

CONCEPT DRIFT AND NON-STATIONARITY IN MACHINE LEARNING - AN INFORMATION-THEORETIC VIEW

Dr. REJI KURIEN THOMAS

CEO, TOL Biotech, USA. Email: founder@technopilot.in

Abstract

Classical learning theory is predominantly formulated under stationarity assumptions, wherein observations are drawn from a fixed probability measure. In many practical settings, however, the data-generating process evolves over time, inducing non-stationarity that fundamentally alters the limits of learnability. This paper presents an information-theoretic review of machine learning under non-stationary conditions, examining how temporal variation in the underlying distribution constrains achievable performance. We consider stochastic processes with time-indexed probability measures and analyse learning objectives in terms of excess risk, dynamic regret, and stability under distributional drift. Existing results are synthesised using information-theoretic quantities - including entropy rate, mutual information, and Kullback–Leibler divergence - to characterise how rates of change in the source process bound adaptation speed and generalisation accuracy. Across supervised learning, bandit models, and reinforcement learning, we highlight common structural dependencies between drift magnitude, information availability, and attainable error guarantees. Rather than introducing new bounds, the paper consolidates theoretical insights that reveal why persistent non-stationarity imposes irreducible performance degradation. This formulation provides a unified mathematical perspective on learning in time-varying environments and motivates the development of adaptive algorithms whose guarantees are explicitly parameterised by information-theoretic measures of change.

Keywords: Machine Learning, Non- Stationary Environments, Information Theory, Continual Learning, Reinforcement Learning, Dynamic Forecasting, Adaptive Systems, Topological Clustering, Model Falsification, Predictive Coding.

1.0 INTRODUCTION


Machine learning (ML) has revolutionized many fields by making systems able to identify trends in data, come up with predictions and adapt to new conditions. However, in the face of non-stationary settings, where distributions of data, or system dynamics, change over time, more traditional models based on ML face significant difficulties. The information theory provides an effective paradigm of capturing the core constraints of what can be absorbed and what cannot be absorbed by ML in such dynamic environments. This introduction looks at the information-theoretic restrictions that can be applied to ML in non-stationary conditions, the constraints that are caused by dynamical changes, the complications of constant learning, and solutions that can be used to overcome the complications.

1.1 The Problem of Non-stationary Environment

The non-static nature of the environment is characteristic of the dynamic nature of the data, i. e. a relationship between variables varies. This continues to change, bringing several complexities during the process of applying ML models. With changing environment, the models developed using historic data could be outdated which will

prevent generalization and the ability to adapt to new trends. Musaev et al. (2025) emphasize the significance of using metrics-based forecasting in non-stationary environments, which presupposes the use of complex metrics to trace dynamic changes in system dynamics. The rationale under the method is that models that can dynamically adapt to changing conditions are required without depending on past data alone which is an uphill task in predicting non-stationary processes (Musaev et al., 2025).

Graph 1: Non-Stationary vs. Stationary Processes



As shown in the graph above, stationary processes maintain a consistent distribution over time, while non-stationary processes experience significant variations in data distribution, making prediction and adaptation more challenging."

1.2 Never-Ending Learning and Constrained Computation

The main difficulty of non-stationary settings is lifelong learning whereby models need to keep changing with time without forgetting what has already been learned. According to Kumar et al, (2025), there is a correlation between continual learning and computationally constrained reinforcement learning. The authors state that the learning systems should retain the possibility to refresh their knowledge with new information without losing their efficiency in the resource-constrained environments. This poses a computational problem, in that models have to trade off between model complexity, computational need, and real time adaptation. The approaches of lifelong learning, like those examined by Kumar et al. (2025) are the strategies to deal with this problem, focusing on the reinforcement learning algorithms that are capable of handling non-stationary conditions and reduce catastrophic forgetting.

1.3 Falsehood and Experimental Design

In non-stationary learning Model falsification is a vital aspect of learning in which models are evaluated and validated by locating scenarios that give rise to inaccurate predictions. Murari et al. (2019) uses model falsification to the experimental design in non-stationary settings. They suggest that model validation based on systematic testing on real-world data can detect and use corrections on model assumptions to improve predictive behaviour. The methodology is required in applications where the models are robust enough to change dynamically with the behavior of the systems, especially in robotics and autonomous systems (Murari et al., 2019).

1.4 Adaptability in Information Gathering in Robotics

In robotic tasks, adaptive information collection can be taken as a fundamental learning approach to non-stationary settings. Along with non-stationary Gaussian processes, Chen et al. (2023) present a framework of an adaptive robotic information gathering. This form of methodology will allow the robots to revise their informational models based on the varying environmental conditions thus increasing flexibility and effectiveness in activities like exploration and decision-making in unpredictable environments. With the help of the Gaussian processes, robots can be used to model the uncertainty in the environment and constantly update the strategy as they receive new information (Chen et al., 2023).

1.5 The Information Theory and Topological Clustering Role

Topological clustering is another strategy, which has the capability of tackling the challenge of non-stationary data. Masuyama et al. (2019) explore the field of topological clustering based on the principles of adaptive resonance theory (ART) and information-theoretic learning. Their results are used to show how this approach can adapt dynamically to changes in the distribution of data over time, and therefore address concept drift. Combining ART with information-theoretic learning, this method provides a potent process of grouping time-varying data, which is indispensable in terms of pattern recognition and recognition of anomalies in dynamic settings (Masuyama et al., 2019).

1.6 Information Theoretic Measures of Model Optimization

The information theory is crucial in achieving the optimization of ML models to non-stationary environments. The idea suggested by Hilbert (2017) is that the principles of information theory can guide the development of informational-age expanding and adaptable models. Through the emphasis on the amount of information that a system is capable of handling and learning, these principles contribute to the definition of the boundaries of the performance that is possible under the influence of ML. This view is crucial to understand the limitations in non-stationary situations since it provides a framework of analyzing how learning systems may evolve over a period of time (Hilbert, 2017).

1.7 Generalization in the evolving domain

Xie et al. (2024) offer the ways to increase domain generalization with the help of dynamic latent representations. This solution focuses on making ML models more adaptable to new and unknown environments by focusing on their ability to generalize across a variety of, and changing, domains. The authors stress the significance of dynamic latent space, that enables a model to be sensitive and responsive to a change in data distributions over time. Such an approach is especially relevant in non-stationary settings, in which the ability to apply to new circumstances is essential to maintain model functionality (Xie et al., 2024).

1.8 Predictive Coding; the Slowness Principle

Another information-theoretic view of the dynamics of learning in non-stationary environments is the idea of predictive coding, which was explored by Creutzig and Sprekeler (2008). The idea here is that predictive coding postulates that both naturalistic and artificial systems enhance performance by reducing prediction error between anticipated and observed data on senses. This principle is in non-stationary situations where slow learning in the form of gradual modification of predictions in a system over time to prevent the traps of rapid, unsteady adjusting (Creutzig and Sprekeler, 2008).

Non-stationary machine learning also has theoretical and practical problems, especially on the interpretation of information-theoretic limits in time. Since there is continuous learning to adaptive clustering and predictive coding, various methods are available in facing these challenges. However, more research should be done to perfect models that can learn effectively in such a dynamic environment. The presented work represents the basis of realizing the inherent limitations to the learning and adaptation abilities of AI, and in such a way, research on information-theoretic methods of learners to ML in non-stationary conditions continues to be essential.

Table 1: Key Features of Non-Stationary Processes and Machine Learning Challenges

Feature	Stationary Processes	Non-Stationary Processes	Machine Learning Challenge in Non-Stationary Environments
Data Distribution	Invariant over time	Varies over time	Difficult to adjust models to changing data distributions
Predictability	High predictability of data	Low predictability of data due to dynamic changes	Models face challenges in predicting future states due to dynamic data
Adaptation to Changes	Low adaptability needed	Continuous adaptation required	Models must adapt continuously while maintaining previous knowledge
Learning Stability	Stable over time	Highly variable, time-dependent	Requires constant learning to avoid catastrophic forgetting
Computational Resources	Adequate resources for training	High computational resources needed for on-the-fly adaptation	Resource constraints hinder real-time learning and adaptation

Correlation of Data Points	Strong correlation between data points	Weak or evolving correlations	Models must handle evolving correlations between data points
Examples	Controlled static environments (e.g., controlled systems)	Real-world dynamic environments (e.g., financial markets, robotics)	Complex, dynamic environments where training data are not static

2.0 LITERATURE REVIEW

2.1 The Non-Stationary Environments and the Problems of the Machine Learning

Non-stationary environments are characterized by unsteady data distributions, which are time-dependent, and therefore can create a big barrier to ML systems. Such settings experience changes in the statistical characteristics of data with time, which makes traditional models of such changes based on the assumption of stationarity less useful. Musaev et al. (2025) suggest a metric-based forecast system of non-stationary process dynamics, which is a new forecasting metric able to adjust to the changing environment circumstances. Their work throws light on the need of considering time-effects in the process of forecasting in non-stationary conditions.

2.2 On-Going Learning in Non-Stationary Place

Constant learning has become the key to overcoming the constraints of non-stationary environments. Kumar et al. (2025) study the interaction between the continual learning and reinforcement learning in computational limited settings. They highlight that learning systems should be able to change with new data distributions they usually come across with as time passes without forgetting the previously learned information which is a phenomenon referred to as catastrophic forgetting. Their results indicate that reinforcement learning, when computational constraints are imposed on it, is forced to come up with strategies that are responsive to the continual changes in the data.

2.3 Non-stationary environment model falsification

The model falsification proposed by Murari et al. (2019) is an experimental design technique of a non-stationary environment. They also talk about the way that the methodology based on falsification can be used to optimize models according to changing environmental dynamics. This adaptation mechanism is important in dealing with real world data which is unpredictable and changing in nature. The ability to falsify models and continually improve the models is what makes the ML models both robust and flexible over time, despite changing data distributions.

2.4 Dynamic Reinforcement Learning as Self Supervised Learning

Schmidhuber (1990) suggested the use of fully recurrent neural networks (RNNs) to self-supervise such networks to overcome the non-stationary environment. His work formed the basis of the dynamic reinforcement learning systems with the ability to plan and adapt to the constantly changing data. This method will allow a continuous learning process to take place, as recurrent neural networks will allow models to make predictions and decisions using past and current data and adjust to changing situations.

2.5 Adaptive Information Gathering of Robot Systems

Chen et al. (2023) discuss non-stationary Gaussian processes to use in the adaptive robotic information collection. By doing so, robots can learn and change with time, successfully gathering the information, adjusting the strategy, and making decisions within dynamic settings. It is especially essential when it needs to be used in an environment that changes continuously, including autonomous vehicles and industrial robots.

2.6 Topological Clustering and Information theoretic Learning

Masuyama et al. (2019) investigate the topological clustering, presented by the adaptive resonance theory (ART) in the context of information-theoretic learning. Their approach includes the clustering of data which is responsive to changing nature of non stationary environment. Using the combination of ART and information-theoretic learning, the approach manages to update the clusters dynamically over time as new information is presented, which is an essential aspect in the operation of the ML systems that are under the non-stationary conditions.

2.7 Theoretical Informatics of Growth and Learning

Hilbert (2017) suggests that information-theoretic approach to the growth in the information age. This strategy drives the value of information as a learning and decision-making tool in changing environments. Hilbert states that the more information an AI system can process and use, the more adaptable it is and the more it can grow according to the alterations in data. This information is especially applicable in non-stationary conditions, where the inflow of data is continuous and requires effective extraction of information.

2.8 Domain Generalization with Dynamic Latent Representation

Xie et al. (2024) talk about improving the domain generalization through the dynamic latent representations. The representations enable models to make inferences in different environments where dynamic and non-fixed data distributions are learned. The method is particularly effective in non-stationary conditions, where the distributions of data change, and it requires the flexibility of the model.

2.9 Deployment-efficient Reinforcement Learning

In their study, Huang et al. (2022) explore the issue of reinforcement learning deployment efficiency in non-stationary conditions with lower bounds and optimality. According to their work, these systems have to be run within some theoretical limits to be efficient under the real world usage, but their limitations due to the non-stationary character of the data. To achieve success in deployment in the dynamic environment, optimization of learning algorithms in this range is essential.

2.10 Predictive Coding and Slowness Principle in Non-stationary Environment

Creutzig and Sprikeler (2008) use an information-theoretic method of learning in non-stationary environments, namely, predictive coding and the slowness principle. The

slowness principle suggests paying attention to the slowly changing variables to be able to model the dynamic environment. Through predictive coding, systems develop forward states based on the previous knowledge and hence, more apt to new information and stable in the face of environmental fluctuation.

2.11 Information Pruning to Dynamic Learning

Anagnostopoulos and Gramacy (2013) discuss the issue of information-theoretic data discarding of dynamic data streams. They also talk about dealing with massive amounts of incoming data that might not be of use in the decision-making that is taking place. The selective discarding of data of less importance enables the models to focus on relevant data which enhances performance when operating within non-stationary environments where the relevance of certain data varies with time.

2.12 Non-Stationary Bandits of Reinforcement Learning

Liu et al. (2023) introduce a concept of non-stationary bandits that is used to solve the decision-making task when the distribution of rewards varies with time. Whereas in the multi-armed bandits model, the rewards are constant, in the non-stationary bandits, the rewards are changing at any given time which is a challenge to adaptive decision-making. They give the basis of their work to reinforcement learning algorithms that can dynamically adapt to non-stationary rewards.

3.0 METHODOLOGY

The study examines the information-theoretic ingredients of the limitations of ML algorithms under non-stationary conditions. The main aim is to identify the impact of these limitations on the efficiency of learning and to offer the frameworks of learning to understand the ML limitations in dynamic environments. These constraints are systematically analyzed and modeled using the methodology.

3.1 Problem Formulation and Conceptual Framework

The dynamics of underlying processes in the non-stationary environment depend with time and hence making predictive modeling difficult. To address the problem, the methodology starts with the formalisation of the problem as a non-stationary dynamic forecasting task based on the metric-based ML framework of Musaev, Makshanov, and Grigoriev (2025). This includes the task of characterising the environment as dynamically changing metrics which are updated in real time to capture such dynamics, and offer the advantage of enabling ML models to respond well to changes in time.

3.2 Reinforcement Learning Continuous Learning

The key to dealing with non-stationary environments is reinforcement learning (RL). The approach embraces the approach of Kumar et al. (2025) in which continuous learning is considered computational constrained RL. This ensures that agents are able to keep on learning based on the new experiences and are able to effectively deal with the limitation of resources. It is centered on exploration and exploitation balancing algorithms and allows RL agents to evolve over a long time.

Model Falsification and Experimental Design

Model falsification is an important tool used to measure and improve the performance of the learning systems within dynamic environments. The approach of falsification of hypotheses on the environmental behavior proposed by Murari et al. (2019) is introduced, according to which the hypotheses are consistently tested on the real information. This makes sure that ML models are rooted in reality, which ensure that models can be continuously tuned on the basis of real-time feedback, which is an essential aspect in the design of experiments in non-stationary conditions.

3.2.1 Simulation setup and data generation (RL)

To provide empirical evidence for reinforcement learning behavior under changing conditions, reward–learning curves were generated using a Q-learning agent in a standard control benchmark (CartPole). The environment dynamics were treated as a controlled proxy for non-stationary decision-making, and the agent was trained across repeated episodes using an ϵ -greedy exploration policy. At each episode, the cumulative reward was recorded and aggregated across training to produce a learning curve (total reward vs. episode index). This dataset was used to visualize how policy performance evolved over time under continual interaction and adaptation constraints.

Graph 2




Figure 2: Reinforcement learning performance over time (Q-learning, CartPole).

Total episodic reward across training episodes, illustrating learning progression and stabilization as the agent adapts its policy through reward feedback.

3.3 Topographical Clustering of Data Representation

Topological techniques provide the structure of non-stationary data. The topology used in the methodology is the topological clustering of Msuyama et al. (2019) through adaptive resonance theory produced in the information theoretical framework based on entropy measures of cluster relevance. Such clustering makes it easy to manage the increasing amounts of data since it splits the data into meaningful, independently processable windows.

3.4 Latent Variables Dynamic Representations

A modern system of dynamic learning deals with dynamic latent representations, which are explained by Xie et al. (2024). This approach builds a framework which includes latent variables that have the ability to capture the features of data which change with time hence flexibility in representation. Such latent variables are developed, which enables the model to project on changing realms and reduce retraining needs.

Table 2: Review of Machine Learning Methods for Non-Stationary Environments

Technique	Description	Benefits	Difficulties	Reference
Continual Learning	A paradigm of training that builds models continually with incoming data streams, avoiding catastrophic forgetting.	Supports smooth adaptation to changing data distributions in real-time applications.	Requires advanced regularization and memory management to prevent harmful interference between learned concepts.	Kumar et al., 2025
Reinforcement Learning (RL)	A dynamic decision-making model where agents learn from actions and reward feedback in changing environments.	Useful for adaptive policies in real-time strategies and robotics.	High computational cost; slow convergence when environmental statistics change rapidly.	Schmidhuber, 1990
Model Falsification	The process of identifying and removing models that perform poorly in varying environmental conditions.	Enhances reliability by ensuring only empirically validated models are used.	Risk of losing potentially valuable predictive capacity; requires significant computational resources.	Murari et al., 2019
Adaptive Resonance Theory (ART)	A neural network that uses competitive learning to stabilize clusters while preserving the plasticity-stability balance.	Effective for data streams where temporal development doesn't impact cluster integrity.	Poor scalability with high-dimensional or large datasets; requires frequent parameter updates.	Masuyama et al., 2019
Gaussian Processes (GPs)	A Bayesian non-parametric regression model used to model functions as a distribution of latent variables, adaptable to non-stationary effects.	Provides a robust framework for lifelong learning and quantifying uncertainties.	Computational complexity increases cubically with dataset size, making it difficult for large-scale applications.	Chen et al., 2023

Topological Clustering	A clustering algorithm based on algebraic topology, clustering data by inherent topological properties like persistence diagrams.	Well-suited for dynamic distributional structures and slow changes in cluster topology.	High computational cost; scalability issues with large datasets.	Masuyama et al., 2019
Dynamic Latent Representations	A learning strategy that adapts time-varying latent representations for long-term prediction across changing distributions.	Improves domain generalization by aligning latent space with changing environments.	Careful tuning of latent dimensionality and structure is needed to prevent overfitting and under-representation.	Xie et al., 2024
Adaptive Robotic Information Gathering	Autonomous robots use probabilistic models (e.g., Gaussian processes) to iteratively improve their perceptions and make context-sensitive decisions in dynamic environments.	Promotes effective information gathering and flexible discovery in changing environments.	Requires real-time inference processes and continuous learning to maintain performance in highly dynamic environments.	Chen et al., 2023

3.5 Deployment-Efficiency-Constrained Reinforcement Learning

Reinforcement learning when applied to non-stationary environments requires efficiency constraints to be strictly taken into account. The current research employs the methodological framework utilized in Huang et al. (2022) to derive lower limits on the efficiency of reinforcement-learning deployment, hence making sure that the algorithms make their best use of accessible resources as dictated. This involves reduction of computing load and memory use as well as achieving proper real time predictive performance.

The optimality conditions are defined formally to establish the minimum resource allocation to be able to maintain high performance systems at varying environmental conditions. This methodological feature is important to the fact that the machine-learning models can be scaled to real-life application, where both computational and memory resources are often limited (Huang et al., 2022).

3.6 Information-Theoretic Data Discarding

Due to the fact that data streams in non-stationary systems are usually noisy and outdated, an information-theoretic method of data discarding is used, in the direction of Anagnostopoulos and Gramacy (2013). The methodology is based on entropy-based measurements to select and eliminate the irrelevant or redundant data, and therefore only the most informative aspects can be learnt. The strategy makes the computation process less burdensome and increases the overall efficiency of the learning process. The model chooses which data to discard (selectively), therefore focusing on assimilating the most informative experiences, and it has been shown to be quicker and more effective in generalisation in non-stationary environments.

4.0 RESULT

4.1 Essential Limitations to Learning Non-Stationary Processes

Non-stationary processes are challenging to learn since the distribution of data changes as time progresses. Musaev, Makshanov, and Grigoriev (2025) explored the application of metric-based machine-learning forecasting to overcome those issues in the dynamic process dynamics. Their work determines the main limitations that restrict the accuracy and efficiency of models in the environment where there are time-dependent changes in distributions. This is because the failure of traditional models to change without retraining or recalibration is a source of fundamental bottleneck, highlighting the challenge of having the best predictions in environments where process dynamics constantly change.

4.2 Constant Learning and Computational Constraints

The article by Kumar et al. (2025) focuses on the analysis of continual learning as a computationally constrained problem of reinforcement-learning. Their reasoning is that, as much as continuous learning is a promising way to get out of the non-stationary conditions, it is constrained by computational capacity and the complexity of the model. Their results show that a system facing new information has to trade-off between stability and plasticity as it has to maintain previous knowledge and also integrate new information. Such balance is critical to those systems that work in non-stationary conditions, and quick adaptations as well as long-term learning are demanded.




Fig 1: Balancing Stability and Plasticity in Continual Learning

4.3 Model Falsification in Non-stationary Environments

Murari et al. (2019) presented a model falsification strategy of learning in non-stationary conditions. Their approach uses falsification as a methodology to design experiments, thus resolving the issue of measuring the performance of the model when assumptions on the state of the system are unlikely to hold. This method can be used to generate new models which are aligned to new data by using information-theoretic principles and rejecting models that do not fit new data in real-time. The technique plays a significant role in the robustness of models in environments where data distribution is changing on a continuous basis.

4.4 Dynamic Latent Representations and Domain Generalization

As suggested by Xie et al. (2024), dynamic latent representations are used to improve evolving domain generalisation. When applied in the non-stationary environment, they claim that the concept of static representations tend to be ill-adapted to changing data properties which leads to poor performance.

Through their studies, they have proved that the addition of dynamic representations enables the models to generalise over changing domains and, in the process, alleviate distribution shift. It is also in this strategy where the points are made that models in non-stationary environments need to be able to constantly change their internal representations, so as to enhance performance.

4.5 Reinforcement Learning Under Non-Stationary Environment: Lower Bound and Optimality

The article by Huang et al. (2022) is about the implementation of reinforcement learning into the non-stationary environment and the lower limits and optimality of that environment. They emphasize that reinforcement-learning agents can in some cases not achieve optimal performance when facing non-stationary dynamics, unless some special conditions, including the need to have effective exploration-exploitation trade-offs, are in place. According to their research, knowledge of these lower bounds can be used to develop more efficient algorithms to be used in real-time in non-stationary environments, especially in relation to computational efficiency and convergence rates.

4.6 Robotic Information Gathering with Adaptive Gaussian Processes

Chen, Kharden, and Liu (2023) focus on the adaptive robotic information collection of non-stationary Gaussian process. Their work describes the adaptation of Gaussian processes to learning in the case where the underlying data-generation process is non-static. The model is more effective in decision-making within robots by evolving the dynamics through time-varying the kernel functions to improve the decision-making process. This method shows potential in enhancing flexibility and resiliency of robotic systems in the real world where the environmental factors continue to vary.

4.7 Data discarding based on Information theoretic data discarding of dynamic trees

Anagnostopoulos and Gramacy (2013) suggest an information-theoretic data discarding algorithm in dynamic trees of data streams. This approach is specifically relevant to those systems that need to keep decision trees, but that have to change in response to non-stationary data sources.

The approach is able to cut down the information that is not so important as time goes by thus simplifying the model and making sure that the most informative data are being retained to make a decision. This information-shedding plan is essential in dealing with the informational clutter that often comes along with non-stationary setups where the relevance of information changes with the coming of new information.

4.8 Information-Theoretic Dynamic Decision-Making

Schmidhuber (1990) and Creutzig and Sprekeler (2008) give the background theory of self-supervised learning and predictive coding in non-stationary environment. The article of Schmidhuber referencing recurrent neural networks in studying dynamic reinforcement learning highlights the possibility of such models to evolve over time by using the previous experiences to make decisions in the future. The information-theoretic predictive coding by Creutzig and Sprekeler strengthens the arguments of ensuring a minimal amount of surprise or uncertainty during decision-making. The two theories highlight that models are needed to effectively update internal states so as to generate trustworthy forecasts in non-stationary situations.

5.0 DISCUSSION

This part investigates the fundamental issues and knowledge connected with information-theoretic constraints of machine learning in non-stationary contexts. These environments are dynamic and change with time, putting in place basic constraints to AI abilities. In this case, major points like continuous learning, evolutionary prediction, adaptation of model, and strengthening gain are examined in the framework of the non-stationary environment.

5.1 Recurrent Education and Resource Limitation

In non-stationary conditions, continuous learning is a key challenge, and such models need to be able to adapt to changes without forgetting the previous knowledge. Kumar et al. (2025) address computational problems that occur in these cases, especially in the framework of reinforcement-learning. The adaptive capacity of the model with time may be hindered by the limited memory and processing power. Their work draws attention to the fact that as the environment evolves, the ability to retain past information and adjust it with new potentially conflicting data is becoming more complicated that there are fundamental constraints on accuracy and computational efficiency.

5.2 Adaptive Models and Metric Based Forecasting

The paper by Musaev et al. (2025) suggests metric-based forecasting non-stationary process dynamics, using dynamic metrics as the model of changing nature. Nonetheless, non stationary, time varying patterns, which appear in non-stationary environments, are frequently not reflected in traditional measures, especially in weather forecasting or long term stock market projections where external conditions change in an unpredictable manner. Therefore, although metric-based forecasting provides practical estimates, it is bound to be less accurate in the prediction when the underlying systems vary quickly/drastically or in a random manner.

5.3 Theoretical falsification of non-stationary environments

Murari et al. (2019) explore the concept of model falsification as one of the methods of dealing with non-stationary environments and focus on the use of this tool with the help of experimental design. They are concerned with determining when the assumptions of a model become invalid due to the changing environment. Even though we can recalibrate

models in real time with the help of falsification, it shows that there is no guarantee that a model will be perfectly accurate in the long run, particularly when other factors are not included in the models. The unpredictability that is inevitable due to non-stationary conditions poses a structural challenge in the establishment of models that can be optimally adapted.

5.4 The Limitations of Exploration and Reinforcement Learning

Non-stationary conditions in reinforcement learning make the exploration exploitation dilemma worse. Huang et al. (2022) discuss the deployment efficiency, which defines lower limits of performance in dynamic situations. The non-stationary environments cause the disruption of the state-action mappings, often changing the optimal policies and leading to the necessity to explore. However, more exploration needs more computational resources which does not necessarily provide valuable information when the environment is unpredictable which places substantial information-theoretic constraints on learning.

5.5 Information -Theory and Topological Clustering

Masuyama et al. (2019) discuss the adaptive resonance theory of topological clustering, which involves the utilization of information theoretic learning strategies on dynamically changing streams of data. Their publication demonstrates the ability of clustering to handle changes in the environment, and however, there are always limits to such adaptation when the distribution moves too fast relative to the adaptability of the algorithm. The information-theoretic methods, promising as they are, are limited by entropy in more and more complicated data spaces, and the resulting computation requirements are growing exponentially, and may push whole computational systems to unachievable limits.

5.6 Dynamic Representations of Domain generalisation

This method is better in enhancing flexibility, but at the same time, it reveals the boundaries of generalisation. Adaptation of latent representations to different environments makes it harder to be robust in many environments, as the more data changes, the harder it is to balance flexibility and stability.

5.7 Predictive Coding and the Principle of Slowness

Creutzig and Sprekeler (2008) apply the meaning of predictive coding and slowness principle to dynamic systems. Such principles are useful in making predictions using previous experience but when rapid changes take place, the predictability of changes is reduced thereby nullifying the hypothesis of gradual evolution. The principle of slowness comes to grief in unstable systems (financial markets or climate dynamics) the predictive errors become large in such systems. Information-theoretic method to the interpretation of machine learning in non-stationary conditions reveals the inherent constraints of the dynamic conditions. Although the sorts of strategies that could be identified include the continuous learning process, the use of metrics to make predictions, the falsification of models and adaptive clustering, they are limited by the continued complexities that are being experienced in the target environments. These limitations are due to fundamental

information-theoretic concepts that define the trade-off between the computational efficiency, flexibility and predictive power. Future studies need to resolve such limits in order to come up with techniques that are more consistent with the dynamics of the real world that are not stationary.

6.0 CONCLUSION

The exploration of information theoretic limits of machine learning in non-stationary processes explains the inherent limitations of AI systems in further training. With the changing field of machine-learning methods, it is important to understand such shortcomings to create more efficient and adaptable systems that can be applied in the dynamic setting.

6.1 Limitations of Non-stationary Processes

Machine-learning models, particularly models that are intended to predict non-stationary processes, are faced with notable difficulties in predicting and adjusting to variable dynamics. Musaev et al. (2025) illustrate the challenges of the metric-based forecasting approaches that have limitations posed by the requirement to keep up with the changes in the environment. These issues render the need of models that can dynamically update and adapt to changing streams of data.

6.2 On-Going learning and Re-enforcement Learning

It is the constant learning that can help reduce limitations posed by non-stationary environments. Kumar et al. (2025) investigate the interaction of the process of continuous learning with computationally limited reinforcement learning, which makes the process of learning in constantly dynamic environments challenging. Such limitations are dictated by the need to ensure that the models can absorb new information without forgetting the already obtained information and this requires complex architectures that will enable the models to maintain stability as well as flexibility.

6.3 Model Adaptive Techniques and Falsification

In non-stationary settings, model falsification is a crucial way of improving the predictability of learning algorithms. Murari et al. (2019) present a model falsification of experimental design, which aims at invalidating inappropriate models in real-time, thus making the systems resistant to instability in dynamic situations. Such practice is an essential part of AI systems that are to be used in changing environments when current models are not sufficient.

6.4 Reinforcement Learning in Area based Dynamics

The theme of sustaining optimal decision-making through time in non-stationary environments highlighted by Schmidhuber (1990) in his work on recurrent neural networks in reinforcement learning makes it difficult to compare current decision-making with that of the past. Though self-supervised RNNs have shown potential to react to changing conditions, the efficiency of self-supervised RNNs is restricted by the computational needs of the process of constantly learning on time-varying data.

6.5 The Adaptive Robotic Systems and Gaussian Processes

Chen et al. (2023) use non-stationary Gaussian processes to adaptive robotic information obstruction, which gives an indication of how robots can fit in a changing setting. Their study demonstrates the power of Gaussian processes in non-stationary data modelling in addition to weakness of their complex processes as well as the necessity of real-time updates.

6.6 Information Theory and Topological Clustering

Masuyama et al. (2019) suggest that one of the solutions to the problem of non-stationary environments is offered by the topological clustering. Their approach is more efficient in processing and classifying the changing data as they dynamically respond to the changing data environments by using adaptive resonance theory along with information-theoretic learning. However, its real-world implementation is still limited considering it requires large computational capabilities and memory.

Domain Generalisation and Dynamic Latent Representations Domain generalisation is a technique used in the field of neural networks to generalise the knowledge acquired during training to new domains (similar to how the neural network was trained).

Xie et al. (2024) present a new strategy based on the improvement of domain generalisation with the help of dynamic latent representations, which is significant in terms of models that should function in non-stationary conditions. They note in their study that dynamic representations can be used to enhance generalisation in various settings, and that it is also challenging to fine-tune representations in the real world.

6.7 Reinforcement Learning Lower Bounds and Optimality

Huang et al. (2022) offer useful information on the lower bounds and optimality of reinforcement learning in the non-stationary environment. They provide a guideline on understanding the limits of reinforcement learning in real life applications especially in deployment efficiency by defining these boundaries. Nevertheless, even real-world implementation is limited by such lower limits and additional algorithm development is required.

6.8 Data Discarding Information-Theoretic

Anagnostopoulos and Gramacy (2013) investigate the information-theoretic data discarding to simplify processing in dynamic situations. Although eliminating irrelevant information can be very efficient in the process of learning, it is also questionable due to loss of information that might be very important in offering long-term adaptation.

6.9 Theoretical and Real- world Issues of Non-stationary Environments

The analysis of non-stationary setting in terms of information-theoretic perspective provides informative theoretical understanding of what can be done within the boundaries of machine-learning. Yet, practical limitations tend to face the theoretical models in reality as studies by Masuyama et al. (2019) and Ji et al. (2025) both indicate. The ongoing

adaptation and the stability of the models are still one of the issues requiring new solutions.

The specified information-theoretic limitations demonstrate how complicated the implementation of machine-learning systems in non stationary setting can be. Although the model falsification, continual learning, and adaptive techniques do provide the possible solutions, their practical usage is still limited by the computational efforts, along with constant adaptation requirements. The next step in research should be to improve adaptability and efficiency of learning algorithms in order to adapt to the dynamics of real world data. Increased capabilities of machine learning in overcoming the challenges of non-stationary environments can be reached through the development of information-theoretic research into the field, which will contribute to the development of the boundaries of the AI capability over time.

References

- 1) Musaev, A. Makshanov and D. Grigoriev, "Metric-Based Machine-Learning Forecast of Non-Stationary Process Dynamics," 2025 International Russian Automation Conference (RusAutoCon), Sochi, Russian Federation, 2025, pp. 320-325, **DOI:** 10.1109/RusAutoCon65989.2025.11177378
- 2) Kumar, S., Marklund, H., Rao, A., Zhu, Y., Jun Jeon, H., Yueyang, L., & Van Roy, B. (2025). Continual learning as computationally constrained reinforcement learning. *Foundations and Trends in Machine Learning*, 18(5), 913-1053. <https://doi.org/10.1561/2200000116>
- 3) Murari, A., Lungaroni, M., Peluso, E. et al. A Model Falsification Approach to Learning in Non-Stationary Environments for Experimental Design. *Sci Rep* 9, 17880 (2019). <https://doi.org/10.1038/s41598-019-54145-7>
- 4) Schmidhuber, J. (1990). *Making the world differentiable: on using self supervised fully recurrent neural networks for dynamic reinforcement learning and planning in non-stationary environments* (Vol. 126). Inst. für Informatik. <https://people.idsia.ch/~juergen/FKI-126-90ocr.pdf>
- 5) Chen W, Khardon R, Liu L. Adaptive Robotic Information Gathering via non-stationary Gaussian processes. *The International Journal of Robotics Research*. 2023;43(4):405-436. doi:10.1177/02783649231184498
- 6) N. Masuyama, C. K. Loo, H. Ishibuchi, N. Kubota, Y. Nojima and Y. Liu, "Topological Clustering via Adaptive Resonance Theory With Information Theoretic Learning," in IEEE Access, vol. 7, pp. 76920-76936, 2019, doi: 10.1109/ACCESS.2019.2921832.
- 7) Hilbert, M. (2017). The more you know, the more you can grow: an information theoretic approach to growth in the information age. *Entropy*, 19(2), 82. <https://doi.org/10.3390/e19020082>
- 8) Xie, B., Chen, Y., Wang, J., Zhou, K., Han, B., Meng, W., & Cheng, J. (2024, March). Enhancing evolving domain generalization through dynamic latent representations. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 38, No. 14, pp. 16040-16048). **DOI:** <https://doi.org/10.1609/aaai.v38i14.29536>
- 9) Huang, J., Chen, J., Zhao, L., Qin, T., Jiang, N., & Liu, T. Y. (2022). Towards deployment-efficient reinforcement learning: Lower bound and optimality. *arXiv preprint arXiv:2202.06450*. <https://doi.org/10.48550/arXiv.2202.06450>
- 10) F. Creutzig and H. Sprekeler, "Predictive Coding and the Slowness Principle: An Information Theoretic Approach," in *Neural Computation*, vol. 20, no. 4, pp. 1026-1041, April 2008, doi: 10.1162/neco.2008.01-07-455.

- 11) Anagnostopoulos, C., & Gramacy, R. B. (2013). Information-theoretic data discarding for dynamic trees on data streams. *Entropy*, 15(12), 5510-5535. <https://doi.org/10.3390/e15125510>
- 12) Sujiono, Muhammad Ridha Akhsanu, Adaptive Optics for Spacetime: Breaking the Red Noise Floor to Reveal Quantum Gravity Signatures with Deep Reinforcement Learning (July 30, 2025). <https://ssrn.com/abstract=5985514>
- 13) Liu, Y., Kuang, X., & Van Roy, B. (2023). A definition of non-stationary bandits. *arXiv preprint arXiv:2302.12202*. <https://doi.org/10.48550/arXiv.2302.12202>
- 14) Saal, H. (2011). Information theoretic approach to tactile encoding and discrimination. <http://hdl.handle.net/1842/5737>
- 15) E. Ji, Y. Wang, S. Xing and J. Jin, "Hierarchical Reinforcement Learning for Energy-Efficient API Traffic Optimization in Large-Scale Advertising Systems," in IEEE Access, vol. 13, pp. 142493-142516, 2025, doi: 10.1109/ACCESS.2025.3598712.