model $p(z_t \mid x_t, m_{c_t}, c_t)$ in the same way as EKF SLAM. As usual, we approximate the measurement function h by Taylor expansion:

(13.15)
$$\begin{aligned} h(m_{c_t}, x_t^{[k]}) &\approx \underbrace{h(\mu_{c_t,t-1}^{[k]}, x_t^{[k]})}_{=:\,\hat{z}_t^{[k]}} + \underbrace{h'(\mu_{c_t,t-1}^{[k]}, x_t^{[k]})}_{=:\,H_t^{[k]}} (m_{c_t} - \mu_{c_t,t-1}^{[k]}) \\ &= \hat{z}_t^{[k]} + H_t^{[k]}\,(m_{c_t} - \mu_{c_t,t-1}^{[k]}) \end{aligned}$$

Here the derivative h' is taken with respect to the feature coordinates m_{c_t}. This linear approximation is tangent to h at $x_t^{[k]}$ and $\mu_{c_t,t-1}^{[k]}$. Under this approximation, the posterior for the location of feature c_t is indeed Gaussian. The new mean and covariance are obtained using the standard EKF measurement update:

(13.16)
$$K_t^{[k]} = \Sigma_{c_t,t-1}^{[k]} H_t^{[k]T} (H_t^{[k]} \Sigma_{c_t,t-1}^{[k]} H_t^{[k]T} + Q_t)^{-1}$$

(13.17)
$$\mu_{c_t,t}^{[k]} = \mu_{c_t,t-1}^{[k]} + K_t^{[k]}(z_t - \hat{z}_t^{[k]})$$

(13.18)
$$\Sigma_{c_t,t}^{[k]} = (I - K_t^{[k]} H_t^{[k]})\Sigma_{c_t,t-1}^{[k]}$$

Steps 1 and 2 are repeated M times, resulting in a temporary set of M particles.

3. **Resampling.** In a final step, FastSLAM resamples this set of particles. We already encountered resampling in a number of algorithms. FastSLAM draws from its temporary set M particles (with replacement) according to a yet-to-be-defined importance weight. The resulting set of M particles then forms the new and final particle set, Y_t. The necessity to resample arises from the fact that the particles in the temporary set are not distributed according to the desired posterior: Step 1 generates poses x_t only in accordance with the most recent control u_t, paying no attention to the measurement z_t. As the reader should know well by now, resampling is the common technique in particle filtering to correct for such mismatches.

   This situation is once again illustrated in Figure 13.5, for a 1-D example. Here the dashed line symbolizes the *proposal distribution*, which is the distribution at which particles are generated, and the solid line is the target distribution. In FastSLAM, the proposal distribution does not depend on z_t, but the target distribution does. By weighting particles as shown in the bottom of this figure, and resampling according to those weights, the resulting particle set indeed approximates the target distribution.