This paper addresses the problem of distributed detection in multi-agent networks. Agents receive private signals about an unknown state of the world. The underlying state is globally identifiable, yet informative signals may be dispersed throughout the network. Using an optimization-based framework, we develop an iterative local strategy for updating individual beliefs. In contrast to the existing literature which focuses on asymptotic learning, we provide a finite-time analysis. Furthermore, we introduce a Kullback-Leibler cost to compare the efficiency of the algorithm to its centralized counterpart. Our bounds on the cost are expressed in terms of network size, spectral gap, centrality of each agent and relative entropy of agents' signal structures. A key observation is that distributing more informative signals to central agents results in a faster learning rate. Furthermore, optimizing the weights, we can speed up learning by improving the spectral gap. We also quantify the effect of link failures on learning speed in symmetric networks. We finally provide numerical simulations for our method which verify our theoretical results.