Reduced Dimensionality Extended Kalman Filter for SLAM

Dinesh Gamage
http://eng.monash.edu.au/ecse/about/people/profile/dineshg

Tom Drummond
http://eng.monash.edu.au/ecse/about/people/profile/twd

Department of Electrical and Computer Systems Engineering, Monash University, Australia.

Computational complexity of the Kalman filter grows at least quadratically with the number of dimensions in the filter. This is a particular problem for applications like monocular simultaneous localization and mapping (SLAM) where it is not possible to run a single filter on a large map with many thousands of landmarks. The filtering approach for SLAM, maintains only the current camera pose with all landmarks of interest as the state [2].

This paper presents a method for reducing the computational complexity of the Kalman filters by reducing the dimensionality as information is acquired. The method reduces the dimensionality of the extended Kalman filter (EKF) for SLAM by identifying dominant modes of the filter. This can be used in general to reduce the dimensionality of the EKF irrespectively of its application, without being limited to SLAM.

For a filter with zero process noise, the mean of this distribution can be represented as a point in a nD space with its uncertainty as a hyper ellipse. Further information will reduce this uncertainty along some directions. After some time uncertainty along many directions becomes comparatively small, making further information redundant. So the covariance matrix can be decompose into certain and uncertain dimensions.

$$\Sigma_{t_1} = UDU^T \tag{1}$$

where U is the set of singular vectors of Σ_{t_1} and D is the diagonal matrix of singular values. The singular vectors represent de-correlated uncertainty directions with variances proportional to their singular values. We partition D into D_s the significant (large) set and D_i the insignificant (small) set. U is also partitioned into U_s vectors corresponding to D_s and U_i vectors corresponding to D_i. With this Σ_t can now be written in block form:

$$\Sigma_{t_1} = \begin{bmatrix} U_s & | & U_i \end{bmatrix} \begin{bmatrix} D_s & 0 \\ \hline 0 & D_i \end{bmatrix} \begin{bmatrix} U_s^T \\ \hline U_i^T \end{bmatrix} \tag{2}$$

So the original covariance matrix becomes:

$$\Sigma_{t_1} = U_s D_s U_s^T + U_i D_i U_i^T \tag{3}$$

As the second term of the above expression is comparatively small the column space of U_s can be used as the reduced space ignoring small singular vectors. Let the significant vectors U_s extracted at $t = t_1$ be U_{st_1} and x_t' be the reduced state which is Gaussian distributed in the reduced space with a mean μ_t' at $t(\geq t_1)$. The reduced mean μ_t' relates to the original state mean μ_t as:

$$\mu_t = \mu_{t_1} + U_{st_1}\mu_t' \tag{4}$$

This new state x_t' represents the variations of the original state around the point μ_{t_1} along the directions of the column space of U_{st_1}. Initially we start with μ_{t_1} being a zero vector, indicating our knowledge about uncertainty is zero along corresponding singular vectors in U_{st_1}. Subsequent observations can be projected onto the derived reduced space to gather information about the new state x_t'. The projected covariance matrix Σ_{t_1} is obtained as:

$$\Sigma_{t_1}' = U_{st_1}^T \left(\Sigma_{t_1} - U_i D_i U_i^T \right) U_{st_1} \tag{5}$$

The dimensionality of x_t' can be kept quite small, compared to the original state x_t. For all time steps $t(\geq t_1)$, information can be collected to update x_t' by changing its mean and the covariance. This makes the reduced state time dependent. To obtain the prediction equation in the reduced space, the linearized EKF states at time steps t and $t - 1$ can be decomposed according to the equation 4. As we are assuming zeros process noise, the process model becomes the identity. Substituting the decomposed states into the process equation yields:

$$\mu_{t_1} + U_{st_1}\mu_t' = \left(\mu_{t_1} + U_{st_1}\mu_{t-1}' \right) \tag{6}$$

making the predicted state same as the previous state:

$$\mu_t' = \mu_{t-1}' \tag{7}$$

Similarly, if the measurement is z_t with the model Jacobian H_t and measurement noise v_t, update equation, after substitution becomes:

$$z_t = H_t \left(\mu_{t_1} + U_{st_1}\mu_t' \right) + v_t \tag{8}$$

which can be modified as:

$$z_t - H_t\mu_{t_1} = H_t U_{st_1}\mu_t' + v_t \tag{9}$$

Here $z_t - H_t\mu_{t_1}$ becomes the modified observation and $H_t U_{st_1}$ the projected Jacobian. The dimensionality reduction described so far works only when the state is static. In a SLAM setup, the camera pose keeps changing requiring a slightly different approach.

The computational savings achieved by reducing the dimensionality of the filter can be spent by admitting more variables into the filter to be measured (thus increasing its dimensionality again). A new variable l, can be directly added to the reduced state to obtain an augmented state. It has to be noted that this augmentation increases the dimensionality of the original space as well by the same number of dimensions. So the froze mean μ_{t_1} has to be augmented with zeros to get a modified mean $\hat{\mu}_{t_1}$.

With continuous augmentations the reduced state $x_{k|k}'$ also will start growing. To keep the dimensionality of $x_{k|k}'$ manageable, we decompose it continuously by distributing newly learned knowledge over the froze mean μ_{t_1} and select a new reduced basis.

To model the desired variation while retaining standard operations on the camera, here we perform reduction only upon landmark states. As the first step, landmarks observed up to time $t = t_1$ are decomposed, keeping the camera state intact by directly transferring it into the reduced state.

When landmarks are static, the prediction involves estimating the camera and its covariance with cross-covariance [1]. As we transfer the camera state directly in to the reduced state x_t', camera parameters can be predicted in the usual manner while keeping x_{y_t}' unchanged to get the predicted state $x_{t|(t-1)}'$. The covariance block in the reduce covariance Σ_t' corresponding to the camera state has to be added with process noise to get the predicted covariance $\Sigma_{t|(t-1)}'$.

State update is done by projecting the measurement model onto the reduced space. Measurement model Jacobian H_t, can be projected onto reduced space by a right multiplication. If we drop the subscript t and denote the Jacobian as H:

$$H_s = H U_{st_1} \tag{10}$$

Using these construction, in this paper we introduce a dimensionality reduction technique to handle the complexity growth of the extended Kalman filter for SLAM. Though the Kalman filter is not the state of the art any more with current sparse matrix methods, we believe still the Kalman filter has a considerable potential generally in computer vision.

[1] J. Civera, A. Davison, and J. Martínez Montiel. 1-point ransac. *Structure from Motion using the Extended Kalman Filter*, pages 65–97, 2012.

[2] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(6):1052–1067, 2007.