that artificially inflates the amount of noise in the sensor. One can think of this inflation as accommodating not just the measurement uncertainty, but also the uncertainty induced by the approximate nature of the particle filter algorithm.

An alternative, more sound solution involves a modification of the sampling process which we already discussed briefly in Chapter 4.3.4. The idea is that for a small fraction of all particles, the role of the motion model and the measurement model are reversed: Particles are generated according to the measurement model

$$x_t^{[m]} \sim p(z_t \mid x_t) \tag{8.6}$$

and the importance weight is calculated in proportion to

$$w_t^{[m]} = \int p(x_t^{[m]} \mid u_t, x_{t-1}) \, bel(x_{t-1}) \, dx_{t-1} \tag{8.7}$$

This new sampling process is a legitimate alternative to the plain particle filter. It alone will be inefficient since it entirely ignores the history when generating particles. However, it is equally legitimate to generate a fraction of the particles with either of those two mechanisms and merge the two particle sets. The resulting algorithm is called *MCL with mixture proposal distribution*,

MIXTURE MCL

or *Mixture MCL*. In practice, it tends to suffice to generate a small fraction of particles (e.g., 5%) through the new process.

Unfortunately, our idea does not come without challenges. The two main steps—sampling from $p(z_t \mid x_t)$ and calculating the importance weights $w_t^{[m]}$—can be difficult to realize. Sampling from the measurement model is only easy if its inverse possesses a closed form solution from which it is easy to sample. This is usually not the case: imagine sampling from the space of all poses that fit a given laser range scan! Calculating the importance weights is complicated by the integral in (8.7), and by the fact that $bel(x_{t-1})$ is itself represented by a set of particles.

Without delving into too much detail, we note that both steps can be implemented, but only with additional approximations. Figure 8.17 shows comparative results for MCL, MCL augmented with random samples, and Mixture MCL for two real-world data sets. In both cases, $p(z_t \mid x_t)$ was itself learned from data and represented by a density tree—an elaborate procedure whose description is beyond the scope of this book. For calculating the importance weights, the integral was replaced by a stochastic integration, and the prior belief was continued into a space-filling density by convolving each particle with a narrow Gaussian. Details aside, these results illustrate that the mixture idea yields superior results, but it can be challenging to implement.