ASYMPTOTIC LEARNING IN FEEDFORWARD NETWORKS WITH BINARY SYMMETRIC CHANNELS

Zhenliang Zhang⋆*, Edwin K. P. Chong*⋆*, Ali Pezeshki*⋆*, and William Moran*†

⋆ Electrical & Computer Engineering, Colorado State University, Fort Collins, CO 80523, USA
† Defence Science Institute, Melbourne, VIC 3053, Australia

ABSTRACT

Each of a large number of nodes takes a measurement in sequence to decide between two hypotheses about the state of the world. Each node also has available the decisions of some of its immediate predecessors and uses these and its own measurement to make its decision. Each node broadcasts its decision through a binary symmetric channel, which randomly flips the decision. The question treated here is whether there exists a decision strategy consisting of a sequence of likelihood ratio tests such that the decisions approach the true hypothesis as the number of nodes increases. We show that if each node learns from bounded number of predecessors, then the decisions cannot converge to the underlying truth. We show that if each node learns from all predecessors then the decisions converge in probability to the underlying truth when the flipping probabilities are bounded away from $1/2$. We also derive, in the case when the flipping probabilities tend to $1/2$, a condition on the convergence rate of the flipping probabilities that is required for the decisions to converge to the true hypothesis in probability.

***Index Terms*—** Decentralized detection, social learning.

1. INTRODUCTION

We consider a large number of nodes, which sequentially make decisions between two hypotheses H_0 and H_1. At stage k, node a_k takes a measurement X_k (called its *private signal*), receives the decisions of its $m_k < k$ immediate predecessors, and makes a binary decision $d_k = 0$ or 1 about the prevailing hypothesis H_0 or H_1, respectively. It then broadcasts a decision to its successors. Note that m_k is often referred to as the *memory size*. A typical question is this: Can these nodes asymptotically learn the underlying true hypothesis? In other words, does the decision d_k converge (in probability) to the true hypothesis as $k \to \infty$? If so, what is the convergence rate of the error probability?

One application of the sequential hypothesis testing problem is decentralized detection in sensor networks, in which case the set of nodes represents a set of spatially distributed sensors attempting to jointly solve the hypothesis testing problem. Because of limited resources for processing and transmitting data, each sensor aggregates its measurement and the observed decisions from the previous sensors into a much smaller message (e.g., a 1-bit decision) and then sends it to other sensors for further aggregation. A central question is whether we can design a sequence of decision rules to aggregate the spatially distributed information such that the decisions converge to the underlying truth as the number of sensors increases.

Another application is to social learning in multi-agent networks, in which case the set of nodes represents a set of agents trying to learn the underlying truth (also known as the state of the world). Each agent makes a decision based on its own measurement and what it learns from the actions/decisions of the previous agents. In this case, we usually assume that each agent uses a myopic decision rule to minimize a local objective function; for example, the probability of error is locally minimized using the Bayesian likelihood ratio test with a threshold given by the ratio of the prior probabilities. The question in this setting is whether the agents in the social network can asymptotically learn the state of the world.

The research on our problem begins with a seminal paper by Cover [1], which considers the case where each node only observes the decision from its immediate previous node, i.e., $m_k = 1$ for all k. This structure is also known as a serial network or tandem network and has been studied extensively in [1]–[14]. We use $\mathbb{P}_j$ and π_j to denote the probability measure and the prior probability associated with H_j, $j = 0, 1$, respectively. Cover [1] shows that if the (log)-likelihood ratio for each private signal X_k is bounded almost surely, then the (Bayesian) error probability $\mathbb{P}_e^k = \pi_0\mathbb{P}_0(d_k = 1) + \pi_1\mathbb{P}_1(d_k = 0)$, using a sequence of likelihood ratio tests, does not converge in probability to 0 as $k \to \infty$. Conversely, if the likelihood ratio is unbounded, then the error probability converges to 0. In the case of unbounded likelihood ratios for the private signals, Veeravalli [8] shows that the error probability converges sub-exponentially with respect to the number k of nodes in the case where the private signals are independent and identically Gaussian distributed. Tay *et al.* [10] show that the convergence is in general sub-exponential and

This work was supported in part by AFOSR under contract FA9550-09-1-0518, and by NSF under grants CCF-0916314 and CCF-1018472.

derive an upper bound for the convergence rate of the error probability in the tandem network. Lobel *et al.* [11] derive a lower bound for the convergence rate of the error probability in the case where each node learns randomly from one previous node (not necessarily its immediate predecessor). In the case of bounded likelihood ratios, Drakopoulos *et al.* [12] provide a non-Bayesian decision strategy, which results in convergence of the error probability.

At the other extreme, consider the situation when each node can observe *all* the previous decisions, that is, $m_k = k-1$ for all k. This scenario was first studied in the context of social learning [15],[16], where each node uses the Bayesian likelihood ratio test to make its decision. In the case of bounded likelihood ratios for the private signals, the authors of [15] and [16] show that the error probability does not converge to 0, resulting in a wrong decision with positive probability. In [17], we show that in balanced binary trees, the decisions converge to the right decision even if the likelihood ratios of signals converge to 1 as the number of nodes increases. We further studied in [18] the convergence rate of the error probability in more general tree structures. In the case of unbounded likelihood ratios for the private signals, Smith and Sorensen [19] study this problem using martingales and show that the error probability converges to 0. Krishnamurthy [20],[21] studies this problem from the perspective of quickest time change detection. Chamley [22] provides a convergence rate analysis of the error probability in these structures. Acemoglu *et al.* [23] show that the nodes can asymptotically learn the underlying truth in more general network structures.

Most previous work including those reviewed above assume that the communication channels are perfect. We consider the situation where each broadcast decision is flipped with a certain probability, modeled by a binary symmetric channel. This situation was not considered in earlier studies. If the broadcast decision of a node is flipped, then all the successors of that node observe that flipped decision. We assume that each node uses a likelihood ratio test to generate its binary decision. We call the sequence of likelihood ratio tests a *decision strategy*. We want to know whether or not there exists a decision strategy such that the node decisions converge in probability to the underlying true hypothesis.

We use the following notation to characterize the scaling law of the asymptotic rate. Let f and g be positive functions defined on positive integers. We write $f(N) = O(g(N))$ if there exists a positive constant c_1 such that $f(N) \leq c_1 g(N)$ for sufficiently large N. We write $f(N) = \Omega(g(N))$ if there exists a positive constant c_2 such that $f(N) \geq c_2 g(N)$ for sufficiently large N.

We show that if each node can only learn from a bounded number of immediate predecessors, then for any decision strategy, the error probabilities cannot converge to 0. We also show that if each node can learn from *all* previous nodes, i.e., $m_k = k-1$, then the error probabilities converge to 0 using the myopic decision strategy, provided the flipping probabilities are bounded away from $1/2$. In this case, we show that the error probabilities converge to 0 as $\Omega(k^{-2})$. In the case where the flipping probabilities converge to $1/2$, we derive a necessary condition on the convergence rate of the flipping probabilities (that is, how fast they must converge) such that the error probabilities converge to 0. More specifically, we show that if there exists $p > 1$ such that the flipping probabilities converge to $1/2$ as $O(1/k(\log k)^p)$, then it is impossible for the error probability to converges to 0. Therefore, only if the flipping probabilities converge as $\Omega(1/k(\log k)^p)$ for some $p \leq 1$ can we hope for asymptotic learning.

This paper is a summary of results in a more extended paper [24]. We have omitted all the proofs for lack of space. A comprehensive treatment with proofs is presented in [24]. We also study in [24] the case where the broadcast message at each node is subject to random erasure.

2. PRELIMINARIES

We use $\mathbb{P}$ to denote the underlying probability measure. We use π_j to denote the prior probability (assumed nonzero) and $\mathbb{P}_j$ to denote the probability measure associated with H_j, $j = 0, 1$. At stage k, node a_k takes a measurement X_k of the scene and makes a decision $d_k = 0$ or $d_k = 1$ about the prevailing hypothesis H_0 or H_1. It then broadcasts its decision d_k through a binary symmetric channel, which randomly flips d_k into $\hat{d}_k$. The decision d_k of node a_k is made based on the private signal X_k and the sequence of potentially corrupted decisions $\hat{D}_{m_k} = \{\hat{d}_1, \hat{d}_2, \ldots, \hat{d}_{m_k}\}$ received from the m_k immediate predecessor nodes using a likelihood ratio test.

Our aim is to find a sequence of likelihood ratio tests such that the probability of making a wrong decision about the state of the world tends to 0 as $k \to \infty$. Before proceeding, we introduce the following definitions and assumptions:

1. The private signal X_k takes values in a set S, endowed with a σ-algebra $\mathcal{S}$. We assume that X_k is independent of the broadcast history $\hat{D}_{m_k}$. Moreover, the X_ks are mutually independent and identically distributed with distribution $\mathbb{P}_j^X$, under H_j, $j = 0, 1$. (Note that $\mathbb{P}_j^X$ is a probability measure on the σ-algebra $\mathcal{S}$.) We assume that the underlying hypothesis, H_0 or H_1, does not change with k.

2. The two probability measures $\mathbb{P}_0^X$ and $\mathbb{P}_1^X$ are equivalent; i.e., they are absolutely continuous with respect to each other. In other words, if $A \in \mathcal{S}$, then $\mathbb{P}_0^X(A) = 0$ if and only if $\mathbb{P}_1^X(A) = 0$.

3. Let the likelihood ratio of a private signal $s \in S$ be

$$L_X(s) = \frac{d\mathbb{P}_1^X}{d\mathbb{P}_0^X}(s),$$

where $d\mathbb{P}_1^X / d\mathbb{P}_0^X$ denotes the Radon–Nikodym derivative (which is guaranteed to exist because of the assumption that the two measures are equivalent). We assume that the likelihood ratios for the private signals are unbounded; i.e., for any set $S' \subset S$ with probability 1 under the measure $(\mathbb{P}_0^X + \mathbb{P}_1^X)/2$, we have

$$\inf_{s \in S'} \frac{d\mathbb{P}_1^X}{d\mathbb{P}_0^X}(s) = 0 \text{ and } \sup_{s \in S'} \frac{d\mathbb{P}_1^X}{d\mathbb{P}_0^X}(s) = \infty.$$

4. Suppose that θ is the underlying truth. Let $\bar{b}_k = \mathbb{P}(\theta = H_1 | X_k)$, which we call the *private belief* of a_k. By Bayes' rule, we have

$$\bar{b}_k = \left(1 + \frac{\pi_0}{\pi_1} \frac{1}{L_X(X_k)}\right)^{-1}. \tag{1}$$

5. Recall that node a_k observes m_k decisions $\hat{D}_{m_k}$ from its immediate predecessors. Let p_j^k be the conditional probability mass function of $\hat{D}_{m_k}$ under H_j, $j = 0, 1$. The likelihood ratio of a realization $\mathcal{D}_{m_k}$ is

$$L_D^k(\mathcal{D}_{m_k}) = \frac{p_1^k(\mathcal{D}_{m_k})}{p_0^k(\mathcal{D}_{m_k})} = \frac{\mathbb{P}_1(\hat{D}_{m_k} = \mathcal{D}_{m_k})}{\mathbb{P}_0(\hat{D}_{m_k} = \mathcal{D}_{m_k})}.$$

6. Let $b_k = \mathbb{P}(\theta = H_1 | \hat{D}_{m_k})$, which we call the *public belief* of a_k. We have

$$b_k = \left(1 + \frac{\pi_0}{\pi_1} \frac{1}{L_D^k(\hat{D}_{m_k})}\right)^{-1}. \tag{2}$$

7. Each node a_k makes its decision using its own measurement and the observed decisions based on a likelihood ratio test with a threshold $t_k > 0$:

$$d_k = \begin{cases} 1 & \text{if } L_X(X_k) L_D^k(\hat{D}_{m_k}) > t_k, \\ 0 & \text{if } L_X(X_k) L_D^k(\hat{D}_{m_k}) \le t_k. \end{cases}$$

   If $t_k = \pi_0/\pi_1$, then this test becomes the maximum a-posteriori probability (MAP) test, in which case the probability of error is locally minimized for node a_k. If $t_k = 1$, then the test becomes the maximum-likelihood (ML) test. If the prior probabilities are equal, then these two tests are identical. A decision strategy $\mathbb{T}$ is a sequence of likelihood ratio tests with thresholds $\{t_k\}_{k=1}^{\infty}$. Given a decision strategy, the decision sequence $\{d_k\}_{k=1}^{\infty}$ is a well-defined stochastic process.

8. We say that the system *asymptotically learns* the underlying true hypothesis with decision strategy $\mathbb{T}$ if

$$\lim_{k \to \infty} \mathbb{P}(d_k = \theta) = 1.$$

   In other words, the probability of making a wrong decision goes to 0, i.e., $\lim_{k\to\infty} \mathbb{P}_e^k = 0$. The question we are interested in is this: In each of the two classes of failures, is there a decision strategy such that the system asymptotically learns the underlying true hypothesis?

3. MAIN RESULTS

Recall that d_k is the input to a binary symmetric channel and $\hat{d}_k$ is the output, which is either equal to d_k (no flipping) or is equal to its complement $1 - d_k$ (flipping). The channel matrix is given by $\mathbb{P}(\hat{d}_k = i | d_k = j)$, $i, j = 0, 1$. We assume that $\mathbb{P}(\hat{d}_k = 1 | d_k = 0) = \mathbb{P}(\hat{d}_k = 0 | d_k = 1) = q_k$, where q_k denotes the probability of a flip. The assumption of symmetry is for simplicity only, and all results obtained in this section can be generalized easily to a general binary communication channel with unequal flipping probabilities, i.e., $\mathbb{P}(\hat{d}_k = 1 | d_k = 0) \neq \mathbb{P}(\hat{d}_k = 0 | d_k = 1)$.

Theorem 1. *Suppose that there exists C and $\epsilon > 0$ such that for all k, $m_k \le C$ and $q_k \in [\epsilon, 1 - \epsilon]$. Then, there does not exist a decision strategy such that the error probabilities converge to 0.*

Now consider the case where a_k can observe all its predecessors: $m_k = k - 1$. We will show that, using the myopic decision strategy, the error probabilities converge to 0 in the presence of random flipping when the flipping probabilities are bounded away from $1/2$. We further derive, for the case where the flipping probabilities converge to $1/2$, a necessary condition on the convergence rate of the flipping probabilities such that the error probabilities converge to 0.

If we state the conditions on the private signal distributions in a symmetric way, then it suffices to consider the case when the true hypothesis is H_0. In this case, our aim is to show that the Type I error probabilities converge to 0, that is, $\mathbb{P}_0(d_k = 1) \to 0$, or equivalently that the *public likelihood ratio* $L_k = \mathbb{P}_1(\hat{D}_k)/\mathbb{P}_0(\hat{D}_k)$ converges to 0. We consider the myopic decision strategy; that is, the decision made by the kth node is on the basis of the MAP test. From symmetry considerations, there is no loss in assuming that $q_k \le 1/2$. We consider two cases:

1) The flipping probabilities are bounded away from $1/2$ for all k; i.e., there exists $c > 0$ such that $q_k \le 1/2 - c$ for all k. This ensures that the corrupted decision still contains some useful information about the true hypothesis. We call this the case of *uniformly informative nodes*.

2) The flipping probabilities q_k converge to $1/2$ as $k \to \infty$. This means that the broadcast decisions become increasingly uninformative as we move towards the later nodes. We call this the case of *asymptotically uninformative nodes*.

3.1. Uniformly Informative Nodes

We first show that the error probabilities converge to 0. Recall that $\bar{b} = \mathbb{P}(H_1 | X)$ denotes the private belief given by signal X. Let $(\mathbb{G}_0, \mathbb{G}_1)$ be the conditional distributions of the private belief $\bar{b}$: $\mathbb{G}_j(s) = \mathbb{P}_j(\bar{b} \le s)$. These distributions exhibit two important properties:

a) *Proportionality*: This follows easily by Bayes' rule:

$$\frac{d\mathbb{G}_1}{d\mathbb{G}_0}(\bar{b}) = \frac{\bar{b}}{1-\bar{b}}.$$

b) *Dominance*: $\mathbb{G}_1(s) < \mathbb{G}_0(s)$ for all $s \in (0,1)$, and $\mathbb{G}_j(0) = 0$ and $\mathbb{G}_j(1) = 1$ for $j = 0, 1$. Moreover, $\mathbb{G}_1(r)/\mathbb{G}_0(r)$ is monotone non-decreasing as a function of r.

In the case of uniformly informative nodes, we can show that the network asymptotically learn the underlying truth.

Theorem 2. *Suppose that the flipping probabilities are bounded away from $1/2$. Then, $\mathbb{P}_e^k \to 0$ as $k \to \infty$.*

The proof of Theorem 2 involves the fact that the public likelihood ratio L_k is a martingale under H_0. By Doob's martingale convergence theorem, a non-negative martingale converges to a finite limit almost surely.

Suppose that the conditional densities of the private belief exists. By property a), we can write the conditional densities of the private belief as follows:

$$f^1(\bar{b}) = \frac{d\mathbb{G}_1}{d\bar{b}}(\bar{b}) = \bar{b}\rho(\bar{b}),\ f^0(\bar{b}) = \frac{d\mathbb{G}_0}{d\bar{b}}(\bar{b}) = (1-\bar{b})\rho(\bar{b}),$$

where $\rho(\bar{b})$ is a non-negative function. Next we provide a lower bound on the error probability with respect to the number of nodes. For simplicity, we assume that $\rho(1)$ is a non-negative constant. More general cases where $\rho(\bar{b}) \to 0$ as $\bar{b} \to 1$ are studied in [24].

Theorem 3. *Suppose that the flipping probabilities are bounded away from $1/2$ and $\rho(1)$ is a non-negative constant. Then, the Type I error probability converges to 0 as $\Omega(k^{-2})$.*

3.2. Asymptotically Uninformative Nodes

In this section, we consider the case where the node decisions become increasing uninformative; that is, $q_k \to 1/2$. Let $Q_k = (1-2q_k)/(1-q_k)$. Note that $q_k \to 1/2$ if and only if $Q_k \to 0$. This parameter measures how "informative" the corrupted decision is. For example, if $q_k = 0$ (where there is no flipping), then the decision is maximally informative in terms of updating the public belief. However if $q_k = 1/2$, in which case $Q_k = 0$, then the decision is completely uninformative in terms of updating the public belief.

Theorem 4. *Suppose that there exists $p > 1$ such that $Q_k = O\left(\frac{1}{k(\log k)^p}\right)$. Then, the public belief converges to a nonzero limit almost surely.*

It is evident that if the public belief converges to a nonzero limit almost surely, then $\mathbb{P}_0(d_k = 1)$ is bounded away from 0 and $\mathbb{P}_0(d_k = 0)$ is bounded away from 1. In consequence, the system does not asymptotically learn the true situation, and Theorem 4 provides a necessary condition for asymptotic learning.

Theorem 4 implies that for the public belief to tend to zero with positive probability, there must exist a $p \le 1$ such that $Q_k = \Omega(1/k(\log k)^p))$. If the public belief does not converge to zero, then it is impossible for there to be an eventual collective arrival at the true hypothesis. To explain this further, let $\mathcal{H}$ denote the event that there exists a (random) k_0 such that the sequence of decisions $d_k = 0$ (true hypothesis) for all $k \ge k_0$. Occurrence of this event signifies that after a finite number of decisions, the agents arrive at the true underlying state. Such an outcome also means that, eventually, each agent's private signal is overpowered by the past collective true verdict, so that a false decision is never again declared. In the literature on social learning, this phenomenon is called *information cascade* (e.g., [25]) or *herding* (e.g., [19]). We use $\mathcal{L}$ to denote the event $\{b_k \to 0\}$. Notice that $\mathcal{H}$ occurs only if $\mathcal{L}$ occurs. Hence, $\mathcal{H}$ is a subset of the event that $b_k \to 0$; that is, $\mathcal{H} \subset \mathcal{L}$. These leads to the following corollary of Theorem 4.

Corollary 1. *If $Q_k = O(1/k(\log k)^p)$ for some $p > 1$, then $\mathbb{P}(\mathcal{H}) = 0$.*

4. CONCLUSION

We have studied the sequential hypothesis testing problem in a feedforward network in which node decisions experience random flippings after broadcasts. We show that if the memory sizes are bounded, then there does not exist a decision strategy such that the error probabilities converge to 0. If each node learns from all the previous decisions, then with the myopic decision strategy, the error probabilities converge to 0, provided the flipping probabilities are bounded away from $1/2$. In the case where the flipping probabilities converge to $1/2$, we derive a necessary condition on the convergence rate of the flipping probabilities such that the error probabilities converge to 0.

Our analysis leads to several open questions. We wish to study the case where the memory size goes to infinity but each node cannot learn from *all* previous decisions. We also want to generalize the techniques used in this paper to more general network topologies. Moreover, besides erasure and flipping failures, we expect that our techniques can be used in the additive Gaussian noise scenario. With finite signal-to-noise ratios (SNR), the martingale convergence proof in Theorem 2 easily generalizes to this scenario. However, if SNR goes to 0 (e.g., the fading coefficient goes to 0, the noise variance goes to infinity, or the broadcasting signal power goes to 0), it is obvious that the convergence of error probability is not always true. We want to derive necessary and sufficient conditions on the convergence rate of SNR such that the error probability still converges to 0.

5. REFERENCES

[1] T. M. Cover, “Hypothesis testing with finite statistics,” *Ann. Math. Statist.*, vol. 40, no. 3, pp. 828–835, 1969.

[2] M. E. Hellman and T. M. Cover, “Learning with finite memory,” *Ann. Math. Statist.*, vol. 41, no. 3, pp. 765–782, 1970.

[3] P. Swaszek, “On the performance of serial networks in distributed detection,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 29, no. 1, pp. 254–260, Jan. 1993.

[4] Z. B. Tang, K. R. Pattipati, and D. L. Kleinman, “Optimization of detection networks: Part I—Tandem structures,” *IEEE Trans. Syst., Man and Cybern.*, vol. 21, no. 5, pp. 1044–1059, Sept./Oct. 1991.

[5] R. Viswanathan, S. C. A. Thomopoulos, and R. Tumuluri, “Optimal serial distributed decision fusion,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 24, no. 4, pp. 366–376, Jul. 1988.

[6] J. Koplowitz, “Necessary and sufficient memory size for m-hypothesis testing,” *IEEE Trans. Inform. Theory*, vol. IT-21, no. 1, pp. 44–46, Jan. 1975.

[7] J. D. Papastravrou and M. Athans, “Distributed detection by a large team of sensors in tandem,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 28, no. 3, pp. 639–653, Jul. 1992.

[8] V. V. Veeravalli, “Topics in decentralized detection,” Ph.D. dissertation, Univ. Illinois, Urbana-Champaign, 1992.

[9] J. N. Tsitsiklis, “Decentralized detection,” *Adv. Statist. Signal Process.*, vol. 2, pp. 297–344, 1993.

[10] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, “On the sub-exponential decay of detection error probabilities in long tandems,” *IEEE Trans. Inform. Theory*, vol. 54, no. 10, pp. 4767–4771, Oct. 2008.

[11] I. Lobel, D. Acemoglu, M. A. Dahleh, and A. Ozdaglar, “Lower bounds on the rate of learning in social networks,” in *Proc. IEEE American Control Conf.*, Hyatt Regency Riverfront, St. Louis, MO, June 10-12, 2009, pp. 2825–2830.

[12] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis, “On learning with finite memory,” *preprint*, available in arXiv:1209.1122.

[13] P. K. Varshney, *Distributed detection and data fusion*, New York: Springer-Verlag, 1997.

[14] P. M. Djuric and Y. Wang, “Disturbed Bayesian learning in multiagent systems: Improving our understanding of its capabilities and limitations,” *IEEE Signal Process. Magazine*, vol. 29, no. 2, pp. 65–76, Mar. 2012.

[15] A. V. Banarjee, “A simple model for herd behavior,” *Quart. J. Econ.*, vol. 107, no. 3, pp. 797–817, Aug. 1992.

[16] S. Bikchandani, D. Hirshleifer, and I. Welch, “A theory of fads, fashion, custom, and cultural change as information cascades,” *J. Political Econ.*, vol. 100, no. 5, pp. 992–1026, Oct. 1992.

[17] Z. Zhang, A. Pezeshki, W. Moran, S. D. Howard, and E. K. P. Chong, “Error probability bounds in balanced binary relay trees,” *IEEE Tran. Inform. Theory*, vol. 58, no. 6, pp. 3548–3563, Jun. 2012.

[18] Z. Zhang, E. K. P. Chong, A. Pezeshki, W. Moran, and S. D. Howard, “Learning in hierarchical social networks,” *IEEE J. Sel. Topics Signal Process.*, to appear.

[19] L. Smith and P. Sorensen, “Pathological outcomes of observational learning,” *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000.

[20] V. Krishnamurthy, “Bayesian sequential detection with phase-distributed change time and nonlinear penalty–A POMDP lattice programming approach,” *IEEE Trans. Inform. Theory*, vol. 57, no. 10, pp. 7096–7124, Oct. 2011.

[21] V. Krishnamurthy, “Quickest detection POMDPs with social learning: Interaction of local and global decision makers,” *IEEE Trans. Inform. Theory*, vol. 58, no. 8, pp. 5563–5587, Aug. 2012.

[22] C. P. Chamley, *Rational herds*, New York: Cambridge, 2004.

[23] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, “Bayesian learning in social networks,” *Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[24] Z. Zhang, E. K. P. Chong, A. Pezeshki, and W. Moran, “Hypotheses testing in feedforward networks with broadcast failures,” *preprint*, available from arXiv:1211.4518.

[25] J. Surowiecki, *The wisdom of crowds*, New York: Doubleday, 2005.