

DISTRIBUTED HIDDEN MARKOV MODELS FOR SOCIAL NETWORKS

Gokulnath¹, Kanishk², Hareesh³

¹SRM College, India.

ABSTRACT

This article explores the application of a distributed Hidden Markov Model (HMM) over graphs to address the problem of detecting influential and normal nodes within dynamic social networks. We review the state-of-the-art approaches for multi-agent filtering and examine a novel algorithm capable of superior state estimation in decentralized settings. Our findings indicate that the proposed solution not only offers accurate tracking of time-varying states but also outperforms traditional social learning algorithms, making it a valuable tool for various fields including social networks, environmental monitoring, and target tracking.

1. INTRODUCTION

Tracking hidden states within dynamic environments, such as identifying the location of objects or the diffusion of opinions in social networks, requires advanced filtering algorithms. Among the available techniques, Hidden Markov Models (HMMs) provide an ideal framework for these tasks due to their capacity to model sequential data.

In decentralized settings, where data is distributed among multiple agents or nodes, traditional centralized methods can become inefficient. The focus of this work is on developing decentralized algorithms where agents, such as users within a social network, collaborate to infer hidden states based on their local observations and communication with neighbors. This scenario finds applications in numerous areas such as environmental monitoring, target tracking, and notably, analyzing opinion formation in social networks.

In particular, previous studies have focused on decentralized Kalman filters and Bayesian frameworks. These approaches either rely on linear system models or multiple rounds of communication between agents, which may not be practical in fast-changing environments. To overcome these challenges, we propose an efficient distributed HMM filtering algorithm that only requires a single round of communication for each state change.

Algorithm Description

The core idea behind the algorithm is that a network of agents, each observing data from an evolving system, can infer the system's hidden states collaboratively. In our context, each agent represents a node in a social network, tasked with tracking influential users or topics over time. The influence of these users can significantly affect the overall opinions in the network.

Agents communicate and update their beliefs regarding the true hidden state of the system by incorporating their observations into a probabilistic model. In the case of social networks, this allows nodes to adjust their perception of opinion leaders or influencers, based on their neighbors' feedback and real-time information.

The algorithm operates in two main stages. First, agents revise their previous beliefs by calculating a time-adjusted prior based on the Markov transition probabilities. Second, each agent updates its belief using a Bayesian update rule, considering new local observations. Once these steps are completed, agents exchange their updated beliefs with their neighbors to ensure that information diffuses throughout the network. The use of time-adjusted priors allows agents to better track

changes in rapidly evolving systems compared to traditional methods, which may not account for these dynamics effectively.

2. MAIN RESULTS

The performance of the proposed algorithm is evaluated by comparing its deviation from the optimal centralized solution. The algorithm achieves asymptotically bounded performance, meaning that as the number of iterations increases, the difference between the centralized and distributed beliefs remains within an acceptable range. This makes it particularly useful for environments where the hidden states change frequently.

One of the key strengths of the proposed approach lies in its ability to handle rapidly changing states. Unlike previous methods that require slow convergence to account for state transitions, our algorithm leverages geometrically ergodic transition models to provide faster convergence and more accurate tracking. In particular, it performs better in scenarios where the state transition probabilities are fast, which is often the case in real-world social networks.

Furthermore, the algorithm is robust to different network structures. Even in sparse networks, where agents may have limited communication with their neighbors, the performance remains stable. This property is crucial for applications in social networks where connectivity may not always be dense.

3. RESULTS

To validate the theoretical findings, we ran simulations on a network of 10 agents. The agents were tasked with tracking a hidden state that could transition between two hypotheses. Our simulation results demonstrate that the proposed distributed HMM algorithm closely follows the performance of a centralized HMM filter, which has access to all agent observations. In particular, the distributed filter exhibits a remarkable ability to track abrupt changes in the hidden state, which is essential for real-time applications in social networks.

Moreover, we compared the proposed algorithm with a state-of-the-art social learning algorithm, known as Adaptive Social Learning (ASL). While ASL provides satisfactory performance, it is slower to respond to rapid state transitions due to its reliance on fixed state assumptions. The proposed algorithm, by contrast, accounts for the dynamic nature of state transitions, allowing it to respond more quickly to changes in the system.

An additional experiment explored the effect of different network topologies on the algorithm's performance. As expected, networks with more connected agents performed better due to the faster diffusion of information. However, even in less connected networks, the distributed HMM algorithm was able to maintain accurate tracking of the hidden states, highlighting its robustness.

4. DISCUSSION

This article presented a distributed HMM filtering algorithm designed for decentralized environments such as social networks. The algorithm's ability to efficiently track time-varying states with minimal communication between agents makes it well-suited for real-time applications. By leveraging prior information about state dynamics, it outperforms traditional social learning algorithms in terms of both accuracy and speed.

Future research could explore further optimizations to handle more complex scenarios where the number of hidden states is large, or where the state transition probabilities vary significantly over time. Additionally, integrating the proposed algorithm with real-world social network data could offer valuable insights into opinion formation and information diffusion processes.

5. REFERENCES

- [1] Olfati-Saber, R. (2007). "Distributed Kalman filtering for sensor networks." IEEE Conference on Decision and Control, 5492-5498.
- [2] Khan, U. A., & Moura, J. M. (2008). "Distributing the Kalman filter for large-scale systems." IEEE Transactions on Signal Processing, 56(10), 4919-4935.
- [3] Krishnamurthy, V. (2016). Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing. Cambridge University Press.
- [4] Djuric, P. M., et al. (2008). "Distributed particle filtering for dynamic state estimation." IEEE Signal Processing Magazine, 25(4), 20-34.
- [5] Jadbabaie, A., Lin, J., & Morse, A. S. (2003). "Coordination of groups of mobile autonomous agents using nearest neighbor rules." IEEE Transactions on Automatic Control, 48(6), 988-1001.
- [6] Nedic, A., & Ozdaglar, A. (2009). "Distributed subgradient methods for multi-agent optimization." IEEE Transactions on Automatic Control, 54(1), 48-61.
- [1] Kar, S., & Moura, J. M. (2012). "Consensus + innovations distributed inference over networks: Cooperation and sensing in networked systems." IEEE Signal Processing Magazine, 29(5), 15-30.
- [2] Rabbat, M. G., & Nowak, R. D. (2004). "Distributed optimization in sensor networks." IEEE/ACM Transactions on Networking (TON), 13(5), 917-930.