

Erratum

An error appeared on page 20 of Richard S. Sutton's paper, "Learning to Predict by the Methods of Temporal Differences," which appeared in volume 3, number 1, of *Machine Learning*. The correct version of Figure 3 appears below, along with its caption.




Figure 3. Average error on the random-walk problem under repeated presentations. All data are from TD(λ) with different values of λ . The dependent measure used is the RMS error between the ideal predictions and those found by the learning procedure after being repeatedly presented with the training set until convergence of the weight vector. This measure was averaged over 100 training sets to produce the data shown. The $\lambda = 1$ data point is the performance level attained by the Widrow-Hoff procedure. For each data point, the standard error is approximately $\sigma = 0.01$, so the differences between the Widrow-Hoff procedure and the other procedures are highly significant.