Eurographics Symposium on Geometry Processing 2019
D. Bommes and H. Huang
(Guest Editors)

Volume 38 (2019), Number 5

Supplementary Material of Divergence-Free Shape Correspondence by Deformation

M. Eisenbergen, Z. Lähner and D. Cremers

Figure 1: Example of result on the TOPKIDS dataset. The correspondence between the target (left) and our deformation (right) are color-coded. Topological merging appears both on the right hand to face as well as the left hand to the body in the target shape. Due to our volume-preservation constraint we cannot align the surfaces in these parts which leads to squishing in the hands and the belly. Nevertheless, most parts of the body are matched well and even the correspondences on the affected parts are reasonable.

Appendix A: Optimization: Details

In section 5.3 of the paper we outlined our expectation maximization framework. Here, we want to provide a more detailed description of the method. At the same time, we try to keep it as brief as possible, because most parts of this summary are standard techniques when dealing with Gaussian mixture models.

Gaussian mixture model

As outlined before, we interpret the shifted points $f_n = x_n^{(T)}$ as the centers of Gaussian distributions with the covariance matrix $\sigma^2 I_D \in \mathbb{R}^{D\times D}$ which in the end should describe $\mathcal{Y}$ well. Furthermore, each point y_m is assumed to correspond to some point x_n. This relationship is encoded by the correspondence matrix $Z \in \{0,1\}^{(N+1)\times M}$, where $\sum_{n=1}^{N+1} Z_{nm} = 1$. If $Z_{(N+1)m} = 1$ for some m, the point y_m does not correspond to any point x_n and it is assumed to be uniformly sampled from Ω instead. This way we acknowledge the presence of outliers and counteract them by explicitly modeling them.

According to Bayes' theorem, the posterior probability distribution of the desired parameters a_k in (12) given the latent correspondences Z_{nm} and the observed points $\mathcal{Y}$ is defined as follows:

$$p(a|Z,\mathcal{Y}) \propto p(a)p(\mathcal{Y}|Z,a) = p(a)\prod_{m=1}^{M} p(y_m|Z,a)$$
$$= p(a)\prod_{m=1}^{M}\prod_{n=1}^{N+1} p(y_m|Z_{nm}=1,a)^{Z_{nm}}. \quad (21)$$

In Section 5.3 it was mentioned that the prior of the parameters a_k is a Gaussian distribution $a_k \sim \mathcal{N}(0,\lambda_k)$. As a shorthand notation we define the diagonal matrix $L := \mathrm{diag}(\lambda_1,\dots,\lambda_K)$ and set the prior $a \sim \mathcal{N}(0,L)$ for the coefficient vector a. In order to explicitly evaluate the posterior density of a in (21) we have to investigate the data likelihood $p(y_m|Z_{nm}=1,a)$ in detail. For $n < N+1$ it is a Gaussian distribution in the product space of the embedding space Ω and the space of descriptor values:

$$p(y_m|Z_{nm}=1,a) = \frac{1}{(2\pi\sigma^2)^{\frac{D}{2}}}\exp\left(-\frac{1}{2\sigma^2}\mathrm{d}_{nm}^2\right). \quad (22)$$

For the case $n = N+1$ it is simply a uniform distribution because y_m is considered to be an outlier:

$$p(y_m|Z_{(N+1)m}=1,a) = 1. \quad (23)$$

Note that the GMM is not only defined on the D dimensional embedding space but rather on the product space of Ω and the (possibly high dimensional) feature space. However, this only affects how the correspondences are computed and can be considered a theoretical nuance.

Expectation maximization

We want to determine the coefficients a by applying an expectation maximization approach similar to [MS10]. In the E step soft correspondences $W \in [0,1]^{(N+1)\times M}$ are determined as a relaxed version of the latent variables Z:

$$W_{nm} = \mathbb{E}_{Z|\mathcal{Y},a}(Z_{nm}) = p(Z_{nm}=1|y_m,a) = \frac{p(y_m|Z_{nm}=1,a)}{p(y_m|a)} = \frac{p(y_m|Z_{nm}=1,a)}{\sum_{\tilde{n}=1}^{N+1} p(y_m|Z_{\tilde{n}m}=1,a)}. \tag{24}$$

The data likelihood terms $p(y_m|Z_{nm}=1,a)$ are defined in (22) and (23). This leads to the expression proposed in (16).

The M step now consists of minimizing the following energy with respect to a:

$$\mathbb{E}_{Z|\mathcal{Y},a}\big(-\log p(a|Z,\mathcal{Y})\big) = -\log p(a) - \sum_{m=1}^{M}\sum_{n=1}^{N+1} W_{nm}\log p(y_m|Z_{nm}=1,a) \propto \frac{1}{2}a^T L^{-1} a + \frac{1}{2\sigma^2}\sum_{m=1}^{M}\sum_{n=1}^{N} W_{nm}\|y_m - f_n\|_2^2. \tag{25}$$

In this context the shifted points f_n depend on the unknown coefficients a. However, the descriptor distances $\mathrm{d}^{\mathrm{SHOT}}$ are independent of a, therefore they vanish in the last step in (25).

Robust correspondences

As proposed in (17) we will reformulate the correspondence penalization term $\frac{1}{2}\|y_m - f_n\|_2^2$ in (25) to in order to make our method more robust:

$$\rho(\|y_m - f_n\|_2). \tag{26}$$

Note that this has a probabilistic interpretation as a mixture of Huber densities. We choose the outer slope as $r_0 := 0.01$. For values $\|y_m - f_n\|_2 \le r_0$ the term (26) remains exactly the same but for bigger residua the penalization by the Huber loss only grows linearly. Due to this property the Huber norm does not penalize outliers exorbitantly high and is therefore more robust than the standard least squares loss.

Gauss Newton

In order to compute the optimal deformation parameters a we have to minimize the energy $E(a)$ defined in (17). For this purpose, we first derive how to optimize the following simplified energy and later derive how to handle the Huber loss distance penalization:

$$E^{\mathrm{LS}}(a) := \frac{1}{2}a^T L^{-1} a + \frac{1}{2\sigma^2}\sum_{m=1}^{M}\sum_{n=1}^{N} W_{nm}\|y_m - f_n\|_2^2. \tag{27}$$

We assume in this context that the correspondences W are fixed. For the optimization we use a Gauss-Newton type method which yields an iteration scheme minimizing E^{LS}. Like in the Levenberg-Marquardt algorithm [Lev44] the iteration contains an additional damping term L^{-1} which is added to the Hessian of the non-linear least squares term.

Applying the standard Gauss-Newton methodology we get an iterative method to determine the weights a. The general idea of this approach is that the shifted points $f_n(a)$ are linearized around the current iterate $a^{(i)}$ in the energy (27):

$$E^{\mathrm{LS}}(a) \approx \frac{1}{2}a^T L^{-1} a + \frac{1}{2\sigma^2}\sum_{m=1}^{M}\sum_{n=1}^{N} W_{nm}\|y_m - \underbrace{\big(f_n(a^{(i)}) + \mathrm{D}_a f_n(a^{(i)})(a - a^{(i)})\big)}_{\approx f_n(a)}\|_2^2. \tag{28}$$

This approximate energy is linear in the current unknown a so the remaining task is a simple linear least squares problem. The recursion formula to compute the approximate deformation parameters $a^{(i)}$ then admits the following explicit form:

$$a^{(i+1)} := a^{(i)} - \big(J^T\tilde{W}J + \sigma^2 L^{-1}\big)^{-1}\big(J^T r - \sigma^2 L^{-1} a^{(i)}\big). \tag{29}$$

In this context J consists of the Jacobians of f_n, r of the (weighted) distance residuals and $\tilde{W}$ is a diagonal matrix containing the column sums of W. Let $e_D = (1,...,1)^T \in \mathbb{R}^D, e_M = (1,...,1)^T \in \mathbb{R}^M$, then these quantities are explicitly defined as:

$$J = \begin{pmatrix} \mathrm{D}_a f_1 \\ \vdots \\ \mathrm{D}_a f_N \end{pmatrix} \in \mathbb{R}^{ND\times K}. \tag{30a}$$

$$r = \begin{pmatrix} \sum_{m=1}^{M} W_{1m}(f_1 - y_m) \\ \vdots \\ \sum_{m=1}^{M} W_{Nm}(f_N - y_m) \end{pmatrix} \in \mathbb{R}^{ND}. \tag{30b}$$

$$\tilde{W} = \mathrm{diag}\big((We_M)\otimes e_D\big) = \mathrm{diag}\begin{pmatrix} e_D\sum_{m=1}^{M} W_{1m} \\ \vdots \\ e_D\sum_{m=1}^{M} W_{Nm} \end{pmatrix} \in \mathbb{R}^{ND\times ND}. \tag{30c}$$

What is left to specify is how to compute the derivatives $\mathrm{D}_a f_n \in \mathbb{R}^{D\times K}$ in (30a). Note that $f_n = x_n^{(T)}$ is recursively defined in (14), therefore we need to apply the chain rule. As a result the derivative $\mathrm{D}_a x_n^{(t)}$ is passed from the first time step $t = 0$ to the last $t = T$ and gradually modified in each step. Inserting the Karhunen-Loève representation (12) in the definition (14) yields the following recursive formula:

$$x_n^{(t+1)}(a) := x_n^{(t)} + h\sum_{k=1}^{K} v_k\Big(x_n^{(t)} + \frac{h}{2}\sum_{k=1}^{K} v_k\big(x_n^{(t)}\big)a_k\Big)a_k. \tag{31}$$

The dependencies on a are denoted explicitly in order to make it more comprehensible. The quantities $x_n^{(t)}(a)$ can now be differentiated wrt. a:

$$\mathrm{D}_a x_n^{(0)} = 0. \tag{32a}$$

$$\mathrm{D}_a x_n^{(t+1)} =$$

$$\mathrm{D}_a x_n^{(t)} + h \sum_{k=1}^{K} \mathrm{D}_x v_k(\cdot) \left(\left(I_D + \frac{h}{2} \sum_{k=1}^{K} \mathrm{D}_x v_k(x_n^{(t)}) a_k \right) \mathrm{D}_a x_n^{(t)} + \begin{pmatrix} | & & | \\ v_1(x_n^{(t)}) & \dots & v_K(x_n^{(t)}) \\ | & & | \end{pmatrix} \right) a_k + h \begin{pmatrix} | & & | \\ v_1(\cdot) & \dots & v_K(\cdot) \\ | & & | \end{pmatrix}. \tag{32b}$$

Note that the Jacobian $\mathrm{D}_x v_k \in \mathbb{R}^{D \times D}$ can be computed analytically for any basis element v_k. In this context $I_D \in \mathbb{R}^{D \times D}$ is the identity matrix.

The only thing left to discuss is how to extend this approach for the Huber loss penalization of the energy E in (17). For point distances $\|y_m - f_n\|_2 \leq r_0$ the Huber loss and the least squares loss are the same. For residual values $\|y_m - f_n\|_2 > r_0$ the derivate wrt. the deformation parameters a is the following:

$$\mathrm{D}_a \rho(\|y_m - f_n\|_2) = r_0 \frac{(f_n - y_m)^T}{\|f_n - y_m\|_2} \mathrm{D}_a f_n. \tag{33}$$

This eliminates the possibility of a direct Gauss-Newton type optimization which requires non linear least squares terms. We can however incorporate this in our algorithm using a simple heuristic. For this purpose we multiply the respective weights W_{nm} with the factor $r_0 \frac{1}{\|f_n - y_m\|_2}$ for $\|y_m - f_n\|_2 > r_0$ in each iteration.

Appendix B: Karhunen-Loève expansion of the deformation field

Here, we provide some theoretical justification for the particular choice of basis in (11) and the construction of the weights (13). These λ_k can be interpreted to be the eigenvalues of the linear operator $\mathcal{C} := (-\Delta)^{-\frac{D}{2}}$ corresponding to the eigenfunctions ϕ_k. We can then apply the so-called Karhunen-Loève expansion [Sul15, Ch. 11] to our setup. This framework provides us with an alternative representation of the potential field Φ which can in turn be used to define the deformation field v. For further reference concerning the mathematical foundation of this approach the interested reader is referred to [Stu10], [CRSW13], [DS17]. Following this approach one can now derive a construction which enables us to sample arbitrary square integrable scalar fields $\hat{\Phi} : \Omega \to \mathbb{R}$:

$$\hat{\Phi}(x) = \sum_{k=1}^{\infty} \phi_k(x) \sqrt{\lambda_k} \xi_k. \tag{34}$$

According to the Karhunen-Loève expansion the coefficients are samples of the Gaussians $\xi_k \sim \mathcal{N}(0,1)$. This approach can now be applied to get an alternative description of each entry of the potential vector field Φ. Inserting this in (5) we obtain an alternative representation of the deformation field v. In particular, we get the summation (12) for the basis elements (11) in the 3D case. Indeed we can derive the following Gaussian prior distribution for the weights a_k:

$$a_k = \sqrt{\lambda_k} \xi_k \sim \mathcal{N}(0, \lambda_k). \tag{35}$$

Remark The choice of the exponent $\frac{D}{2}$ in the definition of the weights λ_k (13) is not arbitrary. In general it is supposed to be chosen strictly larger than $\frac{D}{2}$ in order for our resulting basis to fulfill certain approximation properties in the limit of infinitely many basis functions, see [DS17, Ch. 2.4]. However, we achieved good results in our experiments by choosing it as small as possible in order to not suppress the high frequencies more severely than necessary. In particular, the expressiveness of or method seems to deteriorate when a large exponent is chosen because then the weights (13) decay too rapidly. Therefore, we typically even set it to $\frac{D}{2}$ which works fine for our purposes, although this is not theoretically justified when the number of basis functions approaches infinity. On the other hand, choosing it smaller than $\frac{D}{2}$ certainly causes the expected value of the velocity series to diverge for $K \to \infty$.

To conclude this section we want to motivate our choice of the Karhunen-Loève framework and the particular linear operator $\mathcal{C}$ to model the deformation fields v. In the context of the Karhunen-Loève expansion the operator $\mathcal{C}$ is called covariance operator. It is typically chosen to incorporate some assumptions about the regularity of the produced sample functions. A natural assumption about the deformation fields v is that they are as uniform as possible. This yields that the resulting correspondence mappings are to some degree spatially continuous. Therefore, we require the Dirichlet energy to be small:

$$\|\nabla v\|_{L_2}^2 = \sum_{d=1}^{D} \int_{\Omega} \|\nabla v_d(x)\|_2^2 \mathrm{d}x. \tag{36}$$

We can achieve this by penalizing the high frequency components of v. These frequencies are strongly related to those of the potential field Φ because according to (5) the basis elements are simply mapped onto the velocity basis elements. This mapping does not change the frequencies:

$$\|\nabla v\|_{L_2} = \|\nabla(\nabla \times \Phi)\|_{L_2} = \|\nabla \Phi\|_{L_2}. \tag{37}$$

If we choose e.g. $D = 2$ one can prove that the Dirichlet energy $\|\nabla v\|_{L_2}^2$ is equivalent to the squared ℓ_2 norm of the weights ξ:

$$\|\xi\|_{\ell_2}^2 = \|\nabla v\|_{L_2}^2, \text{ for } D = 2. \tag{38}$$

A derivation of this property can be found in [DS17, Ch. 7.1.3]. In the case of finitely many parameters $\xi_1, ..., \xi_K$ the norm $\|\xi\|_{\ell_2}$ is equivalent to the Euclidean norm $\|\xi\|_2$ of the vector $\xi = (\xi_1, ..., \xi_K)^T$. The term $\|\xi\|_2^2$ is in turn proportional to the negative log likelihood of the standard normal distributed parameter $\xi \sim \mathcal{N}(0, I_K)$:

$$-\log(p(\xi)) = \frac{K}{2} \log(2\pi) + \frac{1}{2} \|\xi\|_2^2 \propto \frac{1}{2} \|\xi\|_2^2. \tag{39}$$

This indicates that a maximum likelihood approach involving ξ leads to an enforcement of uniformity of the vector field v. This can be extended to the case $D = 3$ in a similar manner but we refrain from providing more details here for the sake of brevity.

References

[CRSW13] COTTER S. L., ROBERTS G. O., STUART A. M., WHITE D.: Mcmc methods for functions: modifying old algorithms to make them faster. *Statistical Science 28*, 3 (2013), 424–446. 3

[DS17] DASHTI M., STUART A. M.: The bayesian approach to inverse problems. *Handbook of Uncertainty Quantification* (2017), 311–428. 3

[Lev44] LEVENBERG K.: A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics 2*, 2 (1944), 164–168. 2

[MS10] MYRONENKO A., SONG X.: Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence 32*, 12 (2010), 2262–2275. 1

[Stu10] STUART A. M.: Inverse problems: a bayesian perspective. *Acta Numerica 19* (2010), 451–559. 3

[Sul15] SULLIVAN T. J.: *Introduction to Uncertainty Quantification*. Springer, 2015. 3