INCREMENTAL B-SPLINE DEFORMATION MODEL FOR GEOMETRIC GRAPH MATCHING

Miguel Amável Pinheiro, Jan Kybic

Faculty of Electrical Engineering, Czech Technical University in Prague

ABSTRACT

We propose a B-spline deformation model, which can be efficiently updated from sequential measurements. Our incremental update is based on the Kalman filtering, is optimal in the least squares sense, and is very fast, an order of magnitude faster than the direct methods, while converging to the same solution. While this method can be applied to any multidimensional scattered data interpolation task, our main application is a geometric graph matching algorithm, which is used for registering large images from the biomedical domain or remote sensing, based on matching linear structures such as blood vessels or roads. The B-spline transformation model needs to be updated incrementally as more points are added to the match and using the proposed Kalman-like update yields substantial speed gains over the previously used bi-Lipschitz model.

Index Terms— B-splines, Kalman filter, image registration, geometrical transformation, graph matching.

1. INTRODUCTION

Generalized geometric graphs consist of vertices in space connected by curves and can represent many real world structures such as roads, rivers, blood vessels (Fig. 3), or neuronal fibers. We have previously described a matching method for such graphs [1], with applications to image registration. The matching is constructed by incrementally adding pairs of matching edges and corresponding vertices. One of the key components is a deformation model which is used to determine whether a new pair being considered for correspondence is geometrically compatible with the current partial matching, i.e. if there is a smooth geometrical transformation consistent with the matching. In each iteration, the model needs to be updated. Since many possible partial matching are considered, the update and evaluation need to be fast in order not to become a bottleneck. In the preceding version [1], only a simple bi-Lipschitz model was used, constraining the maximum relative change of distances to a given ε. Such a model is fast but not very discriminative.

This work was supported by the Czech Science Foundation project 17-15361S and by the Ph.D. grant SFRH/BD/77134/2011 of the Portuguese Fundação para a Ciência e Tecnologia.

To improve the matching efficiency, we propose to use the classical uniform B-splines deformation model [2], which can be evaluated quickly [3]. However, finding B-splines coefficients from scattered data requires solving a linear system of equations, which would be prohibitively expensive to do in each iteration. We therefore propose to update the coefficients incrementally using a Kalman filter approach, with a fast, constant complexity update. The procedure converges quickly to the exact solution. A similar approach has been used to recover a 3D position [4] and shape [5].

More generally, we are solving a scattered data approximation problem with sequentially provided measurements. A scattered data interpolation or approximation problem can be solved by numerous methods including kernel regression, Gaussian process regression (also known as kriging), or thin plate splines [6]. Incrementally updating a scattered point implementation was described for a quadratic approximation [7], thin-plate splines [8], or a triangulation-based approximation [9]. Closely related work includes estimating non-uniform B-spline coefficients using a Kalman filter for 1D uniform measurements [10], for 1D sequences with a sliding time window [11], or for multivariate non-uniform B-splines on a mesh [12].

2. METHOD

2.1. Standard approach

Let us first define and solve the scalar problem of recovering a continuous function $f\colon \mathbb{R}^D \to \mathbb{R}$, represented as a linear combination of $n = n_1 n_2 \dots n_D$ B-spline basis functions φ_j with coefficients c_j:

$$f(\mathbf{x}) = \sum_{j=1}^{n} c_j \varphi_j(\mathbf{x}) = \boldsymbol{\varphi}^T(\mathbf{x})\, \mathbf{c} \tag{1}$$

$$\text{with} \quad \varphi_j(\mathbf{x}) = \prod_{k=1}^{D} \beta^3\Big(\frac{x_k}{h_k} - z_{j,k}\Big) \tag{2}$$

where β^3 are 1D cubic B-splines [3] placed on a uniform grid at positions $\mathbf{z}_j$ with spacing h_k. Given a set of m scattered data points $\mathbf{x}_i \in \mathbb{R}^D$ with $i = 1, \dots, m$, we shall assume to observe $y_i = f(\mathbf{x}_i) + r_i$, where r_i is an i.i.d. Gaussian measurement noise with standard deviation σ_r. In a vectorized

Preprint submitted to 2018 IEEE International Symposium on Biomedical Imaging.
Received October 22, 2017.

form, $\mathbf{y} = \Phi\mathbf{c} + \mathbf{r}$, with $R = I\sigma_r^2$ being the covariance of $\mathbf{r}$. Let us further assume that the a priori p.d.f. of the coefficients $\mathbf{c}$ is also Gaussian, with the mean $\bar{\mathbf{c}}$ and covariance $P = I\sigma_{\mathbf{c}}^2$, uncorrelated with the measurement noise $\mathbf{r}$.

Given measurements $\mathcal{M}^t = \{(y_i, \mathbf{x}_i); \quad i = 1, \ldots, t\}$ at time t, we are interested in calculating the maximum a posteriori probability (MAP) estimate $\hat{\mathbf{c}}^t$ of $\mathbf{c}$. From the Bayes law,

$$\hat{\mathbf{c}}_t = \arg\max_{\mathbf{c}} p(\mathbf{c}|\mathcal{M}^t) = \arg\max_{\mathbf{c}} p(\mathbf{y}|\mathbf{c}, X)p(\mathbf{c}), \quad (3)$$

where X represents all $\mathbf{x}_i$. Taking the negative logarithm leads to $\hat{\mathbf{c}}_t$ being a minimizer of

$$(\mathbf{y} - \Phi\mathbf{c})^T R^{-1}(\mathbf{y} - \Phi\mathbf{c}) + (\mathbf{c} - \bar{\mathbf{c}})^T P^{-1}(\mathbf{c} - \bar{\mathbf{c}}) \quad (4)$$

The first part of (4) corresponds to a standard least squares fitting, while the second term introduces a quadratic regularization for the spline coefficients towards $\bar{\mathbf{c}}$. If no a priori information is known, we use large $\sigma_{\mathbf{c}}$ to avoid bias. Setting the derivative to zero yields the normal equations

$$\left(\Phi^T R^{-1}\Phi + P^{-1}\right)\hat{\mathbf{c}}_t = \Phi^T R^{-1}\mathbf{y} + P^{-1}\bar{\mathbf{c}} \quad (5)$$

In the incremental setting, the normal equations (5) need to be solved for each t, which can be costly. Alternatively, the normal equations can be solved iteratively, using for example the conjugated gradients method, initialized from the previous estimate $\hat{\mathbf{c}}^{t-1}$.

2.2. Kalman update

We use a Kalman-like approach [4, 5] to update the a posteriori Gaussian p.d.f. $p(\mathbf{c}|\mathcal{M}^t) = \mathcal{N}(\hat{\mathbf{c}}_t, P_t)$ given the previous estimate $p(\mathbf{c}|\mathcal{M}^{t-1})$ and the current measurement $(y_t, \mathbf{x}_t)$, with initial values $P_0 = P$ and $\mathbf{c}_0 = \bar{\mathbf{c}}$. The principal differences with respect to a standard Kalman filter are that (i) the (unknown) system state $\mathbf{c}$ is static, with no evolution and no process noise, while (ii) the measurement matrix $\boldsymbol{\varphi}^T(\mathbf{x}_t)$ varies with t. The update procedure to incorporate a measurement $(y_t, \mathbf{x}_t)$ is as follows:

1. Calculate the (scalar) measurement residual

$$\tilde{y}_t = y_t - \boldsymbol{\varphi}^T(\mathbf{x}_t)\,\hat{\mathbf{c}}_{t-1} \quad (6)$$

2. Calculate the scalar variance s_t of the residual

$$s_t = \boldsymbol{\varphi}^T(\mathbf{x}_t)\,P_{t-1}\boldsymbol{\varphi}(\mathbf{x}_t) + \sigma_r^2 \quad (7)$$

3. Calculate the Kalman gain

$$\mathbf{k}_t = P_{t-1}\boldsymbol{\varphi}(\mathbf{x}_t)s_t^{-1} \quad (8)$$

4. Calculate the updated a posteriori mean and covariance

$$\hat{\mathbf{c}}_t = \hat{\mathbf{c}}_{t-1} + \mathbf{k}_t\tilde{y}_t \quad (9)$$

$$P_t = \left(I - \mathbf{k}_t\boldsymbol{\varphi}^T(\mathbf{x}_t)\right)P_{t-1} \quad (10)$$

These steps are repeated for each new match received. The update procedure is very efficient because of the compact support of the B-splines, leading to sparse $\boldsymbol{\varphi}(\mathbf{x}_t)$; no system of linear equations needs to be solved. Because of the Gaussianity, the mean $\hat{\mathbf{c}}_t$ from (9) is equal to the desired MAP estimate (3).

Fig. 1 shows an example of the approximation f determined by the estimate $\hat{\mathbf{c}}^t$ for a 1D function approximation task with uniformly distributed random measurements. Observe that with increasing t the approximation converges to the function being approximated.

2.3. Vector extension

Both the direct method (5) and the Kalman update method (Section 2.2) can be easily extended to vector measurements. In our case $\mathbf{f}\colon \mathbb{R}^D \to \mathbb{R}^D$ represents the geometrical transformation, $\mathbf{f}(\mathbf{x}) = \boldsymbol{\varphi}^T(\mathbf{x}_i)\ \mathbf{C}$, where $\mathbf{C}$ of size $n \times D$ contains the B-spline coefficients. For simplicity, let us consider that both the measurement noise $\mathbf{r}$ and the a priori probability $p(\mathbf{C})$ are independent between coordinates. Then the estimation can be performed for each coordinate separately. If in addition the initial covariance matrices R and P are the same for all coordinates (isotropic) then also all $P_t, \mathbf{k}_t$ are the same and the calculations (7), (8), and (10) can be shared.

2.4. Deformation model

We now have all ingredients of our deformation model. Given the initial guess of the transformation, e.g. identity, we find $\hat{\mathbf{C}}_0$ by least squares fitting, while $\sigma_{\mathbf{c}}$ corresponds to the expected difference between the initial guess and the final transformation. During the matching process, our current belief about $\mathbf{C}$ is described by the mean $\hat{\mathbf{C}}_t$ and correlation P_t (the same for all coordinates). When a new pair $(\mathbf{x}_t, \mathbf{y}_t)$ comes, we calculate the prediction error

$$\tilde{\mathbf{y}}_t = \mathbf{y}_t - \boldsymbol{\varphi}^T(\mathbf{x}_t)\hat{\mathbf{C}}_{t-1} \quad (11)$$

as in (6). The covariance matrix of the residual is diagonal, $\Sigma_{\mathbf{y}} = I_{D\times D}s_t$, with s_t from (7). The new pair is accepted if

$$\tilde{\mathbf{y}}_t^T\Sigma_{\mathbf{y}}^{-1}\tilde{\mathbf{y}}_t \le d_M^2 \quad (12)$$

where the left term is a Mahalonobis distance [5]. We use $d_M = 3$, which corresponds to a 99% confidence region.

In many biomedical applications, the transformation is almost rigid, with a small nonlinear component. We proceed by first fitting a rigid transformation using e.g. a fast SVD-based algorithm [13] to the first few matched edges. Initial $\hat{\mathbf{C}}_0$ is then determined to fit the B-spline approximation to the rigid transformation at knots $\mathbf{z}_i$ by solving the normal equations (5). The B-spline approximation is then updated as described above. The rigid component could be reestimated in later iterations but in practice we found it unnecessary.

Preprint submitted to 2018 IEEE International Symposium on Biomedical Imaging.
Received October 22, 2017.

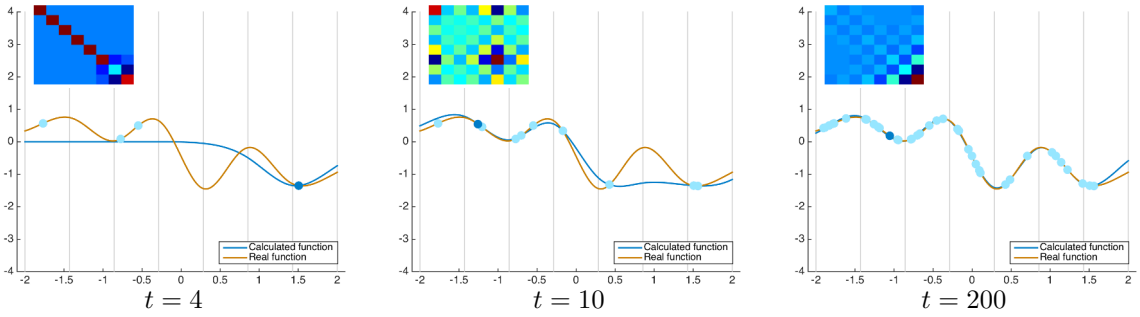

Fig. 1. Evalution of the approximation calculated by the proposed method with the true function in brown, the B-spline approximation f in blue, the randomly sampled data points shown as blue dots and the most current point in dark blue. Vertical gray lines represent the B-spline knots. The covariance matrix P_t of $\mathbf{c}$ is also shown (blue values are small and red are high).

3. EXPERIMENTS

3.1. Synthetic data

On the task of approximating a known function from its random samples, we compare the accuracy and speed of the proposed B-spline Kalman update method (Section 2.2) with four other methods: B-spline approximation by the direct method and by the iterative conjugate gradient method (with at most 20 iterations), Gaussian process regression (GPR) as described in [14], and thin plate splines radial basis functions (RBF) [6]. We set $\sigma_{\mathbf{c}} = 0.1$ and $\sigma_{\mathbf{r}} = 0.0025$ and $n_1, n_2 = 8$ in all experiments. All methods had to update the approximation after each sample. We report the mean time for the update and the mean ℓ_2 approximation error (evaluated on a fine, regular grid). The parameters governing the trade-off between accuracy and speed in all methods were set so that the final approximation error is similar.

Figure 2 shows the average over 100 repetitions of the mean approximation error and update time for the task of approximating an $\mathbb{R}^2 \to \mathbb{R}$ B-spline function with 8×8 knots and random normal coefficients. The Kalman method is at least as accurate as the other methods; note that exact recovery is not possible due to the measurement noise. Moreover, the Kalman method is much faster than all other methods and the update complexity per iteration is constant, independent of the number of measurements.

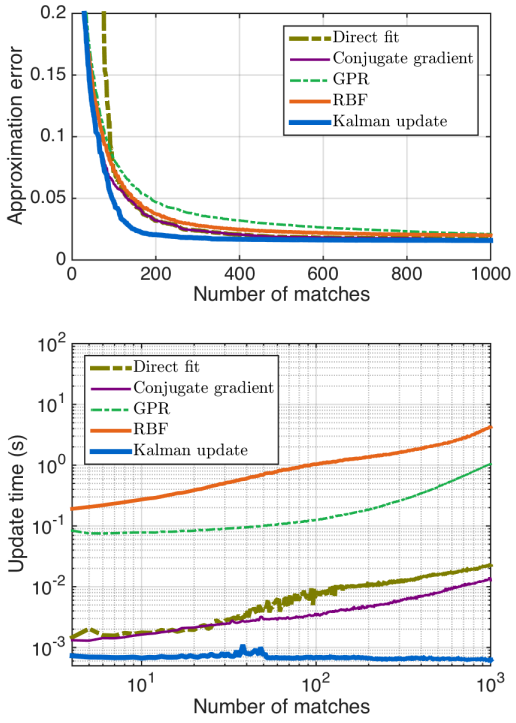

Fig. 2. Mean approximation error and update time per iteration for an $\mathbb{R}^2 \to \mathbb{R}$ function generated with B-splines, as a function of the number of measurements (matches).

3.2. Graph matching

In Table 1, we present the numerical geometric matching results for the Monte Carlo Tree Search algorithm [1]. We compare our B-spline model with Kalman update described here with fitting the same model either directly or via the iterative conjugated gradient method. We compare it further with the original bi-Lipschitz method [1] (denoted GMMC) and the Gaussian process model (denoted GPR) [14].

We use the following datasets (see [1] for example images): *Brain Circuits* (3D 2-photon microscopy volumes of brain networks acquired one week apart), *EM/LM* (3D brain neuronal connections acquired with electron and light microscopy), *Brain Vessels* (3D brain blood vessels acquired with optical and 2-photon microscopy), *Angiography* (heart blood vessels acquired by X-rays at different time instances), *Retina* (overlaping retinal photography images acquired at different positions, see Fig. 3), *Roads* (roads from satellite images with graphs extracted from maps).

In all datasets, the new method works best in terms of precision (fraction of correctly retrieved matches) thanks to the

Preprint submitted to 2018 IEEE International Symposium on Biomedical Imaging.
Received October 22, 2017.

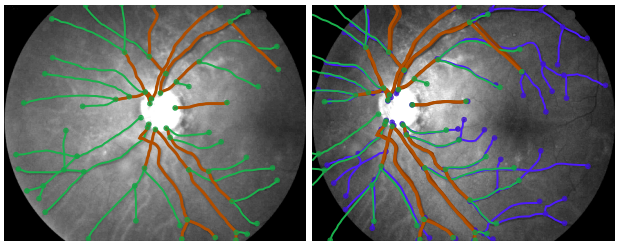

(a) (b)

Fig. 3. Example matching on the retinal dataset. *(a)* First image with extracted edges in green, *(b)* second image with extracted edges in blue and transformed edges from the first image after alignment. Matched edges are shown in red.

more efficient search space pruning. In 5 out of the 6 datasets, the proposed method is the most geometrically accurate. It is also the fastest on the two largest datasets.

4. CONCLUSION

We have presented a novel and computationally efficient method of updating B-spline coefficients for the task of incremental approximation of scattered data, based on Kalman filtering. We have shown that our procedure is by several orders of magnitude faster than alternative methods. We have further integrated the incremental update into a geometric graph matching method [1], to check the compatibility of incrementally built matches with the geometric model. This has turned out to be a very good fit, using the incremental updating procedure for a B-spline model is both fast and efficient in pruning the search tree, further increasing the match quality and the size of the graphs that can be matched.

Dataset		Kalman update	GMMC	Direct Fit	Conjugate Gradients	GPR
Brain	e	**0.026**	**0.026**	0.558	0.926	0.067
Circuits	p	**76.0**	72.0	5.0	0.0	64.7
$\|\mathbf{V}\| \in$	r	31.1	29.5	1.6	0.0	**36.1**
$[124, 130]$	t	0.58	**0.16**	147.56	80.17	0.17
	e	**0.014**	0.016	0.033	0.214	0.017
EM/LM	p	**77.8**	**77.8**	44.4	12.5	62.5
$\|\mathbf{V}\| \in$	r	**70.0**	**70.0**	40.0	10.0	50.0
$[22, 30]$	t	**0.00**	**0.00**	48.97	310.68	0.01
Brain	e	0.053	**0.045**	0.062	0.077	0.062
Vessels	p	**71.4**	62.5	**71.4**	**71.4**	**71.4**
$\|\mathbf{V}\| \in$	r	**41.7**	**41.7**	**41.7**	**41.7**	**41.7**
$[28, 37]$	t	0.02	**0.01**	34.71	399.13	0.02
Angio-	e	**0.024**	0.031	0.040	0.171	0.065
graphy	p	**72.2**	70.1	52.1	46.4	55.6
$\|\mathbf{V}\| \in$	r	72.2	**80.6**	58.3	50.0	69.4
$[18, 33]$	t	**0.01**	0.22	0.98	1.42	0.31
	e	**0.010**	0.017	0.181	0.169	0.083
Retina	p	**94.1**	87.6	64.0	76.8	63.4
$\|\mathbf{V}\| \in$	r	76.2	**76.8**	7.4	11.7	54.2
$[64, 209]$	t	**0.22**	0.50	1.93	10.59	8.66
	e	**0.005**	**0.005**	0.132	0.184	0.423
Roads	p	**90.5**	89.5	49.1	40.4	45.3
$\|\mathbf{V}\| \in$	r	**70.8**	66.0	20.1	19.6	33.7
$[19, 6050]$	t	**4.25**	62.24	83.34	100.52	17.28

Table 1. Average distance between true matches after aligment e (the graph coordinates were normalized to the $[-1, 1]^D$ interval), recall r and precision p in percents, and a processing time t in seconds. Bold denotes the best values. Maximum and minimum number of vertices $|\mathbf{V}|$ in each dataset is also given.

REFERENCES

[1] M. A. Pinheiro, J. Kybic, and P. Fua, "Graph matching using Monte Carlo tree search," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 39, no. 11, pp. 2171–2185, 2017.

[2] J. Kybic and M. Unser, "Fast parametric elastic image registration," *IEEE Trans. Imag. Proc.*, vol. 12, no. 11, pp. 1427–1442, 2003.

[3] M. Unser, "Splines: A perfect fit for medical imaging," *Progress in Biomedical Optics and Imaging*, vol. 3, pp. 225–236, 2002.

[4] F. Moreno-Noguer, V. Lepetit, and P. Fua, "Pose priors for simultaneously solving alignment and correspondence," in *ECCV*, 2008, pp. 405–418.

[5] J. Sanchez-Riera, J. Östlund *et al.*, "Simultaneous pose, correspondence and non-rigid shape," in *CVPR*, 2010, pp. 1189–1196.

[6] M. J. D. Powell, *Algorithms for Approximation*. Clarendon Press, 1987, ch. Radial Basis Functions for Multivariable Interpolation: A Review, pp. 143–167.

[7] M. Powell, "Least Frobenius norm updating of quadratic models that satisfy interpolation conditions," *Mathematical Programming*, vol. 100, no. 1, pp. 183–215, 2004.

[8] V. Skala, "Fast interpolation and approximation of scattered multidimensional and dynamic data using radial basis functions," in *WSEAS Trans.on Mathematics*, vol. 12, no. 5, 2013, pp. 501–511.

[9] J. R. Shewchuk, "Delaunay refinement algorithms for triangular mesh generation," *Computational Geometry: Theory and Applications*, vol. 1–3, no. 22, pp. 21–74, 2002.

[10] M. Harashima, L. A. Ferrari, and P. Sankar, "Spline approximation using Kalman filter state estimation," *IEEE Trans. Circuits and Systems II.*, vol. 44, no. 5, 1997.

[11] J. Jauch, F. Bleimund *et al.*, "Recursive B-spline approximation using the kalman filter," *Engineering Science and Technology*, vol. 20, no. 1, pp. 28 – 34, 2017.

[12] C. de Visser, Q. Chu, and J. Mulder, "Differential constraints for bounded recursive identification with multivariate splines," *Automatica*, vol. 47, no. 9, pp. 2059 – 2066, 2011.

[13] K. Arun, T. Huang, and S. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 9, no. 5, pp. 698–700, 1987.

[14] E. Serradell, P. Głowacki *et al.*, "Robust non-rigid registration of 2D and 3D graphs," in *CVPR*, 2012, pp. 996–1003.

Preprint submitted to 2018 IEEE International Symposium on Biomedical Imaging.
Received October 22, 2017.