

ON THE FUSION STRATEGIES FOR FEDERATED DECISION MAKING

Mert Kayaalp*, Yunus İnan*, Visa Koivunen†, Emre Telatar*, Ali H. Sayed*

*École Polytechnique Fédérale de Lausanne (EPFL)

†Aalto University

ABSTRACT

We consider the problem of information aggregation in federated decision making, where a group of agents collaborate to infer the underlying state of nature without sharing their private data with the central processor or each other. We analyze the non-Bayesian social learning strategy in which agents incorporate their individual observations into their opinions (i.e., soft-decisions) with Bayes rule, and the central processor aggregates these opinions by arithmetic or geometric averaging. Building on our previous work, we establish that both pooling strategies result in asymptotic normality characterization of the system, which, for instance, can be utilized in order to give approximate expressions for the error probability. We verify the theoretical findings with simulations and compare both strategies.

1. INTRODUCTION AND RELATED WORK

The centralized strategy has been the traditional norm for collaborative decision-making, where data is collected from all sources and processed at a central controller, fusion center or cloud. However, this approach raises concerns about data privacy and communication costs. In this work, we study the paradigm of federated decision-making, which enables cooperative decision-making without data collection [1]. For example, in healthcare applications, federated decision-making can be used to test possible hypotheses without gathering private data from siloed institutions [2]. Moreover, in classification tasks, agents (e.g., mobile device users) may have different views of a physical phenomenon such as scenery as in multi-view learning [3], and the server may want to classify the scene or infer the state of environment without collecting the local data of users.

The setting we consider in this work, which is described in detail in Sec. 2, is an instance of locally Bayesian (a.k.a. non-Bayesian) social learning [4–7] that is based on general decentralized networks. Here, we only consider the special case of a star network topology where all agents are connected to a central node. In social learning, the goal is to infer the hypothesis that best describes the observations received from the environment. The true hypothesis is common for all agents and hence the cooperation is of interest for all agents. To this end, at each time instant, (i) distributed agents process their private data locally and send their opinions (i.e., soft-decisions) to the server, (ii) the server pools the information received from agents and broadcasts the aggregated belief back to the agents, (iii) the agents use the broadcasted belief as a prior belief for their subsequent opinions.

Note that this setting has close ties with the distributed detection literature [8–11], which typically involves sharing log-likelihood ratios for each data and assumes spatial independence and/or homo-




Fig. 1: Data types at the edge devices can be highly heterogeneous.¹

geneity across agents. In contrast, in the social learning framework, agents share beliefs (posteriors) formed over time, rather than a statistic about individual data. By this way, they can incorporate their individual priors as well. Furthermore, in social learning, data is not required to be independent between agents nor identically distributed. In fact, agents can possess different data types (e.g., text, video), as in multi-modal machine learning — see Fig. 1. The federated decision-making is also related to federated learning [12, 13] as they both fall under the umbrella of federated analytics [1], which aims to enable cooperation among multiple users while preserving privacy of data. However, federated learning focuses on training machine learning models with labeled data samples, whereas the current setting involves making collaborative inferences without sharing raw observations.

In this paper, we address the following problem: How should the server aggregate information received from multiple distributed agents that process the data locally and share only their local soft decision statistics? Two of the most widely used methods for combining probability vectors are *arithmetic* averaging (AA), a.k.a. linear opinion pooling, and *geometric* averaging (GA), a.k.a. logarithmic opinion pooling. The distinct properties of these procedures are well-established in the literature — see, e.g., [14]. In our previous work [15], we investigated the performances of AA and GA strategies in the general social learning setting. For the special case of federated architectures, we derived an exact expression for the asymptotic convergence rates of beliefs. Different from prior works that only compared AA and GA with limited assumptions, such as assuming only one step of fusion [16–18] or assuming specific/particular distributions like Gaussian or Poisson [16, 18, 19]; our analysis considered repeated fusion of beliefs without any assumption on the distribution of data. In this work, we extend our previous findings and show that the asymptotic behaviors of both AA and GA can be characterized with normal distributions and the

This work was supported in part by grant 205121-184999 from the Swiss National Science Foundation (SNSF). Emails: {mert.kayaalp, yunus.inan, emre.telatar, ali.sayed}@epfl.ch., visa.koivunen@aalto.fi

¹Image credit: freepik.com.

asymptotic error probabilities of both strategies can be established accordingly. The obtained results provide further insight into the arithmetic and geometric fusion rules, and can be utilized for designing federated decision-making systems.

Notation. Random variables are written in boldface letters. We use “proportional to” symbol \propto whenever LHS of an equation is a proper normalization of RHS, and is an entry of a probability mass function. The KL-divergence between two distributions p and q is denoted as $D_{\text{KL}}(p||q)$.

2. FEDERATED DECISION MAKING

We consider a distributed hypothesis testing task in which a group \mathcal{K} of K agents, aided by a central controller or server, perform inferences on a phenomenon of interest in their environment. More formally, the agents seek to choose the true hypothesis θ° from a finite set of H hypotheses $\Theta = \{1, 2, \dots, H\}$. At each time instant i , agent k receives personal observation $\xi_{k,i}$ from the environment that is distributed according to the marginal likelihood function $L_k(\cdot|\theta^\circ)$ (since the true hypothesis is θ°). Notice that the likelihood functions depend on agent index k , i.e., the data distributions across different agents can be heterogeneous. We assume that the environment is stationary, and that the data is independent and identically distributed (i.i.d.) over time at each agent. However, we do not make any assumptions on spatial dependence, i.e., there is no assumption on the dependence among different agents. Each agent k only knows its own local likelihood model $L_k(\xi|\theta)$ for each $\theta \in \Theta$ that represents how likely that an observation ξ is produced by θ . Furthermore, the server does not need to possess any information about the distribution of data on agents.

The agents do not share their private raw data with the server: all data processing and ownership are local. Each agent k only shares its belief (i.e., opinion) $\psi_{k,i}$ about the possible hypotheses, which is a probability mass function over Θ . In other words, $\psi_{k,i}(\theta)$ represents the amount of confidence agent k has on the proposition “ $\theta = \theta^\circ$ ”, at time instant i . Once the server receives the intermediate beliefs from distributed agents, it pools them and broadcasts the averaged belief back to the agents. In this work, we study the widely used linear (AA) and logarithmic (GA) pooling strategies, although there are other variations [14]. We now motivate the setting with two real-world applications.

Example 1 (Crowd Size Estimation). *Crowd counting aims to estimate the number of people in a given area using visual or other sensory data. The possible hypotheses are the size of the crowd, such as “the number of people in the event is between 500-1000”. Agents, i.e., users cooperating with the server, can obtain data in the form of photos and videos of their surroundings, or their mobile devices can measure the number of possible connections over Bluetooth. Due to privacy concerns, users may prefer not to share their raw data with the server. In addition, transmitting raw data can create communication bottlenecks at the server. Therefore, employing the social learning strategy, users can only share their beliefs over the set of hypotheses. They can form their beliefs by using local likelihood models, e.g., by using pre-trained neural network models on their devices that can estimate the crowd size based on images [20].* ■

Example 2 (Flexible Spectrum Use). *Spatially distributed sensors or mobile devices can try to estimate the number of targets or active wireless emitters within a particular area. For example, in the radio spectrum, the mobile devices can measure the local occupancy and*

field levels to distinguish between idle and congested spectrum, or map the radio environment. This practice can enhance the environmental awareness, and allocate the resources optimally in order to improve the wireless communication. Exchanging beliefs only rather than raw data can be advantageous in terms of privacy and communication costs. ■

To form the beliefs, agents execute the non-Bayesian social learning strategy [4–7] under the special case of star topologies. At each time instant i , agents first process their observations in a *locally* Bayesian manner:

$$\psi_{k,i}(\theta) \propto L_k(\xi_{k,i}|\theta)\mu_{i-1}(\theta) \quad (\text{Adapt}), \quad (1)$$

and then share these intermediate beliefs with the server. If the server employs AA for information fusion, the updated belief is a weighted arithmetic average of the intermediate beliefs from different agents:

$$\mu_i(\theta) = \sum_{k \in \mathcal{K}} \pi_k \psi_{k,i}(\theta) \quad (\text{AA}). \quad (2)$$

Here, $\pi = [\pi_1, \dots, \pi_K]^\top$ is a vector of the confidence weights π_k the server assigns to each agent k [9, 21]. We assume that these weights are positive, sum up to 1, and constant over time. Alternatively, if the fusion rule is GA, the updated belief is a weighted geometric average:

$$\mu_i(\theta) \propto \prod_{k \in \mathcal{K}} (\psi_{k,i}(\theta))^{\pi_k} \quad (\text{GA}). \quad (3)$$

The server then sends the updated belief to the agents, and they execute the same steps repeatedly over time by using the updated beliefs and locally observed data.

Remark 1. *An important distinction between AA and GA lie in their handling of initial beliefs and their sensitivity to individual beliefs. In AA, only one agent having a positive initial belief on the true hypothesis is sufficient. In contrast, in GA, all agents must have positive initial beliefs on the truth in order not to discard it from the beginning. Furthermore, GA is more sensitive to small beliefs compared to AA. This is because, if some agents’ beliefs are very small, they impact the product more significantly compared to the sum. In other words, GA grants the agents a veto power: if one agent transmits a belief entry as 0, the combined belief also becomes 0. In general, giving this level of autonomy to individual agents depends on the application at hand.* ■

3. ASYMPTOTIC NORMALITY OF AA AND GA

In this section, we analyze the asymptotic behavior under both fusion strategies. In order to avoid pathological cases, we assume that each observation has finite information about the true hypothesis. That is to say, for each agent k and for each hypothesis $\theta \in \Theta$, we assume $D_{\text{KL}}(L_k(\cdot|\theta^\circ)||L_k(\cdot|\theta)) < \infty$. This implies that likelihood functions have the same support. Moreover, in order to uniquely distinguish the true hypothesis, we need the following condition.

Assumption 1 (Global Identifiability). *For each wrong hypothesis $\theta \neq \theta^\circ$, there exists at least one clear-sighted agent k with $D_{\text{KL}}(L_k(\cdot|\theta^\circ)||L_k(\cdot|\theta)) > 0$.* ■

Note that this assumption does not require *local* identifiability, which is an agent’s ability of inferring θ° without any cooperation. As a result, all agents can benefit from cooperation. Under Assumption 1, it is already known that AA [4, 5, 22] and GA [6, 7] result

in consistent truth learning, that is, the beliefs on the true hypothesis converge to one almost surely. In this work, we further show that under both AA and GA fusion rules, the asymptotic error probabilities can be calculated with Gaussian cumulative distribution functions (c.d.f.s). We start with the AA strategy.

Proposition 1 (Asymptotic normality of AA). *For each wrong hypothesis $\theta \neq \theta^\circ$, we define the mean and variance*

$$\rho_A \triangleq -\mathbb{E}\left[\log\left(\sum_{k \in \mathcal{K}} \pi_k \mathbf{r}_{k,i}\right)\right], \quad \sigma_A^2 \triangleq \text{Var}\left(\log\left(\sum_{k \in \mathcal{K}} \pi_k \mathbf{r}_{k,i}\right)\right), \quad (4)$$

in terms of the likelihood ratio

$$\mathbf{r}_{k,i} \triangleq \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta)}{L_k(\boldsymbol{\xi}_{k,i}|\theta^\circ)}. \quad (5)$$

If $\sigma_A^2 < \infty$, it holds under Assumption 1 that

$$\lim_{i \rightarrow \infty} \mathbb{P}\left(\frac{\log \boldsymbol{\mu}_i(\theta) + \rho_A i}{\sigma_A \sqrt{i}} \leq t\right) = \Phi(t), \quad (6)$$

where $\Phi(t)$ is the standard Gaussian cumulative distribution function.

Proof Sketch. For the complete proof, see Appendix A. Here, we only give a sketch of the proof. Consider an arbitrary $\theta \neq \theta^\circ$. For notational simplicity, define

$$\mathbf{x}_i \triangleq \log \boldsymbol{\mu}_i(\theta), \quad \mathbf{L}_i \triangleq \log\left(\sum_{k \in \mathcal{K}} \pi_k \mathbf{r}_{k,i}\right). \quad (7)$$

Since \mathbf{L}_i 's are assumed to be i.i.d. and of finite variance, the central limit theorem [23] yields:

$$\lim_{i \rightarrow \infty} \mathbb{P}\left(\frac{\sum_{j=1}^i \mathbf{L}_j + \rho_A i}{\sigma_A \sqrt{i}} \leq t\right) = \Phi(t). \quad (8)$$

Therefore, if we show that the distance between $\frac{1}{\sqrt{i}} \mathbf{x}_i$ and $\frac{1}{\sqrt{i}} \sum_{j=1}^i \mathbf{L}_j$ goes to zero with high probability, we are done. To that end, we utilize the result from [15] which states

$$\frac{1}{\sqrt{i}} \mathbf{x}_i \rightarrow -\rho_A \quad \text{as } i \rightarrow \infty, \text{ almost surely.} \quad (9)$$

In particular, we use the following bootstrapping argument on (9): There exists an i_0 after which (i.e., $\forall i \geq i_0$)

$$\boldsymbol{\mu}_i(\theta) \leq \exp\{-i(\rho_A - \epsilon)\}, \quad \forall \theta \neq \theta^\circ \quad (10)$$

with high probability. This result enables a finer study of the evolution of \mathbf{x}_i as $i \rightarrow \infty$. ■

Note that Proposition 1 can equivalently be given in terms of the log-belief ratios:

Corollary 1. *If we define the log-belief ratio $\boldsymbol{\lambda}_i(\theta) \triangleq \log \frac{\boldsymbol{\mu}_i(\theta^\circ)}{\boldsymbol{\mu}_i(\theta)}$, for each $\theta \neq \theta^\circ$, Proposition 1 implies that*

$$\lim_{i \rightarrow \infty} \mathbb{P}\left(\frac{\boldsymbol{\lambda}_i(\theta) - \rho_A i}{\sigma_A \sqrt{i}} \leq t\right) = \Phi(t). \quad (11)$$

Proof Sketch. For the complete proof, see Appendix B. Essentially, the proof follows from consistent truth learning, i.e., $\boldsymbol{\mu}_i(\theta^\circ) \rightarrow 1$ almost surely as $i \rightarrow \infty$. This implies that, as $i \rightarrow \infty$, $\log \boldsymbol{\mu}_i(\theta)$ and $\log \boldsymbol{\mu}_i(\theta) - \log \boldsymbol{\mu}_i(\theta^\circ)$ behaves similarly. ■

Next, we investigate the behavior of beliefs under the GA fusion rule.

Proposition 2 (Asymptotic normality of GA). *For each wrong hypothesis $\theta \neq \theta^\circ$, we define the mean*

$$\rho_G \triangleq -\mathbb{E}\left[\sum_{k \in \mathcal{K}} \pi_k \log \mathbf{r}_{k,i}\right] = \sum_{k \in \mathcal{K}} \pi_k D_{\text{KL}}(L_k(\cdot|\theta^\circ) || L_k(\cdot|\theta)) \quad (12)$$

and the variance

$$\sigma_G^2 \triangleq \text{Var}\left(\sum_{k \in \mathcal{K}} \pi_k \log \mathbf{r}_{k,i}\right). \quad (13)$$

If $\sigma_G^2 < \infty$, then under Assumption 1 it holds that

$$\lim_{i \rightarrow \infty} \mathbb{P}\left(\frac{\boldsymbol{\lambda}_i(\theta) - \rho_G i}{\sigma_G \sqrt{i}} \leq t\right) = \Phi(t). \quad (14)$$

Proof. The proofs from [6, 7] for the convergence rate analysis can be easily extended to establish asymptotic normality. From (1) and (3), the log-belief ratio $\boldsymbol{\lambda}_i(\theta)$ evolves according to the recursion:

$$\boldsymbol{\lambda}_i(\theta) = \boldsymbol{\lambda}_{i-1}(\theta) - \sum_{k \in \mathcal{K}} \pi_k \log \mathbf{r}_{k,i}. \quad (15)$$

Expanding the recursion over time, we obtain

$$\boldsymbol{\lambda}_i(\theta) = -\sum_{j=1}^i \left(\sum_{k \in \mathcal{K}} \pi_k \log \mathbf{r}_{k,j}\right) + \boldsymbol{\lambda}_0(\theta). \quad (16)$$

Since (i) $\boldsymbol{\lambda}_0(\theta) < \infty$ due to initial beliefs being positive, (ii) $\mathbf{r}_{k,i}$'s are i.i.d., and (iii) $\sigma_G^2 < \infty$, an application of the central limit theorem [23] concludes the proof. ■

In Propositions 1 and 2, we have established the asymptotic normality for each wrong hypothesis, which allows approximating the error probabilities when agents or the central processor employs a maximum a-posteriori (MAP) rule to estimate θ° . It is important to note that under spatial independence, GA matches the optimal Bayes posterior, which attains minimum error probability, when the confidence weight π_k on each agent k is uniform. Additionally, the convergence rates ρ_G and ρ_A were already established in previous works [6, 7] and [15] respectively. Notice that $\rho_G \geq \rho_A$ — applying Jensen's inequality to the equation on the left in (4), we obtain (12). However, this does not directly imply that GA has a better performance in probability of error criterion since it depends on the variances as well. Our simulations indicate, however, that GA has lower error probability in general. We leave it to a future work to determine whether GA is superior to AA in terms of asymptotic error probability under any data distribution or confidence weight configuration — such as in the convergence rate case.

4. NUMERICAL RESULTS

We consider a binary hypothesis testing problem with $K = 10$ distributed agents connected to a fusion center. In the first experiment, agents receive i.i.d. Gaussian observations with mean 0 under the true hypothesis θ° , and with mean 1 under $\theta \neq \theta^\circ$. For both hypotheses, the correlation matrix is the identity matrix I_K , i.e., data is independent across agents. A possible application of this experiment is in flexible spectrum use, for example in cognitive radios. The possible hypotheses are whether the spectrum is idle or not. If there is noise only, the signal distributions would be central (i.e., zero




Fig. 2: (a) For the first experiment with Gaussian data, the histograms of $\tilde{\lambda}_i^{(A)}(\theta)$ and $\tilde{\lambda}_i^{(G)}(\theta)$ are drawn as shaded bar plots with colors blue and red respectively. The standard normal density $\mathcal{G}(t)$ is shown as solid black line. (b) For the second experiment with exponentially distributed data, the histograms of $\tilde{\lambda}_i^{(A)}(\theta)$ and $\tilde{\lambda}_i^{(G)}(\theta)$ are drawn as shaded bar plots with colors blue and red respectively. The standard normal density $\mathcal{G}(t)$ is shown as solid black line.




Fig. 3: (a) For the uncorrelated Gaussian case, histograms of $\lambda_i^{(A)}(\theta)$ and $\lambda_i^{(G)}(\theta)$ are drawn as shaded bar plots with colors blue and red respectively. (b) For the correlated Gaussian case, histograms of $\lambda_i^{(A)}(\theta)$ and $\lambda_i^{(G)}(\theta)$ are drawn as shaded bar plots with colors blue and red respectively.

mean). Otherwise, if there is a signal present, it would be non-zero mean. We choose the combination vector as

$$\pi = [0.13, 0.2, 0.09, 0.15, 0.08, 0.05, 0.1, 0.05, 0.1, 0.05]. \quad (17)$$

If we take the time horizon $i = 5000$ and repeat the experiment for 500 realizations, we obtain the histograms given in Figure 2a. Specifically, we obtain the normalized log-belief ratios

$$\tilde{\lambda}_i^{(A)}(\theta) \triangleq \frac{\lambda_i^{(A)}(\theta) - \rho_A i}{\sqrt{i}\sigma_A}, \quad \tilde{\lambda}_i^{(G)}(\theta) \triangleq \frac{\lambda_i^{(G)}(\theta) - \rho_G i}{\sqrt{i}\sigma_G} \quad (18)$$

and plot their histograms in density forms (denoted as $f^{(A)}(t)$ and $f^{(G)}(t)$ respectively) to compare them with the standard normal density $\mathcal{G}(t) \triangleq \frac{1}{\sqrt{2\pi}} \exp(-t^2/2)$. In addition to the visual matching supporting our claims, we performed a Shapiro-Wilk test for testing the normality of $\tilde{\lambda}_i^{(A)}(\theta)$ and $\tilde{\lambda}_i^{(G)}(\theta)$. The resulting p -values turn out to be 0.66 and 0.54, which strongly supports the null hypothesis of normal distribution.

For the second experiment, we assume that agents receive i.i.d. exponential observations with mean 1 under θ° and with mean 0.5 under $\theta \neq \theta^\circ$. The other problem parameters are the same as the first experiment. The resulting histograms are given in Figure 2b. The Shapiro-Wilk test yields p -values 0.93 and 0.53 respectively, again, strongly supporting the null hypothesis of normal distribution.

We also conducted a third experiment to verify that correlations between the agents affect ρ_A but do not affect ρ_G (see (12)). To

compare correlated versus uncorrelated data across agents, we simulate the setting from the first experiment, except that this time we set the correlation matrix as $\Sigma = 0.95\mathbb{1}_K\mathbb{1}_K^\top + 0.5I_K$. Note that this choice makes the data highly correlated. Then, we plot the histograms of $\lambda_i^{(A)}(\theta)$ and $\lambda_i^{(G)}(\theta)$ (without normalization) for both uncorrelated and correlated cases in Figure 3a and 3b. We can observe that $\rho_G = 0.5$ is unchanged by the correlations — the observation that $f^{(G)}$'s are concentrated around $i \times \rho_G = 2500$ supports this fact. In contrast, ρ_A increases drastically under the presence of positive correlations. Nevertheless, in both cases considered, GA outperforms AA in terms of the error probability.

5. CONCLUSION

In this work, we considered two pooling algorithms for federated decision making, namely AA and GA. The choice between AA and GA may depend on the specific task at hand. We study their performances in terms of error probability, and establish asymptotic normality which is a valuable tool that can help avoid the need to simulate the algorithm for large time horizons. Note that the parameters in the normal approximations can be obtained either analytically or through Monte Carlo simulations, which is more efficient compared to the exact simulation of the entire system. We leave the study for asymptotic normality of general decentralized networks to a future work.

6. REFERENCES

- [1] A. R. Elkordy, Y. H. Ezzeldin, S. Han, S. Sharma, C. He, S. Mehrotra, and S. Avestimehr, "Federated analytics: A survey," *APSIPA Transactions on Signal and Information Processing*, vol. 12, no. 1, pp. 1–33, 2023.
- [2] H. Hallock, S. E. Marshall, P. A. C. 't Hoen, J. F. Nygård, B. Hoorne, C. Fox, and S. Alagaratnam, "Federated networks for distributed analysis of health data," *Frontiers in Public Health*, vol. 9, 2021.
- [3] V. Bordinon, S. Vlaski, V. Matta, and A. H. Sayed, "Learning from heterogeneous data based on social interactions over graphs," *IEEE Transactions on Information Theory*, pp. 1–25, Dec. 2022.
- [4] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.
- [5] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, 2012, pp. 709–713.
- [6] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-Bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.
- [7] A. Lalitha, T. Javidi, and A. D. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp. 6161–6179, 2018.
- [8] J. N. Tsitsiklis, "Decentralized detection," in *In Advances in Statistical Signal Processing*. JAI Press, 1993, pp. 297–344.
- [9] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer, 2012.
- [10] D. Bajovic, D. Jakovetic, J. M. Moura, J. Xavier, and B. Sinopoli, "Large deviations performance of consensus+ innovations distributed detection with non-Gaussian observations," *IEEE Transactions on Signal Processing*, vol. 60, no. 11, pp. 5987–6002, 2012.
- [11] Y. Inan, M. Kayaalp, A. H. Sayed, and E. Telatar, "A fundamental limit of distributed hypothesis testing under memoryless quantization," in *Proc. IEEE International Conference on Communications*, 2022, pp. 4824–4829.
- [12] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. International Conference on Artificial Intelligence and Statistics*, vol. 54, April 2017, pp. 1273–1282.
- [13] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [14] G. Koliander, Y. El-Laham, P. M. Djurić, and F. Hlawatsch, "Fusion of probability density functions," *Proc. IEEE*, vol. 110, no. 4, pp. 404–453, April 2022.
- [15] M. Kayaalp, Y. Inan, E. Telatar, and A. H. Sayed, "On the arithmetic and geometric fusion of beliefs for distributed inference," *arXiv:2204.13741*, 2022.
- [16] T. Li, H. Fan, J. García, and J. M. Corchado, "Second-order statistics analysis and comparison between arithmetic and geometric average fusion: Application to multi-sensor target tracking," *Information Fusion*, vol. 51, pp. 233–243, 2019.
- [17] T. Li, Y. Xin, Y. Song, E. Song, and H. Fan, "Some statistic and information-theoretic results on arithmetic average density fusion," *arXiv:2110.01440*, 2021.
- [18] L. Gao, G. Battistelli, and L. Chisci, "Multiobject fusion with minimum information loss," *IEEE Signal Processing Letters*, vol. 27, pp. 201–205, Jan. 2020.
- [19] K. C. Chang, C.-Y. Chong, and S. Mori, "Analytical and computational evaluation of scalable distributed fusion algorithms," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 4, pp. 2022–2034, 2010.
- [20] W. Liu, M. Salzmann, and P. Fua, "Context-aware crowd counting," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5099–5108.
- [21] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, July 2014.
- [22] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, no. 13-28, pp. 1–38, May 2013.
- [23] A. H. Sayed, *Inference and Learning from Data*. Cambridge University Press, 2022, 3 vols.
- [24] D. Williams, *Probability with Martingales*. Cambridge University Press, 1991.

A. PROOF OF PROPOSITION 1

Consider an arbitrary $\theta \neq \theta^\circ$. We denote the log-belief on θ as $\mathbf{x}_i \triangleq \log \boldsymbol{\mu}_i(\theta)$. It was shown in [15] that $\frac{1}{i} \mathbf{x}_i \rightarrow -\rho_A$ almost surely as $i \rightarrow \infty$. This implies that for any $\epsilon > 0$, there exists an i_0 such that

$$\mathcal{A}_0 \triangleq \{\boldsymbol{\mu}_i(\theta) \leq \exp\{-i(\rho_A - \epsilon)\}, \forall i \geq i_0, \forall \theta \neq \theta^\circ\} \quad (19)$$

is a high-probability event. After this observation, we split the proof into two parts:

- (a) $\liminf_{i \rightarrow \infty} \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} < t\right) \geq \Phi(t)$
- (b) $\limsup_{i \rightarrow \infty} \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \leq t\right) \leq \Phi(t).$

Then, it readily follows from (a) and (b) that

$$\lim_{i \rightarrow \infty} \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \leq t\right) = \Phi(t). \quad (20)$$

A.1. Proof of (a)

Using (1) and (2), the log-belief \mathbf{x}_i can be upper bounded under the event \mathcal{A}_0 as

$$\mathbf{x}_i \leq \mathbf{x}_{i_0} + \sum_{j=i_0+1}^i \mathbf{L}_j - \sum_{j=i_0+1}^i \log(1 - H \exp\{-j(\rho_A - \epsilon)\}) \quad (21)$$

where we defined $\mathbf{L}_i \triangleq \log(\sum_k \pi_k \mathbf{r}_{k,i})$ and recall that H is the number of hypotheses. Since

$$-\sum_{j=i_0+1}^{\infty} \log(1 - H \exp\{-j(\rho_A - \epsilon)\}) < \infty, \quad (22)$$

for every $\delta > 0$ there must exist an $i_1^+ > i_0$ such that for all $i \geq i_1^+$:

$$-\frac{1}{\sqrt{i}} \sum_{j=i_0+1}^i \log(1 - H \exp\{-j(\rho_A - \epsilon)\}) \leq \delta. \quad (23)$$

Consequently, the event

$$\mathcal{A}_1^+ \triangleq \left\{ \frac{1}{\sqrt{i}} \mathbf{x}_i \leq \frac{1}{\sqrt{i}} \mathbf{x}_{i_0} + \frac{1}{\sqrt{i}} \sum_{j=i_0+1}^i \mathbf{L}_j + \delta, \forall i \geq i_1^+ \right\} \quad (24)$$

is also a high-probability event. Furthermore, since \mathbf{x}_{i_0} is almost surely finite, there is an $i_2^+ > i_0$ such that

$$\mathcal{A}_2^+ \triangleq \left\{ \frac{1}{\sqrt{i}} \mathbf{x}_{i_0} \leq \delta, \forall i \geq i_2^+ \right\} \quad (25)$$

is a high-probability event as well. Accordingly, for any outcome in $\mathcal{A}^+ \triangleq \mathcal{A}_1^+ \cap \mathcal{A}_2^+$, and for $i \geq \max\{i_1^+, i_2^+\}$, we have

$$\frac{1}{\sqrt{i}} \mathbf{x}_i \leq \frac{1}{\sqrt{i}} \sum_{j=i_0+1}^i \mathbf{L}_j + 2\delta, \quad (26)$$

which directly implies

$$\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \leq \frac{\sum_{j=i_0+1}^i \mathbf{L}_j + \rho_A i}{\sigma_A \sqrt{i}} + \frac{2\delta}{\sigma_A}. \quad (27)$$

Finally, with $\overline{\mathcal{A}^+}$ denoting the complement of \mathcal{A}^+ , we obtain that for $i \geq \max\{i_1^+, i_2^+\}$:

$$\begin{aligned} & \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \geq t\right) \\ &= \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \geq t \cap \mathcal{A}^+\right) + \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \geq t \cap \overline{\mathcal{A}^+}\right) \\ &\leq \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \geq t \cap \mathcal{A}^+\right) + \mathbb{P}(\overline{\mathcal{A}^+}) \\ &\leq \mathbb{P}\left(\frac{\sum_{j=i_0+1}^i \mathbf{L}_j + \rho_A i}{\sigma_A \sqrt{i}} \geq t - \frac{2\delta}{\sigma_A}\right) + \mathbb{P}(\overline{\mathcal{A}^+}), \end{aligned} \quad (28)$$

where (28) follows from (27). Since \mathbf{L}_i 's are i.i.d. random variables with finite variance, we can use the central limit theorem (CLT) [23] and arrive at

$$\limsup_{i \rightarrow \infty} \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \geq t\right) \leq 1 - \Phi(t) \quad (29)$$

by using the facts that δ and $\mathbb{P}(\overline{\mathcal{A}^+})$ can be made arbitrarily small and $\Phi(\cdot)$ is continuous. Note that (29) is equivalent to

$$\liminf_{i \rightarrow \infty} \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} < t\right) \geq \Phi(t). \quad (30)$$

A.2. Proof of (b)

Similar to the proof of (a), we obtain a lower bound for \mathbf{x}_i under the event \mathcal{A}_0 by

$$\begin{aligned} \mathbf{x}_i &\geq \mathbf{x}_{i_0} + \sum_{j=i_0+1}^i \log\left(\frac{\pi_k \mathbf{r}_{k,j}}{1 + \exp\{-j(\rho_A - \epsilon)\} \sum_{\theta \neq \theta^\circ} \mathbf{r}_{k,j}(\theta)}\right) \\ &\geq \mathbf{x}_{i_0} + \sum_{j=i_0+1}^i \log\left(\frac{\pi_k \mathbf{r}_{k,j}}{1 + \exp\{-j(\rho_A - \epsilon)\} \sum_{k \in \mathcal{K}} \sum_{\theta \neq \theta^\circ} \mathbf{r}_{k,j}(\theta)}\right) \\ &\geq \mathbf{x}_{i_0} + \sum_{j=i_0+1}^i \mathbf{L}_j - \sum_{j=i_0+1}^i \log\left(1 + \exp\{-j(\rho_A - \epsilon)\} \sum_{k \in \mathcal{K}} \sum_{\theta \neq \theta^\circ} \mathbf{r}_{k,j}(\theta)\right) \\ &\geq \mathbf{x}_{i_0} + \sum_{j=i_0+1}^i \mathbf{L}_j - \sum_{j=i_0+1}^i \exp\{-j(\rho_A - \epsilon)\} \sum_{k \in \mathcal{K}} \sum_{\theta \neq \theta^\circ} \mathbf{r}_{k,j}(\theta). \end{aligned} \quad (31)$$

Comparing with (21), we aim to show that the rightmost term in (31), i.e.,

$$\mathbf{S}_i \triangleq \sum_{j=i_0+1}^i \exp\{-j(\rho_A - \epsilon)\} \sum_{k \in \mathcal{K}} \sum_{\theta \neq \theta^\circ} \mathbf{r}_{k,j}(\theta) \quad (32)$$

tends to zero almost surely when divided by \sqrt{i} . Observe that \mathbf{S}_i is non-decreasing. Hence,

$$\mathbf{S}_\infty \triangleq \lim_{i \rightarrow \infty} \mathbf{S}_i \quad (33)$$

exists almost surely, i.e.,

$$\limsup_{i \rightarrow \infty} \mathbf{S}_i = \liminf_{i \rightarrow \infty} \mathbf{S}_i. \quad (34)$$

Moreover, by the monotone convergence theorem [24],

$$\mathbb{E}[\mathbf{S}_\infty] = \lim_{i \rightarrow \infty} \mathbb{E}[\mathbf{S}_i] < \infty. \quad (35)$$

Consequently, \mathbf{S}_∞ must be finite almost surely. Then, by combining the inequality

$$0 \leq \frac{\mathbf{S}_i}{\sqrt{i}} \leq \frac{\mathbf{S}_\infty}{\sqrt{i}} \quad (36)$$

and the limit $\frac{\mathbf{S}_\infty}{\sqrt{i}} \rightarrow 0$ as $i \rightarrow \infty$ (both of which hold almost surely), we conclude that $\frac{\mathbf{S}_i}{\sqrt{i}} \rightarrow 0$ almost surely as well. Therefore, for every $\delta > 0$ there must exist an $i_1^- > i_0$ such that for all $i \geq i_1^-$:

$$\frac{1}{\sqrt{i}} \sum_{j=i_0+1}^i \exp\{-j(\rho_A - \epsilon)\} \sum_{k \in \mathcal{K}} \sum_{\theta \neq \theta^\circ} \mathbf{r}_{k,j}(\theta) \leq \delta. \quad (37)$$

Then, proceeding similarly to the part (a), we define the high-probability sets:

$$\mathcal{A}_1^- \triangleq \left\{ \frac{1}{\sqrt{i}} \mathbf{x}_i \geq \frac{1}{\sqrt{i}} \mathbf{x}_{i_0} + \frac{1}{\sqrt{i}} \sum_{j=i_0+1}^i \mathbf{L}_j - \delta, \forall i \geq i_1^- \right\} \quad (38)$$

and

$$\mathcal{A}_2^- \triangleq \left\{ \frac{1}{\sqrt{i}} \mathbf{x}_{i_0} \geq -\delta, \forall i \geq i_2^- \right\}. \quad (39)$$

By defining $\mathcal{A}^- \triangleq \mathcal{A}_1^- \cap \mathcal{A}_2^-$, we arrive at the following analog to (28) for $i \geq \max\{i_1^-, i_2^-\}$:

$$\mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \leq t\right) \leq \mathbb{P}\left(\frac{\sum_{j=i_0+1}^i \mathbf{L}_j + \rho_A i}{\sigma_A \sqrt{i}} \leq t + \frac{2\delta}{\sigma_A}\right) + \mathbb{P}(\overline{\mathcal{A}^-}), \quad (40)$$

which implies

$$\limsup_{i \rightarrow \infty} \mathbb{P}\left(\frac{\mathbf{x}_i + \rho_A i}{\sigma_A \sqrt{i}} \leq t\right) \leq \Phi(t), \quad (41)$$

and the proof is complete.

B. PROOF OF COROLLARY 1

Notice that for any outcome in \mathcal{A}_0 , defined in (19), it is true that

$$1 - H \exp\{-i(\rho_A - \epsilon)\} \leq \boldsymbol{\mu}_i(\theta^\circ) \leq 1, \quad (42)$$

which implies

$$\log(1 - H \exp\{-i(\rho_A - \epsilon)\}) \leq \log \boldsymbol{\mu}_i(\theta^\circ) \leq 0. \quad (43)$$

Since

$$\frac{1}{\sqrt{i}} \log(1 - H \exp\{-i(\rho_A - \epsilon)\}) \rightarrow 0, \quad (44)$$

the result of Proposition 1 applies here as well.