



Figure 12.6 Comparison of (a) SEIF without sparsification with (b) SEIF using the sparsification step with 6 active landmarks. The comparison is carried out in a simulated environment with 50 landmarks. In each row, the left panel shows the set of links in the filter, the center panel the correlation matrix, and the right panel the normalized information matrix. Obviously, the sparsified SEIF maintains many fewer links, but its result is less confident as indicated by its less-expressed correlation matrix.

12.7 How Sparse Should SEIFs Be?

A key question pertains to the degree of sparseness one should enforce in a SEIF. In particular, the number of active features in SEIFs determines the degree of sparseness. The sparseness trades off two factors: the computational efficiency of the SEIF, and the accuracy of the result. When implementing a SEIF algorithm, it is therefore advisable to get a feeling for this trade-off.

The “gold standard” for a SEIF is the EKF, which avoids sparsification and also does not rely on relaxation techniques for recovering the state estimate. The following comparison characterizes the three key performance measures that set sparse SEIFs apart from EKFs. Our comparison is based on a simulated robot world, in which the robot senses the range, proximity, and identity of nearby landmarks.