On Word-of-Mouth and Private-Prior Sequential Social Learning

Andrea Da Col, Cristian R. Rojas and Vikram Krishnamurthy

I. Introduction

The growing prominence of generative machine learning models [1] has raised significant concerns about data incest-induced model collapse [2]. As synthetic content proliferates on public platforms, it is increasingly likely to be incorporated into future training sets, leading to model degradation [3]. This paper examines the online learning dynamics of a group of decision-makers whose outputs are incorporated in future training data.

We approach this problem through the lens of social learning, a fundamental mathematical framework for modeling interactions between social sensors [4, 5]. In classical social learning models, agents form estimates of an unknown state by combining private beliefs with the observed actions of others. These models typically assume that each agent updates her *private prior* (PP) belief as new actions are revealed. A second architecture is considered, inspired by *Word-of-Mouth* (WoM) learning [6, 7], in which only the initial agent has access to external state information, while all subsequent agents rely solely on the propagated beliefs of their predecessors. The final agent broadcasts her belief to the entire network, thus creating a shared public dataset.

II. Problem Formulation

A set $\mathcal{I} = \{1, \dots, m\}$ of agents aim to learn a time-dependent state of nature from private online measurements. The agents are only allowed to operate sequentially, i.e., one cannot take actions before her predecessor.

A. Model and Agent Description

We consider a data-generating mechanism governed by the asymptotically stable first-order autoregressive dynamics:

$$x_k = a x_{k-1} + w_k, \qquad k \in \mathbb{N}, \tag{1}$$

where $x_k, w_k \in \mathbb{R}$ are the state of the system and the process noise at time k, respectively. We assume that the agents $i \in \mathcal{I}$ can measure changes in the state according to:

$$y_k^i = x_k + n_k^i, \qquad k \in \mathbb{N},\ i \in \mathcal{I}, \tag{2}$$

where $y_k^i, n_k^i \in \mathbb{R}$ are the observed output and the measurement noise of agent i at time k, respectively. The random variables x_0, (w_k) and (n_k^i) are assumed mutually independent, with $x_0 \sim \mathcal{N}(\hat{x}_0, p_0)$, $w_k \sim \mathcal{N}(0, q)$, and $n_k^i \sim \mathcal{N}(0, r_k^i)$. An agent i implementing (1) and (2) can optimally estimate x_k from her own local observations $\{\tilde{y}_1^i, \dots, \tilde{y}_k^i\}$ by using the Kalman filter update rule

$$\begin{aligned}
p_{k|k-1}^i &= a^2 p_{k-1|k-1}^i + q, \\
\hat{x}_{k|k-1}^i &= a \hat{x}_{k-1|k-1}^i, \\
\alpha_k^i &= p_{k|k-1}^i / (p_{k|k-1}^i + r_k^i), \\
p_{k|k}^i &= p_{k|k-1}^i (1 - \alpha_k^i), \\
\hat{x}_{k|k}^i &= \hat{x}_{k|k-1}^i + \alpha_k^i (\tilde{y}_k^i - \hat{x}_{k|k-1}^i),
\end{aligned} \tag{3}$$

where the variances $p_{k|k-1}^i$ and $p_{k|k}^i$, as well as the Kalman gains α_k^i, do not depend on data, and are assumed to be publicly available to all agents. We will denote by (v_k^i) the additive white Gaussian noise injected immediately before i, where $v_k^i \sim \mathcal{N}(0, s^i)$ and $s^i \in \mathbb{R}_{>0}$. We also assume that (v_k^i) is independent of the other random variables, namely x_0, (w_k), and (v_k^j) for $i \neq j$.

B. Private-Prior Setup

The prediction step does not depend on the most recently collected measurement $\tilde{y}_k^i$, and hence

$$\begin{aligned}
p_{k|k-1}^i &= a^2 p_{k-1|k-1}^i + q, \\
\hat{x}_{k|k-1}^i &= a \hat{x}_{k-1|k-1}^i
\end{aligned} \tag{4}$$

can be updated independently by each $i \in \mathcal{I}$ using her private prior information $\hat{x}_{k-1|k-1}^i$ of x_{k-1}. On the other hand, the a-posteriori update depends on the estimate of other decision-makers. At each time step k, the leftmost agent can observe the unknown state directly through $y_k^1 = x_k + v_k^1$. This observation model is the same as in (2), but with $n_k^1 = v_k^1$. Thus, the estimates for agent 1 can be obtained from (3). The Kalman filters of subsequent agents $i \in \mathcal{I} \setminus \{1\}$ can see the estimate of its predecessor, after the processing operation

$$y_k^i = x_k + n_k^{i-1} + v_k^i / \alpha_k^{i-1}.$$

Using a recursive expression, we are able to define the equivalent noise $n_k^i := n_k^{i-1} + v_k^i / \alpha_k^{i-1}$, which is a zero-mean Gaussian random variable. The posterior mean $\hat{x}_{k|k}^i$ is obtained with the following measurement update:

$$\begin{aligned}
\alpha_k^i &= p_{k|k-1}^i \Big/ \Big[p_{k|k-1}^i + s^1 + \textstyle\sum_{j=2}^i s^j \big/ \big(\alpha_k^{j-1}\big)^{\mathbf{2}} \Big], \\
p_{k|k}^i &= p_{k|k-1}^i (1 - \alpha_k^i), \\
\hat{x}_{k|k}^i &= \hat{x}_{k|k-1}^i + \alpha_k^i (\tilde{y}_k^i - \hat{x}_{k|k-1}^i), \qquad\qquad i \in \mathcal{I}.
\end{aligned} \tag{5}$$

C. Word-of-Mouth Setup

Our previous derivations can be used as the starting point for describing WoM social learning. After substituting the

Andrea Da Col and Cristian R. Rojas are with the Division of Decision and Control Systems, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden. Vikram Krishnamurthy is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, 14853, USA.

quantity $x^i_{k-1|k-1} = x^m_{k-1|k-1}$ into (3) we obtain

$$\begin{aligned} p^i_{k|k-1} &= a^2 p^m_{k-1|k-1} + q, \\ \hat{x}^i_{k|k-1} &= a\hat{x}^m_{k-1|k-1}, \\ \alpha^i_k &= p^m_{k|k-1} \Big/ \Big[p^m_{k|k-1} + s^1 + \textstyle\sum_{j=2}^i s^j / \big(\alpha^{j-1}_k\big)^{\mathbf{2}} \Big], \\ p^i_{k|k} &= p^m_{k|k-1}(1-\alpha^i_k), \\ \hat{x}^i_{k|k} &= \hat{x}^m_{k|k-1} + \alpha^i_k(\tilde{y}^i_k - \hat{x}^m_{k|k-1}), \end{aligned} \tag{6}$$

that is, the update rule for WoM social learning.

III. Analysis of PP and WoM Learning

A. Private-Prior Setup

We expect the following cascade behavior to manifest for a cascade of agents embedded in a PP setup: after the Kalman Filter of agent i converges to her steady-state, so will do the Kalman filter of agent $i+1$, for every $i \in \mathcal{I} \setminus \{m\}$.

Theorem 1: *Consider the PP interconnection of Kalman filters implementing* (4) *and* (5). *Then, there exists a unique positive solution* p^i_∞ *to*

$$p^i_\infty = a^{\mathbf{2}} p^i_\infty - \frac{(a p^i_\infty)^{\mathbf{2}}}{p^i_\infty + r^i_\infty} + q, \qquad i \in \mathcal{I},$$

where $r^i_\infty = s^1 + \sum_{j=2}^i s^j / (\alpha^{j-1}_\infty)^{\mathbf{2}}$. *Furthermore,* $(1-\alpha^i_\infty)a \in (-1,1)$, *and* $p^i_{k|k-1} \to p^i_\infty$ *for any* $p^i_{1|0} \in \mathbb{R}_{>0}$. *This result holds for every* $i \in \mathcal{I}$.

B. Word-of-Mouth Setup

To characterize the asymptotic behavior of the predicted variance when the agents implement the WoM setup, we study the following update rule:

$$T(p^m_{k|k-1}) = a^{\mathbf{2}} \left(\frac{p^m_{k|k-1} r^m_k}{p^m_{k|k-1} + r^m_k} \right) + q. \tag{7}$$

The following theorem formalizes the existence of a unique positive fixed point for (7).

Theorem 2: *Consider the WoM interconnection of Kalman filters implementing* (6). *Then,* (7) *has a unique positive fixed point* p^m_∞.

Theorem 2 tells us that, if the positive sequence $(p^m_{k|k-1})$ is convergent, then it must converge to p^m_∞. The following theorem establishes the convergence of $(p^m_{k|k-1})$ to p^m_∞ for the case $m = 2$.

Theorem 3: *If* $m = 2$, *then* (7) *admits a unique positive fixed point* p^2_∞, *and* $p^2_{k|k-1} \to p^2_\infty$ *for any* $p^2_{1|0} \in \mathbb{R}_{>0}$.

For more details the reader is referred to [8].

IV. Numerical Examples

A key finding of our study is that some agents, particularly those with large $i \in \mathcal{I}$ attain smaller $p^i_{k|k-1}$, benefit from using the public belief (WoM) rather than relying on their private knowledge (PP). Intuitively, as i increases, the prediction and posterior variances increase, due to the larger variance of the equivalent noise. Thus, in the WoM setup, all the agents (save for the last one) are forced to use a more spread-out prior than their own. Inevitably, agent 1, who alone can enjoy measurement noise with constant variance, will suffer a performance degradation, and output a larger Kalman gain compared to the (optimal) private-prior case. In other words, it will put more trust on the measurements than on its prior. As a result of this, the next agent will enjoy an equivalent noise with smaller variance. Together with being fed worse prior knowledge, she will also trust her measurements (i.e., the estimate of her predecessor) more than her prior belief, resulting in a larger Kalman gain. One can apply this reasoning inductively to all agents with $i > 1$, and when it comes to the last agent, the variance of her equivalent noise has decreased considerably compared to the PP case. This, plus the fact that she has been using her own prior (which corresponds to the public prior in the WoM case), makes the Kalman filter of agent m converge to a steady-state filter with a smaller prediction error variance p^m_∞ than in the PP setup. This is evident in Figure 1.

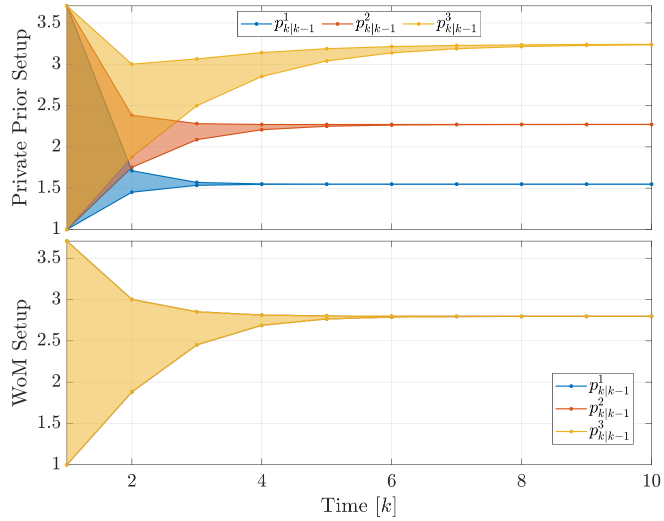

Fig. 1: Asymptotic behavior of the prediction error variances in the PP (upper panel) and WoM cases (lower panel).

References

[1] H. Cao, C. Tan, Z. Gao, Y. Xu, G. Chen, P.-A. Heng, and S. Z. Li, "A survey on generative diffusion models," *IEEE Transactions on Knowledge and Data Engineering*, 2024.

[2] M. Marchi, S. Soatto, P. Pratik, and P. Tabuada, "Heat death of generative models in closed-loop learning," *arXiv preprint arXiv:2404.02325*, 2024.

[3] B. Biggio, B. Nelson, and P. Laskov, "Poisoning attacks against support vector machines," *arXiv preprint arXiv:1206.6389*, 2012.

[4] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, pp. 3–49, 2011.

[5] A. V. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, 1992.

[6] X. Vives, "How fast do rational agents learn?," *The Review of Economic Studies*, vol. 60, no. 2, pp. 329–347, 1993.

[7] X. Vives, "Learning from others: A welfare analysis," *Games and Economic Behavior*, vol. 20, no. 2, pp. 177–200, 1997.

[8] A. Da Col, C. R. Rojas, and V. Krishnamurthy, "On word-of-mouth and private-prior sequential social learning," *In Proceedings of the 64th IEEE Conference on Decision and Control (CDC) (submitted for publication)*, 2025.