



Here you will find experiments and results obtained when performing  $n$ -step TD learning on the random walk examples from this chapter. I ran the experiments suggested in the book with one thousand episodes rather than the suggested one hundred. I also ran these experiments for one hundred alphas uniformly spaced between the limits given. For the online version of  $n$ -step TD learning we obtain



This result looks quite similar to that presented in the book. For the offline version of the same algorithm we obtained



This looks qualitatively the same as that found in the book but has a more "jagged" appearance of the graphs for small  $n$ . This I believe is due to the more stringent requirements that offline TD learning methods require of their alpha parameter, i.e. they must be small enough so that convergence is guaranteed.

---

[John Weatherwax](#)

Last modified: Sun May 15 08:46:34 EDT 2005