters, so that $\mathbf{s} = t_{\mathbf{q}}(\mathbf{s}_0 + \mathbf{B}_s \mathbf{c})$, where t is a Procrustes transformation parameterised by $\mathbf{q}$. $\mathbf{B}_s \in \mathbb{R}^{2n \times m}$ and $\mathbf{s}_0 \in \mathbb{R}^{2n}$ are learnt during training and represent the linear subspace of flexible shape variations. This way, the Continuous Regression formulation would be simply transformed into:

$$\sum_{j=1}^{M} \int_{\delta \mathbf{p}} p(\delta \mathbf{p}) \|\delta \mathbf{p} - \mathbf{R} f(\mathbf{I}_j, \mathbf{p}_j^* + \delta \mathbf{p})\|_2^2 d\delta \mathbf{p}, \quad (32)$$

where the Jacobians now need to be computed with respect to the PDM. To do so, it suffices to apply the chain rule:

$$\frac{\partial f(\mathbf{I}, \mathbf{p} + \delta \mathbf{p})}{\partial \mathbf{p}} = \frac{\partial f(\mathbf{I}, \mathbf{p} + \delta \mathbf{p})}{\partial \mathbf{s}} \frac{\partial \mathbf{s}}{\partial \mathbf{p}}, \quad (33)$$

where $\frac{\partial \mathbf{s}}{\partial \mathbf{p}}$ can be analytically computed from the shape model, and $\frac{\partial f(\mathbf{I}, \mathbf{p} + \delta \mathbf{p})}{\partial \mathbf{s}}$ is computed as in Eq. (7). This

way, the Continuous Regression solution is exactly the same as shown in Eq. (20), with the Jacobians now being computed with respect to the PDM. Now, the matrix $\mathbf{W}_\lambda \mathbf{B}^{-1} + \mathbf{D}_\mathcal{S}^{*T} \mathbf{V}_\mathcal{T}^{-1} \mathbf{D}_\mathcal{S}^*$ is m dimensional, and thus its inversion is $\mathcal{O}(m^3)$. In our experimental set-up, $2n = 132$, and $m = 24$. The computational saving is obvious.

Now, we can further analyse the update complexity, noting that the most expensive step in the updating process is not computing the inverse of $\mathbf{W}_\lambda \mathbf{B}^{-1} + \mathbf{D}_\mathcal{S}^{*T} \mathbf{V}_\mathcal{T}^{-1} \mathbf{D}_\mathcal{S}^*$, but rather computing $\mathbf{D}_\mathcal{S}^T \mathbf{V}_\mathcal{T}^{\mathcal{C}^{-1}}$, which has a complexity cost of $\mathcal{O}(d^2 m)$ (assuming a PDM is used). The overall computational cost of updating a regressor is

$$\mathcal{O}(3md^2) + \mathcal{O}(3m^2 d) + \mathcal{O}(m^3). \tag{34}$$

We compare this updating cost with the **sampling-based** SDM update proposed in [2], which was formulated as:

$$\begin{aligned}
\mathbf{R}_{\mathcal{T} \cup \mathcal{S}} &= \mathbf{R}_\mathcal{T} - \mathbf{R}_\mathcal{T} \mathbf{Q} + \mathbf{Y}_\mathcal{S} \mathbf{X}_\mathcal{S}^T \mathbf{V}_{\mathcal{T} \cup \mathcal{S}} \\
\mathbf{Q} &= \mathbf{X}_\mathcal{S} \mathbf{U} \mathbf{X}_\mathcal{S}^T \mathbf{V}_\mathcal{T} \\
\mathbf{U} &= \left(\mathbb{I}_K + \mathbf{X}_\mathcal{S}^T \mathbf{V}_\mathcal{T} \mathbf{X}_\mathcal{S} \right)^{-1} \\
\mathbf{V}_{\mathcal{T} \cup \mathcal{S}} &= \mathbf{V}_\mathcal{T} - \mathbf{V}_\mathcal{T} \mathbf{Q}
\end{aligned} \tag{35}$$

where $\mathbb{I}_K$ is the K-dimensional identity matrix. We can readily see that the main bottleneck in this update comes from computing $\mathbf{V}_\mathcal{T} \mathbf{Q}$, which is $\mathcal{O}(d^3)$, i.e., the *iCCR update is an order of magnitude faster than the sampling-based SDM update* proposed in [2]. This brings the cost of updating from the ~ 4 seconds reported in [2] down to ~ 0.2 seconds, which is the total update time in our sequential Matlab implementation of iCCR.

7 EXPERIMENTAL RESULTS

To evaluate our proposed approach, we have tested our CCR and iCCR methods on the most extensive benchmark that exists to date: 300VW [41]. To compare our proposed methods against top participants, we develop a fully automated system, which includes automated initialisation, as well as a tool to detect that tracking is lost. This way, we can compare our methods in exactly the same conditions that were set up for the benchmark. We will show that our fully automated system achieves state of the art results, as well as the importance of incremental learning. Code for training and testing is available on the lead author's website.

7.1 Experimental set-up

Test Data: All methods are tested on the 300VW [41] dataset, which, to the best of our knowledge, is the most extensive and recent benchmark in Face Tracking. The dataset consists of 114 videos, each ~ 1 minute long, divided into different categories. 50 videos are used for training, whereas the 64 remaining videos are subdivided into three categories, intended to represent increasingly unconstrained scenarios. Category 1 contains 31 videos recorded in controlled conditions, whereas Category 2 includes 19 videos recorded under severe changes in illumination. Category 3 contains 14 videos captured in totally unconstrained scenarios. All the videos are of a single-person, and have been annotated in a semi-supervised manner using two different methods [8],

[46]. All frames in which the face appears beyond profile-view have been removed from the challenge evaluation, and therefore are not considered in our experiments either.

Error measure: The error measure is that defined for 300VW. It is computed for each frame by dividing the average point-to-point Euclidean error by the inter-ocular distance, understood to be the distance between the two outer eye corners. More specifically, if $\hat{\mathbf{s}}$ is the estimated shape, and $\mathbf{s}^*$ is the ground-truth shape, the RMSE is given as:

$$RMSE = \frac{\sum_{i=1}^n \sqrt{(\hat{x}_i - x_i^*)^2 + (\hat{y}_i - y_i^*)^2}}{d_{outer} n}, \tag{36}$$

where d_{outer} is the Euclidean distance between the points defined for the outer corner of the eyes, measured on the ground-truth. The results are summarised in the form of Cumulative Error Distribution curves (CED), along with the Area Under the Curve (AUC).

7.2 Training

Data: We use data from different datasets of static images to construct our training set. Specifically, we use a total of ~ 7000 images taken from Helen [20], LFPW [3], AFW [60], IBUG [36], and a subset of MultiPIE [18]. We have used the facial landmark annotations provided by the 300 faces in the wild challenge [35], [41], as they are consistent across datasets. Our models are trained for a 66-point configuration. In order to ensure the consistency of the annotations with respect to the test set, the Shape Model is constructed using the training partition of 300VW, and has 20 non-rigid parameters and 4 rigid parameters [27].

Training: As shown in Algorithm 1, the training of CCR starts from a set of given statistics ($\boldsymbol{\mu}^0$ and $\boldsymbol{\Sigma}^0$). In our experiments, the statistics are computed across the training sequences, by computing the differences between consecutive frames. That is to say, $\boldsymbol{\mu}^0$ and $\boldsymbol{\Sigma}^0$ are meant to model how the shapes vary from frame to frame. The main idea of this approach is to replicate a real scenario when "perturbing" the ground-truth. Given that the training set shows smoother variations w.r.t. the test set, we compute the differences between frames separated by two and three time steps. This way, higher displacements are also captured.

Features: We use a HOG [14] implementation, with a block-size of 24 pixels, subdivided in 4 blocks, and 9 bins, resulting in a 9504-dimensional vector. We apply PCA on the output, retaining 2000 dimensions, i.e., $d = 2000$. To generate a PCA model per cascade level, we compute the functional covariance shown in Eq. (27).

Cascade levels: In our experimental set-up, we have fixed the number of Cascade levels to $L = 4$.

7.3 Testing

Initialisation: The tracker is initialised at the beginning of each sequence, as well as each time the tracker is detected to have lost a fitting (see details below). The initialisation utilises the open-source dlib face detection to locate the face bounding box (dlib.net), and then predicts the shape with a Context-based SDM [38]. Then, a single CCR iteration is carried out.
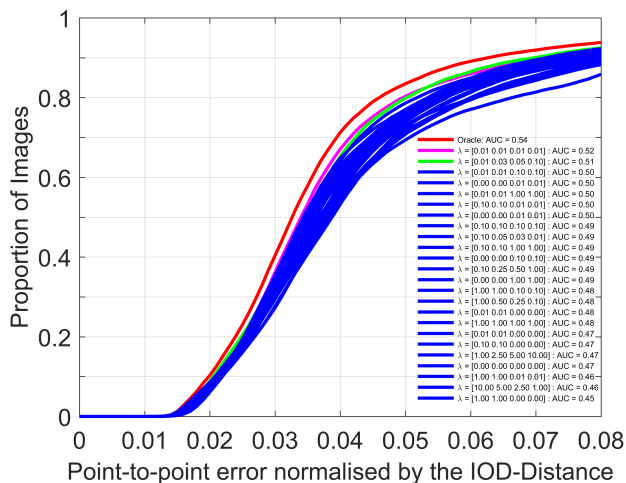
Fig. 7. CED's for a set of different forgetting factors, along with the corresponding AUCs (Category 3). Red curve shows the Oracle. Magenta shows the best factor found for that Category. Green curve shows the results attained by the forgetting factor found on the training set. Blue curves are results attained by other learning rates

Incremental learning update: The incremental learning needs to filter frames to decide whether a fitting is suitable or harmful to be used for the model updates. We use a simple heuristic with a linear SVM, trained to decide whether a particular fitting is "correct", understood as being below a threshold error. In order to avoid redundancy, once a frame has been used to update the models, we consider no more frames within a window of 3 consecutive frames. Given that videos are of a single person, we do not need to perform any face recognition to guarantee that updated models keep tracking the same person with updated, and thus personalised, models. For the iCCR update, we collect person-specific statistics on a sliding window of 50 frames, so that the person-specific variations are encoded in the statistics used to update the model. The forgetting factor λ, has been found through validation on the training set, and is set to $\lambda^1 = 0.01$, $\lambda^2 = 0.025$, $\lambda^3 = 0.05$ and $\lambda^4 = 0.1$, where λ^l is the learning rate for the l-th cascade level, and the lower the λ, the more influence a frame has on the update. We evaluate the sensitiveness of iCCR w.r.t. λ in Section 7.4.
Loss of tracking detection: Sometimes the tracker fails to track a face. This occurs in particular when there are significant occlusions or head-poses beyond profile. In these cases, we need to detect that the tracker has lost a face, in order to reinitialise it in the next frame from the face detector output. We used the same SVM learnt for the incremental learning update to also detect whether the tracker is lost or not, by empirically selecting a suitable threshold for the score returned by the SVM model.

7.4 Influence of the forgetting factor

The forgetting factor has been found through validation on the training partition, and set up as shown above. In general, it has been found that the top levels of the cascade are less sensitive to an error in the update, and therefore a smaller factor can be used (meaning a higher impact of the given frame). However, for the last levels of the cascade, a low

factor can have a big impact if the frame to be updated is not "entirely" correct. Since the score of an SVM is used to detect whether a frame is suitable or not for the update, it is likely that eventually a partially correct frame will be used. To measure how well the found values generalise to the test set, we have included an experiment in Category 3 showing the performance of other forgetting factors. The overall results are shown in Fig. 7. It can be seen that other values would produce similar results. In addition, we include the error given by an "Oracle" (red curve), understood as the aggregated error for a scenario in which the best forgetting factor is chosen for each video independently. That is to say, red curve represents the results that would be attained if we could choose the best forgetting factor per video (fixed for the whole video). It can be seen that the red curve does not show a significant improvement over the best blue curve, meaning that in general the learning rate might have a limited impact in the results.

7.5 Equivalence with SDM

In order to demonstrate the performance of our CCR method compared to SDM, we have trained an SDM model under the same conditions as for CCR. More specifically, we have used the same training set, and the same statistics measured for the first level. The training was then done in a sequential manner, in a similar fashion to that of Algorithm 1. The number of cascade levels was again set to 4. The PCA models were directly computed over the sampled images, and comprise 2000 components. The SDM model was tested in the test partition of 300-VW under exactly the same conditions than those of CCR and iCCR. The results attained by our SDM model are shown in Fig. 8.

7.6 Comparison with state of the art

We evaluate our fully automated tracking system in exactly the same conditions as those used for the 300VW benchmark. We compare our method with respect to the top two 300VW results [50], [56]. We report results for the configurations of 49 points. In order to compare our method against [50], [56], we asked these authors for their results, which were kindly provided. For the fairest comparison with respect to the 300VW evaluation, we removed the frames that were not considered for the evaluation (i.e. those including faces beyond profile). To show the benefits of incremental learning, we include the results of both our CCR and iCCR systems in all curves, as well as our trained SDM. CED results are shown in Fig. 8, and the AUC for each method is shown in Table 1. It is interesting to remark that the methods we compare with made contributions to the initialisation of the tracker in each frame, or to failure detection. In our case, even with a simple heuristic for both initialisation and failure detection, we are capable of attaining comparable or even superior performance over [50], [56]. The initialisation for each frame in our tracking system is simply the output of the previous frame. We do not include any multi-view models or progressive initialisation, but instead ensure that the statistics used to generate our models generalise well to unseen scenarios.

It is unlikely that algorithms with slower running times than iCCR, such as [50], [56], would attain their high accuracies if forced to process challenging videos in real-time. This

has important consequences for tracking streaming video, in which frames must be tracked as fast as they arrive in order to make use of the limited movement of facial points between adjacent frames, something that is beneficial to trackers of any ilk.

It is worth highlighting that CCR and SDM models were reinitialised in $\sim 0.51\%$ and $\sim 0.56\%$ of total frames (> 121000), respectively, whilst the iCCR was reinitialised in only $\sim 0.34\%$ of the frames. This illustrates the stability of our models, resulting in a simple yet effective tracking system. Specially remarkable is the importance of the Incremental Learning in challenging sequences, such as those shown in Category 3, to achieve state of the art results. As shown throughout the paper, the complexity of our incremental learning allows for real-time implementation, something that could not be achieved by previous works on Cascaded Regression. Some visual examples are depicted in Fig. 9. Code and videos can be found at http://continuousregression.wordpress.com.

Method	Category 1	Category 2	Category 3
iCCR	0.5978	0.5918	**0.5141**
CCR	0.5657	0.5539	0.4410
SDM	0.5617	0.5522	0.4380
Yang et al. [56]	**0.5981**	0.6025	0.4996
Xiao et al. [50]	0.5814	**0.6093**	0.4865

TABLE 1
AUC for 49 points configuration for the different categories.

8 CONCLUSION

We have proposed a novel formulation for the problem of Continuous Least Squares in spaces of correlated variables, namely Continuous Regression. We have shown the relation that exists between the sampling distribution and the space within which samples are taken, and have proposed a solution for general feature descriptors. We then incorporated the Continuous Regression within the Cascaded Regression framework, and showed its performance is similar to sampling-based methods. We devised an incremental learning approach using Continuous Regression, and showed its complexity allows for real-time performance. Results demonstrate the state-of-the-art performance in challenging sequences, whilst being able to run in real-time. We also want to note that while we tackle the facial landmark tracking problem, cascaded regression has also been applied to a wider range of problems such as pose estimation [16], model-free tracking [48] or object localisation [58], thus meaning the contributions of this paper can be extended to other problems. Future work will include also the optimisation of the learning process, given that a better classifier than a SVM would probably select better frames, and we will also investigate the effect of the forgetting factor under different configurations.
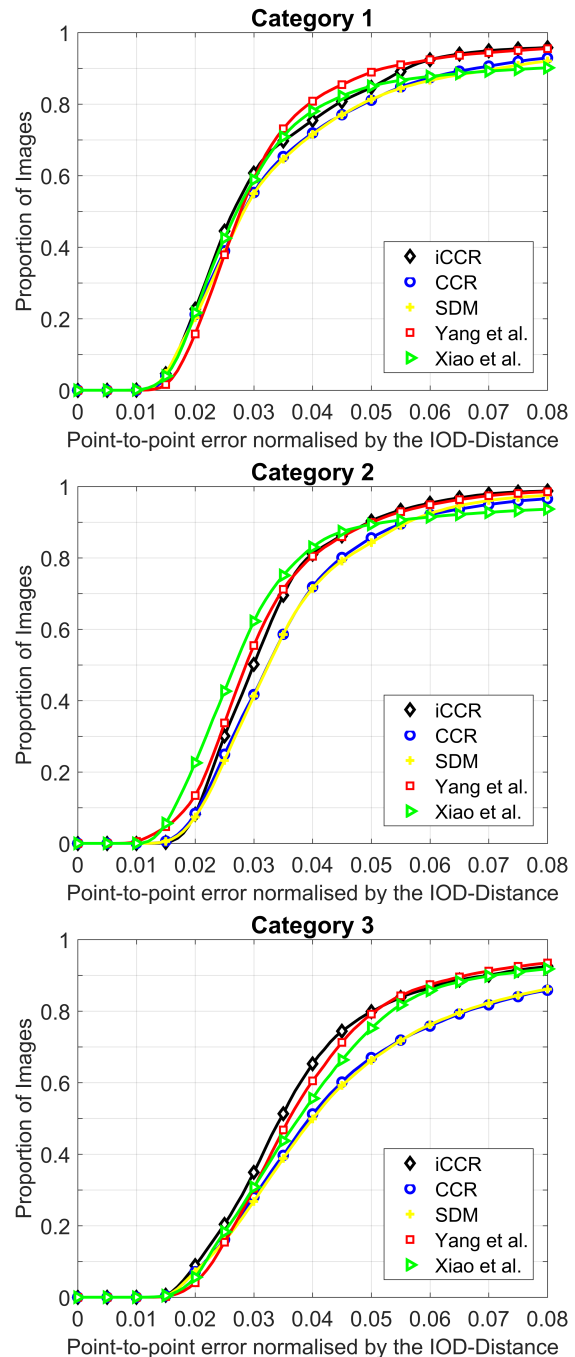
Fig. 8. CED's for the 49-points configuration. Better seen in colour

REFERENCES

[1] Amari, S.: Differential-geometrical methods in statistics Springer Science and Business Media (1985)
[2] Asthana, A., Zafeiriou, S., Cheng, S., Pantic, M.: Incremental face alignment in the wild. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2014)
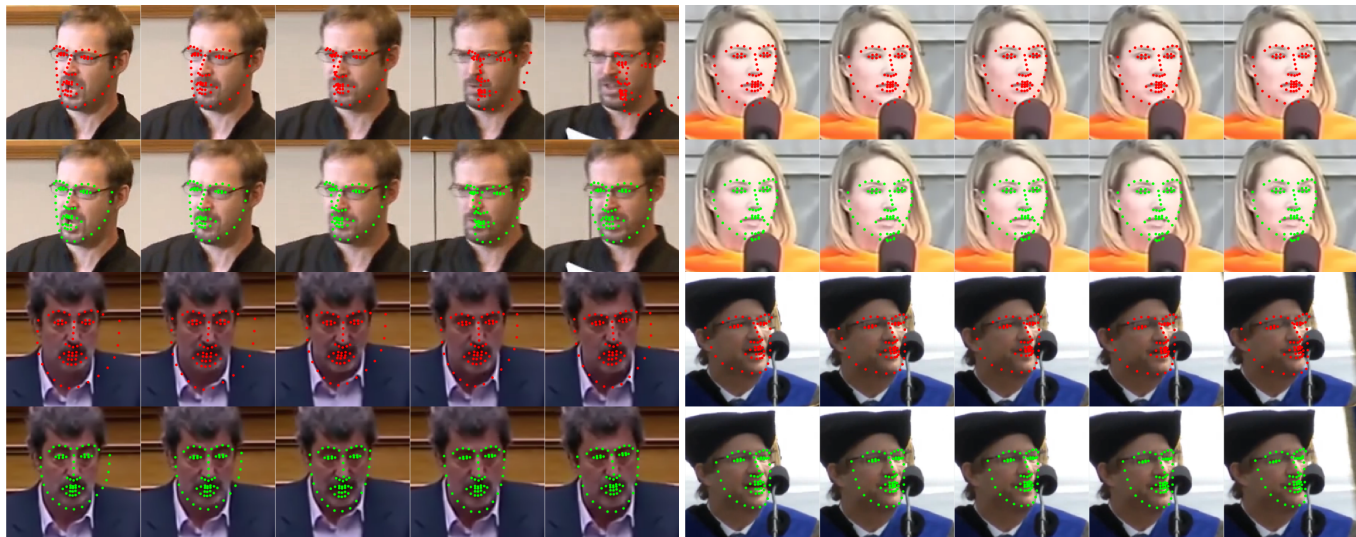
Fig. 9. Visual examples from Category 3 on 300VW. Upper rows show the CCR fitting (red points), whereas the lower rows show the iCCR fitting (green points). The sequences on the left show a clear case where iCCR outperforms CCR thanks to its incremental learning capabilities (fast movement in the upper sequence, low resolution in the sequence at the bottom). The sequences on the right show a failure case where the learned features are not capable to outperform CCR. All qualitative videos and results can be found at https://continuousregression.wordpress.com/code.

[3] Belhumeur, P.N., Jacobs, D.W., Kriegman, D.J., Kumar, N.: Localizing parts of faces using a consensus of exemplars. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2011) 545–552

[4] Brookes, M.: The matrix reference manual (2011) [online] http://www.ee.imperial.ac.uk/hp/staff/dmb/matrix/intro.html (2011)

[5] Bulat, A., Tzimiropoulos, G.: Convolutional aggregation of local evidence for large pose face alignment. In: British Machine Vision Conf. (2016)

[6] Cao, X., Wei, Y., Wen, F., Sun, J.: Face alignment by explicit shape regression. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2012) 2887–2894

[7] Cao, X., Wei, Y., Wen, F., Sun, J.: Face alignment by explicit shape regression. Int'l Journal of Computer Vision **107**(2) (2014) 177–190

[8] Chrysos, G.S., Antonakos, E., Zafeiriou, S., Snape, P.: Offline deformable face tracking in arbitrary videos. In: Int'l Conf. Computer Vision - Workshop. (2015)

[9] Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models-their training and application. Comp. Vision and Image Understanding **61**(1) (1995) 38–59

[10] Cootes, T.F., Taylor, C.J.: Statistical models of appearance for computer vision (2004) Technical Report. University of Manchester

[11] Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. Trans. on Pattern Analysis and Machine Intelligence **23**(6) (2001) 681–685

[12] Cootes, T.F., Ionita, M.C., Lindner, C., Sauer, P.: Robust and accurate shape model fitting using random forest regression voting. In: European Conf. on Computer Vision. (2012) 278–291

[13] Cristinacce, D., Cootes, T.: Feature detection and tracking with constrained local models. In: British Machine Vision Conf. (2006) 929–938

[14] Dalal, N., Triggs, B.: Histograms of Oriented Gradients for human detection In: IEEE Conf. on Computer Vision and Pattern Recognition. (2005)

[15] Dhall, A., Goecke, R., Joshi, J., Sikka, K., Gedeon, T.: Emotion recognition in the wild challenge 2014: Baseline, data and protocol. In: Int'l Conf. on Multimodal Interaction. (2014) 461–466

[16] Dollár, P., Welinder, P., Perona, P.: Cascaded pose regression. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2010) 1078–1085

[17] Gross, R., Matthews, I., Baker, S.: Generic vs. person specific active appearance models. Image and Vision Computing **23**(11) (2005) 1080–1093

[18] Gross, R., Matthews, I., Cohn, J., Kanade, T., Baker, S.: Multi-pie. Image and Vision Computing **28**(5) (2010) 807–813

[19] Kazemi, V., Sullivan, J.: One millisecond face alignment with an ensemble of regression trees. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2014)

[20] Le, V., Brandt, J., Lin, Z., Bourdev, L.D., Huang, T.S.: Interactive facial feature localization. In: European Conf. on Computer Vision. (2012) 679–692

[21] Levin, A., Shashua A.: Principal Component Analysis Over Continuous Subspaces and Intersection of Half-spaces In: European Conference on Computer Vision (2002)

[22] Liu, L., Hu, J., Zhang, S., Deng, W.: Extended supervised descent method for robust face alignment In: Asian Conf. on Computer Vision Workshops. (2014) 71–84

[23] Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**(2) (2004) 91–110

[24] Martinez, B., Valstar, M., Binefa, X., Pantic, M.: Local Evidence Aggregation for Regression Based Facial Point Detection IEEE Trans.Pattern Analysis and Machine Intelligence **35**(5) (2013) 1149–1163

[25] Martins, P., Henriques, J.F., Caseiro, R., Batista, J.: Bayesian constrained local models revisited IEEE Trans. Pattern Analysis and Machine Intelligence **38**(4) (2016) 704–716

[26] Marx, B.D., Eilers, P.H.C.: Generalized Linear Regression on Sampled Signals and Curves: A P-spline approach Techno. **41**(1) 1–13

[27] Matthews, I., Baker, S.: Active appearance models revisited. Int'l Journal of Computer Vision **60**(2) (2004) 135–164

[28] Morris, J.S.: Functional Regression Annual Review of Statistics and Its Applications **2** (2015) 321–359

[29] Peng, X., Feris, R., Wang, X., Metaxas, D.: A recurrent encoder-decoder for sequential face alignment In: European Conf. on Computer Vision. (2016) 38–56

[30] Pennec, X.: Intrinsic statistics on Riemannian Manifolds: basic tools for geometrical measurements J. Math Imagin Vis. **25**, 127–154. (2006)

[31] Quak, E., Sivakumar, N., Ward, J.D.: Least squares approximation by radial functions SIAM J. Math. Anal. **24**(4) (1993) 1043–1066

[32] Ramsay, J., Silverman, B.: Functional Data Analysis Springer (2005).

[33] Ratliffe, S.J., Heller, G.Z., Leader, L.R. : Functional data analysis with application to periodically stimulated foetal heart rate data. I: Functional regression. Stat. in Med. **21**(8) 1115–1127

[34] Ren, S., Cao, X., Wei, Y., Sun, J.: Face alignment at 3000 FPS via regressing local binary features. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2014) 1685–1692

[35] Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: 300 faces in-the-wild challenge: the first facial landmark localization challenge. In: Int'l Conf. Computer Vision - Workshop. (2013)

[36] Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: A semi-automatic methodology for facial landmark annotation. In: IEEE

Conf. on Computer Vision and Pattern Recognition - Workshops. (2013)

[37] Sánchez-Lozano, E., De la Torre, F., González-Jiménez, D.: Continuous regression for non-rigid image alignment. In: European Conf. on Computer Vision. (2012) 250–263

[38] Sánchez-Lozano, E., Martinez, B., Valstar, M.F.: Cascaded regression with sparsified feature covariance matrix for facial landmark detection Pattern Recognition Letters 73 (2016) 19–25

[39] Sánchez-Lozano, E., Martinez, B., Tzimiropoulos, G., Valstar, M.: Cascaded continuous regression for real-time incremental face tracking In: European Conf. on Computer Vision. (2016) 645–661

[40] Saragih, J.M., Lucey, S., Cohn, J.F.: Deformable model fitting by regularized landmark mean-shift. Int'l Journal of Computer Vision 91(2) (2011) 200–215

[41] Shen, J., Zafeiriou, S., Chrysos, G.S., Kossaifi, J., Tzimiropoulos, G., Pantic, M.: The first facial landmark tracking in-the-wild challenge: Benchmark and results. In: Int'l Conf. Computer Vision - Workshop. (2015)

[42] Tresadern, P., Sauer, P., Cootes, T.: Additive update predictors in Active Appearance Models In: British Machine Vision Conf. (2010)

[43] Tresadern, P., Ionita, M., Cootes, T.: Real-time facial feature tracking on a mobile device. Int'l Journal of Computer Vision 96(3) (2012) 280–289

[44] Tzimiropoulos, G., Pantic, M.: Optimization problems for fast AAM fitting in-the-wild. In: Int'l Conf. Computer Vision. (2013)

[45] Tzimiropoulos, G., Pantic, M.: Gauss-newton deformable part models for face alignment in-the-wild. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2014) 1851–1858

[46] Tzimiropoulos, G.: Project-out cascaded regression with an application to face alignment. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2015) 3659–3667

[47] Tzimiropoulos, G., Pantic, M.: Fast algorithms for fitting Active Appearance Models to unconstrained images Int'l Journal of Computer Vision (in press)

[48] Wang, X., Valstar, M., Martinez, B., Khan, M.H., Pridmore, T.: Tric-track: Tracking by regression with incrementally learned cascades. In: Int'l Conf. Computer Vision. (2015)

[49] Wang, J.L., Chiou, J.M., Muller, H.G.: Functional data analysis. In: Annual Review of Statistics and its Application (2016) 3:1 257–295

[50] Xiao, S., Yan, S., Kassim, A.: Facial landmark detection via progressive initialization. In: ICCV-W 1st Facial Landmark Tracking in-the-wild challenge and workshop (2015)

[51] Xiao, S., Feng, J., Xing, J., Lai, H., Yan, S., Kassim, A.: Robust facial landmark detection via recurrent attentive-refinement networks In: European Conf. on Computer Vision. (2016)

[52] Xiong, X., De la Torre, F.: Supervised descent method and its applications to face alignment. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2013)

[53] Xiong, X., la Torre, F.D.: Supervised descent method for solving nonlinear least squares problems in computer vision. CoRR abs/1405.0601 (2014)

[54] Xiong, X., De la Torre, F.: Global supervised descent method In: IEEE Conf. on Computer Vision and Pattern Recognition (2015)

[55] Yan, J., Lei, Z., Yi, D., Li, S.: Learn to combine multiple hypotheses for accurate face alignment. In: Int'l Conf. Computer Vision - Workshop. (2013) 392–396

[56] Yang, J., Deng, J., Zhang, K., Liu, Q.: Facial shape tracking via spatio-temporal cascade shape regression In: ICCV-W, 1st Facial Landmark Tracking in-the-wild challenge and workshop (2015)

[57] Yao, F., Muller, H.G., Wang, J.L.: Functional Linear regression analysis for longitudinal data The Annals of Statistics 33(2) (2005) 2873–2903

[58] Yan, J., Lei, Z., Yang, Y., Li, S.: Stacked deformable part model with shape regression for object part localization. In: European Conf. on Computer Vision. (2014) 568–583

[59] Zhou, S., Krueger, V., Chellappa, R.: Probabilistic recognition of human faces from video. Comp. Vision and Image Understanding 91(12) (2003) 214 – 245

[60] Zhu, X., Ramanan, D.: Face detection, pose estimation, and landmark localization in the wild. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2012) 2879–2886
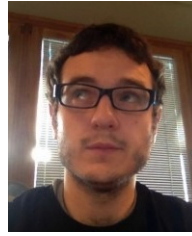
Enrique Sánchez-Lozano is a Research Fellow in the Computer Vision Lab at the University of Nottingham, where he received his PhD in Computer Science (2017). He received his M.Eng. in Telecommunication Engineering (2009) and M.Sc. in Signal Theory and Communications (2011) from the University of Vigo (Spain). He spent 6 months as a visiting scholar at the Human Sensing Lab, with Fernando De la Torre, between 2011 and 2012. He was given a Fundación Pedro Barrié Scholarship to do a 5 month research stay at the Computer Vision Lab (CVL), University of Nottingham, with Michel Valstar, between 2013 and 2014.
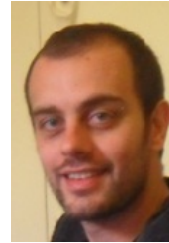
Georgios (Yorgos) Tzimiropoulos received the M.Sc. and Ph.D. degrees in Signal Processing and Computer Vision from Imperial College London, U.K. He is Assistant Professor with the School of Computer Science at the University of Nottingham, U.K. Prior to this, he was Assistant Professor with the University of Lincoln, U.K. and Senior Researcher in the iBUG group, Department of Computing, Imperial College London. He is currently Associate Editor of the Image and Vision Computing Journal. His main research interests are in the areas of face and object recognition, alignment and tracking, and facial expression analysis.
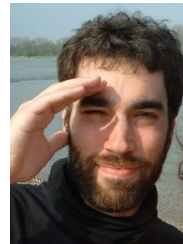
Brais Martinez (M'10) is a Research Fellow with the Computer Vision Laboratory at the University of Nottingham, UK. He has previously obtained his PhD in Computer Science from the Universitat Autonoma de Barcelona, and held a postdoctoral research position at the Imperial College London. He has worked on the topics of Computer Vision and Pattern Recognition, mainly focusing on model-free tracking and on face analysis problems such as face alignment, face detection and facial expression recognition, publishing his findings on these topics on a variety of authoritative venues such as CVPR, ICCV, ECCV or TPAMI.

Fernando De la Torre is an Associate Research Professor in the Robotics Institute at Carnegie Mellon University. He received his B.Sc. degree in Telecommunications, as well as his M.Sc. and Ph. D degrees in Electronic Engineering from La Salle School of Engineering at Ramon Llull University, Barcelona, Spain in 1994, 1996, and 2002, respectively. His research interests are in the fields of Computer Vision and Machine Learning. Currently, he is directing the Component Analysis Laboratory (http://ca.cs.cmu.edu) and the Human Sensing Laboratory (http://humansensing.cs.cmu.edu) at Carnegie Mellon University. He has over 150 publications in referred journals and conferences and is Associate Editor at IEEE TPAMI. He has organized and co-organized several workshops and has given tutorials at international conferences on component analysis.

Michel Valstar (M'09) is an Associate Professor in the Computer Vision and Mixed Reality Labs at the University of Nottingham. He received his masters degree in Electrical Engineering at Delft University of Technology in 2005 and his PhD in computer science with the intelligent Behaviour Understanding Group (iBUG) at Imperial College London in 2008. His main interest is in automatic recognition of human behaviour. In 2011 he was the main organiser of the first facial expression recognition challenge, FERA 2011. In 2007 he won the BCS British Machine Intelligence Prize for part of his PhD work. He has published technical papers at authoritative conferences including CVPR, ICCV and SMC-B and his work has received popular press coverage in New Scientist and on BBC Radio.