

Appendix A

Performing DR in the Parameter Space Instead of the Task Space

In this appendix, we consider the cases where the dimensionality reduction of the MPs is performed not on the task space, but rather in the parameter space by basically switching the order in which the kernel function matrix and the linear projection matrix defined in Chap. 7 are used for the representation of the state vector. A principal component approach is introduced for DMPs, while an EM approach is used for ProMPs.

A.1 DR in the Parameter Space for DMPs

Dimensionality reduction (DR) of DMPs, which in Sect. 7.2 was performed in the robot joint space, can also be performed in the parameter space with the expression:

$$\mathbf{f}(x_t) = \Phi_t^T \Omega \mathbf{v}, \tag{A.1}$$

with Φ_t^T being a $(d \times dN_f)$ matrix of Gaussian basis functions and Ω being a $(dN_f \times M_f)$ matrix with the mappings from a parameter space of dimension M_f to the whole DMP parameter space. The DMP weight vector \mathbf{v} now has dimension M_f . In order to initialize the projection matrix Ω , we have to take the data matrix in Eq. (7.26), $\bar{\mathbf{F}}$, knowing that:

$$\bar{\mathbf{F}} = [\mathbf{f}_1, \dots, \mathbf{f}_{N_t}] = [\Phi_1^T \Omega \mathbf{v}, \dots, \Phi_{N_t}^T \Omega \mathbf{v}], \tag{A.2}$$

which can be expressed as a least-squares minimization problem as:

$$[\Phi_1^{\dagger,T} \mathbf{f}_1, \dots, \Phi_{N_t}^{\dagger,T} \mathbf{f}_{N_t}] \simeq \Omega \mathbf{v}, \tag{A.3}$$

where \dagger represents the pseudoinverse matrix. We can then obtain the projection matrix $\mathbf{\Omega}$ by performing PCA in Eq.(A.3). Note that, in order to fit the matrix $\mathbf{\Omega}$ ($dN_f \times M_f$), we need at least a number of timesteps $N_t > M_f$.

A.2 DR in the Parameter Space for ProMPs

In Chap.7, we devised an Expectation-Maximization based method for finding a linear Dimensionality Reduction (DR) technique to represent ProMPs in a latent space of the DoF of the robot, using the expression $\mathbf{y}_t = \mathbf{\Omega} \mathbf{\Phi}_t^T \mathbf{v}$ for the state space, where $\mathbf{\Omega}$ was a ($d \times r$) matrix, d being the DoF of the trajectory and r a latent space dimension. In this paper, we used the reverse order between the two terms $\mathbf{\Omega}$ and $\mathbf{\Phi}_t$, resulting in:

$$\mathbf{y}_t = \mathbf{\Phi}_t^T \mathbf{\Omega} \mathbf{v}, \quad (\text{A.4})$$

where $\mathbf{\Omega}$ is now a ($dN_f \times M_f$) matrix and we use $\mathbf{v} \sim \mathcal{N}(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v)$ to represent the latent space parameters. In this case, the linear DR is performed directly in the Gaussian weights space of the ProMP, instead of the DoF' space, resulting in the probability distribution for the state space:

$$\begin{aligned} p(\mathbf{y}_t) &= \int_{\mathbf{v}} \mathcal{N}(\mathbf{y}_t | \mathbf{\Phi}_t^T \mathbf{\Omega} \mathbf{v}, \boldsymbol{\Sigma}_y) \mathcal{N}(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v) d\mathbf{v} \\ &= \mathcal{N}(\mathbf{y}_t | \mathbf{\Phi}_t^T \mathbf{\Omega} \boldsymbol{\mu}_v, \boldsymbol{\Sigma}_y + \mathbf{\Phi}_t^T \mathbf{\Omega} \boldsymbol{\Sigma}_v \mathbf{\Omega} \mathbf{\Phi}_t^T) \end{aligned} \quad (\text{A.5})$$

We then describe the Expectation-Maximization step for obtaining the parameters $\boldsymbol{\theta} = \{\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v, \mathbf{\Omega}, \boldsymbol{\Sigma}_y\}$ by using the complete-data log-likelihood. $\mathbf{\Phi}_t^T$ will be a ($d \times M_f$) matrix, and \mathbf{y}_t^j will be a d -dimensional vector. We use the vector \mathbf{Y}^j to denote the concatenated position vectors of a single trajectory j ,

$$\mathbf{Y}^j = \left[\mathbf{y}_1^{jT}, \dots, \mathbf{y}_{N_t}^{jT} \right]^T.$$

Now, we want to use our model estimation algorithm for policy search algorithms that are based on data re-weighting. These algorithms introduce a weighting d_k for each trajectory. Hence, we also have to consider such a weighting in our EM algorithm. We will maximize the weighted marginal log-likelihood $\sum_k d_k \log p(\mathbf{y}_t, \boldsymbol{\theta})$ of the data and the latent space representation, thus we have to derive the equations with the marginalized \mathbf{v} with the difficulties it entails. The following sections explain how to obtain the log-likelihood function and differentiate it.

Expectation Step

In the expectation step, we must find the probabilities for each demonstration j with the old parameters $\boldsymbol{\theta}^{old}$:

$$p(\mathbf{v} | \mathbf{Y}^j) = \frac{p(\mathbf{Y}^j | \mathbf{v}) p(\mathbf{v})}{p(\mathbf{Y}^j)} \propto p(\mathbf{Y}^j | \mathbf{v}) p(\mathbf{v}), \quad (\text{A.6})$$

where:

$$p(\mathbf{Y}^j | \mathbf{v}) = \mathcal{N}(\mathbf{Y}^j | \Phi^T \Omega \mathbf{v}, \mathbf{I}_{N_t} \otimes \Sigma_y). \quad (\text{A.7})$$

Note that in the contextual case, we have omitted the conditioning on the context variables for simplicity of the equations. Using the Bayes rule for Gaussian distributions, we obtain $p(\mathbf{v} | \mathbf{Y}^j) = \mathcal{N}(\mathbf{v} | \mu_j, \Sigma_j)$, where

$$\begin{aligned} \mu_j &= \mu_v + \Sigma_v \Omega_{N_t}^T \Phi \Gamma^{-1} (\mathbf{Y}^j - \Phi^T \Omega_{N_t} \mu_v) \\ \Sigma_j &= \Sigma_v - \Sigma_v \Omega_{N_t}^T \Phi \Gamma^{-1} \Phi^T \Omega_{N_t} \Sigma_v, \end{aligned}$$

and Γ is given by $\Gamma = \mathbf{I}_{N_t} \otimes \Sigma_y + \Phi^T \Omega_{N_t} \Sigma_v \Omega_{N_t}^T \Phi$.

Maximization Step

Given the probabilities $p(\mathbf{v} | \mathbf{Y}^j)$ for each demonstration, we maximize the weighted expectation of the log-likelihood function, where d_j is used as weight for each trajectory, i.e.,

$$L = \sum_{j=1}^{N_d} d_j \mathbb{E}_{\mathbf{v} | \mathbf{Y}^j; \theta^{old}} [\log (p(\mathbf{v}, \mathbf{Y}^j; \theta))]$$

After some calculations, analogously to Sect. 7.1.1, we obtain:

$$\begin{aligned} L &= -\frac{1}{2} \left[\left(\sum_{j=1}^{N_d} d_j \right) \log |2\pi \Sigma_v| + N_t \log |2\pi \Sigma_y| \right] \\ &\quad - \frac{1}{2} \sum_{j=1}^{N_d} d_j (\mu_v - \mu_j) \Sigma_v^{-1} (\mu_v - \mu_j) \\ &\quad - \frac{1}{2} \sum_{j=1}^{N_d} d_j \sum_{t=1}^{N_t} [(y_t^j - \Phi_t^T \Omega \mu_j)^T \Sigma_y^{-1} (y_t^j - \Phi_t^T \Omega \mu_j) \\ &\quad + \text{tr}(\Omega^T \Phi_t \Sigma_y^{-1} \Phi_t^T \Omega \Sigma_j)] - \frac{1}{2} \sum_{j=1}^{N_d} d_j \text{tr}(\Sigma_v^{-1} \Sigma_j). \end{aligned}$$

Then, we derivate L w.r.t. $\theta = \{\mu_v, \Sigma_v, \Omega, \Sigma_y^{-1}\}$:

$$\mu_v = \left(\sum_{j=1}^{N_d} d_j \right)^{-1} \sum_{j=1}^{N_d} d_j (\mu_j), \quad (\text{A.8})$$

$$\Sigma_v = \left(\sum_{j=1}^{N_d} d_j \right)^{-1} \sum_{j=1}^{N_d} d_j [\Sigma_j + (\mu_v - \mu_j)(\mu_v - \mu_j)^T], \quad (\text{A.9})$$

$$\mathbf{\Omega} = \left[\sum_{t=1}^{N_t} d_j \mathbf{\Phi}_t^T \mathbf{\Sigma}_y^{-1} \mathbf{\Phi}_t \right]^\dagger \left[\sum_{t=1}^{N_t} \mathbf{\Phi}_t \mathbf{\Sigma}_y^{-1} \sum_{j=1}^{N_d} d_j y_t^j \boldsymbol{\mu}_j^T \right] \cdot \left[\sum_{j=1}^{N_d} d_j (\boldsymbol{\Sigma}_j + \boldsymbol{\mu}_j \boldsymbol{\mu}_j^T) \right]^\dagger \quad (\text{A.10})$$

$$\begin{aligned} \mathbf{\Sigma}_y = & \left(\sum_{j=1}^{N_d} d_j \right)^{-1} \frac{1}{N_t} \sum_{j=1}^{N_d} \sum_{t=1}^{N_t} d_j [\mathbf{y}_t^j (\mathbf{y}_t^j - \mathbf{\Phi}_t^T \boldsymbol{\Omega} \boldsymbol{\mu}_j)^T + \mathbf{\Phi}_t^T \\ & + \boldsymbol{\Omega} \boldsymbol{\mu}_j (\mathbf{\Phi}_t^T \boldsymbol{\Omega} \boldsymbol{\mu}_j - \mathbf{y}_t^j)^T + \mathbf{\Phi}_t^T \boldsymbol{\Omega} \boldsymbol{\Sigma}_j \boldsymbol{\Omega} \mathbf{\Phi}_t^T] \end{aligned} \quad (\text{A.11})$$

Appendix B

External Resources

- 1- Closed-loop inverse kinematics for redundant robots
<http://www.iri.upc.edu/groups/perception/#inverseKinematics>
- 2- Friction model applications
<http://www.iri.upc.edu/groups/perception/#ScarfTask>
- 3- Realtime tracking and grasping of a moving object from range video
<http://www.iri.upc.edu/groups/perception/#trackGrasp>
- 4- Human-guided compliant control
<http://www.iri.upc.edu/groups/perception/#adapt>
<http://www.iri.upc.edu/groups/perception/#adapt2>
- 5- Dual Relative Entropy Policy Search
<http://www.iri.upc.edu/groups/perception/#dualReps>
- 6- Dimensionality Reduction for Dynamic Movement Primitives
<http://www.iri.upc.edu/groups/perception/#drdmp>
- 7- Contextualized ProMPs with obstacle avoidance
<http://www.iri.upc.edu/groups/perception/#ProMPobstacles>
- 8- Mixture models of adaptable ProMPs
<http://www.iri.upc.edu/groups/perception/#DRGMM>

Glossary

J	Geometric Jacobian
J[†]	Jacobian pseudoinverse
q = [q ₁ , . . . , q _m]	Joint state
Δq	joint state variation
x	Task space variable
e	Positioning error of the robot
κ(·)	Condition number of a matrix
FK(·)	Forward kinematics function
α	Gain in a CLIK algorithm
σ₁, . . . , σ_n	Jacobian singular values
m	Number of joints
n	Task space dimension
u_c	Robot control torque
u_e	External torque acting on the robot
u_{PD}	Proportional-derivative torque
K_P	Proportional gain matrix
K_D	Derivative gain matrix
M(q)	Robot joint inertia matrix
C(q, q̇)	Coriolis and centripetal forces
F_f	Robot joint friction
G(q)	Robot's gravity forces
n(q, q̇)	Coriolis, centripetal, friction and gravity forces
π	Policy
Θ	Policy parameter space
θ ∈ Θ	Policy parameters
τ	Trajectory
y	Robot state variable
x	Robot's latent space state variable
s	Contextual variable

u	Action taken by the policy π
R	Reward function
J_{π_θ}	Expected reward under policy π
$\omega \sim \mathcal{N}(\boldsymbol{\mu}_\omega, \boldsymbol{\Sigma}_\omega)$	Sample from a normally distributed policy with mean $\boldsymbol{\mu}_\omega$ and covariance $\boldsymbol{\Sigma}_\omega$
λ, η	Weighting temperatures for PI2 and REPS, respectively
N_t	Number of timesteps of a trajectory
N_f	Number of Gaussian kernels used per DoF
N_k	Number of trajectories (rollouts) used per policy update
y_g	DMP goal position
$\tau, \alpha_z, \beta_z, \alpha_x$	DMP fixed parameters
x	DMP phase variable
f(x) = f_t	DMP excitation function
g(x)	N_f -dimensional vector of DMP Gaussian kernels
$\boldsymbol{\psi} = \mathbf{I}_d \otimes \mathbf{g}(x)$	$d \times dN_f$ matrix of Gaussian kernels
$\boldsymbol{\Omega}$	Robot's DoF coupling matrix
$\boldsymbol{\Sigma}_y$	ProMP fitting and system noise
$\boldsymbol{\Sigma}_f$	DMP fitting and system noise