

Predictive Goal-Oriented Communication Policy for Remote Tracking of Markovian Sources

Isaac Andrés Balbuca Guachizaca and Paulo Cardieri

Abstract—We investigate the remote tracking of a Markovian source in an energy-harvesting wireless system within the goal-oriented communication framework. We propose a Predictive transmission policy that exploits the source dynamics to suppress redundant updates while preserving energy for correcting persistent errors. The proposed policy employs prediction at the transmitter to anticipate the receiver’s self-synchronization, thereby avoiding unnecessary transmissions. Simulations across different source dynamics and channel conditions show that, under severe energy constraints, the Predictive policy outperforms a reactive baseline that transmits whenever source and receiver states diverge, achieving low reconstruction error with reduced energy consumption.

Keywords—Goal-oriented communication, energy harvesting, Prediction, remote tracking, Markovian source.

I. INTRODUCTION

Over time, telecommunications have advanced alongside how humans interact with information. Initially, networks were designed for voice transmission. Later, they evolved to transport multimedia content and to connect a wide range of devices, creating a hyperconnected society. The development of wireless communications, driven by sixth-generation (6G) networks and the proliferation of Internet of Things (IoT) systems and autonomous devices, demands a paradigm shift, as networks are no longer mere data transport channels but platforms for interconnected intelligence [1].

The traditional communication paradigm faces limitations when confronted with emerging applications that impose critical requirements for latency and reliability, in addition to severe energy constraints. Such demands push conventional systems to operate close to the physical capacity limits defined by Shannon [2].

Given this scenario, it is necessary to revisit the foundations established by Shannon and Weaver, whose Mathematical Theory of Communication defined three fundamental problems related to the communication process: (i) the *technical problem*, which deals with the accuracy of symbol transmission; (ii) the *semantic problem*, focused on conveying the meaning of the message; and (iii) the *effectiveness problem*, which evaluates the impact of the message in accomplishing a task [3].

Historically, traditional communication systems have concentrated on the technical problem. However, with advances in artificial intelligence and computational power, we can

revisit the other two problems, leading to two new research areas: **Semantic Communication**, which explores the semantic problem, and **Goal-Oriented Communications**, which focuses on the effectiveness problem [4].

In goal-oriented communication, the objective is to optimize transmission strategies to ensure the successful completion of a predefined task while using fewer communication resources. Unlike conventional systems, which prioritize delivering all raw information regardless of its relevance, goal-oriented systems transmit only the data most relevant to the task. This selective approach, however, requires cooperation between source and destination terminals, as well as knowledge of the task objectives at the transmitter side.

Goal-oriented transmission policies have been introduced to exploit the semantic relevance of the message in several problems. In this work, we focus on the remote process tracking problem, in which a signal generated by a source terminal must be reconstructed at the destination terminal. Pappas et al. [5] and later Salimnejad et al. [6] investigated this problem, considering a Markov process. More recently, studies on energy-harvesting IoT systems derived optimal threshold-based policies that jointly account for source state, age of information, and battery levels [7].

This study investigates a transmission strategy for the remote tracking problem under energy constraints. We consider the same scenario adopted by Pappas et al. in [5], i.e., the process to be tracked is modeled as a Markov process X_t and through the use of ACK and NACK feedback messages, the transmitter is aware of the reconstructed process at the destination $\hat{X}_t$. We combine a semantic-aware strategy with a predictive technique to reduce the number of transmissions and, consequently, the transmitter’s energy consumption. We show that, under severe energy constraints, the proposed approach can preserve energy for future transmissions, thereby reducing the signal reconstruction error at the destination.

The remainder of this paper is organized as follows. Section II describes the system model adopted in the analysis of the transmission policies. Section III presents the transmission policies evaluated in this work: the Semantics-Aware baseline, proposed by Pappas et al. in [5], the proposed Predictive policy, and the Ideal Predictive policy. Numerical results are presented and discussed in Section IV, and conclusions are drawn in Section V.

II. SYSTEM MODEL

We consider the system shown in Fig. 1, where a transmitter observes a discrete-time stochastic process X_t at time slots $t \in \{0, 1, \dots\}$ and sends status updates to a remote receiver

Isaac Andrés Balbuca Guachizaca (i290889@dac.unicamp.br) and Paulo Cardieri (cardieri@unicamp.br) are with the School of Electrical and Computer Engineering (FEEC), State University of Campinas (UNICAMP), Campinas-SP, Brazil. This work was financed by RNP/MCTI Brasil 6G (01245.020548/2021-07), Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) - Finance Code 001, INCT-STREAM 409179/2024-8, CAPES (88887.003087/2024-00), and CNPq (311344/2025-9).




Fig. 1: System model for goal-oriented semantic communication with energy harvesting.

following the selected transmission policy. The receiver reconstructs the source process, denoted as $\hat{X}_t$.

The source process X_t is modeled as an N -state Birth-Death Markov Process (BDMP) with a finite state space $\mathcal{S} = \{0, 1, \dots, N-1\}$. State transitions are restricted to the current state or its immediate neighbors, simulating a process that evolves gradually over time. Specifically, the source transitions from state i to $i+1$ with birth probability p and to state $i-1$ with death probability q , while remaining in the current state with probability $1-p-q$, subject to the boundary conditions.

The system operation follows the model adopted by Pappas et al. in [5], as described next. At each time slot t , the transmitter decides whether to send an update according to the binary decision variable $a_t \in \{0, 1\}$, where $a_t = 1$ denotes a transmission attempt, otherwise $a_t = 0$. The status of a_t is determined by the selected transmission strategy. Transmissions occur over a channel with a success probability p_s , and the receiver sends an acknowledgment (ACK) for a successful transmission or a negative-ACK (NACK) for a failure, assumed to be delivered instantaneously and error-free. Through this feedback, the transmitter maintains perfect knowledge of the reconstructed state $\hat{X}_t$ at any given time. Consequently, the transmitter can compare the actual source state X_t with the reconstructed state at the receiver $\hat{X}_t$ at any time slot. This allows it to detect discrepancies and make decisions based on its transmission policy.

The transmitter is powered by an Energy Harvesting module equipped with a finite capacity battery, denoted by $E_{\max}$. Similarly to [8], we assume that energy arrivals follow a Bernoulli process with parameter μ . We define a binary indicator $u_t \in \{0, 1\}$ to represent the energy arrival process at the beginning of slot t , where $\Pr[u_t = 1] = \mu$. When $u_t = 1$, one unit of energy is harvested and added to the battery. Consequently, the battery level e_t evolves according to the following dynamics:

$$e_{t+1} = \min\{e_t + u_t - a_t c_t, E_{\max}\}, \quad (1)$$

where c_t denotes the energy cost associated with a transmission attempt.

A transmission is permitted at slot t only if the battery has sufficient energy, imposing the energy constraint $e_t - a_t c_t \geq 0$, $\forall t$.

The energy consumed by the transmitter depends on the number of transmission attempts, which, in turn, is determined by both the adopted transmission strategy and the channel

conditions, since failed transmissions may trigger different transmitter behavior. Therefore, in order to enable a fair comparison among different transmission strategies, the energy cost must be properly associated to the target transmission success probability. For tractability, we consider an AWGN wireless channel with BPSK modulation. Under this assumption, the energy required for the successful transmission of a packet containing L bits with probability p_s is given by¹:

$$E = \frac{L N_0}{2} \left[Q^{-1} \left(1 - p_s^{1/L} \right) \right]^2, \quad (2)$$

where N_0 denotes the noise power spectral density and $Q(\cdot)$ denotes the Gaussian Q -function.

III. TRANSMISSION POLICIES

In this section, we describe the transmission policies evaluated in our study. First, we present the *Semantics-Aware* policy, which serves as the performance baseline established in the literature. Then, we introduce our proposed *Predictive* policy and its ideal version.

A. Semantics-Aware Policy (SAP)

The *Semantics-Aware* policy, proposed by Pappas and Kountouris [5], relies on two key assumptions: (a) the transmitter has perfect knowledge of the reconstructed state $\hat{X}_t$, and (b) the receiver updates its reconstructed state whenever a message carrying a new state update is successfully received. Otherwise, the receiver retains its previously reconstructed state.

Consider the system operation at time slot $t+1$. If the system is synchronized, i.e., $X_{t+1} = \hat{X}_t$, no transmission is required, since, in the absence of a new update, the receiver simply keeps its current estimate, that is, $\hat{X}_{t+1} \leftarrow \hat{X}_t$. On the other hand, if a mismatch exists, i.e., $X_{t+1} \neq \hat{X}_t$, the transmitter sends an update packet to the receiver. If the transmission is successful (ACK), the receiver updates its estimate according to $\hat{X}_{t+1} = X_{t+1}$. Otherwise, if the transmission fails (NACK), no retransmission is attempted during that time slot, and the receiver maintains its previous estimate, i.e., $\hat{X}_{t+1} = \hat{X}_t$. The *Semantics-Aware* policy is goal-oriented, as it transmits only when the receiver's state diverges from the actual source state, ensuring that every transmission directly contributes to reducing the reconstruction error.

B. Predictive Policy (PP)

In the *Semantics-Aware* policy, a transmission is triggered whenever the system becomes unsynchronized. Although this strategy is effective at keeping the receiver updated, it may lead to high energy consumption. Motivated by this limitation, and aiming to simultaneously reduce energy consumption while maintaining the reconstruction error at acceptable levels, we propose the *Predictive Policy*, which can be interpreted as an extension of the *Semantics-Aware* policy.

The proposed *Predictive Policy* operates as follows. When the system is synchronized ($X_t = \hat{X}_t$), no update is needed.

¹We assume that a packet is successfully received only if all L bits are correctly decoded.

However, if a discrepancy exists, a predictor estimates the next state of the source, denoted by $\tilde{X}_{t+1}$, and compares it with the current state at the receiver $\hat{X}_t$. If the prediction indicates that the system will naturally resynchronize (i.e., $\tilde{X}_{t+1} = \hat{X}_t$), no transmission is performed, thereby saving energy at the expense of a temporary increase in reconstruction error, as the error at time t is not corrected. On the other hand, if the mismatch is expected to persist, i.e., $\tilde{X}_{t+1} \neq \hat{X}_t$, a transmission is triggered to restore synchronization at the receiver.

Any prediction technique can be employed to estimate $\tilde{X}_{t+1}$. In this work, we adopt a simple probabilistic prediction approach in which the transmitter recursively estimates the source transition probabilities, denoted by $\hat{p}$ and $\hat{q}$, from the observed birth and death events. Based on these estimated transition probabilities, the predicted next state $\tilde{X}_{t+1}$ is chosen as the most likely future state of the source.

C. Ideal Predictive Policy (IPP)

To establish a performance upper bound for the Predictive policy, we introduce the *Ideal Predictive Policy*. In this policy, the transmitter has perfect knowledge of the next source state X_{t+1} at the time of decision, eliminating any prediction error. Although physically unrealizable, this policy serves as a benchmark that isolates the contribution of prediction errors to the overall reconstruction performance. The residual error of the IPP arises only from channel losses and intentionally suppressed transmissions.

IV. NUMERICAL RESULTS

The performance of the transmission policies presented in Sec. III is evaluated through simulations under different representative operating scenarios.

A. Simulation Setup and Scenarios

The information source is modeled as a BDMP with $N = 10$ states, considering two distinct scenarios regarding source variability:

- **Rare Event:** Represents a slowly changing source, characterized by transition probabilities $p = 0.01$ and $q = 0.50$.
- **Fast Source:** Represents a rapidly changing source, characterized by $p = 0.20$ and $q = 0.70$.

Both scenarios are evaluated under different energy constraints by varying the energy buffer capacity ($E_{\max} \in \{2, 10\}$) and the probability μ that one unit of energy is harvested during a time-slot. These energy-constrained scenarios allow us to evaluate the potential benefits of the Predictive policies, which seek to preserve energy resources at the expense of a temporary increase in reconstruction error.

Two representative channel conditions are considered: $p_s = 0.5$, corresponding to a *non-reliable channel*, and $p_s = 0.9$, corresponding to a *reliable channel*.

Regarding the transmission cost c_t , we assume that achieving a transmission success probability of $p_s = 0.5$ requires $c_t = 1$ unit of energy. Using the modeling presented in

Section II (expression (2)), increasing the channel reliability to $p_s = 0.9$, the transmitter requires approximately $c_t = 1.56$ unit of energy per transmission, assuming a packet size of $L = 100$ bits.

In all figures presented below, the following acronyms are used: SAP = Semantics-Aware Policy, IPP = Ideal Predictive Policy, and PP = Predictive Policy.

In Sections IV-C and IV-D, we compare the performance of the *Semantics-Aware* and Predictive policies. The results obtained with the Ideal Predictive Policy are discussed in Section IV-E, together with general discussions.

B. Performance Metrics

To quantify the efficiency of both policies, we consider the following performance metrics.

Reconstruction error

This metric, as defined in [5], measures the discrepancy between the original source process X_t and the reconstructed state at the receiver $\hat{X}_t$. The time-averaged reconstruction error over a large observation interval T (expressed in number of time slots) is given as:

$$\bar{E} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}(X_t \neq \hat{X}_t), \quad (3)$$

where $\mathbb{1}(\cdot)$ is the indicator function.

Energy consumption rate

To evaluate the energy efficiency of the transmission policies, we use the *energy consumption rate*, defined as

$$\bar{C} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T a_t c_t, \quad (4)$$

where $a_t \in \{0, 1\}$ is the binary transmission decision at time slot t and c_t is the energy cost associated with a transmission attempt, as introduced in Sec. II. This metric represents the average energy consumed per time slot over a large observation interval T .

Both metrics were evaluated by averaging over $T = 100,000$ slots.

C. Results for the Rare Event Scenario

1) **Non-reliable channel** ($p_s = 0.5$): Fig. 2a shows the reconstruction error achieved by the three policies as a function of the energy harvesting probability, for different energy buffer capacities.




Fig. 2: Reconstruction error (a) and energy consumption rate (b) as functions of energy harvesting probability, for energy buffer capacity $E_{\max} \in \{2, 10\}$, under the *Rare Event scenario* with $p_s = 0.50$.

Under severe energy constraints ($E_{\max} = 2$ and $\mu < 0.5$), the Semantics-Aware policy exhibits the highest reconstruction error. This occurs because, for every detected discrepancy, the transmitter attempts a correction over a channel that fails half the time, rapidly depleting the energy buffer, preventing future transmissions. On the other hand, the Predictive Policy, anticipating self-corrections, tolerates transient errors and triggers transmissions only when the mismatch is likely to persist. In this way, energy is preserved for future transmissions, leading to better performance than the Semantics-Aware policy. However, for $\mu > 0.5$, the battery is recharged frequently enough for the Semantics-Aware policy to effectively correct discrepancies, thereby outperforming the Predictive Policy.

For $E_{\max} = 10$, the larger battery capacity enables the Semantics-Aware policy to transmit whenever required, thereby achieving the lowest reconstruction error.

Fig. 2b shows that the Predictive policies are more energy-efficient. By estimating the future state of the source, these policies trigger transmissions only when the mismatch is likely to persist. As a result, they achieve significantly lower energy consumption.

Overall, while Semantics-Aware can achieve lower reconstruction error when energy is abundant, it operates at the highest energy consumption, whereas the Predictive Policy maintains better energy efficiency by accepting a slight residual error.

2) **Reliable channel** ($p_s = 0.9$): Operating over a reliable channel requires a higher per-transmission energy cost ($c_t \approx 1.56$). Consequently, under severe energy constraints ($E_{\max} = 2$ and $\mu < 0.5$), a single transmission consumes nearly the entire battery capacity, leading to degraded performance for all policies (see Fig. 3a), particularly for the *Semantics-Aware* policy. On the other hand, the higher channel reliability increases the probability of successful packet delivery, which partially mitigates the adverse effects of the limited battery capacity.

For a larger battery capacity ($E_{\max} = 10$), the *Semantics-Aware* policy has sufficient energy to transmit nearly every update packet, as in the non-reliable channel scenario. In the present case, however, it achieves a substantially lower reconstruction error than the Predictive policies, since 90% of the transmitted update packets are successfully received on average.




Fig. 3: Reconstruction error (a) and energy consumption rate (b) as functions of energy harvesting probability, for energy buffer capacity $E_{\max} \in \{2, 10\}$, under the *Rare Event scenario* with $p_s = 0.90$.




Fig. 4: Reconstruction error (a) and energy consumption rate (b) as functions of energy harvesting probability, for energy buffer capacity $E_{\max} \in \{2, 10\}$ under the *Fast Source scenario* with $p_s = 0.50$.

As in the non-reliable channel case, the predictive policies are more energy-efficient, as shown in Fig. 3b.

D. Results for Fast Source Scenario

1) **Non-reliable channel** ($p_s = 0.5$): In this scenario involving a rapidly varying source, the demand for status updates becomes significantly higher, further accentuating the adverse effects of energy constraints, as illustrated in Fig. 4a, which shows reconstruction errors ranging from 0.2 to 0.5. As before, under severe energy constraints ($E_{\max} = 2$ and $\mu < 0.4$), the *Semantics-Aware* policy exhibits the worst performance. However, as these constraints are relaxed and more energy becomes available, the *Semantics-Aware* policy achieves the lowest reconstruction error.

Although the resulting reconstruction errors are too high for practical applications, these results highlight how the combination of severe energy constraints, poor channel quality, and a rapidly varying source demands appropriate transmission policies to cope with such adverse operating conditions.

As expected, all policies consume more energy in the Fast Source scenario than in the Rare-Event scenario, as can be observed by comparing Figs. 2b and 4b. Moreover, in the Fast Source scenario, the energy consumption rate increases with the energy harvesting probability μ . This behavior indicates that the highly volatile nature of the source frequently depletes




Fig. 5: Reconstruction error (a) and energy consumption rate (b) as functions of energy harvesting probability, for energy buffer capacity $E_{\max} \in \{2, 10\}$ under the *Fast Source* scenario with $p_s = 0.90$.

the battery, making the energy availability strongly dependent on the harvesting process.

2) **Reliable channel** ($p_s = 0.9$): In the reliable channel case, which requires a higher transmission cost ($c_t \approx 1.56$), the policies exhibit behavior similar to that observed in the non-reliable channel scenario. However, when energy is abundant, the results show that the Predictive policies reach an error floor significantly higher than that of the *Semantics-Aware* policy, even in the case of ideal prediction (see Fig. 5). This indicates that the intentional errors introduced by the Predictive policies account for a substantial portion of the overall reconstruction error when the source varies rapidly.

E. Discussion

The results highlight the benefits of the Predictive policy in scenarios with limited energy availability (low μ or small $E_{\max}$). Under such conditions, the transmission logic adopted by the *Semantics-Aware* policy, which attempts to immediately correct every discrepancy without distinguishing transient errors from those that would naturally self-correct in the next source transition, may lead to rapid battery depletion. As a consequence, future transmissions become impossible, increasing the average reconstruction error. In contrast, the Predictive policy seeks to preserve energy, even at the cost of temporarily tolerating a higher reconstruction error.

The benefits of the Predictive policy are also observed in scenarios with less severe energy constraints (high μ or large $E_{\max}$), although now from the perspective of energy consumption. In these situations, the Predictive policy is able to achieve low reconstruction error levels (below 1%, compared to approximately 0.2% achieved by the *Semantics-Aware* policy, as illustrated in Fig. 3), while requiring lower energy consumption.

The gap between the Predictive and the Ideal Predictive policies provides further insight into the limits of the predictive approach. Since the Ideal Predictive Policy completely eliminates prediction errors, the difference between the reconstruction errors achieved by the Predictive Policy and the Ideal Predictive Policy represents the performance degradation caused solely by the stochastic uncertainty of the source process, which cannot be overcome by any predictor, regardless

of how accurately p and q are estimated. The gap is more pronounced in the *Fast Source* scenario (see Fig. 5a), where the higher transition probabilities increase the unpredictability of the source, making predictions less reliable. On the other hand, in the *Rare Event* scenario, the source evolves slowly, and the gap narrows considerably (see Figs. 3a), indicating that the Predictive policy operates closer to the ideal limit when the source dynamics are more predictable. These observations suggest that the Predictive policy is most useful in a regime where the source is slow enough for predictions to be reliable.

V. CONCLUSIONS AND FUTURE WORKS

This study investigated the remote tracking problem of a Markovian source under energy constraints, proposing a Predictive Policy that balances the reconstruction error and the energy consumption. By predicting the source's state, the policy anticipates the source's natural self-correction, preserving energy for persistent errors rather than depleting it on transient discrepancies, unlike the *Semantics-Aware* policy. The results reveal that while the *Semantics-Aware* policy minimizes error under high energy availability, the Predictive policy offers a more sustainable strategy, achieving substantial energy savings by accepting a residual error.

Future works include formulating the transmission scheduling problem as a Markov Decision Process to derive an optimal policy that jointly minimizes reconstruction error and energy consumption, incorporating the principle of suppressing transmissions when self-synchronization is anticipated.

REFERENCES

- [1] W. Saad, C. Chaccour, C. K. Thomas, and M. Debbah, Eds., *Foundations of Semantic Communication Networks*. Hoboken, NJ, USA: John Wiley & Sons, 2025.
- [2] H. Zhou, Y. Deng, X. Liu, N. Pappas, and A. Nallanathan, "Goal-oriented semantic communications for 6g networks," *IEEE Internet of Things Magazine*, vol. 7, no. 5, pp. 104–110, 2024.
- [3] C. Shannon and W. Weaver, *The Mathematical Theory of Communication*, ser. Illini books. University of Illinois Press, 1949, no. v. 1. [Online]. Available: <https://books.google.com.br/books?id=OfxtvWEACAAJ>
- [4] W. Yang, H. Du, Z. Q. Liew, W. Y. B. Lim, Z. Xiong, D. Niyato, X. Chi, X. Shen, and C. Miao, "Semantic communications for future internet: Fundamentals, applications, and challenges," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 1, pp. 213–250, 2023.
- [5] N. Pappas and M. Kountouris, "Goal-oriented communication for real-time tracking in autonomous systems," in *2021 IEEE International Conference on Autonomous Systems (ICAS)*, 2021, pp. 1–5.
- [6] M. Salimnejad, M. Kountouris, and N. Pappas, "Real-time reconstruction of markov sources and remote actuation over wireless channels," *IEEE Transactions on Communications*, vol. 72, no. 5, pp. 2701–2715, 2024.
- [7] E. Delfani, G. J. Stamatakis, and N. Pappas, "State-aware timeliness in energy harvesting iot systems monitoring a markovian source," *IEEE Transactions on Green Communications and Networking*, vol. 9, no. 3, pp. 977–990, 2025.
- [8] A. Zakeri, M. Moltafet, and M. Codreanu, "Semantic-aware real-time tracking of a markov source under sampling and transmission costs," in *2023 57th Asilomar Conference on Signals, Systems, and Computers*, 2023, pp. 694–698.