

How Does Data Freshness Affect Real-time Supervised Learning?

Md Kamran Chowdhury Shisher

Department of ECE

Auburn University

Auburn, AL, USA

mzs0153@auburn.edu

Yin Sun

Department of ECE

Auburn University

Auburn, AL, USA

yzs0078@auburn.edu

ABSTRACT

In this paper, we analyze the impact of data freshness on real-time supervised learning, where a neural network is trained to infer a time-varying target (e.g., the position of the vehicle in front) based on features (e.g., video frames) observed at a sensing node (e.g., camera or lidar). One might expect that the performance of real-time supervised learning degrades monotonically as the feature becomes stale. Using an information-theoretic analysis, we show that this is true if the feature and target data sequence can be closely approximated as a Markov chain; it is not true if the data sequence is far from Markovian. Hence, the prediction error of real-time supervised learning is a function of the Age of Information (AoI), where the function could be non-monotonic. Several experiments are conducted to illustrate the monotonic and non-monotonic behaviors of the prediction error. To minimize the inference error in real-time, we propose a new “selection-from-buffer” model for sending the features, which is more general than the “generate-at-will” model used in earlier studies. By using Gittins and Whittle indices, low-complexity scheduling strategies are developed to minimize the inference error, where a new connection between the Gittins index theory and Age of Information (AoI) minimization is discovered. These scheduling results hold (i) for minimizing general AoI functions (monotonic or non-monotonic) and (ii) for general feature transmission time distributions. Data-driven evaluations are presented to illustrate the benefits of the proposed scheduling algorithms.

CCS CONCEPTS

• Networks → Network management; • Theory of computation → Mathematical optimization;

KEYWORDS

Age of Information; supervised learning; scheduling; Markov chain; buffer management

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiHoc '22, October 17–20, 2022, Seoul, Republic of Korea

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9165-8/22/10...\$15.00

<https://doi.org/10.1145/3492866.3549711>

ACM Reference Format:

Md Kamran Chowdhury Shisher and Yin Sun. 2022. How Does Data Freshness Affect Real-time Supervised Learning?. In *The Twenty-third International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '22), October 17–20, 2022, Seoul, Republic of Korea*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3492866.3549711>

1 INTRODUCTION

In recent years, the proliferation of networked control and cyber-physical systems such as autonomous vehicle, UAV navigation, remote surgery, industrial control system has significantly boosted the need for real-time prediction. For example, an autonomous vehicle infers the trajectories of nearby vehicles and the intention of pedestrians based on lidars and cameras installed on the vehicle [19]. In remote surgery, the movement of a surgical robot is predicted in real-time. These prediction problems can be solved by real-time supervised learning, where a neural network is trained to predict a time varying target based on feature observations that are collected from a sensing node. Due to data processing time, transmission errors, and queueing delay, the features delivered to the neural predictor may not be fresh. The performance of networked intelligent systems depends heavily on the accuracy of real-time prediction. Hence, it is important to understand how data freshness affects the performance of real-time supervised learning.

To evaluate data freshness, a metric *Age of information* (AoI) was introduced in [15]. Let U_t be the generation time of the freshest feature received by the neural predictor at time t . Then, the AoI of the features, as a function of time t , is defined as $\Delta(t) = t - U_t$, which is the time difference between the current time t and the generation time U_t of the freshest received feature. The age of information concept has gained a lot of attention from the research communities. Analysis and optimization of AoI were studied in various networked systems, including remote estimation, control system, and edge computing. In these studies, it is commonly assumed that the system performance degrades monotonically as the AoI grows. Nonetheless, this is not always true in real-time supervised learning. For example, it was observed that the predictor error of day-ahead solar power forecasting is not a monotonic function of the AoI, because there exists an inherent daily periodic changing pattern in the solar power time-series data [24].

In this study, we carry out several experiments and present an information-theoretic analysis to interpret the impact of data freshness in real-time supervised learning. In addition, we design buffer management and transmission scheduling strategies to improve the accuracy of real-time supervised learning. The key contributions of this paper are summarized as follows:

- We develop an information-theoretic approach to analyze how the AoI affects the performance of real-time supervised learning. It is shown that the prediction errors (training error and inference error) are functions of AoI, whereas they could be non-monotonic AoI functions — this is a key difference from previous studies on AoI functions, e.g., [17, 26, 28, 29]. When the target and feature data sequence can be closely approximated as a Markov chain, the prediction errors are non-decreasing functions of the AoI. When the target and feature data sequence is far from Markovian, the prediction errors could be non-monotonic in the AoI (see Sections 2-3).
- We conduct several experiments and observe that, due to long-range dependence, response delay, and/or communication delay, the target and feature data sequence can be far from Markovian and the corresponding prediction errors are non-monotonic AoI functions. In certain scenarios, even a fresh feature ($\text{AoI}=0$) may generate larger prediction errors than stale features ($\text{AoI} > 0$), i.e., the freshest feature may not be the best feature; see Fig. 2 for an illustration.
- We propose buffer management and transmission scheduling strategies to minimize the inference error. Because the inference error could be a non-monotonic AoI function, we introduce a novel “selection-from-buffer” model for feature transmissions, which is more general than the “generate-at-will” model used in many earlier studies, e.g., [26, 28, 32]. If the AoI function is non-decreasing, the “selection-from-buffer” model achieves same performance as the “generate-at-will” model; if the AoI function is non-monotonic, the “selection-from-buffer” model can potentially achieve better performance.
- In the single-source case, an optimal scheduling policy is devised to minimize the long-term average inference error. By exploiting a new connection with the Gittins index theory [9], the optimal scheduling policy is proven to be a threshold policy on the Gittins index (Theorems 4.1-4.2), where the threshold can be computed by using a low complexity algorithm like bisection search. This scheduling policy is more general than the scheduling policies proposed in [26, 28].
- In the multi-source case, a Whittle index scheduling policy is designed to reduce the weighted sum of the inference errors of the sources. By using the Gittins index obtained in the single-source case, a semi-analytical expression of the Whittle index is obtained (Theorems 5.1-5.2), which is more general than the Whittle index formula in [29, Equation (7)].
- The above scheduling results hold (i) for minimizing general AoI functions (monotonic or non-monotonic) and (ii) for general feature transmission time distributions. Data driven evaluations show that “selection-from-buffer” with optimal scheduler achieves up to 3 times smaller inference error compared to “generate-at-will,” and 8 times smaller inference error compared to periodic feature updating (see Fig. 5). Whittle index policy achieves up to 2 times performance gain compared to maximum age first (MAF) policy (see Fig. 6).

1.1 Related Works

In recent years, AoI has become a popular research topic [33]. Average AoI and average peak AoI are studied in many queueing systems [15, 28, 32]. As surveyed in [26], there exist a number of applications of non-linear AoI functions, such as auto-correlation function [17], estimation error [16, 20, 27], and Shannon’s mutual information and conditional entropy [26]. In existing studies on AoI, it was usually assumed that the observed data sequence is Markovian and the performance degradation caused by information aging was modeled as a monotonic AoI function. However, practical data sequence may not be Markovian [12, 26, 30]. In the present paper, theoretical results and experimental studies are provided to analyze the performance of real-time supervised learning for both Markovian and non-Markovian time-series data. In [34], impact of peak-AoI on the convergence speed of online training was analyzed. Unlike online training in [34], our work considers offline training and online inference.

Moreover, there are significant research efforts on the optimization of AoI functions by designing sampling and scheduling policies. Previous studies [1, 14, 20, 26, 28, 29] focused on non-decreasing AoI functions. Recently, a Whittle index based multi-source scheduling policy was derived in [2] to minimize Shannon’s conditional entropy that could be a non-monotonic function of the AoI. The Whittle index policy in [2] requires that (i) the state of each source evolves as binary Markov process, (ii) the AoI function is concave with respect to the belief state of the Markov process, and (iii) the packet transmission time is constant. The results in [1, 2, 14, 20, 26, 28, 29] are not appropriate for minimizing general (potentially non-monotonic) AoI functions, as considered in the present paper.

2 INFORMATION-THEORETIC MEASURES FOR REAL-TIME SUPERVISED LEARNING

2.1 Freshness-aware Learning Model

Consider the real-time supervised learning system illustrated in Fig. 1, where the goal is to predict a label $Y_t \in \mathcal{Y}$ (e.g., the location of the car in front) at each time t based on a feature $X_{t-\Delta(t)}$ (e.g., a video clip) that was generated $\Delta(t)$ seconds ago. The feature, $X_{t-\Delta(t)} = (V_{t-\Delta(t)}, \dots, V_{t-\Delta(t)-u+1})$ is a time sequence with length u (e.g., each video clip consisting of u consecutive video frames). We consider a class of popular supervised learning algorithms called *Empirical Risk Minimization (ERM)* [10]. In freshness-aware ERM algorithms, a neural network is trained to construct an action $a = \phi(X_{t-\Delta(t)}, \Delta(t)) \in \mathcal{A}$ where $\phi : \mathcal{X} \times \mathcal{D} \mapsto \mathcal{A}$ is a function of feature $X_{t-\Delta(t)} \in \mathcal{X}$ and its AoI $\Delta(t) \in \mathcal{D}$. The performance of learning is measured by a loss function $L : \mathcal{Y} \times \mathcal{A} \mapsto \mathbb{R}$, where $L(y, a)$ is the incurred loss if action a is chosen by the neural network when $Y_t = y$. We assume that \mathcal{Y} , \mathcal{X} , and \mathcal{D} are discrete and finite sets. The loss function L is determined by the *targeted application* of the system. For example, in neural network based estimation, the loss function is usually chosen as the square estimation error $L_2(y, \hat{y}) = \|y - \hat{y}\|^2$, where the action $a = \hat{y}$ is an estimate of $Y_t = y$. In softmax regression (i.e., neural network based maximum likelihood classification), the action $a = Q_Y$ is a distribution of Y_t and the loss function $L_{\log}(y, Q_Y) = -\log Q_Y(y)$ is the negative

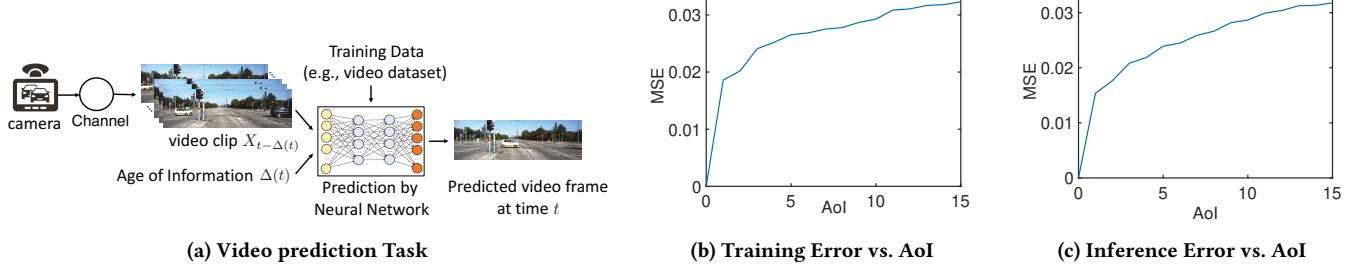


Figure 1: Performance of supervised learning based video prediction. The experimental results in (b) and (c) are regenerated from [18]. The training and inference errors are non-decreasing functions of the AoI.

log-likelihood of the label value $Y_t = y$. Therefore, the loss function L characterizes the goal and purpose of a specific application.

2.2 Offline Training Error

The real-time supervised learning system that we consider consists of two phases: *offline training* and *online inference*. In the offline training phase, the neural network is trained using a training dataset. Let $P_{\tilde{Y}_0, \tilde{X}_{-\Theta}, \Theta}$ denote the empirical distribution of the label \tilde{Y}_0 , feature $\tilde{X}_{-\Theta}$, and AoI Θ in the training dataset, where the AoI $\Theta \geq 0$ of the feature $\tilde{X}_{-\Theta}$ is the time difference between \tilde{Y}_0 and $\tilde{X}_{-\Theta}$. In ERM algorithms, the training problem is formulated as

$$\text{err}_{\text{training}} = \min_{\phi \in \Lambda} \mathbb{E}_{Y, X, \Theta \sim P_{\tilde{Y}_0, \tilde{X}_{-\Theta}, \Theta}} [L(Y, \phi(X, \Theta))], \quad (1)$$

where Λ is the set of functions that can be constructed by the neural network, and $\text{err}_{\text{training}}$ is the minimum training error. The optimal solution to (1) is denoted by $\phi^*_{P_{\tilde{Y}_0, \tilde{X}_{-\Theta}, \Theta}}$.

Let $\Phi = \{f : \mathcal{X} \times \mathcal{D} \mapsto \mathcal{A}\}$ be the set of all functions mapping from $\mathcal{X} \times \mathcal{D}$ to \mathcal{A} . Any action $\phi(x, \theta)$ constructed by the neural network belongs to Φ , whereas the neural network cannot produce some functions in Φ . Hence, $\Lambda \subset \Phi$. By relaxing the feasible set Λ in (1) as Φ , we obtain a lower bound of $\text{err}_{\text{training}}$, i.e.,

$$H_L(\tilde{Y}_0 | \tilde{X}_{-\Theta}, \Theta) = \min_{\phi \in \Phi} \mathbb{E}_{Y, X, \Theta \sim P_{\tilde{Y}_0, \tilde{X}_{-\Theta}, \Theta}} [L(Y, \phi(X, \Theta))], \quad (2)$$

where $H_L(\tilde{Y}_0 | \tilde{X}_{-\Theta}, \Theta)$ is a generalized conditional entropy of \tilde{Y}_0 given $(\tilde{X}_{-\Theta}, \Theta)$ [5, 8, 11]. Compared to $\text{err}_{\text{training}}$, its information-theoretic lower bound $H_L(\tilde{Y}_0 | \tilde{X}_{-\Theta}, \Theta)$ is mathematically more convenient to analyze. The gap between $\text{err}_{\text{training}}$ and the lower bound $H_L(\tilde{Y}_0 | \tilde{X}_{-\Theta}, \Theta)$ was studied recently in [23], where the gap is small if the function spaces Λ and Φ are close to each other, e.g., when the neural network is sufficiently wide and deep [10].

For notational convenience, we refer to $H_L(\tilde{Y}_0 | \tilde{X}_{-\Theta}, \Theta)$ as an *L-conditional entropy*, because it is associated with a loss function L . The *L-entropy* of a random variable Y is defined as [8, 11]

$$H_L(Y) = \min_{a \in \mathcal{A}} \mathbb{E}_{Y \sim P_Y} [L(Y, a)]. \quad (3)$$

Let a_{P_Y} denote an optimal solution to (3), which is called a *Bayes action* [11]. The *L-conditional entropy* of Y given $X = x$ is

$$H_L(Y | X = x) = \min_{a \in \mathcal{A}} \mathbb{E}_{Y \sim P_{Y|X=x}} [L(Y, a)]. \quad (4)$$

Using (4), we can get the *L-conditional entropy* of Y given X [8, 11]

$$H_L(Y | X) = \sum_{x \in \mathcal{X}} P_X(x) H_L(Y | X = x). \quad (5)$$

Similar to (5), (2) can be decomposed as

$$H_L(\tilde{Y}_0 | \tilde{X}_{-\Theta}, \Theta) = \sum_{x \in \mathcal{X}, \theta \in \mathcal{D}} P_{\tilde{X}_{-\Theta}, \Theta}(x, \theta) H_L(\tilde{Y}_0 | \tilde{X}_{-\theta} = x, \Theta = \theta). \quad (6)$$

We assume that in the training dataset, the AoI Θ is independent of the label \tilde{Y}_0 and feature $\tilde{X}_{-\mu}$ for all $\mu \geq 0$. By this assumption and (6), one can get (see our technical report [25] for its proof)

$$H_L(\tilde{Y}_0 | \tilde{X}_{-\Theta}, \Theta) = \sum_{\theta \in \mathcal{D}} P_{\Theta}(\theta) H_L(\tilde{Y}_0 | \tilde{X}_{-\theta}). \quad (7)$$

By choosing different L in (3), a broad class of *L*-entropies is obtained. In particular, Shannon's entropy is derived if L is the logarithmic loss function $L_{\log}(y, Q_Y) = -\log Q_Y(y)$. More examples of the loss function L , the definitions of *L*-divergence $D_L(P_Y || Q_Y)$, *L*-mutual information $I_L(Y; X)$, and *L*-conditional mutual information $I_L(Y; X | Z)$ are provided in [25]. In general, $I_L(X; Y) \neq I_L(Y; X)$, which is different from *f*-mutual information. Moreover, a comparison among the *L*-divergence, Bregman divergence [6], and the *f*-divergence [4] is provided in [25].

2.3 Online Inference Error

In the online inference phase, the neural predictor trained by (1) is used to predict the target in real-time. We assume that $\{(Y_t, X_t), t \in \mathbb{Z}\}$ is a stationary process that is independent of the AoI process $\{\Delta(t), t \in \mathbb{Z}\}$. Using this assumption, the time-average expected inference error during the time slots $t = 0, 1, \dots, T-1$ is given by

$$\text{err}_{\text{inference}}(T) = \frac{1}{T} \mathbb{E} \left[\sum_{t=0}^{T-1} p(\Delta(t)) \right], \quad (8)$$

where

$$p(\delta) = \mathbb{E}_{Y, X \sim P_{Y_t, X_{t-\delta}}} \left[L \left(Y, \phi^*_{P_{\tilde{Y}_0, \tilde{X}_{-\Theta}, \Theta}}(X, \delta) \right) \right], \quad (9)$$

$p(\Delta(t))$ is the expected inference error in time slot t , and $\Delta(t)$ is the inference AoI at time t , i.e., the time difference between label Y_t and feature $X_{t-\Delta(t)}$. The proof of (8) is provided in [25].

Let us define *L-cross entropy* between Y and \tilde{Y} as

$$H_L(Y; \tilde{Y}) = \mathbb{E}_{Y \sim P_Y} \left[L \left(Y, a_{P_{\tilde{Y}}} \right) \right], \quad (10)$$

and L -conditional cross entropy between Y and \tilde{Y} given X as

$$H_L(Y; \tilde{Y}|X) = \sum_{x \in \mathcal{X}} P_X(x) \mathbb{E}_{Y \sim P_{Y|X=x}} \left[L \left(Y, a_{P_{\tilde{Y}|\tilde{X}=x}} \right) \right], \quad (11)$$

where $a_{P_{\tilde{Y}}}$ and $a_{P_{\tilde{Y}|\tilde{X}=x}}$ are the Bayes actions associated with $P_{\tilde{Y}}$ and $P_{\tilde{Y}|\tilde{X}=x}$, respectively. If the neural predictor in (9) is replaced by the Bayes action $a_{\tilde{Y}_0|\tilde{X}_{-\delta}=x}$, i.e., the optimal solution to (2), then $p(\delta)$ becomes an L -conditional cross entropy

$$H_L(Y_t; \tilde{Y}_0|X_{t-\delta}) = \sum_{x \in \mathcal{X}} P_{X_{t-\delta}}(x) \mathbb{E}_{Y \sim P_{Y_t|X_{t-\delta}=x}} \left[L \left(Y, a_{\tilde{Y}_0|\tilde{X}_{-\delta}=x} \right) \right]. \quad (12)$$

If the function spaces Λ and Φ are close to each other, the difference between $p(\delta)$ and $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ is small.

3 INTERPRETATION OF FRESHNESS IN REAL-TIME SUPERVISED LEARNING

In this section, we study how the training AoI Θ and the inference AoI $\Delta(t)$ affect the performance of real-time supervised learning.

3.1 Training Error vs. Training AoI

We first consider the case of deterministic training AoI $\Theta = \theta$. Given $\Theta = \theta$, $H_L(\tilde{Y}_0|\tilde{X}_{-\theta}, \Theta)$ in (7) becomes simply $H_L(\tilde{Y}_0|\tilde{X}_{-\theta})$, which is a function of θ . One may expect that the training error would grow with the AoI θ . If $\tilde{Y}_0 \xleftrightarrow{\epsilon} \tilde{X}_{-\mu} \xleftrightarrow{\epsilon} \tilde{X}_{-\mu-\nu}$ is a Markov chain for all $\mu, \nu \geq 0$, by the data processing inequality for L -conditional entropy [5, Lemma 12.1], one can show that $H_L(\tilde{Y}_0|\tilde{X}_{-\theta})$ is a non-decreasing function of θ . Nevertheless, the experimental results in Figs. 1-2 and [24, 25] show that the training error is a growing function of the training AoI θ in some applications (e.g., video prediction), whereas it is a non-monotonic function of θ in other applications (e.g., temperature prediction and actuator state prediction with delay). As we will explain below, a fundamental reason behind these phenomena is that practical time-series data could be either Markovian or non-Markovian. For non-Markovian $(\tilde{Y}_0, \tilde{X}_{-\mu}, \tilde{X}_{-\mu-\nu})$, $H_L(\tilde{Y}_0|\tilde{X}_{-\theta})$ is not necessarily monotonic in θ .

Next, we develop an ϵ -data processing inequality to analyze information freshness for both Markovian and non-Markovian time-series data. To that end, the following relaxation of the standard Markov chain model is needed, which is motivated by [13]:

Definition 3.1 (ϵ -Markov Chain). Given $\epsilon \geq 0$, a sequence of three random variables Z, X , and Y is said to be an ϵ -Markov chain, denoted as $Z \xrightarrow{\epsilon} X \xrightarrow{\epsilon} Y$, if

$$I_{\chi^2}(Y; Z|X) = \mathbb{E}_{X, Z \sim P_{X, Z}} \left[D_{\chi^2} \left(P_{Y|X, Z} || P_{Y|X} \right) \right] \leq \epsilon^2, \quad (13)$$

where

$$D_{\chi^2}(P_Y || Q_Y) = \sum_{y \in \mathcal{Y}} \frac{(P_Y(y) - Q_Y(y))^2}{Q_Y(y)} \quad (14)$$

is Neyman's χ^2 -divergence and $I_{\chi^2}(Y; Z|X)$ is χ^2 -conditional mutual information.

A Markov chain is an ϵ -Markov chain with $\epsilon = 0$. If $Z \rightarrow X \rightarrow Y$ is a Markov chain, then $Y \rightarrow X \rightarrow Z$ is also a Markov chain [3, p. 34]. A similar property holds for the ϵ -Markov chain.

LEMMA 3.2. If $Z \xrightarrow{\epsilon} X \xrightarrow{\epsilon} Y$, then $Y \xrightarrow{\epsilon} X \xrightarrow{\epsilon} Z$.

Due to space limitation, all the proofs are relegated to our technical report [25]. By Lemma 3.2, the ϵ -Markov chain can be denoted as $Y \xleftrightarrow{\epsilon} X \xleftrightarrow{\epsilon} Z$. In the following lemma, we provide a relaxation of the data processing inequality for ϵ -Markov chain, which is called an ϵ -data processing inequality.

LEMMA 3.3 (ϵ -DATA PROCESSING INEQUALITY). If $Y \xleftrightarrow{\epsilon} X \xleftrightarrow{\epsilon} Z$ is an ϵ -Markov chain, then

$$H_L(Y|X) \leq H_L(Y|Z) + O(\epsilon). \quad (15)$$

If, in addition, $H_L(Y)$ is twice differentiable in P_Y , then

$$H_L(Y|X) \leq H_L(Y|Z) + O(\epsilon^2). \quad (16)$$

Lemma 3.3(b) was mentioned in [24] without proof. Lemma 3.3(a) is new to the best of our knowledge. Now, we are ready to characterize how $H(\tilde{Y}_0|\tilde{X}_{-\theta})$ varies with the AoI θ .

THEOREM 3.4. The L -conditional entropy

$$H_L(\tilde{Y}_0|\tilde{X}_{-\theta}) = g_1(\theta) - g_2(\theta) \quad (17)$$

is a function of θ , where $g_1(\theta)$ and $g_2(\theta)$ are two non-decreasing functions of θ , given by

$$\begin{aligned} g_1(\theta) &= H_L(\tilde{Y}_0|\tilde{X}_0) + \sum_{k=0}^{\theta-1} I_L(\tilde{Y}_0; \tilde{X}_{-k}|\tilde{X}_{-k-1}), \\ g_2(\theta) &= \sum_{k=0}^{\theta-1} I_L(\tilde{Y}_0; \tilde{X}_{-k-1}|\tilde{X}_{-k}). \end{aligned} \quad (18)$$

If $\tilde{Y}_0 \xleftrightarrow{\epsilon} \tilde{X}_{-\mu} \xleftrightarrow{\epsilon} \tilde{X}_{-\mu-\nu}$ is an ϵ -Markov chain for every $\mu, \nu \geq 0$, then $g_2(\theta) = O(\epsilon)$ and

$$H_L(\tilde{Y}_0|\tilde{X}_{-\theta}) = g_1(\theta) + O(\epsilon). \quad (19)$$

According to Theorem 3.4, the monotonicity of $H_L(\tilde{Y}_0|\tilde{X}_{-\theta})$ in θ is characterized by the parameter $\epsilon \geq 0$ in the ϵ -Markov chain model. If ϵ is small, then $\tilde{Y}_0 \xleftrightarrow{\epsilon} \tilde{X}_{-\mu} \xleftrightarrow{\epsilon} \tilde{X}_{-\mu-\nu}$ is close to a Markov chain, and $H_L(\tilde{Y}_0|\tilde{X}_{-\theta})$ is nearly non-decreasing in θ . If ϵ is large, then $\tilde{Y}_0 \xleftrightarrow{\epsilon} \tilde{X}_{-\mu} \xleftrightarrow{\epsilon} \tilde{X}_{-\mu-\nu}$ is far from a Markov chain, and $H_L(\tilde{Y}_0|\tilde{X}_{-\theta})$ could be non-monotonic in θ . Theorem 3.4 can be readily extended to random AoI Θ by using stochastic orders [22].

Definition 3.5 (Univariate Stochastic Ordering). [22] A random variable X is said to be stochastically smaller than another random variable Z , denoted as $X \leq_{st} Z$, if

$$P(X > x) \leq P(Z > x), \quad \forall x \in \mathbb{R}. \quad (20)$$

THEOREM 3.6. If $\tilde{Y}_0 \xleftrightarrow{\epsilon} \tilde{X}_{-\mu} \xleftrightarrow{\epsilon} \tilde{X}_{-\mu-\nu}$ is an ϵ -Markov chain for all $\mu, \nu \geq 0$, and the training AoIs in two experiments 1 and 2 satisfy $\Theta_1 \leq_{st} \Theta_2$, then

$$H_L(\tilde{Y}_0|\tilde{X}_{-\Theta_1}, \Theta_1) \leq H_L(\tilde{Y}_0|\tilde{X}_{-\Theta_2}, \Theta_2) + O(\epsilon). \quad (21)$$

According to Theorem 3.6, if Θ_1 is stochastically smaller than Θ_2 , then the training error in Experiment 1 is approximately smaller than that in Experiment 2. If, in addition to the conditions in Theorems 3.4 and 3.6, $H_L(\tilde{Y}_0)$ is twice differentiable in $P_{\tilde{Y}_0}$, then the last term $O(\epsilon)$ in (19) and (21) becomes $O(\epsilon^2)$.

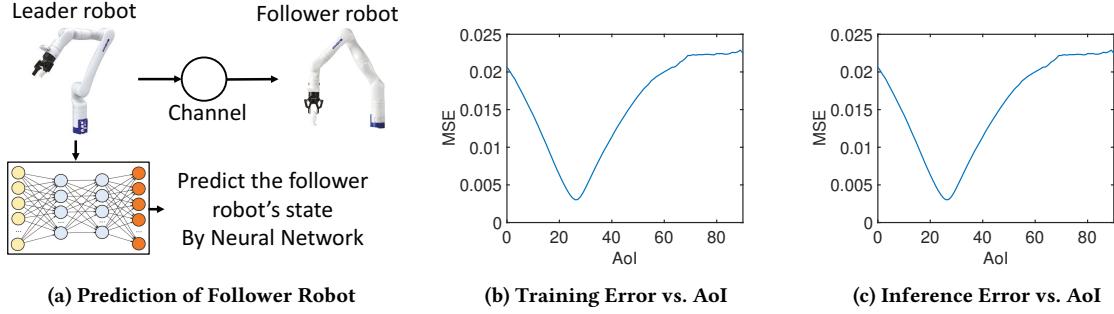


Figure 2: Robot state prediction in a leader-follower robotic system. The leader robot uses a neural network to predict the follower robot’s state. The training and inference errors decrease in the $\text{AoI} \leq 25$ and increase when $\text{AoI} \geq 25$.

3.2 Inference Error vs. Inference AoI

According to (4), (5), and (11), $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ is lower bounded by $H_L(Y_t|X_{t-\delta})$. In addition, $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ is close to its lower bound $H_L(Y_t|X_{t-\delta})$, if the conditional distributions $P_{Y_t|X_{t-\delta}}$ and $P_{\tilde{Y}_0|\tilde{X}_{-\delta}}$ are close to each other, as shown by the following lemma.

LEMMA 3.7. If for all $x \in \mathcal{X}$

$$D_{\chi^2} \left(P_{Y_t|X_{t-\delta}=x} || P_{\tilde{Y}_0|\tilde{X}_{-\delta}=x} \right) \leq \beta^2, \quad (22)$$

then

$$H_L(Y_t; \tilde{Y}_0|X_{t-\delta}) = H_L(Y_t|X_{t-\delta}) + O(\beta). \quad (23)$$

By combining Theorem 3.4 and Lemma 3.7, the monotonicity of $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ versus δ is characterized in the next theorem.

THEOREM 3.8. The following assertions are true:

- (a) If $\{(Y_t, X_t), t \in \mathbb{Z}\}$ is a stationary process, then $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ is a function of the inference AoI δ .
- (b) If, in addition, $Y_t \xrightarrow{\epsilon} X_{t-\mu} \xrightarrow{\epsilon} X_{t-\mu-v}$ is an ϵ -Markov chain for all $\mu, v \geq 0$ and (22) holds for all $x \in \mathcal{X}$ and $\delta \in \mathcal{D}$, then for all $0 \leq \delta_1 \leq \delta_2$

$$H_L(Y_t; \tilde{Y}_0|X_{t-\delta_1}) \leq H_L(Y_t; \tilde{Y}_0|X_{t-\delta_2}) + O(\max\{\epsilon, \beta\}). \quad (24)$$

According to Theorem 3.8, $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ is a function of the inference AoI δ . If ϵ and β are close to zero, $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ is nearly a non-decreasing function of δ ; otherwise, $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ can be far from a monotonic function of δ .

3.3 Interpretation of Experimental Results

We conduct several experiments to study how the training and inference errors of real-time supervised learning vary with the AoI. The code of these experiments is provided in an open-source Github repository.¹

Fig. 1 illustrates the experimental results of supervised learning based video prediction, which are regenerated from [18]. In this experiment, the video frame V_t at time t is predicted based on a feature $X_{t-\delta} = (V_{t-\delta}, V_{t-\delta-1})$ that is composed of two consecutive video frames, where $\Delta(t) = \delta$ is the AoI. A pre-trained neural network model called “SAVP” [18] is used to evaluate on 256 samples of “BAIR” dataset [7], which contains video frames of a randomly

moving robotic arm. The pre-trained neural network model can be downloaded from the Github repository of [18]. One can observe from Fig. 1(b)-(c) that the training and inference errors are non-decreasing functions of the AoI, because the video clips V_t are approximately a Markov chain.

Fig. 2 depicts the performance of robot state prediction in a leader-follower robotic system. As illustrated in a Youtube video ², the leader robot sends its state (joint angles) X_t to the follower robot through a channel. One packet for updating the leader robot’s state is sent periodically to the follower robot every 20 time-slots. The transmission time of each updating packet is 20 time-slots. The follower robot moves towards the leader’s most recent state and locally controls its robotic fingers to grab an object. We constructed a robot simulation environment using the Robotics System Toolbox in MATLAB. In each episode, a can is randomly generated on a table in front of the follower robot. The leader robot observes the position of the can and illustrates to the follower robot how to grab the can and place it on another table, without colliding with other objects in the environment. The rapidly-exploring random tree (RRT) algorithm is used to control the leader robot. Collision avoidance algorithm and trajectory generation algorithm are used for local control of the follower robot. The leader robot uses a neural network to predict the follower robot’s state Y_t . The neural network consists of one input layer, one hidden layer with 256 ReLU activation nodes, and one fully connected (dense) output layer. The dataset contains the leader and follower robots’ states in 300 episodes of continue operation. The first 80% of the dataset is used for the training and the other 20% of the dataset is used for the inference. In Fig. 2, the training and the inference error decreases in AoI, when $\text{AoI} \leq 25$ and increases in AoI when $\text{AoI} \geq 25$. In this case, even a fresh feature with $\text{AoI}=0$ is not good for prediction. In this experiment, $(Y_t, X_{t-\mu}, X_{t-\mu-v})$ is not a Markov chain for all $\mu, v \geq 0$. Hence, the training and the inference error are not non-decreasing functions of AoI.

To facilitate understanding the experimental results in Fig. 2, we provide a toy example to interpret it: Let X_t be a Markov chain and $Y_t = f(X_{t-d})$. One can view X_t as the input of a causal system with delay $d \geq 0$, and Y_t as the system output. Because $Y_t = f(X_{t-d})$, a stale system input X_{t-d} at time $t-d$ is informative for inferring the current output Y_t at time t . If the training and inference datasets

¹<https://github.com/Kamran0153/Impact-of-Data-Freshness-in-Learning>

²https://youtu.be/_z4FHuu3-ag

have similar empirical distributions, we can use data processing inequality to show that $H_L(\tilde{Y}_0|\tilde{X}_{-\delta})$ and $H_L(Y_t; \tilde{Y}_0|X_{t-\delta})$ decrease with δ when $0 \leq \delta \leq d$ and increase with δ when $\delta \geq d$, which is similar to Fig. 2. Moreover, $H_L(\tilde{Y}_0|\tilde{X}_{-d})$ is close to zero if the function space Λ is sufficiently large. It is equal to zero if $\Lambda = \Phi$. The leader-follower robotic system in Fig. 2 can be viewed as a causal system, where the system input is the leader robot's state, and the system output is the follower robot's state. Non-monotonicity occurs in Fig. 2 because the input of a causal system is used to predict the system output in this experiment, which is similar to the toy example. However, the relationship between the system input and output in Fig. 2 is more complicated than the toy example, due to the control algorithms used by the follower robot.

Besides these experiments, if there exists response delay, long-range dependence, and periodic patterns in the target and feature data sequence, the training and inference errors could also be non-monotonic functions of the AoI. This phenomenon is observed in actuator state prediction, temperature prediction, wireless channel state information prediction, and solar power prediction [24, 25].

4 SINGLE-SOURCE SCHEDULING FOR INFERENCE ERROR MINIMIZATION

As shown in Section 3, the inference error is a function of the AoI $\Delta(t)$, whereas the function is not necessarily monotonic. To reduce the inference error, we devise a new scheduling algorithm that can minimize general functions of the AoI, no matter whether the function is monotonic or not.

4.1 System Model

We consider the networked supervised learning system in Fig. 3, where a source progressively sends features through a channel to a receiver. The channel is modeled as a non-preemptive server with i.i.d. service times. At any time t , the receiver uses the latest received feature to predict the current label Y_t . To minimize the inference error, we propose a new “selection-from-buffer” model for feature transmissions, which is more general than the “generate-at-will” model [32]. Specifically, at the beginning of time slot t , the source generates a fresh feature X_t and appends it to a buffer that stores the B most recent features $(X_t, X_{t-1}, \dots, X_{t-B+1})$; meanwhile, the oldest feature X_{t-B} is removed from the buffer. The transmitter can pick any feature from the buffer and submit it to the channel when the channel is idle. A transmission scheduler determines (i) when to submit features to the channel and (ii) which feature in the buffer to submit. When $B = 1$, the “selection-from-buffer” model reduces to the “generate-at-will” model.

We assume that the system starts to operate in time slot $t = 0$ with B features $(X_0, X_1, \dots, X_{B-1})$ in the buffer. Hence, the feature buffer is full at all time $t \geq 0$. The i -th feature sent over the channel is generated in time slot G_i , is submitted to the channel in time slot S_i , is delivered and available for inference in time slot $D_i = S_i + T_i$, where $T_i \geq 1$ is the feature transmission time, $G_i \leq S_i < D_i$, and $D_i \leq S_{i+1} < D_{i+1}$. The feature transmission times T_i could be random due to time-varying channel conditions, congestion, random packet sizes, etc. We assume that the T_i 's are i.i.d. with a finite mean $1 \leq \mathbb{E}[T_i] < \infty$. In time slot $t = S_i$, the $(b_i + 1)$ -th freshest feature in the buffer is submitted to the channel,

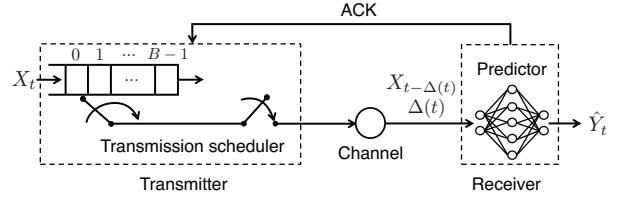


Figure 3: A networked real-time supervised learning system. At each time slot t , the transmitter generates a feature X_t and keeps it in a buffer that stores B most recent features $(X_t, X_{t-1}, \dots, X_{t-B+1})$. The scheduler decides when to submit features to the channel and which feature in the buffer to submit.

where $b_i \in \{0, 1, \dots, B-1\}$. Hence, the submitted feature is $X_{S_i-b_i}$ that was generated at time $G_i = S_i - b_i$. Once a feature is delivered, an acknowledgment (ACK) is fed back to the transmitter in the same time slot. Thus, the idle/busy state of the channel is known at the transmitter.

4.2 Scheduling Problem

Let $U(t) = \max_i \{G_i : D_i \leq t\}$ be the generation time of the latest received feature in time slot t . The age of information (AoI) at time t is given by [15]

$$\Delta(t) = t - U(t) = t - \max_i \{G_i : D_i \leq t\}. \quad (25)$$

Because $D_i < D_{i+1}$, $\Delta(t)$ can be also written as

$$\Delta(t) = t - G_i = t - S_i + b_i, \text{ if } D_i \leq t < D_{i+1}. \quad (26)$$

The initial state of the system is assumed to be $S_0 = 0$, $D_0 = T_0$, and $\Delta(0)$ is a finite constant.

A scheduling policy is denoted by a 2-tuple (f, g) , where $g = (S_1, S_2, \dots)$ determines when to submit the features and $f = (b_1, b_2, \dots)$ specifies which feature in the buffer to submit. We consider the class of *causal scheduling policies* in which each decision is made by using the current and historical information available at the transmitter. Let Π denote the set of all causal scheduling policies. We assume that the scheduler has access to the distribution of $\{(Y_t, X_t), t \in \mathbb{Z}\}$ but not its realization, and the T_i 's are not affected by the adopted scheduling policy.

Our goal is to find an optimal scheduling policy that minimizes the time-average expected inference error among all causal scheduling policies in Π :

$$\bar{p}_{opt} = \inf_{(f,g) \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{(f,g)} \left[\sum_{t=0}^{T-1} p(\Delta(t)) \right]. \quad (27)$$

where $p(\Delta(t))$ is the inference error at time slot t , defined in (9), and \bar{p}_{opt} is the optimum value of (27). Because $p(\cdot)$ is not necessarily a non-decreasing function, (27) is more challenging than the scheduling problems in [26, 28].

4.3 Optimal Single-source Scheduling

We solve (27) in two steps: (i) Given a fixed feature selection policy $f_b = (b, b, \dots)$ with $b_i = b$ for all i , find the optimal feature

submission times $g = (S_1, S_2, \dots)$ that solves

$$\bar{p}_b = \inf_{(f_b, g) \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{(f_b, g)} \left[\sum_{t=0}^{T-1} p(\Delta(t)) \right], \quad (28)$$

(ii) Use the solution to (28) to describe an optimal solution to (27).

It turns out that optimal solution to (28) can be obtained by using the Gittins index of the following *AoI bandit process with a random termination delay T_1* : A bandit process $\Delta(t)$ is controlled by a decision-maker that chooses between two actions **CONTINUE** and **STOP** in each time slot. If the bandit process is not terminated in time slot t , its state evolves according to

$$\Delta(t) = \Delta(t-1) + 1, \quad (29)$$

and a reward $[r - p(\Delta(t))]$ is collected, where $p(\cdot)$ is defined in (9) and r is a constant reward. If the **CONTINUE** action is selected, the bandit process continues to evolve. If the **STOP** action is selected, the bandit process will terminate after a random delay T_1 and no more action is taken. Once the bandit process terminates, its state and reward remain zero. The total profit of the bandit process starting from time t is maximized by solving the following optimal stopping problem:

$$\sup_{v \in \mathfrak{M}} \mathbb{E} \left[\sum_{k=0}^{v+T_1-1} [r - p(\Delta(t+k))] \middle| \Delta(t) = \delta \right], \quad (30)$$

where $v \geq 0$ is a history-dependent stopping time and \mathfrak{M} is the set of all stopping times of the bandit process $\{\Delta(t+k), k = 0, 1, \dots\}$. Following the derivation of the Gittins index in [9, Chapter 2.5], the decision-maker should choose the **STOP** action at time

$$\min_{t \in \mathbb{Z}} \{t \geq 0 : \gamma(\Delta(t)) \geq r\}, \quad (31)$$

where

$$\gamma(\delta) = \inf_{v \in \mathfrak{M}, v \neq 0} \frac{\mathbb{E} \left[\sum_{k=0}^{v-1} p(\Delta(t+k+T_1)) \middle| \Delta(t) = \delta \right]}{\mathbb{E}[v \mid \Delta(t) = \delta]} \quad (32)$$

is the Gittins index, i.e., the value of reward r for which the **CONTINUE** and **STOP** actions are equally profitable at state $\Delta(t) = \delta$. As shown in [25], (32) can be simplified as

$$\gamma(\delta) = \inf_{\tau \in \{1, 2, \dots\}} \frac{1}{\tau} \sum_{k=0}^{\tau-1} \mathbb{E}[p(\delta + k + T_1)], \quad (33)$$

where τ is a positive integer.

THEOREM 4.1. *If $|p(\delta)| \leq M$ for all δ and the T_i 's are i.i.d. with a finite mean $\mathbb{E}[T_i]$, then $g = (S_1(\beta_b), S_2(\beta_b), \dots)$ is an optimal solution to (28), where*

$$S_{i+1}(\beta_b) = \min_{t \in \mathbb{Z}} \{t \geq D_i(\beta_b) : \gamma(\Delta(t)) \geq \beta_b\}, \quad (34)$$

$D_i(\beta_b) = S_i(\beta_b) + T_i$ is the delivery time of the i -th feature submitted to the channel, $\Delta(t) = t - S_i(\beta_b) + b$ is the AoI at time t , $\gamma(\delta)$ is the Gittins index in (33), and β_b is the unique root of

$$\mathbb{E} \left[\sum_{t=D_i(\beta_b)}^{D_{i+1}(\beta_b)-1} p(\Delta(t)) \right] - \beta_b \mathbb{E}[D_{i+1}(\beta_b) - D_i(\beta_b)] = 0. \quad (35)$$

The optimal objective value to (28) is given by

$$\bar{p}_b = \frac{\mathbb{E} \left[\sum_{t=D_i(\beta_b)}^{D_{i+1}(\beta_b)-1} p(\Delta(t)) \right]}{\mathbb{E}[D_{i+1}(\beta_b) - D_i(\beta_b)]}. \quad (36)$$

Furthermore, β_b is exactly the optimal value to (28), i.e., $\beta_b = \bar{p}_b$.

The optimal scheduling policy in Theorem 4.1 has an intuitive structure. Specifically, a feature is transmitted in time-slot t if two conditions are satisfied: (i) The channel is idle in time-slot t , (ii) the Gittins index $\gamma(\Delta(t))$ exceeds a threshold β_b (i.e., $\gamma(\Delta(t)) \geq \beta_b$), where the threshold β_b is exactly equal to the minimum time-averaged inference error \bar{p}_b . The optimal objective value \bar{p}_b is computed by solving (35). Three low-complexity algorithms for solving (35) were provided in [20, Algorithms 1-3]. In practical supervised learning algorithms, the features are shifted, rescaled, and clipped during the data preprocessing step, which can improve the convergence speed. Because of these operations, the inference error is finite in practice (See Figures 1-2 for a few example), and the condition $|p(\delta)| \leq M$ for all δ in Theorem 4.1 is not restrictive in practice.

Theorem 4.1 is proven by directly solving the Bellman optimality equation of the Markov decision process (28), whereas the techniques for minimizing non-decreasing AoI functions in, e.g., [26, 28], could not solve (28). We remark that if $p(\delta)$ is non-monotonic, then $\gamma(\delta)$ is not necessarily monotonic. Hence, (34) in general could not be rewritten as a threshold policy of the AoI $\Delta(t)$ in the form of $\Delta(t) \geq \beta$. This is a key difference from the minimization of non-decreasing AoI functions, e.g., [26, Eq. (48)]. The adoption of the Gittins index $\gamma(\delta)$ as a tool for solving (28) is motivated by a similarity between (28) and the restart-in-state formulation of the Gittins index [9, Chapter 2.6.4]. This connection between the Gittins index theory and AoI minimization was unknown before.

Next, we present an optimal solution to (27).

THEOREM 4.2. *If the conditions of Theorem 4.1 hold, then there exists an optimal solution (f^*, g^*) to (27) that satisfies:*

(a) $f^* = (b^*, b^*, \dots)$, where b^* is obtained by solving

$$b^* = \arg \min_{b \in \{0, 1, \dots, B-1\}} \beta_b, \quad (37)$$

and β_b is the unique root to (35).

(b) $g^* = (S_1^*, S_2^*, \dots)$, where

$$S_i^* = \min_{t \in \mathbb{Z}} \{t \geq S_i^* + T_i : \gamma(\Delta(t)) \geq \bar{p}_{opt}\}, \quad (38)$$

$S_i^* + T_i$ is the delivery time of the i -th feature, $\gamma(\delta)$ is the Gittins index in (33), and \bar{p}_{opt} is the optimal objective value of (27), determined by

$$\bar{p}_{opt} = \min_{b \in \{0, 1, \dots, B-1\}} \beta_b. \quad (39)$$

Theorem 4.2 tells us that, to solve (27), a feature is transmitted in time-slot t if two conditions are satisfied: (i) The channel is idle in time-slot t , (ii) the Gittins index $\gamma(\Delta(t))$ exceeds a threshold \bar{p}_{opt} (i.e., $\gamma(\Delta(t)) \geq \bar{p}_{opt}$), where the threshold \bar{p}_{opt} is the optimal objective value of (27). The optimal objective value \bar{p}_{opt} is determined by (39).

In the special case of non-decreasing $p(\cdot)$ studied in [26, 28], the Gittins index in (33) can be simplified as $\gamma(\delta) = \mathbb{E}[p(\delta + T_1)]$ and the optimal solution to (37) is $b^* = 0$ such that it is optimal to always

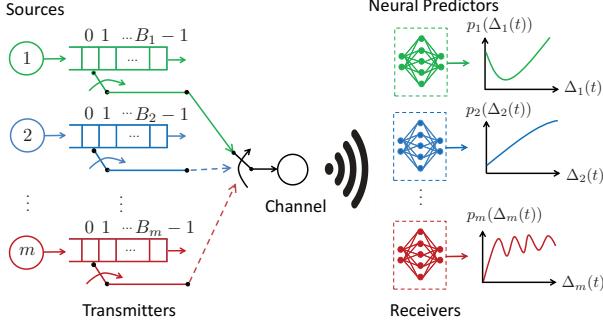


Figure 4: A networked intelligent system, where m sources send features over a shared channel to the corresponding neural predictors. At any time, at most one source can occupy the channel.

select the freshest feature from the buffer. Hence, Theorem 3 in [26] is recovered from Theorem 4.2, and the “generate-at-will” model can achieve the minimum inference error in this special case.

If $p(\cdot)$ is non-monotonic, as in the cases of Fig. 2, the “selection-from-buffer” model could achieve better performance than the “generate-at-will” model, and the optimal scheduler is provided by Theorem 4.2.

5 MULTIPLE-SOURCE SCHEDULING

5.1 System Model and Scheduling Problem

Consider the networked intelligent system in Fig. 4, where m sources send features over a shared channel to the corresponding neural predictors at the receivers. At time slot t , each source l maintains a buffer that stores the B_l most recent features $(X_{l,t}, \dots, X_{l,t-B_l+1})$. When the channel is free, at most one source can select a feature from its buffer and submit the selected feature to the channel.

A centralized scheduler makes two decisions in each time slot: (i) which source should submit a feature to the shared channel and (ii) which feature in the selected source’s buffer to submit. A scheduling policy is denoted by $\pi = (\pi_{l,b_l})_{l=1,2,\dots,m, b_l=0,1,\dots,B_l-1}$, where $\pi_{l,b_l} = (d_{l,b_l}(0), d_{l,b_l}(1), \dots)$ and $d_{l,b_l}(t) \in \{0, 1\}$ represents the scheduling decision for the $(b_l + 1)$ -th freshest feature $X_{l,t-b_l}$ of source l in time slot t . If source l submits the feature $X_{l,t-b_l}$ in its buffer to the channel in time slot t , then $d_{l,b_l}(t) = 1$; otherwise, $d_{l,b_l}(t) = 0$. Let $c_{l,b_l}(t) \in \{0, 1\}$ denote the channel occupation status of the $(b_l + 1)$ -th freshest feature $X_{l,t-b_l}$ of source l in time slot t . If source l submits the feature $X_{l,t-b_l}$ in its buffer to the channel in time slot t , then the value of $c_{l,b_l}(t)$ becomes 1 and remains 1 until it is delivered; otherwise, $c_{l,b_l}(t) = 0$. It is required that $\sum_{l=1}^m \sum_{b_l=0}^{B_l-1} c_{l,b_l}(t) \leq 1$ for all t . Let Π denote the set of all causal scheduling policies.

Let $G_{l,i}$, $S_{l,i}$, $D_{l,i}$, and $T_{l,i}$ denote the generation time, channel submission time, delivery time, and transmission time duration of the i -th feature sent by source l , respectively. The feature transmission times $T_{l,i} \geq 1$ are independent across the sources and i.i.d. among the features from the same source. We assume that the $T_{l,i}$ ’s are not affected by the adopted scheduling policy. The age of information (AoI) of source l at time slot t is given by

$$\Delta_l(t) = t - \max_i \{G_{l,i} : D_{l,i} \leq t\}. \quad (40)$$

Algorithm 1 Whittle Index Policy with Selection-from-Buffer

```

1: Do forever:
2: Update  $\Delta_l(t)$  for all  $l \in \{1, 2, \dots, m\}$ .
3: Calculate the Whittle index  $W_{l,b_l}(\Delta_l(t))$  for all  $l \in \{1, 2, \dots, m\}$ 
   and  $b_l \in \{0, 1, \dots, B_l - 1\}$  using (45)-(47).
4: if the channel is idle and  $\max_{l,b_l} W_{l,b_l}(\Delta_l(t)) \geq 0$  then
5:    $(l^*, b_{l^*}^*) \leftarrow \arg \max_{l,b_l} W_{l,b_l}(\Delta_l(t))$ .
6:   Source  $l^*$  submits its feature  $X_{l^*,t-b_{l^*}^*}$  to the channel.
7: else
8:   No source is scheduled, even if the channel is idle.
9: end if

```

Our goal is to minimize the time-average weighted sum of the inference errors of the m sources, which is formulated by

$$\inf_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{l=1}^m w_l \mathbb{E}_{\pi} \left[\sum_{t=0}^{T-1} p_l(\Delta_l(t)) \right], \quad (41)$$

$$\text{s.t. } \sum_{l=1}^m \sum_{b_l=0}^{B_l-1} c_{l,b_l}(t) \leq 1, \quad t = 0, 1, 2, \dots, \quad (42)$$

where $p_l(\Delta_l(t))$ is the inference error of source l at time slot t and $w_l > 0$ is the weight of source l .

5.2 Multiple-source Scheduling

Problem (41) can be cast as a Restless Multi-arm Bandit (RMAB) problem by viewing the features stored in the source buffers as arms, where (l, b_l) is an arm associated with the $(b_l + 1)$ -th freshest feature of the source l and the state of the arm (l, b_l) is the AoI $\Delta_l(t)$ in (40). Finding the optimal solution for RMAB is generally PSPACE hard [21]. Next, we develop a low-complexity scheduling policy by using both Gittins and Whittle indices.

By relaxing the per-slot channel constraint (42) as the following time-average expected channel constraint

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^m \sum_{b_l=0}^{B_l-1} \mathbb{E}[c_{l,b_l}(t)] \leq 1, \quad (43)$$

and taking the Lagrangian dual decomposition of the relaxed scheduling problem (41) and (43), we obtain following per-arm scheduling problem:

$$\inf_{\pi_{l,b_l} \in \Pi_{l,b_l}} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{\pi_{l,b_l}} \left[\sum_{t=0}^{T-1} w_l p_l(\Delta_l(t)) + \lambda c_{l,b_l}(t) \right], \quad (44)$$

where Π_{l,b_l} is the set of all causal scheduling policies of arm (l, b_l) .

Definition 5.1 (Indexability). [31] Let $\Omega_{l,b_l}(\lambda)$ be the set of all AoI values δ such that if the channel is idle and $\Delta_l(t) = \delta$, the optimal action to (44) is $d_{l,b_l}(t) = 0$. Then, the arm (l, b_l) is *indexable* if $\lambda_1 \leq \lambda_2$ implies $\Omega_{l,b_l}(\lambda_1) \subseteq \Omega_{l,b_l}(\lambda_2)$.

THEOREM 5.2. If $|p_l(\delta)| \leq M$ for all δ and the $T_{l,i}$ ’s are independent across the sources and i.i.d. among the features from the same source with a finite mean $\mathbb{E}[T_{l,i}]$, then all arms are indexable.

Given indexability, the Whittle index $W_{l,b_l}(\delta)$ [31] of the arm (l, b_l) at state δ is $W_{l,b_l}(\delta) = \inf\{\lambda \in \mathbb{R} : \delta \in \Omega_{l,b_l}(\lambda)\}$.

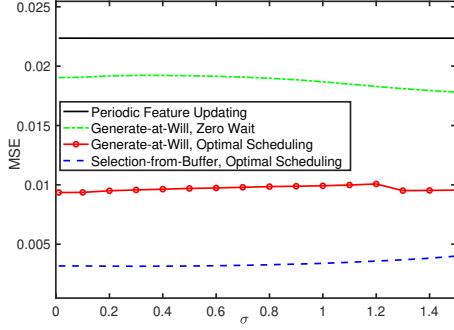


Figure 5: Time average inference error (MSE) vs. the scale parameter σ of discretized i.i.d. log-normal transmission time distribution for single-source scheduling (in robot state prediction task).

THEOREM 5.3. *If the conditions of Theorem 5.2 hold, then the Whittle index $W_{l,b_l}(\delta)$ is given by*

$$W_{l,b_l}(\delta) = \frac{w_l}{\mathbb{E}[T_{l,1}]} \mathbb{E} [z(T_{l,1}, b_l, \delta) + T_{l,2}] \gamma_l(\delta) - \frac{w_l}{\mathbb{E}[T_{l,1}]} \mathbb{E} \left[\sum_{t=T_{l,1}}^{T_{l,1}+z(T_{l,1}, b_l, \delta)+T_{l,2}-1} p_l(t+b_l) \right], \quad (45)$$

where $\gamma_l(\delta)$ is the Gittins index of an AoI bandit process for source l , determined by

$$\gamma_l(\delta) = \inf_{\tau \in \{1, 2, \dots\}} \frac{1}{\tau} \sum_{k=0}^{\tau-1} \mathbb{E} [p_l(\delta + k + T_{l,2})], \quad (46)$$

and

$$z(T_{l,1}, b_l, \delta) = \inf_{z \in \mathbb{Z}} \{z \geq 0 : \gamma_l(T_{l,1} + b_l + z) \geq \gamma_l(\delta)\}. \quad (47)$$

Finding a (semi-)analytical expression of the Whittle index for minimizing non-monotonic AoI functions is in a challenging task. In Theorem 5.3, this challenge is resolved by using the Gittins index $\gamma_l(\delta)$ to solve (44), where the solution techniques of (28) are employed. The Whittle index scheduling policy for reducing the weighted-sum inference error is described in Algorithm 1, where all sources remain silent when the channel is idle, if $W_{l,b_l}(\Delta_l(t)) < 0$ for all l and b_l .

In the special case that (i) the AoI function $p(\cdot)$ is non-decreasing and (ii) the transmission time is fixed as $T_{l,i} = 1$, it holds that $\gamma_l(\delta) = p_l(\delta + 1)$ and $z(T_{l,1}, b_l, \delta) = \max\{\delta - b_l - 1, 0\}$. Hence,

$$W_{l,0}(\delta) = w_l \left[\delta p_l(\delta + 1) - \sum_{t=1}^{\delta} p_l(t) \right] \quad (48)$$

for $\delta \geq 1$ and $b_l = 0$. By this, the Whittle index in Section IV of [29, Equation (7)] is recovered from Theorem 5.3.

6 DATA DRIVEN EVALUATIONS

In this section, we illustrate the performance of our scheduling policies, where the inference error function $p(\delta)$ is collected from the data driven experiments in Section 3.3.

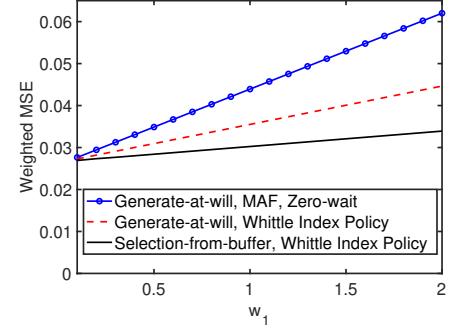


Figure 6: Time-average weighted sum of the inference errors (Normalized MSE) vs. the weight w_1 of Source 1 for multi-source scheduling, where the number of sources is $m = 2$ and the weight of Source 2 is $w_2 = 1$.

6.1 Single-source Scheduling Policies

We evaluate the following four single-source scheduling policies:

1. Generate-at-will, zero wait: The $(i+1)$ -th feature sending time S_{i+1} is given by $S_{i+1} = D_i = S_i + T_i$ and the feature selection policy is $f = (0, 0, \dots)$, i.e., $b_i = 0$ for all i .
2. Generate-at-will, optimal scheduling: The policy is given by Theorem 4.1 with $b_i = 0$ for all i .
3. Selection-from-buffer, optimal scheduling: The policy is given by Theorem 4.2.
4. Periodic feature updating: Features are generated periodically with a period T_p and appended to a queue with buffer size B . When the buffer is full, no new feature is admitted to the buffer. Features in the buffer are sent over the channel in a first-come, first-served order.

Fig. 5 illustrates the time-average inference error achieved by the four single-source scheduling policies defined above. The inference error function $p(\delta)$ used in this evaluation is illustrated in Fig. 2(c), which is generated by using the leader-follower robotic dataset and the trained neural network as explained in Section 3.3. The i -th feature transmission time T_i is assumed to follow a discretized i.i.d. log-normal distribution. In particular, T_i can be expressed as $T_i = \lceil \alpha e^{\sigma Z_i} / \mathbb{E}[e^{\sigma Z_i}] \rceil$, where Z_i 's are i.i.d. Gaussian random variables with zero mean and unit variance. In Fig. 5, we plot the time average inference error versus the scale parameter σ of discretized i.i.d. log-normal distribution, where $\alpha = 1.2$, the buffer size is $B = 30$, and the period of uniform sampling is $T_p = 3$. The randomness of the transmission time increases with the growth of σ . Data-driven evaluations in Fig. 5 show that “selection-from-buffer” with optimal scheduler achieves 3 times performance gain compared to “generate-at-will,” and 8 times performance gain compared to periodic feature updating.

6.2 Multiple-source Scheduling Policies

Now, we evaluate the following three multiple-source scheduling policies:

1. Generate-at-will, maximum age first (MAF), zero wait: At time slot t , if the channel is free, this policy will schedule

the freshest generated feature from source $\arg \max_l \Delta_l(t)$; otherwise no source is scheduled.

2. Generate-at-will, Whittle index policy: Denote

$$W_0(t) = \max_l W_{l,0}(\Delta_l(t)), \quad l_0^* = \arg \max_l W_{l,0}(\Delta_l(t)). \quad (49)$$

If the channel is free and $W_0(t) \geq 0$, the freshest feature of the source l_0^* is scheduled; otherwise no source is scheduled.

3. Selection-from-buffer, Whittle index policy: The policy is described in Algorithm 1.

In Fig. 6, we plot the time average weighted sum of inference errors versus weight w_1 , where the number of sources is $m = 2$ and weight $w_2 = 1$. The inference error function $p_1(\delta)$ is illustrated in Fig. 2(c). The inference error function $p_2(\delta)$ is illustrated in Fig. 1(c), which is generated by using the pre-trained neural network on “BAIR” dataset from [18]. The transmission times for Source 1 and Source 2 are $T_{1,i} = 1$ and $T_{2,i} = 4$ for all i , respectively. The buffer sizes are $B_1 = B_2 = 30$. The weight w_1 is associated with a non-monotonic AoI function. The performance gain of “selection-from-buffer, Whittle index policy” increases as w_1 grows. Due to limited space, more numerical results are provided in [25].

7 CONCLUSIONS

In this paper, we interpreted the impact of data freshness on the performance of real-time supervised learning. We showed that the training error and the inference error of real-time supervised learning could be non-monotonic AoI functions if the target and feature data sequence is far from a Markov model. Our experimental results suggested that the data sequence can be far from Markovian due to response delay, communication delay, and/or long-range dependence. To minimize the time-average inference error, we adopted a new feature transmission model called “selection-from-buffer” and designed an optimal single-source scheduling policy. The optimal single-source scheduling policy is found to be a threshold policy on the Gittins index. Moreover, we developed a Whittle index policy for multiple-source scheduling and provided a semi-analytical expression for the Whittle index. Our numerical results validated the efficacy of the proposed scheduling policies.

ACKNOWLEDGEMENT

The authors are grateful to Vijay Subramanian for one suggestion, to John Hung for useful discussions on this work, and to Shaoyi Li for his help on Fig. 1(b)-(c). This work was supported in part by the NSF grant CCF-1813078 and the ARO grant W911NF-21-1-0244.

REFERENCES

- [1] Ahmed M Bedewy, Yin Sun, Sastry Kompella, and Ness B Shroff. 2021. Optimal sampling and scheduling for timely status updates in multi-source networks. *IEEE Trans. Inf. Theory* 67, 6 (2021), 4019–4034.
- [2] Gongpu Chen, Soung Chang Liew, and Yulin Shao. 2021. Uncertainty-of-information scheduling: A restless multi-armed bandit framework. *arXiv:2102.06384* (2021).
- [3] Thomas M Cover. 1999. *Elements of information theory*. John Wiley & Sons.
- [4] Imre Csiszár and Paul C Shields. 2004. Information theory and statistics: A tutorial. (2004).
- [5] A. Philip Dawid. 1998. Coherent measures of discrepancy, uncertainty and dependence, with applications to Bayesian predictive experimental design. *Technical Report* 139 (1998).
- [6] Inderjit S Dhillon and Joel A Tropp. 2008. Matrix nearness problems with Bregman divergences. *SIAM J. Matrix Anal. Appl.* 29, 4 (2008), 1120–1146.
- [7] Frederik Ebert, Chelsea Finn, Alex Lee, and Sergey Levine. 2017. Self-Supervised Visual Planning with Temporal Skip Connections. In *Conference on Robot Learning (CoRL)*.
- [8] Farzan Farnia and David Tse. 2016. A minimax approach to supervised learning. *NIPS* 29 (2016), 4240–4248.
- [9] John Gittins, Kevin Glazebrook, and Richard Weber. 2011. *Multi-armed bandit allocation indices*. John Wiley & Sons.
- [10] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [11] Peter D. Grünwald and A. Philip Dawid. 2004. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Annals of Statistics* 32, 4 (08 2004), 1367–1433.
- [12] Daojing Guo and I-Hong Hou. 2019. On the Credibility of Information Flows in Real-time Wireless Networks. In *WiOPT*. 1–8.
- [13] Shao-Lun Huang, Anuran Makur, Gregory W Wornell, and Lizhong Zheng. 2019. On universal features for high-dimensional learning and inference. *accepted to Foundations and Trends in Communications and Information Theory: Now Publishers* (2019). Available in <https://arxiv.org/abs/1911.09105>.
- [14] Igor Kadota, Abhishek Sinha, Elif Uysal-Biyikoglu, Rahul Singh, and Eytan Modiano. 2018. Scheduling Policies for Minimizing Age of Information in Broadcast Wireless Networks. *IEEE/ACM Trans. Netw.* 26, 6 (2018), 2637–2650.
- [15] Sanjit Kaul, Roy Yates, and Marco Gruteser. 2012. Real-time status: How often should one update?. In *IEEE INFOCOM*. 2731–2735.
- [16] Markus Klügel, Mohammad H Mamduhi, Sandra Hirche, and Wolfgang Kellerer. 2019. AoI-penalty minimization for networked control systems with packet loss. In *IEEE INFOCOM Age of Information Workshop*. 189–196.
- [17] Antzela Kosta, Nikolaos Pappas, Anthony Ephremides, and Vangelis Angelakis. 2017. Age and value of information: Non-linear age case. In *IEEE ISIT*. 326–330.
- [18] Alex X Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. 2018. Stochastic adversarial video prediction. *arXiv:1804.01523* (2018).
- [19] Sajjad Mozaffari, Omar Y Al-Jarrah, Mehrdad Dianati, Paul Jennings, and Alexandros Mouzakitis. 2020. Deep learning-based vehicle behavior prediction for autonomous driving applications: A review. *IEEE Trans. Intell. Transp. Syst.* 23, 1 (2020), 33–47.
- [20] Tasmeen Zaman Ornee and Yin Sun. 2021. Sampling and remote estimation for the Ornstein-Uhlenbeck process through queues: Age of information and beyond. *IEEE/ACM Trans. on Netw.* 29, 5 (2021), 1962–1975.
- [21] Christos H Papadimitriou and John N Tsitsiklis. 1994. The complexity of optimal queueing network control. In *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory*. 318–322.
- [22] Moshe Shaked and J George Shanthikumar. 2007. *Stochastic orders*. Springer Science & Business Media.
- [23] Md Kamran Chowdhury Shisher, Tasmeen Zaman Ornee, and Yin Sun. 2022. A Local Geometric Interpretation of Feature Extraction in Deep Feedforward Neural Networks. *arXiv:2202.04632* (2022).
- [24] Md Kamran Chowdhury Shisher, Heyang Qin, Lei Yang, Feng Yan, and Yin Sun. 2021. The age of correlated features in supervised learning based forecasting. In *IEEE INFOCOM Age of Information Workshop*.
- [25] MD Kamran Chowdhury Shisher and Yin Sun. 2022. How Does Data Freshness Affect Real-time Supervised Learning? *Technical Report* (2022). Available in <https://arxiv.org/abs/2208.06948>.
- [26] Yin Sun and Benjamin Cyr. 2019. Sampling for Data Freshness Optimization: Non-linear Age Functions. *J. Commun. Netw.* 21, 3 (2019), 204–219.
- [27] Yin Sun, Yury Polyanskiy, and Elif Uysal. 2020. Sampling of the Wiener Process for Remote Estimation Over a Channel With Random Delay. *IEEE Trans. Inf. Theory* 66, 2 (2020), 1118–1135.
- [28] Yin Sun, Elif Uysal-Biyikoglu, Roy D Yates, C Emre Koksal, and Ness B Shroff. 2017. Update or wait: How to keep your data fresh. *IEEE Trans. Inf. Theory* 63, 11 (2017), 7492–7508.
- [29] Vishrant Tripathi and Eytan Modiano. 2019. A Whittle Index Approach to Minimizing Functions of Age of Information. In *IEEE Allerton*. 1160–1167.
- [30] Zijing Wang, Mihai-Alin Badiu, and Justin P Coon. 2022. A Framework for Characterising the Value of Information in Hidden Markov Models. *IEEE Transactions on Information Theory* (2022).
- [31] Peter Whittle. 1988. Restless bandits: Activity allocation in a changing world. *Journal of applied probability* 25, A (1988), 287–298.
- [32] Roy D Yates. 2015. Lazy is timely: Status updates by an energy harvesting source. In *IEEE ISIT*. 3008–3012.
- [33] Roy D Yates, Yin Sun, D Richard Brown, Sanjit K Kaul, Eytan Modiano, and Sennur Ulukus. 2021. Age of information: An introduction and survey. *IEEE Journal on Selected Areas in Communications* 39, 5 (2021), 1183–1210.
- [34] Xin Zhang, Jia Liu, and Zhengyuan Zhu. 2020. Taming convergence for asynchronous stochastic gradient descent with unbounded delay in non-convex learning. In *2020 59th IEEE Conference on Decision and Control (CDC)*. 3580–3585.