Assessing Model Proficiency in Autonomous Agents: A Signal Processing Perspective

Published Version:
Guerra, A., Guidi, F., Dardari, D., Djuric, P.M. (2025). Assessing Model Proficiency in Autonomous Agents: A Signal Processing Perspective. 345 E 47TH ST, NEW YORK, NY 10017 USA : Institute of Electrical and Electronics Engineers Inc. [10.1109/ICASSPW65056.2025.11011249].

Availability:
This version is available at: https://hdl.handle.net/11585/1041311 since: 2026-02-02

Published:
DOI: http://doi.org/10.1109/ICASSPW65056.2025.11011249

Assessing Model Proficiency in Autonomous Agents: A Signal Processing Perspective

A. Guerra[1], F. Guidi[2], D. Dardari[1], and P. M. Djurić[3]
[1]*University of Bologna, DEI, WiLab, Cesena, Italy*
[2]*National Research Council of Italy (CNR), IEIIT, WiLab, Bologna, Italy*
[3]*Stony Brook University, Stony Brook, USA*

***Abstract*—Autonomous agents play a crucial role in applications such as emergency response and urban security, where they are often required to operate independently in critical situations without direct human supervision. A key aspect of autonomy is the agent's ability to assess its proficiency in carrying out tasks and making decisions based on this assessment. This paper introduces a state-space model to present a novel framework for assessing the proficiency of autonomous agents. The proposed metric is based on statistical model assessment, enabling agents to evaluate their model performance in real-time. More specifically, we focus on the proficiency of the model used for measurement predictions. We validate the effectiveness of our metric through simulations with synthetic data. Future work will explore the potential of this framework to enhance decision-making accuracy and improve task performance.**

***Index Terms*—Model proficiency assessment, state-space models, autonomous systems**

I. Introduction

Autonomous agents, such as drones and ground-based robots, are becoming increasingly important today. They are used in various operations, including emergency response, search-and-rescue missions, and urban security [1], [2]. As these agents advance, their capability to operate without direct human supervision is crucial. Autonomy requires them to make independent decisions and adapt to the environment without relying on centralized supervision or leadership [3].

An essential aspect of making autonomous decisions is trustworthiness and effective communication with other team agents [4]. When multiple agents are deployed for the same mission, they should interact to share important information for coordinating and carrying out tasks [5]–[7]. However, as the complexity of their tasks increases, the agents need to evaluate their proficiency in performing a given task and adjust their behavior based on their self- and team-assessment.

From a signal processing perspective, we assume that each agent has a library of statistical models designed to estimate unknown features of its surrounding environment. These models allow agents to make more informed decisions, such as determining optimal navigation paths or selecting other agents for collaboration. The interactions between learning, proficiency assessment, and decision-making by an agent are depicted in Fig. 1. However, statistical models of real-world phenomena often represent only approximations. Thus, it becomes vital for agents to assess the performance of their models, for example, by assigning scores based on their accuracy in predicting new observations and states.

In autonomous systems, model assessment has gained significant attention in recent years [3], [5], [8], [9]. Traditional methods rely on supervised learning techniques, where model performance is evaluated through cross-validation or prediction error metrics [10]. In reinforcement learning-based agents, model assessment often focuses on reward maximization or policy performance under predefined conditions [11]. However, these approaches are usually task-specific and depend heavily on external feedback mechanisms, which may not be suitable for fully autonomous agents operating in dynamic and unpredictable environments.

Recent advances in self-assessment techniques for autonomous agents have explored more adaptive methods. Bayesian frameworks and ensemble learning approaches allow agents to estimate uncertainty in their models, enabling better decision-making in real-time [12], [13]. However, these methods still face limitations in scalability and adaptability, especially when applied to distributed multi-agent systems where communication is crucial.

In this domain, the signal processing community traditionally uses classical approaches to model assessment based on standard test statistics [8], [14]–[16]. For example, in [8], the authors propose a method for assessing dynamic nonlinear models based on empirical and predictive cumulative distributions of data and the Kolmogorov-Smirnov (KS) statistics. Then, in [5], the authors evaluate model proficiency using a Cramer-Rao Lower Bound (CRLB) measure based on the predictive observation distribution.

This paper presents a real-time proficiency assessment mechanism that enables agents to dynamically evaluate their model's performance and adapt their behavior accordingly. The assessment provides a score that could allow agents to make informed decisions about whether to collaborate or act independently based on their model's limitations and the shared knowledge of others. As a further step, agents

This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on "Telecommunications of the Future" (PE00000001 - program "RESTART"), under the NRPP of NextGenerationEU (Mision 4 – Component 2 -Investment 1.1) Prin 2022 (No. 104, 2/2/2022, CUP J53C24002790006), under ERC Grant no. 101116257 (project CUE-GO: Contextual Radio Cues for Enhancing Decision Making in Networks of Autonomous Agents), and the National Science Foundation under Award 2212506. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or ERCEA. Neither the European Union nor the granting authority can be held responsible for them.

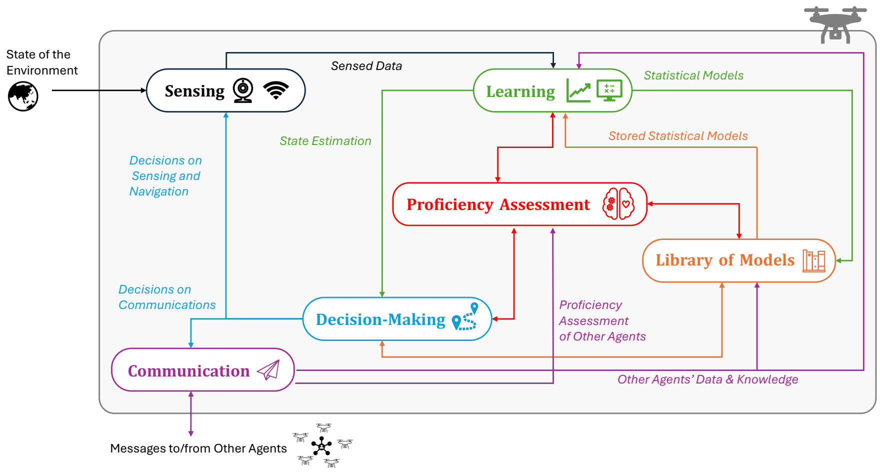

Fig. 1. Interaction between autonomous agents, proficiency assessment, and communication.

can better coordinate their actions and optimize collective outcomes by incorporating this proficiency-based evaluation into a communication framework.

In this context, we provide a signal-processing perspective on model proficiency in autonomous agents by introducing a new metric to assess the fitness of their statistical models. Our approach is validated through simulations on synthetic data. These experiments provide significant insights and allow for meaningful conclusions.

II. Proficiency Assessment

We recall the definition of proficiency self-assessment (PSA), which is applicable to a single model $\mathcal{M}$ and a single agent collecting a scalar observation [5]. We denote with y_t a scalar observation and with $\hat{y}_t$ its prediction at time instant t, as an instantiation of an estimator $\hat{\mathsf{y}}_t$ given by $\hat{y}_t = g_{\mathcal{M}}(y_{1:t-1})$, where $g_{\mathcal{M}}(\cdot)$ is the predictive function according to the model $\mathcal{M}$ and taking as an input the history of measurements $y_{1:t-1}$ up to time instant $t-1$. According to [5], the proficiency of a model $\mathcal{M}$ at time t can be defined as

$$\mathcal{P}_t(\mathcal{M}) = -\mathbb{E}\left[\frac{\partial^2 \ln f_{\mathsf{y}_t,\mathsf{y}_{1:t-1}}(y_t, y_{1:t-1}; \mathcal{M})}{\partial y_t^2}\right], \quad (1)$$

where the expectation $\mathbb{E}[\cdot]$ is computed according to the joint distribution of all the unseen observations that are predicted, $\mathsf{y}_t, \mathsf{y}_{1:t-1}$, and the proficiency is of estimating y_t.

Model proficiency can also be defined for a *specific sequence* of observations. We refer to this measure as the *in situ* proficiency, and it is written as [5, Eq. 6]

$$\mathcal{P}_t(\mathcal{M}|y_{1:t-1}) = -\mathbb{E}\left[\frac{\partial^2 \ln f_{\mathsf{y}_t|\mathsf{y}_{1:t-1}}(y_t|y_{1:t-1}; \mathcal{M})}{\partial\, y_t^2}\right], \quad (2)$$

where now the proficiency is also of estimating y_t but conditioned on seeing $y_{1:t-1}$. We also define the cumulative proficiency of the agent for model $\mathcal{M}$ as

$$\mathcal{P}_{1:t}(\mathcal{M}|y_{1:t-1}) = \sum_{k=1}^{t} \mathcal{P}_k(\mathcal{M}|y_{1:k-1}). \quad (3)$$

By assuming an unbiased estimator, it holds [5, Eq. 7] [17]

$$\begin{aligned}\mathcal{P}_t(\mathcal{M}|y_{1:t-1}) &\geq \left(\mathbb{E}\left[(\hat{\mathsf{y}}_t - \mathsf{y}_t)^2\right]\right)^{-1} \\ &= \left\{\int (\hat{y}_t - y_t)^2 f_{\mathsf{y}_t|\mathsf{y}_{1:t-1}}(y_t|y_{1:t-1}; \mathcal{M})\, dy_t\right\}^{-1}, \end{aligned} \quad (4)$$

meaning that the PSA is an upper bound for the *inverse* of the mean square error (MSE) computed on the measurement prediction error. Extending the concept of the *in situ* proficiency of (2) to the vector case, namely $\boldsymbol{\mathcal{P}}_t(\mathcal{M}|\boldsymbol{y}_{1:t-1})$, we have

$$\begin{aligned}\boldsymbol{\mathcal{P}}_t(\mathcal{M}|\boldsymbol{y}_{1:t-1}) = &\mathbb{E}\Big[\nabla^T_{\boldsymbol{y}_t} \ln f_{\mathbf{y}_t|\mathbf{y}_{1:t-1}}(\boldsymbol{y}_t|\boldsymbol{y}_{1:t-1}; \mathcal{M}) \\ &\times \nabla_{\boldsymbol{y}_t} \ln f_{\mathbf{y}_t|\mathbf{y}_{1:t-1}}(\boldsymbol{y}_t|\boldsymbol{y}_{1:t-1}; \mathcal{M})\Big], \end{aligned} \quad (5)$$

with $\nabla_{\boldsymbol{y}_t} = \frac{\partial}{\partial \boldsymbol{y}_t} = \left[\frac{\partial}{\partial y_{1,t}}, \frac{\partial}{\partial y_{2,t}}, \ldots, \frac{\partial}{\partial y_{N_D,t}}\right]$ being the gradient computed with respect to the measurement vector, and where N_{D} is the size of the observation vector $\boldsymbol{y}_t$.

Equation (5) represents the upper bound on the inverse covariance matrix of the mean squared prediction errors of the observations. By considering that the measurements are homogeneous, we can consider the following two alternatives to obtain a scalar measure of model proficiency

1. $\mathcal{P}_t(\mathcal{M}|\boldsymbol{y}_{1:t-1}) = \mathsf{tr}\left\{\boldsymbol{\mathcal{P}}_t(\mathcal{M}|\boldsymbol{y}_{1:t-1})\right\}, \quad (6)$
2. $\mathcal{P}_t(\mathcal{M}|\boldsymbol{y}_{1:t-1}) = \mathsf{det}\left\{\boldsymbol{\mathcal{P}}_t(\mathcal{M}|\boldsymbol{y}_{1:t-1})\right\}, \quad (7)$

where $\mathsf{tr}(\cdot)$ and $\mathsf{det}(\cdot)$ indicate the trace and the determinant of the matrix, respectively.

III. State-Space Model

We now provide a solution for performing model assessment according to the chosen metric, using a linear Gaussian scalar State-Space Model (SSM) scenario. Let us recall the proficiency in linear scalar Gaussian SSM presented in [5, Sec. 6]. Consider the following first-order auto-regressive state model

$$\mathcal{M}: \begin{cases} \mathsf{x}_t = a\,\mathsf{x}_{t-1} + \mathsf{u}_t, \\ \mathsf{y}_t = b\,\mathsf{x}_t + \mathsf{v}_t, \end{cases} \quad (8)$$

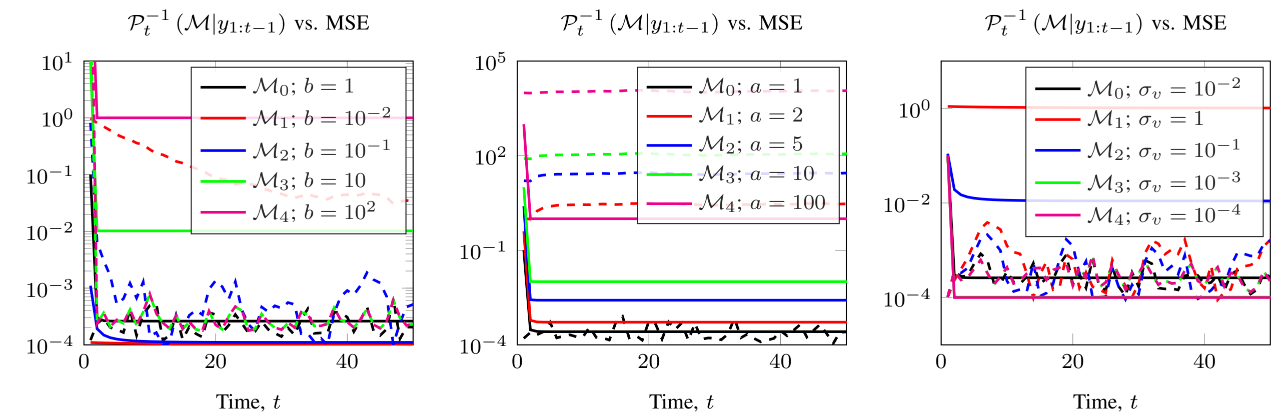

Fig. 2. Solid lines represent the inverse of the proficiency according to the selected models, whereas dashed lines represent the MSE.

where a and b are known scalars, and $\mathsf{u}_t \sim \mathcal{N}(\mathsf{u}_t; 0, \sigma_u^2)$ and $\mathsf{v}_t \sim \mathcal{N}(\mathsf{v}_t; 0, \sigma_v^2)$ are state and measurement noise processes, respectively, with known variances. Hence, we have

$$\mathcal{M}: \begin{cases} \mathsf{x}_t \sim f_{\mathsf{x}_t|x_t}(x_t|x_{t-1};\, \mathcal{M}_x) = \mathcal{N}(\mathsf{x}_t; a\, x_{t-1}, \sigma_u^2), \\ \mathsf{y}_t \sim f_{\mathsf{y}_t\,|\,x_t}(y_t\,|\,x_t;\, \mathcal{M}_y) = \mathcal{N}(\mathsf{y}_t;\, b\, x_t, \sigma_v^2), \end{cases} \tag{9}$$

where $\mathcal{M} = \{\mathcal{M}_x, \mathcal{M}_y\}$, and $\mathcal{M}_x$, $\mathcal{M}_y$ are the models describing the state and observation dynamics. The model in (9) is characterized by the parameters $\boldsymbol{\theta} = \{a,\, b,\, \sigma_u^2,\, \sigma_v^2\}$. The corresponding predictive distribution is given by

$$\begin{aligned} & f_{\mathsf{y}_t|y_{1:t-1}}(y_t|y_{1:t-1}; \mathcal{M}) \\ & = \int f_{\mathsf{y}_t|x_t}(y_t|x_t; \mathcal{M}_y)\, f_{\mathsf{x}_t|y_{1:t-1}}(x_t|y_{1:t-1};\, \mathcal{M})\, dx_t, \end{aligned} \tag{10}$$

where $f_{\mathsf{x}_t|y_{1:t-1}}(x_t|y_{1:t-1};\, \mathcal{M}) = \mathcal{N}(\mathsf{x}_t;\, m_{t|t-1}, \sigma^2_{t|t-1})$ is the predictive distribution of the state with mean and variance given by

$$\begin{aligned} m_{t|t-1} &= \mathbb{E}\left[\mathsf{x}_t \,|\, y_{1:t-1}\right] = a\, m_{t-1|t-1}, \\ \sigma^2_{t|t-1} &= \mathbb{E}\left[(\mathsf{x}_t - m_{t|t-1})^2 \,|\, y_{1:t-1}\right] = a^2 \sigma^2_{t-1|t-1} + \sigma_u^2, \end{aligned} \tag{11}$$

and where the parameters of the posterior distribution of x_t are

$$\begin{aligned} m_{t|t} &= \mathbb{E}\left[\mathsf{x}_t|y_{1:t}\right] = m_{t|t-1} + \kappa_t\, (y_t - b\, m_{t|t-1}), \\ \sigma^2_{t|t} &= \mathbb{E}\left[(\mathsf{x}_t - m_{t|t})^2|y_{1:t}\right] = \sigma^2_{t|t-1} - \kappa_t^2 \left(b^2\, \sigma^2_{t|t-1} + \sigma_v^2\right), \end{aligned} \tag{12}$$

with κ_t being the Kalman gain. The integral in (10) results in [18, Eq. 4.19]

$$f_{\mathsf{y}_t|y_{1:t-1}}(y_t|y_{1:t-1}; \mathcal{M}) = \mathcal{N}\left(\mathsf{y}_t;\, b\, m_{t|t-1}, b^2\, \sigma^2_{t|t-1} + \sigma_v^2\right). \tag{13}$$

By taking the logarithm of (13) we have

$$\ln f_{\mathsf{y}_t|y_{1:t-1}}(y_t|y_{1:t-1}; \mathcal{M}) \propto -\frac{(y_t - b\, m_{t|t-1})^2}{2\, s^2_{t|t-1}}, \tag{14}$$

where the term $(y_t - b\, m_{t|t-1})$ is referred to as the innovation or measurement residual and where we set $s^2_{t|t-1} = \mathbb{E}\left[(\mathsf{y}_t - b\, m_{t|t-1})^2|y_{1:t-1}\right] = b^2\, \sigma^2_{t|t-1} + \sigma_v^2$. The square of the first derivative of (14) results in the following random variable (RV):

$$\mathsf{z}_t = \left(\frac{\partial \ln f_{\mathsf{y}_t|y_{1:t-1}}(y_t|y_{1:t-1}; \mathcal{M})}{\partial\, y_t}\right)^2 = \frac{\bar{\mathsf{y}}_t^2}{s^2_{t|t-1}}, \tag{15}$$

where $\bar{\mathsf{y}}_t = \frac{\mathsf{y}_t - b\, m_{t|t-1}}{\sqrt{s^2_{t|t-1}}} \sim \mathcal{N}(\mathsf{y}_t; 0, 1)$, $\bar{\mathsf{y}}_t^2 \sim \chi_1^2$ is a central chi-squared RV with one degree of freedom, and $\mathbb{E}[\bar{\mathsf{y}}_t^2] = 1$. According to (2), the model proficiency is given by

$$\mathcal{P}_t\left(\mathcal{M} \,|\, y_{1:t-1}\right) = \mathbb{E}\left[\mathsf{z}_t\right] = \frac{1}{s^2_{t|t-1}} = \frac{1}{b^2\, \sigma^2_{t|t-1} + \sigma_v^2}. \tag{16}$$

The PSA alone cannot help in deciding whether a model is correct because, in the presence of model mismatch, the PSA does not represent the bound for the chosen estimator or filtering approach. To determine if $\mathcal{P}_t\left(\mathcal{M}|y_{1:t-1}\right)$ is a correct bound for the inverse of the squared error, we propose the following metric (normalized PSA distance):

$$\bar{d}_{\mathsf{PSA},t}\left(\mathcal{M}\right) = \frac{\left|\mathcal{P}_t^{-1}\left(\mathcal{M}|y_{1:t-1}\right) - e_t^2\left(\mathcal{M}\right)\right|}{\left|\mathcal{P}_t^{-1}\left(\mathcal{M}|y_{1:t-1}\right)\right|}, \tag{17}$$

where $e_t^2\left(\mathcal{M}\right) = (y_t - \hat{y}_t\left(\mathcal{M}\right))^2$ is the squared error for measurement prediction. This metric measures the agreement between the inverse of model proficiency and the squared error in measurement prediction as per (4). When more samples can be collected for each time instant, the squared error is substituted with the MSE. A lower distance indicates that the model is more effective at predicting incoming data. In the extreme case, when $\bar{d}_{\mathsf{PSA},t}\left(\mathcal{M}\right)$ is close to zero, it means that the MSE attains the bound. Due to the normalization, (17) provides a relative measure, making it easier to compare models with different scales of model proficiencies.

IV. Case Studies

We consider two case studies. In the first, the state is described by the SSM in (9) and estimated by a Kalman filter

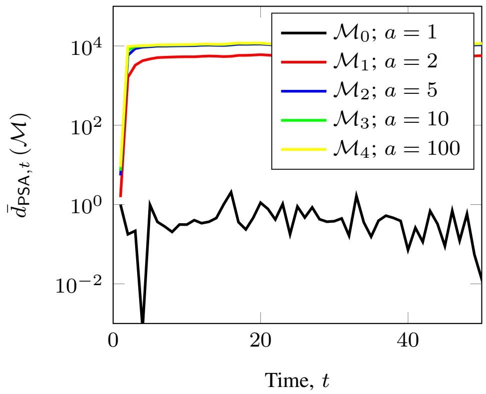

Fig. 3. Normalized PSA distance (17) as a function of time.

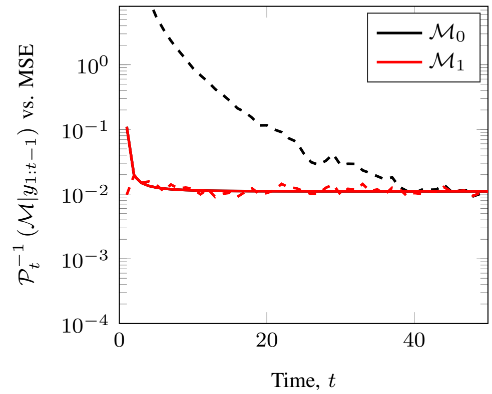

Fig. 4. Inverse of the proficiency (solid lines) and the MSE (dashed lines) of the SSM in (18).

(KF). To generate the actual state trajectory and observations, we consider the following true model parameters $\mathcal{M}_0 : \boldsymbol{\theta}_0 = \{a,\, b,\, \sigma_u^2,\, \sigma_v^2\} = \{0.999,\, 1,\, 10^{-4},\, 10^{-4}\}$ [8]. We observe the evolution of the state over $N_\mathrm{T} = 50$ time instants and consider $N_\mathrm{MC} = 100$ Monte Carlo runs, with results averaged over the trials. Then, we account for $N_\mathrm{M} = 5$ different models for state estimation and measurement prediction, which are $\mathcal{M}_i : \boldsymbol{\theta}_i = \{a_i,\, b_i,\, \sigma_{u,i}^2,\, \sigma_{v,i}^2\}$ with $i = 0, 1, \ldots, N_\mathrm{M}$. The state was initialized as $\mathsf{x}_0 \sim \mathcal{N}\left(\mathsf{x}_0; m_{0|0} = 1,\, \sigma_{0|0}^2 = 0.1\right)$.

In the first case study, shown in Fig. 2-left, we only modified the parameters b as indicated in the legend, and $\{a_i, \sigma_{u,i}^2, \sigma_{v,i}^2\} = \{0.999, 10^{-4}, 10^{-4}\}$, $\forall i$. The inverse of the proficiency is plotted against the considered model and compared with the corresponding mean squared error (MSE). By analyzing (16), it is evident that the model with the lowest b corresponds to the highest model proficiency, and it is denoted as $\mathcal{M}_1$. However, the MSE for this model is not distributed around the inverse of the model proficiency but is two orders of magnitude higher due to model mismatch. This emphasizes that relying solely on model proficiency for model validation results in significant inaccuracies in proficiency assessment. One can notice that using models $\mathcal{M}_{2:4}$ results in estimation errors for measurement predictions comparable to those obtained with the actual model $\mathcal{M}_0$, with the distinction of having a deviation from the bound. At the same time, the MSE on the state estimation errors for the same models (i.e., $\mathcal{M}_{2:4}$) is high. This reinforces the idea that choosing a model exclusively based on the MSE on measurement prediction or solely on proficiency could lead to inaccuracies.

In Fig. 2-middle and Fig. 2-right, we compare the same metrics by analysing model mismatches in the transition parameter a and in the measurement noise variance σ_v^2, respectively. In Fig. 2-middle, it is clear that using an incorrect value for the transition constant a results in significant estimation errors in all cases except the actual model. Unlike the measurement constant b, which amplifies the true state, a controls the dependence of the previous state on the current one. Therefore, while the actual model with $a \approx 1$ indicates that the state is almost constant over time except for the transition noise, other models with $a > 1$ violate this rule and also affect the measurement prediction capability. The same considerations hold when analyzing Fig. 2-right where we changed the measurement noise variance as indicated in the legend. Since all the other transition and measurement parameters are correct, the MSE for all models is low, at around 10^{-4}. However, as observed in the other cases plotted in Fig. 2, it is possible to determine the correct model by measuring the distance of the inverse of the proficiency from the MSE. To this end, to select the correct model, it is possible to use the normalized distance defined in (17). Figure 3 shows this distance as a function of the transition parameter a, showing that the correct model has the lowest value of $\bar{d}_{\mathsf{PSA},t}(\mathcal{M})$.

In the second case study, we consider two models that are given by

$$\mathcal{M}_i : \begin{cases} \mathsf{x}_t = a\,\mathsf{x}_{t-1} + \mathsf{u}_t, \\ \mathsf{y}_t = b\,\mathsf{x}_t + c_i + \mathsf{v}_t, \end{cases} \tag{18}$$

with $i = 0,\, 1$, and where c_i is a measurement bias. In our simulations, we set c_0 to 0 and c_1 to 100. All the other model parameters were set as before. The model $\mathcal{M}_1$ was used to generate the actual measurements. In this case, both models have the same proficiency, but their MSE differs. This is evident in Fig. 4, where the dashed black curve representing the MSE for the biased model $\mathcal{M}_0$ is significantly different from the inverse of its proficiency, which is overlapped with the red continuous curve of model $\mathcal{M}_1$. However, using the metric proposed in (17), we can detect the correct model and, hence, help the agent in decision making.

V. Conclusions

This paper presented a proficiency assessment mechanism for autonomous agents. This metric can be used for improving their decision-making and task performance. The proposed framework allows agents to evaluate the accuracy of their statistical models in real-time, enabling them to adapt their behavior based on model performance. By incorporating this proficiency-based metric, agents can better assess their capabilities, which leads to more informed decisions in complex and dynamic environments.

References

[1] A. Guerra, F. Guidi, D. Dardari, and P. M. Djurić, "Networks of UAVs of low complexity for time-critical localization," *IEEE Aerosp. Elect. Sys. Mag.*, vol. 37, no. 10, pp. 22–38, 2022.

[2] G. Fontanesi, A. Guerra, F. Guidi, J. A. Vásquez-Peralvo, N. Shlezinger, A. Zanella, E. Lagunas, S. Chatzinotas, D. Dardari, and P. M. Djurić, "A deep-NN beamforming approach for dual function radar-communication THz UAV," *IEEE Trans. Veh. Technol.*, To Appear, 2024.

[3] X. Cao, A. Gautam, T. Whiting, S. Smith, M. A. Goodrich, and J. W. Crandall, "Robot proficiency self-assessment using assumption-alignment tracking," *IEEE Trans. Robot.*, 2023.

[4] N. Conlon, D. Szafir, and N. Ahmed, "Investigating the effects of robot proficiency self-assessment on trust and performance," *arXiv preprint arXiv:2203.10407*, 2022.

[5] P. M. Djurić' and P. Closas, "On self-assessment of proficiency of autonomous systems," in *Proc. Int. Conf. Acoustics, Speech, Signal Process.* IEEE, 2019, pp. 5072–5076.

[6] A. Guerra, D. Dardari, and P. M. Djurić, "Dynamic radar network of UAVs: A joint navigation and tracking approach," *IEEE Access*, vol. 8, pp. 116 454–116 469, 2020.

[7] A. Guerra, F. Guidi, D. Dardari, and P. M. Djurić, "Reinforcement learning for joint detection and mapping using dynamic UAV networks," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 60, no. 3, pp. 2586–2601, 2023.

[8] P. M. Djurić and J. Míguez, "Assessment of nonlinear dynamic models by Kolmogorov–Smirnov statistics," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 5069–5079, 2010.

[9] N. Conlon, N. R. Ahmed, and D. Szafir, "A survey of algorithmic methods for competency self-assessments in human-autonomy teaming," *ACM Comput. Surveys*, vol. 56, no. 7, pp. 1–31, 2024.

[10] A. Acharya, R. Russell, and N. R. Ahmed, "Competency assessment for autonomous agents using deep generative models," in *Proc. Int. Conf. Intell. Robots Syst. (IROS)*. IEEE, 2022, pp. 8211–8218.

[11] N. Conlon, A. Acharya, J. McGinley, T. Slack, C. A. Hirst, M. D'Alonzo, M. R. Hebert, C. Reale, E. W. Frew, R. Russell *et al.*, "Generalizing competency self-assessment for autonomous vehicles using deep reinforcement learning," in *AIAA SCITECH 2022 Forum*, 2022, p. 2496.

[12] M. Zambelli and Y. Demirisy, "Online multimodal ensemble learning using self-learned sensorimotor representations," *IEEE Trans. Cogn. Develop. Syst.*, vol. 9, no. 2, pp. 113–126, 2017.

[13] A. S. Alemaw, P. Zontone, L. Marcenaro, P. Marin, D. M. Gomez, and C. Regazzoni, "Integrated learning and decision making for autonomous agents through energy based Bayesian models," in *Proc. Int. Conf. Information Fusion (FUSION)*, 2024, pp. 1–8.

[14] J. M. Bernardo and A. F. Smith, *Bayesian Theory*. John Wiley & Sons, 2009, vol. 405.

[15] P. Closas, M. F. Bugallo, and P. M. Djurić, "Assessing robustness of particle filtering by the Kolmogorov-Smirnov statistics," in *Proc. Int. Conf. Acoustics, Speech, Signal Process.* IEEE, 2009, pp. 2917–2920.

[16] P. M. Djurić and J. Miguez, "Model assessment with Kolmogorov-Smirnov statistics," in *Proc. Int. Conf. Acoustics, Speech, Signal Process.* IEEE, 2009, pp. 2973–2976.

[17] H. L. Van Trees and K. L. Bell, "Bayesian bounds for parameter estimation and nonlinear filtering/tracking," *AMC*, vol. 10, no. 12, pp. 10–1109, 2007.

[18] S. Särkkä and L. Svensson, *Bayesian Filtering and Smoothing*. Cambridge university press, 2023, vol. 17.