Near-optimal Moving Average Estimation at Characteristic Timescales: An Allan Variance Approach

Hossein Haeri, *Student Member, IEEE*, Craig E. Beal, *Member, IEEE*, and Kshitij Jerath, *Member, IEEE*

***Abstract*— A major challenge in moving average (MA) estimation is the selection of an appropriate averaging window length or timescale over which measurements remain relevant to the estimation task. Prior works typically perform timescale selection by examining multiple window lengths (or models) before selecting the 'optimal' one using heuristics, domain knowledge expertise, goodness-of-fit, or information criterion (e.g. AIC, BIC etc.). In the presented work, we propose an alternative mechanism based on Allan Variance (AVAR) that obviates the need for assessing multiple models and systematically reduces reliance on heuristics or rules-of-thumb. The Allan Variance approach is used to identify the timescale that minimizes bias, thus determining the timescale over which past information remains most relevant. We also introduce an alternative method to obtain AVAR for unevenly spaced timeseries. The results from moving average estimation using an Allan Variance-determined window length are compared to the optimal moving average estimator that minimizes mean square error (MSE) for a variety of signals corrupted with Gaussian white noise. While the relevant timescales determined through AVAR tend to be longer than those associated with minimum MSE (i.e. AVAR-based MA estimation requires more measurements spread over a longer period of time), the AVAR-based moving average approach provides a valuable, systematic technique for near-optimal simple moving average estimation.**

***Index Terms*— Moving average estimation, Allan Variance, data relevance, timescale, noise.**

I. Introduction

MOVING Average (MA) estimation has been used effectively across a wide range of domains including hydrology [1], neuroscience [2], econometrics [3], and robotics [4] [5] to name a few. Its widespread adoption is primarily driven by its ease of use, and often by the inability to use model-based techniques such as the Kalman filter, due to lack of reliable, high-fidelity dynamical models in specific domains. For example, in the field of safe autonomous vehicles, knowledge of the friction coefficient between the vehicle tires and roadway may be crucial for safe operation. However, it is extremely difficult to generate high-fidelity models of how the friction coefficient varies over time and space [6]. Thus, moving average estimation or similar data-driven methods may be the only recourse available to make such predictions. While there are many powerful data-driven parameter estimation techniques such as deep neural networks, they often require a large training data set and are unable to respond to sudden unsampled changes in the dynamics of the parameter.

Moving average estimation is itself a broad field of study that includes various techniques such as Auto-regressive Moving Average (ARMA), locally weighted moving average [7], and kernelized ARMA estimation [8]. While the concept of moving averages and MA estimation has been around since the early twentieth century, the challenge of selecting an appropriate or characteristic timescale over which to average measurements persists to this day, encompassing the many techniques mentioned above. For the purposes of this paper, we refer to the *characteristic timescale* as the window length that includes the most relevant measurements or information for a given estimation task, where maximizing relevancy is associated with minimizing bias.

Practically, the research community has addressed this issue by examining several different timescales (or window lengths or orders). By fitting a model for each of these window lengths, we can determine goodness-of-fit (or use information criteria such as AIC) to select an appropriate timescale that best suits the modeling or estimation task. Depending on the complexity of the data and problem, this may be a computationally inefficient and time-intensive approach.

While there exist several model selection methods which help in efficiently searching for an appropriate or characteristic timescale [9] [10] [11] the process of finding an appropriate or characteristic timescale requires multiple iterations and a significant degree of trial and error.

To overcome these limitations, we leverage the concept of Allan Variance (AVAR), and propose a novel and more systematic on-the-fly method to determine the time horizon over which measurements remain relevant to the estimation task. Allan Variance was originally developed to ascertain the frequency of atomic clocks, but has since been adapted to study a large variety of timeseries signals in a diverse range of applications [12] [13]. In this paper, we demon-

Submitted on 11 Sep 2020. "This work was supported by the National Science Foundation (NSF) through the Cyber-Physical Systems (CPS) program award CNS-1932138"

Hossein Haeri is a PhD student at the Department of Mechanical Engineering, University of Massachusetts, Lowell, MA 01854 USA (e-mail: hossein_haeri@uml.edu).

Craig E. Beal is with the Department of Mechanical Engineering, Bucknell University, Lewisburg, PA 17837 USA (e-mail: cbeal@bucknell.edu).

Kshitij Jerath is with the Department of Mechanical Engineering, University of Massachusetts, Lowell, MA 01854 USA (e-mail: kshitij_jerath@uml.edu).

strate that using the AVAR-determined timescale results in near-optimal moving average estimation. While we focus on simple moving average estimation, where measurements are weighted equally across the observing window, the presented methodology may be generalized to other moving average formulations. Moreover, due to the original intent and design of Allan Variance, most studies incorporating AVAR have been restricted to analysis of regularly-sampled timeseries. We extend AVAR to irregularly-sampled timeseries, enabling the identification of characteristic timescales and performing simple MA estimation for such data.

II. Allan Variance

While Allan Variance was originally developed to address the issue of frequency stability and time synchronization between atomic clocks [12], it has quickly become a useful tool for modeling and de-noising inertial sensors [14] [15]. Detailed analysis on noise modeling and characterization using Allan Variance can be found in a previous work by one of the authors [13]. In the presented work, we assume that the underlying signals (deterministic or stochastic) are corrupted by Gaussian white noise, though the AVAR approach extends easily to other noise types as well. Allan Variance helps determine a trade-off between signal components whose variance increases with time versus those whose variance decreases over time, to obtain a characteristic timescale that minimizes bias.

Mathematically, the local moving average $\bar{\theta}_k$ of a regularly-sampled timeseries $\Theta = \{\theta_1, \theta_2, ..., \theta_n\}$ at time instant k with a window length (or timescale) of m is defined as:

$$\bar{\theta}_k = \frac{1}{m} \sum_{i=k-m+1}^{k} \theta_i \tag{1}$$

where m represents the window length or the number of samples in each averaging window $(m < n/2)$, n represents the number of measurements, and θ denotes a measured variable or parameter (such as road-tire friction coefficient in studies on safe vehicle autonomy).

The 2-sample Allan Variance for a regularly-sampled time series Θ as a function of timescale m is defined as [12]:

$$\sigma_A^2(m) = \frac{1}{2}\mathbb{E}\left[(\bar{\theta}_k - \bar{\theta}_{k-m})^2\right] \tag{2}$$

where the expectation operator in (2) may be approximated, so that the Allan Variance can be evaluated using measured timeseries data with the following expression [16]:

$$\hat{\sigma}_A^2(m) = \frac{1}{2(n-2m)} \sum_{k=2m+1}^{n} \left(\bar{\theta}_k - \bar{\theta}_{k-m}\right)^2 \tag{3}$$

Figure 1 shows the Allan Variance for a random walk signal corrupted by white noise as a function of the averaging window length or timescale, plotted on a logarithmic scale (as is typical). While this approach for calculating Allan Variance works well for regularly-sampled data, many real-world applications do not readily generate such data. As a result, we present alternative approaches for scenarios where the timeseries data are not regularly sampled.

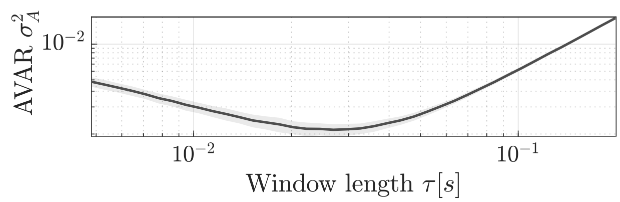

Fig. 1: Allan Variance of a random walk signal corrupted with Gaussian white noise which has been sampled regularly. The Allan Variance is calculated using the expression (3) and the graph depicts mean values and standard deviations obtained from 100 Monte Carlo simulations. The Allan Variance trend indicates that in this particular case, bias is minimized at averaging timescales of approximately $0.03s$.

A. Allan Variance for Irregularly Sampled Data

Irregularly sampled data may come in various forms, such as intermittent data with large intervals between sets of rapid bursts of measurements, or sampling intervals that follow the properties of a known distribution. Several different strategies may be employed to evaluate Allan Variance for such irregularly sampled timeseries data. For example, naïvely speaking, we may choose to sample the data at a higher frequency to obtain regularly sampled data, but this is not always possible or feasible. Other researchers have proposed interpolation to overcome irregular sampling of data [17], but Sesia and Tavella have showed that this approach would produce considerable bias in the evaluation of Allan Variance [16]. Instead, they propose to evaluate Allan Variance by using available data and simply disregarding the absence of missing measurements or data [16].

In this paper, we expand upon this approach and propose an alternative method that evaluates the Allan Variance for irregularly-sampled data without any prior knowledge about the signal characteristics. Nonetheless, readers may use another AVAR estimator, such as the ones included in [18], [19], or [20], that better suits their particular application of interest. Specifically in this work, we first notice that the expression in (1) is no longer applicable for irregularly-sampled data as the number of measurements in any given averaging window are not identical across different windows. To address this limitation, we adopt the expression τ to represent the duration of the averaging window in continuous-time notation. Specifically, as in (1), the local average of the parameter θ may be evaluated in continuous time as the average of recent measurements that are no older than the timescale τ. Mathematically, we can express this as:

$$\bar{\theta}_t = \frac{1}{|C_t|} \cdot \sum_{\theta_{t_i} \in C_t} \theta_{t_i} \tag{4}$$

where $C_t = \{\theta_{t_i} : t - \tau < t_i < t\}$ is the set of recent measurements within a window of length τ, and $|C_t|$ denotes the cardinality of set C_t, i.e. the number of measurements contained in $C_t \subset \Theta$. Clearly, $\bar{\theta}_t$ is undefined for $|C_t| = 0$.

In irregularly-sampled data, it is likely that the number of measurements in different windows of equal length τ is not constant. Throughout this paper, in order to generate irregularly sampled measurements, we create a measurement time vector by randomly sampling timestamps from a uniform distribution. This results in uneven spacing between successive measurements. Further, the modification proposed in (4) enables us to evaluate the Allan Variance using adjacent windows of equal width τ, even if these windows contain a non-identical number of measurements. The original definition of Allan Variance in (2) may now be re-written in continuous time as follows:

$$\sigma_A^2(\tau) = \frac{1}{2}\mathbb{E}\left[(\bar{\theta}_t - \bar{\theta}_{t-\tau})^2\right] \tag{5}$$

To approximate the expected value in (5), we propose an approach similar to (3), but which uses *weighted* averaging. Specifically, each term of $(\bar{\theta}_t - \bar{\theta}_{t-\tau})^2$ may be associated with a joint weight value w_t to account for irregular or missing data. Subsequently, the Allan Variance may be evaluated as a weighted average, with continuous-time windows with more measurements given larger weights, and those with fewer measurements given smaller weights. To evaluate the Allan Variance for irregularly-sampled data we use this joint weight for successive windows as follows:

$$\hat{\sigma}_A^2(\tau) = \frac{1}{2\sum_S \mathrm{w}_t}\sum\nolimits_S \mathrm{w}_t(\bar{\theta}_t - \bar{\theta}_{t-\tau})^2 \tag{6}$$

where both summations are performed over a finite set of time instances $t \in S$, which governs the coarseness of estimation while evaluating the expectation operator in expression (5).

The joint weight of the two adjacent windows may be assumed to be the product of the number of measurements in each window, i.e.:

$$\mathrm{w}_t = |C_t||C_{t-\tau}| \tag{7}$$

though other weighting schemes may be chosen based on domain knowledge expertise or signal processing considerations.

Also, the coarseness of set S, i.e. the time increments used to slide the windows, may depend on how timestamps are distributed. It is suggested that practitioners begin with the mean value of the time interval between measurements, examining the evaluated Allan Variance, and using finer sliding window step size if needed. Alternatively, to make the simulation more time efficient, it is also possible to set the sliding step to the particular window length at which the AVAR is being estimated. A caveat of the proposed approach is that the joint weight w_t becomes zero if either $|C_t|$ or $|C_t - \tau|$ is zero. In such cases, we disregard any time windows during which *no* measurements were obtained.

Figure 2 shows the performance of our proposed method for calculating AVAR in a comparison with equally weighted (unweighted) AVAR evaluation. The performance of the classical formulation with regularly sampled data in the expression (3) is also provided as the reference. As is evident from the figure, the weighted AVAR outperforms the unweighted AVAR evaluation specially across the longer window lengths. Next, we discuss how to leverage the bias-minimizing characteristic timescales obtained from Allan Variance to inform moving average estimation.

III. Moving Average Estimation with Allan Variance-Informed Characteristic Timescale

Simple moving average estimation refers to the process of recovering a time-varying reference or determining a parameter $\theta_r(t)$ by taking the moving average of the given time series of historical noisy measurements Θ over a known temporal horizon τ. For regularly-sampled measurements, the simple moving average estimate of the parameter θ at any arbitrary time step $k \in \{1, 2, ..., n\}$ can be written as:

$$\hat{\theta}_k = \frac{1}{m}\sum_{i=k-m}^{k}\theta_i \tag{8}$$

Similarly, for irregularly-sampled data, the SMAE at any arbitrary time t can be calculated as the average of recent measurements that are no older than the time scale τ. Unsurprisingly, using an expression similar to (1), the SMAE for irregularly-sampled data may be evaluated as the local average at the time t, i.e. $\hat{\theta}_t = \bar{\theta}_t^\tau$ where the superscript τ denotes the window length for local averaging. Throughout this study, we considered the last available estimate $\hat{\theta}_{t^-}$ for undefined values of $\bar{\theta}_t$, i.e. when $|C_t| = 0$.

The choice of the characteristic timescale τ, however, plays an important role in the performance of the moving average estimator. Several aspects govern the choice of the parameter τ, among them being the signal dynamics, noise characteristics, and the distribution of timestamps. As mentioned earlier, the choice of window length or timescale has traditionally been dictated by domain expertise or evaluation using trial and error. The primary contribution of the presented work is to reduce dependence on such heuristics and provide a more systematic method to choose the characteristic timescales at which measurements remain relevant to the estimation task. The method of choice for selection of the characteristic timescale is through identifying the averaging time at which Allan Variance (evaluated based on (5)) is minimized. The timescale at which Allan Variance is minimized corresponds to the minimum bias in the moving average estimate (evaluated using (8)) given the signal dynamics and measurement noise characteristics. This process produces a near-optimal averaging time scale τ^* without having any prior knowledge about the parameter dynamics, i.e. $\hat{\theta}_t = \bar{\theta}_t^{\tau^*}$ where

$$\tau^* = \arg\min_\tau \sigma_A^2(\tau) \tag{9}$$

Next, we discuss results for near-optimal moving average estimation using AVAR-informed characteristic timescales.

IV. Results

In this section, we examine the performance of simple moving average estimation of a time-varying parameter θ, where the parameter may vary according to different deterministic or stochastic signals and is corrupted by white noise. The characteristic timescale at which the moving average estimation is performed is determined through the Allan Variance approach outlined in previous sections. Specifically, the performance of the AVAR-informed moving average estimation task is determined in three steps. First, the Allan Variance of the

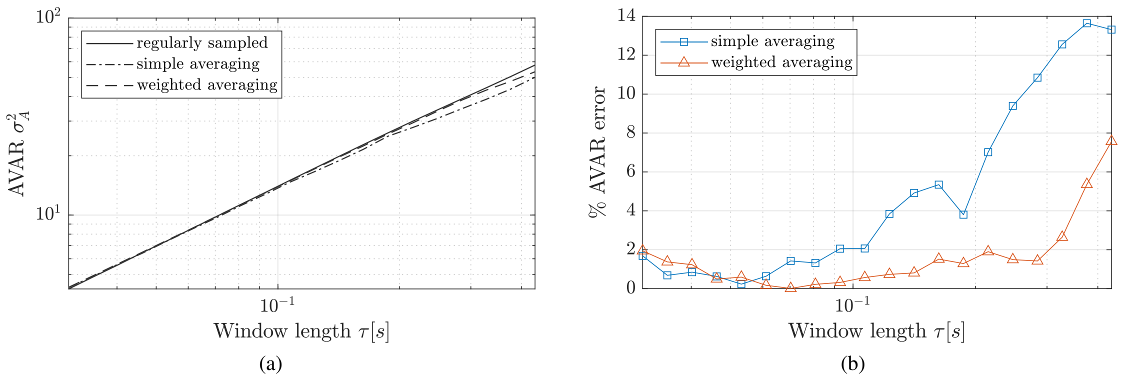

Fig. 2: (a) Allan Variance of random walk signal. Solid line represents the AVAR estimated by expression (3) using the original 5000 regularly sampled timestamps, dashed line represents AVAR estimated using expression (6), and dash-dotted line denotes AVAR obtained by setting all the weights to one (unweighted averaging). All lines represent the mean of 1000 Monte Carlo simulations and for both latter simulations, 1000 random time stamps are uniformly realized from $[0, 1]$. (b) The percentage error between irregularly sampled AVAR (simple and weighted averaging) and the reference regularly sampled AVAR.

provided timeseries data is evaluated using the approximation used in (6). Second, the timescale τ^* associated with the minimum value of the Allan Variance is determined. The moving average estimation is performed using this AVAR-informed timescale for several different signals, such as ramp, sine wave, square wave and random walk, all corrupted by Gaussian white noise. Finally, to validate our results we perform moving average estimation using several different window lengths spread over a large temporal range. The optimal timescale is then determined as the one which minimizes the mean square error between the estimate and the actual value of the signal over a given time period. Finally, the performances of the optimal MSE estimator and the Allan Variance-informed MA estimator are compared.

Figure 3(a) shows the AVAR-informed simple moving average estimation of a ramp signal with an additive noise drawn from $\mathcal{N}(0, 0.1)$. In addition, the figure also includes the true signal (thin black line), noisy measurements (gray dots), and the optimal simple moving average estimate (faint blue line) corresponding to the timescale that minimizes the MSE. Figure 3(b) shows the Allan Variance as a function of the averaging time or window length (top), and the mean-square error for moving average estimates as a function of window length (bottom). The red and blue square markers indicate the window lengths or timescales for which the Allan Variance and the moving average MSE are minimized, respectively. Figures 4, 5, and 6 show these results for random walk, sinusoid, square wave signals, respectively. In each scenario, the reference signal θ_r is corrupted with a specified amount of zero mean white Gaussian noise as indicated in the figures. A simple moving average estimator is then implemented to determine θ_r from a set of irregularly-sampled temporal measurements. In all the simulations, the SMAE returns the previous estimate if there is no measurement across the window.

Table I shows the results of performance comparisons between the optimal MSE estimator and the near-optimal AVAR-informed MA estimator by using the mean square error over the signal length as a measure of accuracy. As is evident, the AVAR-informed MA estimator performs quite well with relatively moderate errors as compared to MSE of the optimal MA estimate, across a variety of noisy deterministic and stochastic signals. Figures 3, 4, 5, and 6 show the corresponding analysis for each signal type. Moreover, the timescale at which AVAR is minimized is also compared against the optimal window length, indicating good agreement between optimal and AVAR-informed characteristic timescales.

The reader should note that AVAR evaluates this characteristic timescale in the absence of any prior knowledge about the signal or the noise model, whereas the MSE estimator possesses knowledge of the true signal. The ability to obtain a near-optimal characteristic timescale may be tremendously valuable in several applications where the original signal characteristics are unknown or change with time.

TABLE I: Near-optimal timescale evaluation using AVAR in a comparison with the optimal timescale calculated based on MSE of the MA estimation for unit ramp, random walk, and two different periodic signals. Asterisks in the last column indicate association with Figs. 3, 4, 5, and 6, respectively.

Signal	Domain [s]	ω	$\text{MSE}_{\text{optimal}}$	MSE_{AVAR}	Error [%]
Unit ramp	$[0, 1]$	-	1.05×10^{-3}	1.12×10^{-3}	6.97*
Random walk	$[0, 1]$	-	6.4238	6.5858	2.52*
$\sin(\omega t)$	$[0, 10]$	1	5.91×10^{-3}	6.48×10^{-3}	9.60
		2	9.57×10^{-3}	11.22×10^{-3}	17.19*
		4	16.19×10^{-3}	21.54×10^{-3}	33.03
square(ωt)	$[0, 10^3]$	0.01	70.64×10^{-3}	70.66×10^{-3}	0.03
		0.02	94.76×10^{-3}	96.82×10^{-3}	2.17*
		0.04	151.27×10^{-3}	162.46×10^{-3}	7.39

V. Concluding Remarks and Future Works

The presented work has demonstrated an Allan Variance-based technique for systematically identifying a characteristic timescale at which to perform simple moving average estimation. The primary insight of this work is in associating

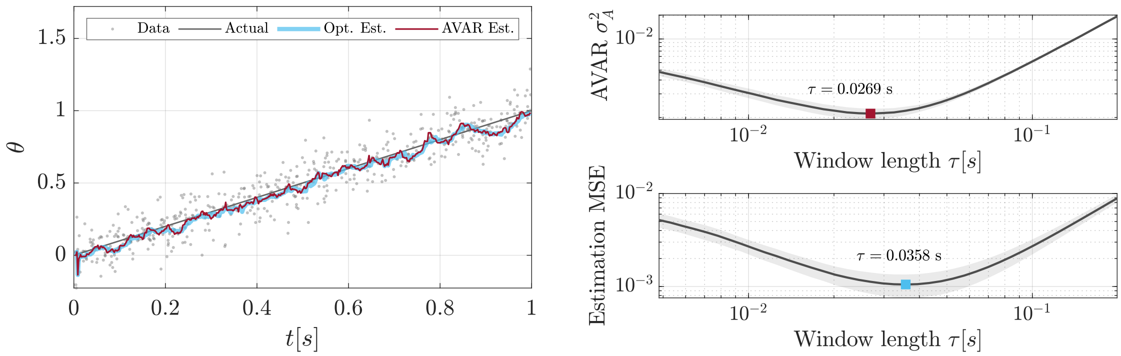

(a) (b)

Fig. 3: (a) A unit ramp reference signal $\theta_r = t$ (black narrow line) and corresponding 500 noisy measurements (dots). The irregular measurement timestamps are randomly generated from a uniform distribution and the reference signal is corrupted by Gaussian white noise with $\sigma = 0.1$ unit. (b) Allan Variance of the measurement time series and the mean square error of the moving average estimation calculated at various window lengths (time scales). Both graphs show mean values and the corresponding standard deviation (shaded area) obtained from 500 Monte Carlo simulations.

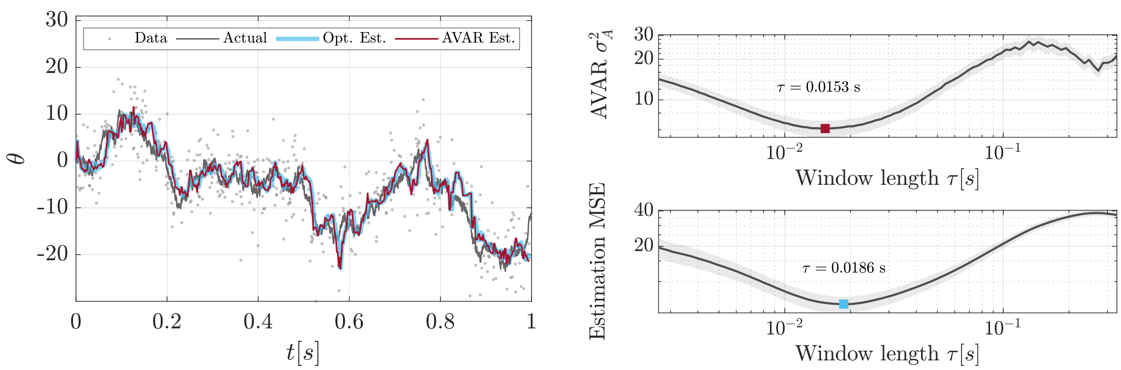

(a) (b)

Fig. 4: (a) A random walk reference signal (black line) and corresponding 500 noisy measurements (gray dots). The random walk signal is generated using 5000 samples and the measurement timestamps are randomly generated from a uniform distribution. The signal is corrupted by zero mean Gaussian white noise with $\sigma = 5$ unit. (b) Allan Variance of the time series and the mean square error of the moving average estimation calculated at various window lengths (time scales). Both graphs show mean values and the corresponding standard deviation (shaded area) obtained from 500 Monte Carlo simulations.

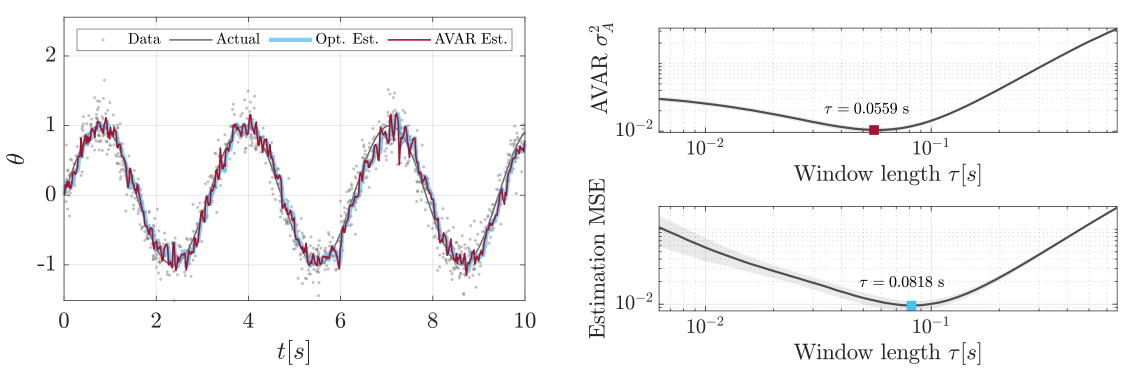

(a) (b)

Fig. 5: (a) A sinusoidal reference signal $\theta_r = \sin(2t)$ (black line) and corresponding 1000 noisy measurements (gray dots). The irregular measurement timestamps are randomly generated from a uniform distribution and the reference signal is corrupted by Gaussian white noise with $\sigma = 0.2$ unit. (b) Allan Variance of the measurement time series and the mean square error of the moving average estimation calculated at various window lengths (time scales). Both graphs show mean values and the corresponding standard deviation (shaded area) obtained from 500 Monte Carlo simulations.

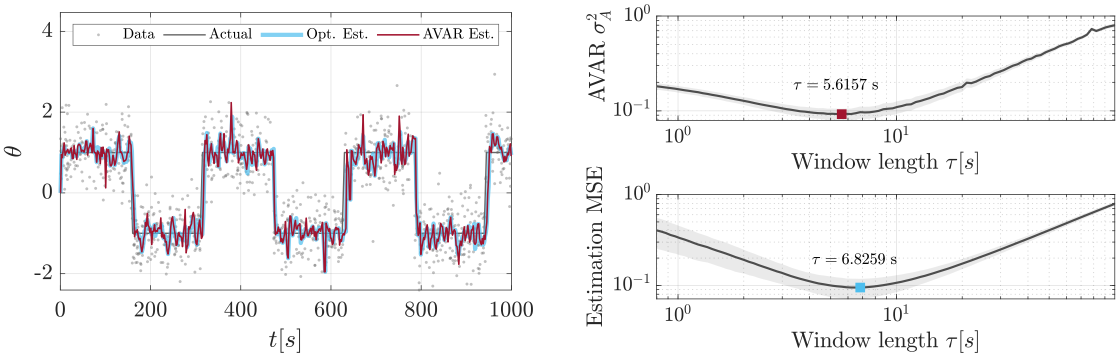

(a) (b)

Fig. 6: (a) A square reference signal $\theta_r = \text{square}(0.02t)$ (black line) and corresponding 1000 noisy measurements (gray dots). The irregular measurement timestamps are randomly generated from a uniform distribution and signal is corrupted by Gaussian white noise with $\sigma = 0.5$ unit. (b) Allan Variance of the measurement time series and the mean square error of the moving average estimation calculated at various window lengths (time scales). Both graphs show mean values and the corresponding standard deviation (shaded area) obtained from 500 Monte Carlo simulations.

the AVAR-informed characteristic timescales to the notion of duration over which data remains relevant. As evinced by the Allan Variance plots, there is an optimal timescale at which bias in the signal is minimized. Performing moving average estimation at shorter time scales produces higher bias estimates due to signal and/or noise characteristics. On the other hand, performing MA estimation at longer timescales also produces higher bias estimates on account of the data losing its relevancy towards the estimation task. The AVAR-informed MA estimation thus balances these two trade-offs of accuracy and relevancy in a near-optimal manner, without relying on iterations or heuristics. Perhaps, equally as importantly, the AVAR-informed moving average estimation is able to perform the estimation task in a near-optimal manner *without* a priori knowledge of the signal or noise model. This makes the presented approach widely applicable to a large number of scenarios where high fidelity models are cumbersome or difficult to obtain, such as determining road-tire friction in safe autonomy applications. In future works, we will address the theoretical aspects of the presented method and expand it to include real-time adaptation to discontinuous signals through the use of Dynamic Allan Variance (DAVAR) which will help improve the moving average estimate. Moreover, the notion of characteristic timescales (or data scales) allows for systematically dealing with under- and over-fitting trade-offs, thus finding applicability in locally weighted estimation methods that utilize more sophisticated kernel functions [7].

References

[1] K. Mohammadi, H. Eslami, and R. Kahawita, "Parameter estimation of an ARMA model for river flow forecasting using goal programming," *Journal of Hydrology*, vol. 331, no. 1-2, pp. 293–299, 2006.

[2] B. Soleimani, P. Das, J. Kulasingham, J. Z. Simon, and B. Babadi, "Granger Causal Inference from Indirect Low-Dimensional Measurements with Application to MEG Functional Connectivity Analysis," *Annual Conf. on Information Sci. and Sys., CISS 2020*, pp. 0–4, 2020.

[3] "Is smarter better? A comparison of adaptive, and simple moving average trading strategies," *Research in International Business and Finance*, vol. 19, no. 3, pp. 399–411, 2005.

[4] N. Soans, E. Asali, Y. Hong, and P. Doshi, "SA-net: Robust state-action recognition for learning from observations," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2020, pp. 2153–2159.

[5] F. Najafi, M. Dadvar, S. Habibian, A. Hosseini, H. Haeri, M. Arvan, B. Peykari, and H. Bagheri, "Robocup rescue 2016 team description paper," in *Robocup Rescue 2016 TDP Collection*, 2016.

[6] C. E. Beal, "Rapid road friction estimation using independent left/right steering torque measurements," *Vehicle System Dynamics*, 2019.

[7] C. G. Atkeson, A. W. Moore, and S. Schaal, "Locally Weighted Learning," *Artificial Intelligence Review*, vol. 11, no. 1-5, pp. 11–73, 1997.

[8] L. Shpigelman, H. Lalazar, and E. Vaadia, "Kernel-ARMA for hand tracking and Brain-Machine Interfacing during 3D motor control," *Advances in Neural Information Processing Systems 21 - Proceedings of the 2008 Conference*, pp. 1489–1496, 2009.

[9] B. Aksasse and L. Radouane, "Two-dimensional autoregressive (2-D AR) model order estimation," *IEEE Transactions on Signal Processing*, vol. 47, no. 7, pp. 2072–2077, 1999.

[10] M. S. Voss and X. Feng, "ARMA model selection using particle swarm optimization and aic criteria," *IFAC Proc. Vol.*, vol. 35, no. 1, pp. 349–354, 2002.

[11] S. S. Pappas, V. C. Moussas, and S. K. Katsikas, "Multivariate ARMA order estimation via multi-model partition theory," *a a*, vol. 11, p. 1m.

[12] D. W. Allan, "Statistics of atomic frequency standards," *Proceedings of the IEEE*, vol. 54, no. 2, pp. 221–230, 1966.

[13] K. Jerath, S. Brennan, and C. Lagoa, "Bridging the gap between sensor noise modeling and sensor characterization," *Measurement*, vol. 116, pp. 350 – 366, 2018.

[14] K. Jerath and S. N. Brennan, "GPS-Free Terrain-Based Vehicle Tracking Performance as a Function of Inertial Sensor Characteristics," in *Proc. of ASME 2011 Dynamic Sys. and Control Conf.*, 2011, pp. 367–374.

[15] M. Kok, J. D. Hol, and T. B. Schön, "Using inertial sensors for position and orientation estimation," *CoRR*, vol. abs/1704.06053, 2017.

[16] I. Sesia and P. Tavella, "Estimating the allan variance in the presence of long periods of missing data and outliers," *Metrologia*, vol. 45, no. 6, p. S134, 2008.

[17] C. Hackman and T. E. Parker, "Noise analysis of unevenly spaced time series data," *Metrologia*, vol. 33, no. 5, pp. 457–466, Oct 1996.

[18] J. Li and J. Fang, "Sliding average allan variance for inertial sensor stochastic error analysis," *IEEE Transactions on Instrumentation and Measurement*, vol. 62, no. 12, pp. 3291–3300, 2013.

[19] D. Howe, "An extension of the allan variance with increased confidence at long term," in *Proc. of the 1995 IEEE International Frequency Control Symposium (49th Annual Symposium)*. IEEE, 1995, pp. 321–329.

[20] C. A. Greenhall, "Estimating the modified Allan variance," in *Proc. of the International Frequency Control Symp.* IEEE, 1995, pp. 346–353.