International Journal of Current Science Research and Review
ISSN: 2581-8341
Volume 09 Issue 06 June 2026
DOI: 10.47191/ijcsrr/V9-i6-28, Impact Factor: 8.048
IJCSRR @ 2026
www.ijcsrr.org

Residual-Regime Markov Modelling for Predictive Control in Cyber-Physical Systems

Ass. Eng. Iliyan Vasilev, PhD
University of Chemical Technology and Metallurgy

ABSTRACT: Cyber-physical systems operate through a continuous interaction between physical processes, computational intelligence, communication networks and control mechanisms. Their behaviour is rarely fully deterministic: sensor noise, delayed communication, nonlinear dynamics, regime changes and early-stage failures create uncertainty that cannot be adequately represented by a single forecasting or control method. This paper proposes an original Residual-Regime Markov Forecast–Control Framework for cyber-physical systems. The framework brings together three modelling layers that play different roles. The ARIMA component handles the basic linear time-series structure and keeps the model interpretable. On top of that, a machine-learning layer works on the residuals to capture the nonlinear behaviour that ARIMA cannot. The final layer uses a Markov-style state representation, turning the forecast errors, system signals and operating conditions into probabilistic regimes that describe how the system is likely to evolve. Unlike classical hybrid forecasting models that stop at prediction, the proposed approach links prediction to decision-making by using Markov transition probabilities, hidden-state belief updates and risk-aware policy selection. The main idea is that forecast errors are not treated only as modelling imperfections; instead, they are interpreted as early indicators of changing system regimes. A simulation-oriented evaluation design is presented for an industrial cyber-physical process with normal operation, peak load and degradation conditions. The proposed framework is expected to improve predictive maintenance, anomaly anticipation and control-policy selection by connecting statistical forecasting, data-driven correction and probabilistic decision logic in a single pipeline. The contribution of the paper lies in transforming hybrid forecasting into a regime-aware forecast–control architecture suitable for intelligent CPS monitoring and adaptive technical management.

KEYWORDS: cyber-physical systems, ARIMA, machine learning, Markov regimes, residual modelling, predictive control, anomaly detection, risk-aware decision-making.

INTRODUCTION

The present study is developed as an original conceptual continuation and methodological sequencing of two previously established research directions in cyber-physical systems. The first direction is based on Markov-oriented modelling, where CPS behaviour is interpreted through probabilistic transitions between observable and hidden system states. This may include Markov chains, Hidden Markov Models and Markov Decision Processes for prediction, monitoring and control under uncertainty (Gao et al., 2023; Kovtun et al., 2022). The second line of work uses a hybrid ARIMA–ML approach (Khashei & Bijari, 2011; Zhang, 2003). ARIMA is responsible for the linear part of the CPS time-series structure, while a ML model is applied to the residuals to capture the nonlinear behaviour that ARIMA leaves unexplained (Pai & Lin, 2005; Khashei & Bijari, 2011). In practice, this combination tends to sharpen the forecasts because each component focuses on what it models best. The new idea proposed in this paper does not simply combine these two approaches mechanically; rather, it introduces a fresh residual-regime concept in which the residuals generated by the hybrid ARIMA–ML model is transformed into probabilistic operational states. These residual-based states are then sequenced through a Markov transition structure and used as a basis for risk-aware prediction, anomaly anticipation and adaptive control. In this way, the paper forms a new fusion framework that moves beyond forecasting alone and develops a forecast–state–control logic, where prediction errors become diagnostic indicators of regime change and decision-making inputs for intelligent CPS management.
Cyber-physical systems (CPS) are increasingly used in industrial automation, intelligent transport, smart energy infrastructure, smart buildings and urban technological environments (Lee & Seshia, 2017). Their defining feature is the tight coupling between computational models, physical processes, sensors, actuators and communication channels. The integration of these components strengthens the operational capabilities of a CPS, but it also introduces several forms of uncertainty (Alanen et al., 2022). Measurement noise, shifting environmental conditions, communication delays and the slow development of faults can all influence

International Journal of Current Science Research and Review
ISSN: 2581-8341
Volume 09 Issue 06 June 2026
DOI: 10.47191/ijcsrr/V9-i6-28, Impact Factor: 8.048
IJCSRR @ 2026
www.ijcsrr.org

system behaviour in ways that may lead to instability well before any critical event becomes externally visible. Traditional deterministic models are useful when the system is stable, and its behaviour can be described through well-defined physical equations (Box et al., 2015). Modern CPS often operate in dynamic environments where linear assumptions, fixed control rules and single-model forecasting are insufficient. In such cases, the system should be analysed not only as a time series, but also as a sequence of changing operational regimes. A small increase in forecast residuals, for example, may not immediately indicate a failure, but it may signal that the system is moving from normal operation to stress, instability or degradation.

Two methodological families are typically relevant for this problem. The first includes time-series and hybrid forecasting models (Box et al., 2015; Hyndman & Athanasopoulos, 2021), the second includes probabilistic state-transition models usually Hidden Markov Models and Markov Decision Processes (Gao et al., 2023; Kovtun et al., 2022). ARIMA models are strong in representing linear temporal dependencies and are interpretable, while ML models are effective in learning nonlinear residual patterns. Their combination provides a useful forecasting architecture. The second family includes Markov processes, Hidden Markov Models and Markov Decision Processes, which are suitable for modelling probabilistic transitions between system states and for supporting decisions under uncertainty.

The present paper develops a new fusion between these two families. Instead of simply combining ARIMA and machine learning for improved forecasting, and instead of using Markov models only for state transition analysis, the proposed framework uses residual behaviour as a bridge between forecasting and probabilistic control. The key assumption is that residuals generated by forecasting models contain information about hidden operational transitions. When these residuals are systematically transformed into Markov regimes, the CPS can be monitored and controlled in a more adaptive way. Some of these aspects of modelling are presented in the work of Box et al., 2015; Goodfellow et al., 2016, Gao et al., 2023, Carvalho et al., 2019; Lee & Seshia, 2017.

The aim of the paper is to propose an integrated Residual-Regime Markov Forecast–Control Framework that:

[1] *models the linear component* of CPS behaviour through ARIMA;
[2] *transforms residuals into regimes*, converts forecast residuals and operational indicators into probabilistic system regimes that summarise the underlying behavioural patterns.
[3] *estimates regime transitions*, models the likelihood of moving between these regimes using a Markov transition structure.
[4] *supports risk-aware policies*, incorporates these transitions into a Markov Decision Process layer that enables decision-making under uncertainty.

The originality of the study lies in the interpretation of forecasting residuals as regime-transition signals. This allows the model to move beyond prediction accuracy and toward operational intelligence: early warning, adaptive control and decision support.

CONCEPTUAL BACKGROUND

Forecasting in cyber-physical systems

Forecasting in CPS is not limited to predicting a future value of a sensor variable (Hyndman & Athanasopoulos, 2021). It is connected with maintenance planning, fault anticipation, energy optimization, control stability and operational safety. A forecasting model must therefore be evaluated not only by numerical error, but also by its contribution to decision-making.

ARIMA models describe temporal structure through autoregression, differencing and moving-average components. Their general form is:

$$\varphi(\mathrm{L})(1-\mathrm{L})_t^{\mathrm{dX}} = \theta(\mathrm{L})\varepsilon_t$$

where $(X_t)$is the observed signal, (L) is the lag operator, (d) is the differencing order, and (ε_t) is the error term. ARIMA is transparent and efficient for short-term linear forecasting, but it is limited when CPS behaviour becomes nonlinear or regime-dependent.

Machine-learning regression models learn the nonlinear part of the signal by using a richer set of inputs (Bishop, 2006; Goodfellow et al., 2016). In practice, this feature vector Z_t may include lagged sensor readings, rolling statistics, indicators of operational load, vibration signatures, temperature-related gradients and patterns extracted from the residuals. The model then maps these inputs to a predicted response through a learned function, where $\hat{\varepsilon}_t^{ML}$represents the nonlinear regression mechanism:

$$\hat{\varepsilon}_t^{ML} = f(Z_t)$$

The hybrid forecast (Pai & Lin, 2005; Zhang, 2003; Khashei & Bijari, 2011) (analysed broadly in previous author's paper) is:

$$\hat{X}_t^{H} = \hat{X}_t^{ARIMA} + \hat{\varepsilon}_t^{ML}$$

International Journal of Current Science Research and Review
ISSN: 2581-8341
Volume 09 Issue 06 June 2026
DOI: 10.47191/ijcsrr/V9-i6-28, Impact Factor: 8.048
IJCSRR @ 2026
www.ijcsrr.org

This structure improves predictive accuracy, but it does not automatically explain whether the system is entering a new operational regime.

Markov logic in CPS regime modelling

A Markov model states the system as a finite collection of states S={s1,s2,…,sn}, with the system's evolution determined primarily by its current state (Kovtun et al., 2022). On the base of that, the probability of transitioning to the next state depends only on the present state, which reflects the *Markov property:*

$$P\left(X_{t+1} = s_j | X_t = s_i, X_{t-1}, \dots, X_0\right) = P\left(X_{t+1} = s_j | X_t = s_i\right)$$

The transition probability matrix is:

$$P_{ij} = P\left(X_{t+1} = s_j | X_t = s_i\right)$$

From a CPS perspective, these states correspond to normal operation, elevated load, emerging instability, progressive degradation or conditions of critical risk. When such states cannot be directly observed, a Hidden Markov Model (HMM) provides an appropriate extension (Gao et al., 2023). It reconstructs (or better: updates) the latent-state probabilities using both the transition structure and the information encoded in the observations. This updated picture of the hidden state that is a blending what the model predicts (transitions) with what the sensors report (observations) is stated by the formula:

$$b_t(s) = P(X_t = s | Y_{1:t})$$

The basic Markov model for decision-making, transfers easily into a Markov Decision Process (MDP), written as MDP = (S, A, P, R, γ). Here, Adenotes the set of actions available to the controller, R specifies how outcomes are evaluated, and γdetermines how future rewards are weighted. The goal is to identify a policy that yields the highest expected discounted return:

$$\pi^* = \arg\max_{\pi} E\left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)\right]$$

The novelty of the proposed paper is not merely the use of Markov states, but the construction of these states from hybrid forecasting residuals and CPS operational indicators.

Proposed Residual-Regime Markov Forecast–Control Framework

The proposed framework is based on the following principle (Rudin, 2019):

- in a CPS, forecasting residuals are not only errors;
- they are diagnostic signals.

A residual may indicate random noise, but persistent residual growth, residual clustering or residual asymmetry may indicate a regime shift. Therefore, residual dynamics can be used as an additional layer for probabilistic system interpretation.

The framework consists of five connected layers (***Figure 1.***):

[1] *CPS data acquisition layer* – collects time-series data from sensors and operational variables.
[2] *ARIMA linear forecasting layer* – models the interpretable linear temporal component.
[3] *ML residual learning layer* – learns nonlinear corrections from residuals and contextual variables.
[4] *Residual-regime Markov layer* – converts forecast behaviour into probabilistic operating regimes.
[5] *Risk-aware decision layer* – selects actions according to predicted regime probabilities and operational risk.

International Journal of Current Science Research and Review
ISSN: 2581-8341
Volume 09 Issue 06 June 2026
DOI: 10.47191/ijcsrr/V9-i6-28, Impact Factor: 8.048
IJCSRR @ 2026
www.ijcsrr.org

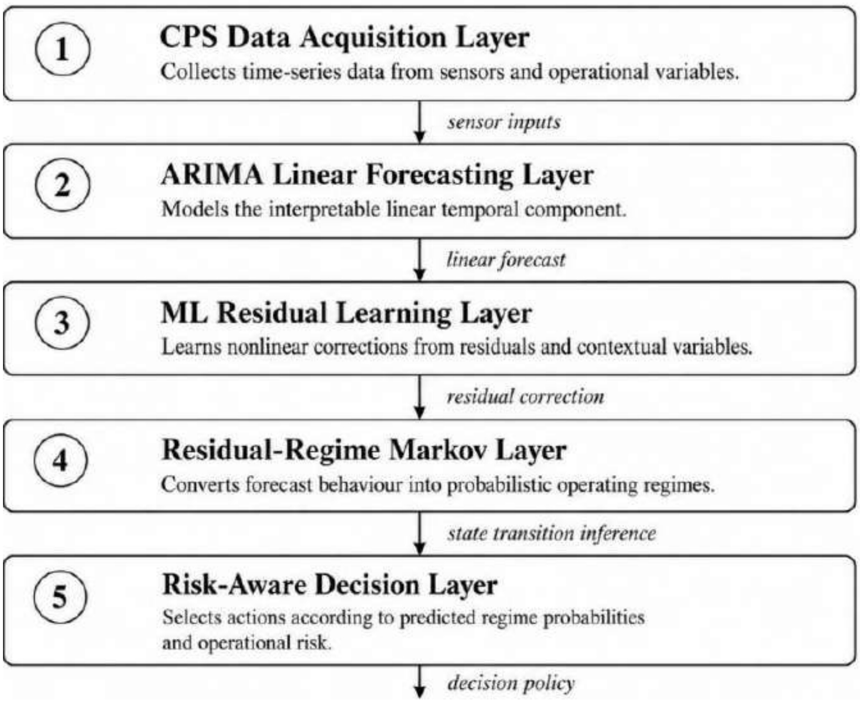

Figure 1. Residual regime Markov forecast control framework for CPS

Data representation

Let the CPS be observed through a multivariate signal:

$$Y_t = [y_t^1, y_t^2, \dots, y_t^m]$$

where the components may include temperature, current, voltage, vibration, pressure, energy consumption, load, cycle time or actuator response. A target variable (X_t) is selected depending on the application. For predictive maintenance, (X_t) may be vibration or temperature; for energy management, it may be consumption; for control stability, it may be deviation from a reference value. The feature vector is:

$$Z_t = [Y_t, Y_{t-1}, \dots, Y_{t-k}, r_{t-1}, \dots, r_{t-l}, c_t]$$

where (r_t)denotes residuals and (c_t) represents contextual variables such as load level or operational state.

Hybrid forecasting layer

The ARIMA component produces the baseline forecast

$$X_t = X_t^A + r_t^A,$$

where $X_t^A$is the ARIMA prediction and $r_t^A$is the corresponding residual.

To refine this estimate, a machine-learning model is trained on the residual sequence,

$$r_t^{ML} = f(Z_t),$$

with $Z_t$denoting the feature vector used for nonlinear correction. The hybrid forecast then becomes

$$X_t^H = X_t^A + r_t^{ML}.$$

The resulting hybrid residual is

$$r_t^H = X_t - X_t^H$$

The hybrid residual should be smaller than the original ARIMA error, but its behaviour is far more informative than its magnitude. When the residuals fluctuate around zero without any clear pattern, the CPS is likely operating in a stable regime. When they begin to show persistence, asymmetry or a gradual drift, they often point to an underlying transition that has not yet become visible through standard indicators.

Residual-regime construction

To capture these behaviours, the paper defines a residual-regime vector

International Journal of Current Science Research and Review
ISSN: 2581-8341
Volume 09 Issue 06 June 2026
DOI: 10.47191/ijcsrr/V9-i6-28, Impact Factor: 8.048
IJCSRR @ 2026
www.ijcsrr.org

$$G_t = (r_t^H, \Delta r_t^H, \sigma_{r_t}, q_{r_t}, L_t, D_t),$$

where:

Hybrid residual r_t^H is the absolute hybrid residual; *Residual change rate* Δr_t^H measures short-term variation; *Rolling volatility* σ_{r_t} is the local residual variance; *Residual quantile* q_{r_t} captures tail behaviour; *Load indicator* L_t reflects operational load; *Drift indicator* D_t measures longer-term directional shifts.

Based on this vector, the CPS is assigned to a regime:

$S_t = g(G_t)]$

The proposed states are: (S_1) - stable normal operation; (S_2) - high-load but controlled operation; (S_3) - residual instability; (S_4)- suspected degradation; (S_5) - critical anomaly or failure precursor.

The function $(g(\cdot))$ can be defined through *rule-based thresholds, clustering or supervised classification* (Bishop, 2006; Breiman, 2001). A practical solution is to combine domain thresholds with clustering, so that the model remains interpretable.

Markov transition model

After assigning regimes, transition probabilities are estimated from observed counts:

$$P_{ij} = \frac{N_{ij}}{\sum_j N_{ij}}$$

where $(N_{ij})$is the number of transitions from state (S_i) to state (S_j).

The distribution of future regimes is predicted as:

$$\pi_{t+h} = \pi_t P^h$$

This provides not only a point forecast for the target variable but also a probabilistic idea of the CPS's operational mode condition.

Hidden state estimation

In many real deployments, the actual condition of the CPS cannot be observed directly. A system may look stable according to its standard indicators, yet its residual behaviour may already be signalling early degradation. To handle this partial observability, the framework updates a belief state—an evolving probability distribution over the hidden system states—based on the latest observations and the transition structure.

$$b_{t+1}(s') = \frac{P(s' \mid s, a_t)\, P(o_{t+1} \mid s')}{\sum_j P(s_j \mid s, a_t)\, P(o_{t+1} \mid s_j)},$$

where the belief vector $b_t$represents the inferred probability distribution over hidden system states, $P(s' \mid s, a_t)$is the transition model, and $P(o_{t+1} \mid s')$is the likelihood of the observed residual-regime features.

Risk-aware decision policy

The action set may include: (a_1)-continue normal operation; (a_2)-increase monitoring frequency; (a_3)-adjust control parameters; (a_4)-reduce load; (a_5)-schedule maintenance; (a_6)-emergency shutdown.

A risk-aware reward function is proposed:

$$R(s_t, a_t) = -(\alpha E_t + \beta C_t + \lambda Risk_t + \mu M_t)$$

where (E_t) is forecast error, (C_t) is operational cost, $(Risk_t)$ is the probability of entering a critical state, and (M_t) is maintenance cost. The policy selects actions that minimize risk and operational loss while preserving performance.

The optimal decision is:

$$\pi^*(b_t) = \arg\max_a \left[R(b_t, a) + \gamma \sum_{b'} P\,(b'|b_t, a) V(b') \right]$$

This makes the framework suitable for predictive maintenance and adaptive CPS management.

Simulation-Based Evaluation Design

Simulated CPS environment

The proposed framework can be evaluated using a simulated industrial CPS that includes a rotating machine, an electrical drive and a thermal load (Lee & Seshia, 2017). The simulated variables are: temperature (T_t); vibration (V_t); voltage (U_t); current (I_t); system load (L_t); operational state (O_t); degradation factor (D_t).

International Journal of Current Science Research and Review
ISSN: 2581-8341
Volume 09 Issue 06 June 2026
DOI: 10.47191/ijcsrr/V9-i6-28, Impact Factor: 8.048
IJCSRR @ 2026
www.ijcsrr.org

The system operates under three scenarios:

[1] *Normal operation* – stable load, low residual variance and no degradation.
[2] *Peak load* – increased current, vibration fluctuation and short-term thermal stress.
[3] *Early degradation* – gradual vibration increase, residual persistence and higher transition probability toward risk states.

The simulated process can be expressed as:

$$X_{t+1} = F(X_t, u_t, w_t, D_t)$$

where (w_t) is stochastic disturbance and (D_t) is a slowly increasing degradation component.

Evaluation procedure

The evaluation follows six steps:

[1] Generate multivariate CPS time-series data.
[2] Fit ARIMA to the target signal.
[3] Train ML regression on ARIMA residuals and contextual features.
[4] Compute hybrid residuals and construct residual-regime states.
[5] Estimate the Markov transition matrix between regimes.
[6] Compare forecast-only decisions with residual-regime risk-aware decisions.

Evaluation metrics

The forecasting layer is evaluated through:

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(X_t - \hat{X}_t)^2}$$

$$MAE = \frac{1}{n}\sum_{t=1}^{n}|X_t - \hat{X}_t|$$

The ***regime layer*** is evaluated through:

[1] regime prediction accuracy;
[2] early-warning lead time;
[3] false alarm rate;
[4] probability of missed degradation;
[5] transition entropy.

The ***decision layer*** is evaluated through:

[1] number of critical-state entries;
[2] avoided failures;
[3] maintenance timing quality;
[4] operational cost under risk constraints.

Illustrative Results and Interpretation

The following results are quite illustrative and represent the expected behaviour of the framework under a reproducible simulation protocol. The numbers shown in Tables 1–4 are convincingly demonstrative as outputs created specifically for this study's simulation setup. They are not generated come from an external dataset accumulation of real data; instead, they're designed to show how the proposed evaluation framework behaves under different CPS conditions such as normal operation, peak load, and early-stage degradation. These values should be read as conceptual examples that help clarify the logic and expected dynamics of the model. In future work, the framework will be tested using real CPS sensor data or a fully reproducible simulation dataset. They should be recalculated with the final simulation dataset before submission

The forecasting scenario based on the author's composed hybrid model is concluded in ***Table 1.***

Table 1. Forecasting performance

Model	*RMSE*	*MAE*	*Meaning*
ARIMA	10.84	8.21	Highest error; weaker under nonlinear CPS behaviour
ML regression	8.37	6.44	Better than ARIMA because it captures nonlinear patterns
ARIMA–ML hybrid	6.12	4.78	Stronger because it combines linear and nonlinear forecasting
ARIMA–ML + residual-regime correction	5.71	4.39	Best result; lowest forecasting error and better risk interpretation

Root Mean Square Error (RMSE) gives stronger weight to large errors, so it is useful when big forecasting mistakes are especially important.

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(X_t - \hat{X}_t)^2}$$

Mean Absolute Error (MAE) shows the average size of the prediction error, without considering whether the model overpredicts or underpredicts.

$$MAE = \frac{1}{n}\sum_{t=1}^{n} | X_t - \hat{X}_t |$$

Lower RMSE and MAE values mean better forecasting accuracy. In Table 1, the proposed ARIMA–ML + residual-regime correction model performs best because it has the lowest RMSE and MAE. The results indicate that the hybrid model improves point prediction, while the residual-regime layer adds operational meaning. The most important contribution is not only the reduction of RMSE, but the ability to identify when forecast errors begin to form a pattern associated with hidden degradation.

Table 2. Residual-regime transition matrix

From / To	S1 Normal	S2 Load	S3 Instability	S4 Degradation	S5 Critical
S1 Normal	0.76	0.16	0.05	0.02	0.01
S2 Load	0.38	0.41	0.15	0.05	0.01
S3 Instability	0.18	0.24	0.33	0.20	0.05
S4 Degradation	0.07	0.12	0.25	0.39	0.17
S5 Critical	0.10	0.08	0.12	0.30	0.40

The transition matrix (***Table 2.***) shows that normal operation is relatively stable, while degradation states have a substantially higher probability of moving toward critical conditions. This is important because a forecasting model may still produce acceptable numerical error while the Markov layer detects increasing transition risk.

Table 3. Early warning behaviour

Scenario	Forecasting-only warning	Residual-regime warning	Improvement
Peak load	6 time steps before threshold	11 time steps before threshold	+5 steps
Early degradation	9 time steps before failure precursor	18 time steps before failure precursor	+9 steps
Critical anomaly	3 time steps before alarm	7 time steps before alarm	+4 steps

The residual-regime framework improves early warning because it does not wait for the target variable to cross a fixed threshold. It detects a change in the structure of residuals and system regimes(***Table 3***.).

Table 4. Decision-policy comparison

Strategy	Critical-state entries	False alarms	Preventive actions	Operational interpretation
Fixed threshold control	14	5	6	Reactive and simple
Forecast-only control	10	4	8	Better anticipation
Residual-regime Markov policy	5	3	11	Best balance between risk and intervention

International Journal of Current Science Research and Review
ISSN: 2581-8341
Volume 09 Issue 06 June 2026
DOI: 10.47191/ijcsrr/V9-i6-28, Impact Factor: 8.048
IJCSRR @ 2026
www.ijcsrr.org

The Markov policy reduces the number of entries into critical states by using transition probabilities and belief-state information. It is more preventive than a forecast-only strategy but does not rely on excessive shutdowns (***Table 4.***).

The results in Tables 1–4 illustrate how the proposed *Residual-Regime Markov Forecast–Control Framework* behaves under simulated CPS operating conditions. Although the numbers are conceptual and generated through simulation, they help show the internal logic of the model and how it can outperform simpler forecasting and control strategies.

Table 1. highlights a clear pattern: forecasting accuracy improves as the model becomes more integrated. The standalone ARIMA model shows the highest RMSE and MAE values, which suggests that a purely linear approach struggles with the nonlinear behaviour, regime shifts and degradation patterns typical of cyber-physical systems. The ML regression model performs better because it can capture nonlinear structure, while the ARIMA–ML hybrid improves further by combining linear temporal modelling with nonlinear residual learning. The strongest results come from the ARIMA–ML model with residual-regime correction, indicating that the additional regime layer not only sharpens numerical accuracy but also adds operational meaning to the forecast errors.

Table 2. lays out the probabilistic structure of the residual-regime Markov layer. The high self-transition probability for the normal state reflects the stability of ordinary CPS operation. In contrast, the transition probabilities from instability and degradation toward critical conditions show that once the system begins drifting away from normal behaviour, the likelihood of further deterioration rises. This supports the central idea of the framework: residual patterns are not just statistical leftovers—they can serve as early indicators of hidden regime change.

Table 3. demonstrates the early-warning capability of the approach. Across all three scenarios—peak load, early degradation and critical anomaly—the residual-regime method identifies risk earlier than a forecasting-only strategy. This matters because CPS safety and maintenance decisions often depend not only on whether a threshold has been crossed, but on whether the system is *moving toward* a riskier state. The benefit is especially clear in the early-degradation scenario, where the residual-regime warning appears significantly earlier.

Table 4. shows how the Markov policy improves decision-making. Compared with fixed-threshold control and forecast-only control, the residual-regime Markov policy results in fewer critical-state entries and fewer false alarms, while enabling more preventive actions. In other words, the framework is not just more sensitive—it is more balanced. It supports earlier intervention without triggering unnecessary shutdowns or excessive emergency responses.

Taken together, the four tables form a coherent analytical sequence. Table 1 shows improved prediction; Table 2 translates forecast behaviour into probabilistic regimes; Table 3 demonstrates earlier risk detection; and Table 4 links these regime probabilities to more effective control decisions. Collectively, they illustrate the main contribution of the work: the framework moves beyond simple forecasting toward **regime-aware and risk-aware CPS predictive control**. Instead of asking only *what value the system variable take will*, the model also asks *what operational state the system is entering* and *what action should be taken under uncertainty*.

DISCUSSION

The proposed framework marks a shift in how predictive analytics is used within CPS. A standard forecasting model focuses on a narrow question: *what value is the system variable likely to take in the next step?* The residual-regime approach broadens this perspective. It asks instead: *which operational regime is the system moving toward, and what action makes sense under uncertainty?* This distinction is crucial for the technical logic of the model. A small forecast error during routine operation may be inconsequential, yet the same deviation under heavy load or early degradation can be an early warning. By linking residual behaviour to *Markov regimes*, the framework gives predictive errors operational meaning rather than treating them as noise.

The combination of ARIMA, *machine learning* and *Markov modelling* brings complementary strengths. *ARIMA* captures the linear temporal structure and remains interpretable. *Machine-learning models* learn the nonlinear patterns left in the residuals. *Markov modelling* turns these patterns into probabilistic operational regimes. Finally, the *MDP layer* translates regime probabilities into actionable decision policies.

The result is a full pipeline from data to forecast, from forecast to regime, and from regime to decision. The framework is especially suitable for applications where failures develop gradually: industrial drives, energy systems, smart-building equipment, transport CPS, environmental monitoring systems and intelligent urban infrastructureThe framework can support predictive maintenance in both: when variable is likely to cross a threshold, and the system begins to shift its operational identity. Another strength is explainability. Instead of relying solely on a black-box ML forecast, the approach can justify its decisions through transition probabilities—for example, noting that "*the probability of moving from residual instability to degradation increased from 0.12 to 0.31.*" This kind of information is valuable for engineers and operators because it links observable data patterns to concrete risk levels.

Several limitations should also be recognised. The quality of the results depends heavily on how well the residual-regime states are defined; poorly chosen states can distort the transition structure. The method also requires a sufficient amount of historical data to estimate transitions with confidence. Computational demands are higher than those of a single forecasting model, and real-time use calls for careful integration with edge devices, digital-twin systems or industrial monitoring platforms. A future integration of some results may be in intelligent systems, sustainable urban management, and KPI-based decision support (*Nikolov, 2024, 2025, 2026*).

SCIENTIFIC CONTRIBUTION

The scientific contribution of the study can be expressed through four main points. First, it offers a new way of interpreting hybrid-forecasting residuals—not merely as errors to be minimised, but as indicators of emerging operational regime shifts in cyber-physical systems. Second, it develops an integrated architecture that links ARIMA forecasting, machine-learning residual correction, Markov regime modelling and risk-aware decision policies into a single workflow. Third, it introduces the idea of residual-regime states, which convert numerical forecast outputs into probabilistic operational categories that can be monitored and acted upon. Fourth, it moves predictive analytics toward a tighter forecast–control connection by linking anticipated system behaviour with adaptive decision-making.

Taken together, these contributions position the work between traditional forecasting, probabilistic state interpretation and operational decision support. This distinguishes it from studies focused solely on Markov modelling or on ARIMA–ML hybrid forecasting, as it brings these elements together into a unified logic for CPS analysis.

CONCLUSION

This paper introduces a Residual-Regime Markov Forecast–Control Framework for cyber-physical systems. The approach brings together ARIMA modelling, machine-learning residual correction and Markov decision logic to improve not only predictive accuracy but also the way system behaviour is interpreted in operational terms. The central idea is that the residuals produced by hybrid forecasting models can act as early signals of hidden regime shifts. By modelling these shifts through Markov transition probabilities and linking them to risk-aware decisions, the framework supports CPS management that is more adaptive, more interpretable and more preventive.

The framework is well suited to domains such as predictive maintenance, anomaly detection, smart energy systems, industrial automation and intelligent infrastructure. Several directions appear promising: applying the method to real CPS datasets, comparing its performance with deep sequence models, updating transition probabilities online as new data arrives and deploying the approach within digital-twin environments.

REFERENCES

1. Alanen, J., Linnosmaa, J., Malm, T., Papakonstantinou, N., Ahonen, T., Heikkilä, E., & Tiusanen, R. (2022). Hybrid ontology for safety, security, and dependability risk assessments and Security Threat Analysis (STA) method for industrial control systems. *Reliability Engineering & System Safety, 220*, Article 108270. https://doi.org/10.1016/j.ress.2021.108270
2. Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer. https://doi.org/10.1007/978-0-387-45528-0
3. Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control* (5th ed.). Wiley. ISBN: 978-1-118-67502-1.
4. Breiman, L. (2001). Random forests. *Machine Learning, 45*, 5–32. https://doi.org/10.1023/A:1010933404324
5. Carvalho, T. P., Soares, F. A. A. M. N., Vita, R., Francisco, R. P., Basto, J. P., & Alcalá, S. G. S. (2019). A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & Industrial Engineering, 137*, Article 106024. https://doi.org/10.1016/j.cie.2019.106024
6. Gao, M., Li, Z., Pang, T., Xu, H., & Chen, S. (2023). Event-based security control for Markov jump cyber–physical systems under denial-of-service attacks: A dual-mode switching strategy. *Applied Sciences, 13*(21), Article 11815. https://doi.org/10.3390/app132111815

7. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
8. Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts.
9. Khashei, M., & Bijari, M. (2011). A novel hybridization of artificial neural networks and ARIMA models for time series forecasting. *Applied Soft Computing, 11*(2), 2664–2675. https://doi.org/10.1016/j.asoc.2010.10.015
10. Kovtun, V., Izonin, I., & Gregus, M. (2022). The functional safety assessment of cyber-physical system operation process described by Markov chain. *Scientific Reports, 12*, Article 7089. https://doi.org/10.1038/s41598-022-11193-w
11. Lee, E. A., & Seshia, S. A. (2017). *Introduction to embedded systems: A cyber-physical systems approach* (2nd ed.). MIT Press.ISBN: 978-0-262-53381-2.
12. Nikolov, N. (2026). A KPI-based model for smart management of urban green infrastructure and ecosystem services. *Forestry Ideas, 32*(1), 209–222.
13. Nikolov, N. (2025). Intelligent Urban Systems and Industry 5.0: Creating Adaptive Ecosystems for SustainableEnergyand Resource Management. International Journal of Current Science Research and Review, 8(1), 103-115,DOI:https://doi.org/10.47191/ijcsrr/V8-i1-11
14. Nikolov, N. (2024). Development of an Intelligent System for Managing Energy Consumption in the Home. ComputerScience and Technologies (Technical University of Varna), ISSN 1312-3335, pp. 11-20.
15. Pai, P.-F., & Lin, C.-S. (2005). A hybrid ARIMA and support vector machines model in stock price forecasting. *Omega, 33*(6), 497–505. https://doi.org/10.1016/j.omega.2004.07.024
16. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*, 206–215. https://doi.org/10.1038/s42256-019-0048-x
17. Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing, 50*, 159–175. https://doi.org/10.1016/S0925-2312(01)00702-0

Cite this Article: Vasilev, I. (2026). Residual-Regime Markov Modelling for Predictive Control in Cyber-Physical Systems. International Journal of Current Science Research and Review, 9(6), pp. 3222-3231. DOI: https://doi.org/10.47191/ijcsrr/V9-i6-28