

06-30-2025

Continual Learning Models in Aviation Systems

Nicki Barari
Drexel University

Ghazal Barari
Embry-Riddle Aeronautical University

This paper explores the vital role of continual lifelong learning in advancing machine learning applications within the aviation industry. Through case studies on predictive maintenance and adaptive flight routing, we demonstrate how human-inspired continual learning enables AI systems to adapt incrementally to evolving conditions while preserving critical prior knowledge. This approach addresses fundamental challenges in dynamic, safety-critical aviation environments, promising improved adaptability, safety, and operational efficiency. The discussion highlights benefits, challenges, current progress, and future directions toward building resilient, intelligent aviation AI systems.

Recommended Citation:

Barari, N. & Barari, G. (2025). Continual Learning Models in Aviation Systems. *Collegiate Aviation Review International*, 43(1), 132-148. Retrieved from <https://ojs.library.okstate.edu/osu/index.php/CARI/article/view/10310/9098>

Introduction

In safety-critical industries such as aviation, machine learning (ML) systems are increasingly being adopted to support and automate complex tasks, from predictive maintenance and real-time flight routing to intelligent crew scheduling and autonomous systems. While conventional ML models have achieved impressive performance under fixed conditions, they typically suffer from a significant limitation: the inability to learn continuously after deployment. Once trained, these systems are rarely designed to adapt to new tasks or data distributions without retraining from scratch. This static learning paradigm is ill-suited to the dynamic nature of the aviation domain, where regulatory frameworks, environmental variables, operational procedures, and technological platforms are in constant flux.

To address this challenge, researchers are turning to human-inspired learning paradigms, particularly continual lifelong learning (Parisi et al., 2019). This approach seeks to equip artificial agents with the ability to learn incrementally, adapt to novel tasks, and refine prior knowledge without experiencing catastrophic forgetting (Barari & Kim, 2021). Continual lifelong learning reflects a key trait of human cognition: the ability to accumulate experience over a lifetime while flexibly incorporating new information (Barari, Lian, & MacLellan, 2024a). In aviation, this capability is essential for ensuring that intelligent systems remain effective as they encounter new aircraft platforms, operational theaters, or emergent failure modes (Barari, Lian, & MacLellan, 2024b).

Central to the design of continual learning systems is the management of the plasticity-stability trade-off. On one hand, plasticity allows a model to integrate new information; on the other, stability ensures that valuable prior knowledge is retained. A system biased too heavily toward plasticity risks overwriting past experience, potentially degrading performance in previously learned domains. Conversely, excessive stability may hinder the model's capacity to adapt, rendering it inflexible in the face of change.

Continual learning models can also contribute to enhanced learner engagement by enabling systems that adapt dynamically to evolving learner behavior and performance. As these models accumulate knowledge over time, they support more personalized and responsive educational experiences (Barari & Sanders, 2024).

To explore how continual learning can be practically implemented in aviation, we build on foundational theory and illustrate two aviation-specific case studies in the following sections. These case studies highlight how continual learning methods address critical challenges in dynamic, safety-critical environments.

Literature Review

The field of Continual Learning, also known as lifelong learning, has emerged in recent years as a response to the limitations of traditional machine learning models, which are generally trained offline on fixed datasets and assume stationary environments. In contrast, continual learning seeks to emulate the incremental and adaptive nature of human cognition, allowing models to learn from a sequence of tasks without experiencing catastrophic forgetting, a

phenomenon in which new learning overwrites previously acquired knowledge (McCloskey & Cohen, 1989).

The issue of catastrophic forgetting was identified early in connectionist models and remains a core challenge in neural networks trained sequentially (Czigler & Winkler, 2010). A model optimized on new data tends to overwrite weights critical to past tasks unless explicitly constrained. Addressing this requires managing the stability-plasticity trade-off: stability ensures retention of old knowledge, while plasticity allows adaptation to new data. Several algorithmic families have emerged to handle this trade-off:

- Regularization-based methods (e.g., Elastic Weight Consolidation) penalize changes to important parameters (Kirkpatrick et al., 2017).
- Replay-based methods store or generate past data to interleave with new training (e.g., Deep Generative Replay, (Shin et al., 2017)).
- Dynamic architectural methods like Progressive Neural Networks (Rusu et al., 2016) expand model capacity with task-specific modules.

Among these, Elastic Weight Consolidation (EWC) has been widely studied and adopted for its simplicity and biological plausibility. It approximates the Fisher Information Matrix (FIM) (Fisher, 1987) to estimate parameter importance and regularizes the loss function to resist changes to critical weights. This method is particularly suited for applications where task boundaries are clear and domain shifts are significant yet structured (Fisher, 1996).

Most continual learning research has focused on benchmark datasets (e.g., Permuted MNIST, Split CIFAR-100), which are designed as academic tools to evaluate a model's ability to avoid catastrophic forgetting in controlled settings. While these datasets are not representative of aviation data, they provide a standardized basis for testing continual learning algorithms before application in real-world domains such as aviation. The growing interest in applying these techniques to real-world, safety-critical domains is now evident, especially in robotics, autonomous vehicles, and aerospace. These fields demand adaptive systems that can update incrementally without complete retraining, while maintaining robustness and trustworthiness.

In the aviation domain, continual learning remains underexplored but highly promising. Aircraft systems evolve over time through new engine types, avionics upgrades, and environmental or regulatory shifts. Each change introduces new "tasks" to which intelligent systems must adapt. This is especially relevant for predictive maintenance, where engine performance and failure patterns differ between models but share underlying principles.

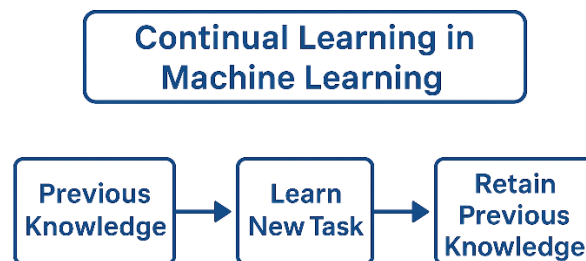
The following sections present aviation-specific examples that illustrate how continual learning can enhance reliability and adaptability in real-world applications. Each case explores how the plasticity-stability trade-off manifests and how it may be mitigated using mechanisms drawn from both neuroscience and contemporary machine learning research.

Theoretical Framework

This work adopts a conceptual case study methodology to explore how continual learning techniques, such as Elastic Weight Consolidation and experience replay, might be applied in real-world aviation scenarios. While we do not present empirical results or formal simulations, the use of theoretical case studies allows us to map established machine learning techniques onto aviation-specific challenges in a structured and illustrative manner. This approach enables early-stage exploration of feasibility, trade-offs, and potential impact. Further validation through implementation, simulation, or field studies will be essential to confirm these findings and guide operational adoption.

As mentioned, the concept of continual lifelong learning emerges from the need to construct machine learning systems that can learn in a manner similar to human cognition, incrementally, adaptively, and without forgetting prior knowledge. Unlike conventional models, which are trained on static datasets and deployed with fixed capabilities, continual learning systems evolve over time. They integrate new information from sequential tasks or data streams while preserving competencies from earlier experiences. This paradigm is essential for building intelligent systems capable of functioning in dynamic, real-world environments such as aviation. A simplified continual learning is illustrated in Figure 1.

Figure 1
Simplified Continual Learning



The theoretical foundations of continual learning are deeply rooted in the stability-plasticity dilemma. Plasticity refers to a model's capacity to acquire new knowledge, whereas stability denotes the ability to retain previously learned information. A system overly biased toward plasticity may suffer from catastrophic forgetting, where new learning disrupts older knowledge. Conversely, a system that is too stable becomes rigid and unable to adapt to new challenges. Effective continual learning involves balancing these forces, ensuring long-term retention of critical skills while enabling flexibility in acquiring new ones.

In the context of aviation, the theoretical constructs are not only intellectually compelling but also practically necessary. Aviation systems operate in non-stationary environments: aircraft fleets evolve, airspace regulations shift, mission parameters change, and sensor technologies advance. AI systems that support maintenance diagnostics, route planning, or autonomous control must therefore learn over time without degrading performance on previously mastered domains.

The objective of integrating continual learning into aviation applications is to enable robust, flexible, and trustworthy AI systems that improve with experience, mirroring the way human experts build domain knowledge. This theoretical framework informs the two case studies that follow. In each, we examine a specific aviation use case where continual learning is applied to real operational challenges:

- The first case explores how EWC can support multi-fleet engine diagnostics, allowing a model to learn about new aircraft engines while preserving knowledge of legacy platforms.
- The second case investigates adaptive flight routing, where a routing model learns to operate effectively across geographically and regulatorily distinct airspaces.

By grounding these applications in continual learning theory, we discuss how conceptual advances in machine learning can directly address pressing needs in safety-critical, data-rich aviation environments.

Case Study 1: Continual Learning in Predictive Maintenance for Multi-Fleet Engine Diagnostics

One of the most promising applications of continual lifelong learning in aviation lies in the domain of predictive maintenance, where machine learning models are used to detect early signs of failure in aircraft engines. Airlines today operate increasingly diverse fleets composed of both legacy and next-generation aircraft, often from multiple manufacturers. As a result, maintenance AI systems must diagnose failures across engine types that differ in design, behavior, and sensor profiles.

Consider a scenario in which an ML-based health monitoring system is initially trained on CFM56 engines used in Boeing 737 aircraft. These engines generate large volumes of telemetry data, including vibration levels, exhaust gas temperature (EGT), fan speeds, and pressure ratios, which are used to build models for anomaly detection and Remaining Useful Life (RUL) estimation. Over time, the airline introduced a new fleet of Airbus A320neo aircraft equipped with LEAP-1A engines. The monitoring system must now learn to interpret and diagnose faults in LEAP-1A engines while retaining its existing diagnostic capabilities for the CFM56 fleet.

This transition highlights the need for continual learning. A naive retraining approach, where the model is updated using only LEAP-1A data, risks catastrophic forgetting, the loss of performance on previously learned CFM56 diagnostics (Kirkpatrick et al., 2017; Goodfellow, Bengio, & Courville, 2016). Conversely, retaining the original model without any adaptation to the new engine data results in suboptimal or even erroneous predictions for the LEAP-1A fleet.

Application of Elastic Weight Consolidation to Multi-Fleet Engine Diagnostics

To address the catastrophic forgetting problem inherent in sequential learning tasks, Elastic Weight Consolidation (EWC) offers an elegant solution grounded in neuroscientific

principles. EWC allows the model to selectively remember important parameters learned from the original task while enabling flexibility to adapt to the new task.

The EWC approach modifies the training objective by introducing a regularization term that penalizes changes to model parameters deemed critical for prior tasks. This term is informed by the FIM, which quantifies the sensitivity of the loss function to each parameter. Parameters with a high FIM value are considered crucial for retaining previous knowledge, and thus deviations from their learned values are heavily penalized.

Formally, the EWC loss function is expressed as:

$$\mathcal{L}(\theta) = \mathcal{L}_{new}(\theta) + \sum_i \frac{\lambda}{2} F_i (\theta_i - \theta_i^*)^2$$


Where:

- $\mathcal{L}_{new}(\theta)$ represents the loss computed on the new dataset, here, the LEAP-1A engine telemetry.
- θ denotes the current parameter set of the diagnostic model.
- θ^* is the parameter set optimized on the prior task, specifically the CFM56 engine data.
- F_i are the diagonal elements of the FIM estimating the importance of each parameter θ_i for the CFM56 task.
- λ is a hyperparameter balancing the trade-off between learning new information and preserving old knowledge.

In the context of multi-fleet predictive maintenance, this approach entails the following steps (as shown in Figure 2):

1. *Training on CFM56 data:*
The model is initially trained to converge on the CFM56 engine dataset, resulting in parameters θ^* . This training phase establishes baseline diagnostic capability on the legacy fleet.
2. *Estimating parameter importance:*
The FIM is approximated by computing the expected squared gradient of the log-likelihood over the CFM56 data. This process identifies parameters essential to maintaining accurate diagnostics for the original fleet.
3. *Training on LEAP-1A data with EWC regularization:*
When adapting the model to the new LEAP-1A data, the loss function is augmented by the EWC penalty term. This regularization discourages significant deviation of critical parameters from their previously learned values, effectively preventing the overwriting of CFM56 knowledge.
4. *Balancing plasticity and stability:*
The hyperparameter λ regulates the balance between plasticity (the ability to learn new LEAP-1A patterns) and stability (retaining CFM56 performance). Careful tuning of λ is necessary, too low a value risks catastrophic forgetting, while too high a value can impede learning new tasks.

Figure 2
Multi-fleet Predictive Maintenance Steps



To implement this approach in practice, we follow a two-phase training loop that reflects the core structure of Elastic Weight Consolidation. The model is first trained on the legacy dataset (CFM56 engine telemetry), allowing it to learn baseline fault detection patterns. Once training on the initial task is complete, the FIM is estimated by computing the average squared gradients of the loss with respect to each model parameter. These gradients quantify the importance of each parameter in performing the original CFM56 task.

The second phase involves fine-tuning the model on the new task, fault prediction for LEAP-1A engines, using a modified loss function that incorporates the EWC regularization term. During this stage, the training loop calculates the new task loss while simultaneously penalizing updates that deviate from the previously learned and important parameters. The strength of this penalty is scaled by the corresponding FIM values and the regularization hyperparameter λ , which controls the balance between retaining legacy diagnostic capabilities and adapting to new engine behaviors.

A simplified pseudocode for this two-phase training procedure is provided within **Appendix A**, outlining the sequential steps of parameter initialization, importance estimation, and regularized fine-tuning. This framework allows predictive maintenance systems to evolve in sync with an expanding aircraft fleet, enabling high diagnostic accuracy across both legacy and emerging engine types without retraining from scratch.

By applying EWC in this way, the predictive maintenance system attains the capacity to learn continuously and adaptively across multiple engine types, supporting a heterogeneous fleet without the need for entirely separate models. This has the potential to reduce maintenance costs, improve fault detection reliability, and align with operational realities where fleets evolve over time.

This mechanism is directly inspired by human cognitive resilience: just as an experienced maintenance engineer draws on prior knowledge while assimilating new technical information, the model learns new engine behaviors without forgetting established diagnostic skills (Czigler & Winkler, 2010).

Case Study 2: Continual Learning in Adaptive Flight Routing Across Global Airspaces

An increasingly important application of machine learning in aviation is the optimization of flight routing in real time, especially in response to changing air traffic, regulatory constraints, and weather conditions. Traditionally, flight plans are generated before departure, with limited adaptability during flight. However, recent advancements in autonomous flight planning systems promise dynamic, onboard route optimization that enhances fuel efficiency, mitigates delays, and improves safety.

Consider an AI-based route optimization system originally trained on operations in North American airspace. This model is familiar with FAA routing structures, jet stream behavior across the Rockies, and regional weather patterns. When deployed in European airspace, the system encounters a different regulatory framework (e.g., Eurocontrol procedures), distinct weather systems (e.g., frequent fog in northern Europe), and denser, more variable airspace due to short-haul flights and temporary military zone activations.

For such a system to be operationally viable on a global scale, it must continually learn from regional data, adapting to new environments while preserving prior capabilities. This requires learning new weather-routing patterns and regulatory constraints without degrading performance in the original domain. A conventional retraining approach, in which the model is updated solely on European data, risks catastrophic forgetting of its North American routing expertise. Conversely, freezing the model to preserve U.S.-specific behaviors leads to rigidity, reducing performance in the new environment.

To strike a balance between adaptability and retention, continual learning techniques must address the plasticity-stability trade-off. In this context, plasticity allows the system to integrate knowledge about Eurocontrol vertical separation rules, active NOTAMs, and regional traffic patterns, while stability ensures that the model retains operational understanding of FAA altitude transitions, NAT (North Atlantic Tracks), and seasonal jet stream routing.

A suitable approach is the use of experience replay, a replay-based method in continual learning. Here, the system maintains a buffer of past flight experiences from North American operations and interleaves them during training on European data. This replay mechanism prevents the loss of critical older knowledge by reintroducing earlier data distributions alongside new ones. Alternatively, dual-model architectures such as Deep Generative Replay (DGR) (Shin et al., 2017) can simulate prior data using generative models, offering a memory-efficient way to maintain performance across domains.

Moreover, modular or architecture-based strategies, such as Progressive Neural Networks (PNNs) (Rusu et al., 2016), could be employed. In this case, the model learns a new set of parameters for European routing while maintaining frozen parameters for North American airspace. These modules are connected via lateral connections that allow for forward knowledge transfer without interfering with past learned features.

This scenario highlights a high-dimensional manifestation of the plasticity-stability trade-off. A system biased toward plasticity may be overfit to European routing peculiarities and fail to

handle U.S.-based transcontinental routes. On the other hand, excessive stability may cause the model to ignore critical changes in the European operational environment, leading to suboptimal or non-compliant flight paths.

In sum, applying continual learning to global route optimization systems enables long-term deployment across diverse airspaces. By mimicking how human pilots and dispatchers accumulate operational experience over multiple regulatory domains, such systems can support safer and more efficient global aviation operations (Rusu et al., 2016; Parisi et al., 2019).

Discussion

The two case studies on predictive maintenance for multi-fleet engine diagnostics and adaptive flight routing across global airspaces serve to illustrate the profound significance of continual lifelong learning within the aviation industry. Their purpose is to concretely demonstrate how machine learning systems, inspired by human cognitive abilities, can address the inherent challenges of dynamic and safety-critical aviation environments. By presenting these examples, the discussion highlights both the practical benefits and complex struggles encountered when attempting to implement continual learning in real-world aviation systems.

At the heart of these case studies lies the principle that aviation AI must not only learn effectively from data but also accumulate knowledge over time, adapting to new information while retaining valuable prior experience. This mirrors the way human operators, be they maintenance technicians or flight planners, build upon their expertise incrementally, integrating lessons from past encounters even as they respond to novel situations. In predictive maintenance, for instance, the system must diagnose faults in both legacy engines and new models without losing diagnostic accuracy on either. Similarly, adaptive flight routing systems must navigate the intricacies of diverse regulatory frameworks and weather patterns, transferring knowledge seamlessly between geographically distinct airspaces.

The benefits of continual lifelong learning in aviation are manifold. Primarily, such learning enables enhanced adaptability, a critical requirement given the fast-evolving nature of aircraft technology, air traffic regulations, and operational contexts. Unlike traditional machine learning models, which are typically trained once and fixed, continual learning systems can update their understanding incrementally, reducing the need for complete retraining. This translates into significant operational cost savings and increased system longevity. Moreover, techniques like EWC and experience replay effectively mitigate the problem of catastrophic forgetting, which would otherwise cause AI models to lose proficiency in previously learned tasks as they absorb new data. Consequently, these systems better support operational efficiency and safety, allowing for more accurate fault predictions, optimized routing decisions, and compliance with ever-changing rules (Li & Hoiem, 2017).

However, implementing continual learning in aviation is not without its challenges. One of the fundamental hurdles is managing the delicate trade-off between plasticity and stability. On one hand, models must be sufficiently plastic to integrate new knowledge, such as unfamiliar engine behaviors or regional airspace regulations. On the other hand, they must maintain enough stability to preserve previously acquired expertise essential for safe operations. Striking this

balance is difficult; excessive plasticity can lead to overwriting important information, resulting in performance degradation, while excessive stability can cause rigidity, preventing the system from adapting to new scenarios.

Beyond algorithmic concerns, continual learning also introduces complexities related to data management. To retain past knowledge, many approaches rely on storing historical data or generating synthetic examples, which can place significant demands on computational and storage resources. Furthermore, in a domain as tightly regulated and safety-critical as aviation, ensuring that adaptive systems can be thoroughly validated and certified is a major obstacle. Machine learning models that change over time challenge traditional verification processes, which assume fixed system behavior. However, the operational benefits of continual learning are significant. In predictive maintenance, continual learning allows earlier detection of emerging fault patterns, potentially reducing aircraft downtime and improving scheduling efficiency. For operators managing evolving fleets, continual learning reduces the need for repeated manual retraining of AI models, enabling seamless adaptation to new engine types or system upgrades. In global flight operations, continual learning supports safer and more flexible cross-border routing by allowing AI systems to adapt to diverse regulatory requirements and regional flight behaviors without sacrificing previously acquired knowledge. These capabilities mirror the adaptive expertise of human operators and offer practical value for maintaining safety and efficiency in a rapidly changing aviation landscape. As a result, regulatory frameworks must evolve to accommodate and oversee adaptive AI technologies.

Currently, the field is making promising strides. Foundational algorithms and architectural designs have been proposed and tested in controlled environments, suggesting the viability of continual learning for aviation tasks. Nonetheless, widespread deployment remains limited. Practical issues such as integration with existing avionics, computational constraints of onboard systems, and the scarcity of aviation-specific continual learning benchmarks slow progress (Olshausen & Field, 1996).

Looking ahead, the future of continual lifelong learning in aviation lies in hybrid approaches that combine multiple learning strategies to leverage their complementary advantages. Emphasis on model explainability and transparency will be critical to gaining operator trust and satisfying certification standards. Additionally, optimizing these learning algorithms for edge deployment, where decisions must be made rapidly on aircraft or unmanned vehicles, will be essential. A promising direction involves developing AI systems that work in partnership with human experts, enhancing situational awareness and decision-making through adaptive human-machine collaboration (Goodfellow et al., 2014).

Moreover, enhancing the robustness of continual learning models against unexpected distribution shifts, such as new failure modes, cyber threats, or sudden regulatory changes, will be crucial for maintaining operational reliability. Finally, advancing standardization and regulatory frameworks tailored to adaptive AI will pave the way for these systems to become integral components of aviation safety and efficiency in the coming decades.

In sum, the case studies underscore that continual lifelong learning is not merely a theoretical concept but a practical necessity for next-generation aviation AI. By enabling models

to learn and adapt in a human-like manner, these approaches promise to transform aviation systems into resilient, intelligent agents capable of safely navigating the complexity and uncertainty of real-world operations.

Conclusion

This paper has explored the vital role of continual lifelong learning in advancing machine learning applications within the aviation industry. Through two detailed case studies, predictive maintenance across diverse engine fleets and adaptive flight routing in global airspaces, we demonstrated how continual learning enables AI systems to dynamically acquire new knowledge while retaining critical prior experience. This human-inspired learning paradigm addresses fundamental challenges posed by the ever-evolving nature of aviation operations, regulatory frameworks, and technological innovations.

The potential benefits of continual lifelong learning are clear: improved adaptability, enhanced safety, and greater operational efficiency. By mitigating catastrophic forgetting through methods such as EWC and experience replay, aviation AI can maintain reliability across changing conditions. However, balancing learning plasticity and memory stability remains a core challenge, compounded by the demands of data management, computational resources, and rigorous certification requirements.

Despite these challenges, recent algorithmic and architectural advances, coupled with growing interest from the aviation community, signal a promising future. Continued research focused on hybrid learning approaches, model interpretability, real-time edge deployment, and human-AI collaboration will be essential. Furthermore, evolving regulatory frameworks must embrace adaptive AI to fully unlock its potential.

An exciting avenue for future research lies in the integration of virtual laboratory environments to support continual learning frameworks in aviation. Virtual labs offer a controlled, scalable, and highly customizable platform for simulating diverse operational scenarios, sensor inputs, and fault conditions that may be rare or difficult to capture in real life. By leveraging these environments, continual learning models can be trained and validated across a broad spectrum of tasks and evolving conditions without risking actual system safety.

Ultimately, continual lifelong learning offers a pathway toward intelligent aviation systems that learn, adapt, and improve in ways akin to human operators. Such systems hold the promise to significantly elevate safety, efficiency, and resilience across the aviation ecosystem as it navigates the complexities of the 21st century.

References

- Barari, G., & Sanders, B. (2024, February). The effectiveness of virtual environments for increasing engagement in engineering technology courses. In the *2024 Conference for Industry and Education Collaboration (CIEC)*.
- Barari, N., & Kim, E. (2021, November). Linking sparse coding dictionaries for representation learning. In *2021 International Conference on Rebooting Computing (ICRC)* (pp. 84–87). IEEE. <https://doi.org/10.1109/ICRC54067.2021.00024>
- Barari, N., Lian, X., & MacLellan, C. J. (2024). Avoiding catastrophic forgetting in visual classification using human concept formation. *arXiv e-prints*, arXiv-2402.
- Barari, N., Lian, X., & MacLellan, C. J. (2024). Incremental concept formation over visual images without catastrophic forgetting. *arXiv Preprint*. <https://arxiv.org/abs/2402.16933>
- Czizler, I., & Winkler, I. (Eds.). (2010). *Unconscious memory representations in perception: Processes and mechanisms in the brain* (Vol. 78). John Benjamins Publishing.
- Fisher, D. (1996). Iterative optimization and simplification of hierarchical clusterings. *Journal of Artificial Intelligence Research*, 4, 147–178. <https://doi.org/10.1613/jair.265>
- Fisher, D. H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning*, 2, 139–172. <https://doi.org/10.1007/BF00169894>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning* (Vol. 1). MIT Press.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 27. https://papers.nips.cc/paper_files/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., ... & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526. <https://doi.org/10.1073/pnas.1611835114>
- Li, Z., & Hoiem, D. (2017). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12), 2935–2947. <https://doi.org/10.1109/TPAMI.2017.2773081>
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (Ed.), *Psychology of learning and motivation* (Vol. 24, pp. 109–165). Academic Press. [https://doi.org/10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8)

- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583), 607–609.
<https://doi.org/10.1038/381607a0>
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54–71.
<https://doi.org/10.1016/j.neunet.2019.01.012>
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., ... & Hadsell, R. (2016). Progressive neural networks. *arXiv Preprint*.
<https://arxiv.org/abs/1606.04671>
- Shin, H., Lee, J. K., Kim, J., & Kim, J. (2017). Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, 30.
- Wang, Z., Haarer, E. L., Barari, N., & MacLellan, C. J. (2025). Taxonomic Networks: A Representation for Neuro-Symbolic Pairing. arXiv preprint arXiv:2505.24601.

Appendix A. EWC Training Procedure for Continual Learning in Predictive Maintenance

The following pseudocode outlines the implementation of Elastic Weight Consolidation (EWC) to enable continual learning in a neural network model applied to multi-fleet engine diagnostics. The method mitigates catastrophic forgetting by selectively preserving important parameters from the original task (e.g., CFM56 engine fault prediction) during training on a new task (e.g., LEAP-1A engine diagnostics).

```
# Inputs:
# D_old: Dataset from old engine type (CFM56)
# D_new: Dataset from new engine type (LEAP-1A)
# Model: Predictive maintenance neural network
#  $\lambda$ : Regularization strength
# epochs: Number of training epochs

# Phase 1: Train on the old task
 $\theta_{\text{star}}$  = Model.initialize_weights()
Model.train(D_old)
 $\theta_{\text{star}}$  = Model.get_weights() # Save trained weights

# Phase 2: Estimate Fisher Information Matrix (FIM)
F = {}
Model.eval()
for x, y in D_old:
    Model.zero_grad()
    loss = compute_loss(Model(x), y)
    loss.backward()
    for param in Model.parameters():
        if param.grad is not None:
            F[param] += (param.grad ** 2) / len(D_old)

# Phase 3: Train on the new task with EWC regularization
for epoch in range(epochs):
    for x, y in D_new:
        Model.zero_grad()

        # New task loss
        prediction = Model(x)
        L_new = compute_loss(prediction, y)

        # EWC regularization loss
        L_ewc = 0
        for param in Model.parameters():
             $\theta_i$  = param
             $\theta_{i\_star}$  =  $\theta_{\text{star}}$ [param]
            F_i = F[param]
            L_ewc += (F_i * ( $\theta_i$  -  $\theta_{i\_star}$ ).pow(2)).sum()

        # Combined loss
        L_total = L_new + ( $\lambda$  / 2) * L_ewc

        # Update parameters
        L_total.backward()
        optimizer.step()
```

Notes:

This pseudocode is adaptable to frameworks such as PyTorch or TensorFlow.

The $F[\text{param}]$ entries approximate the diagonal of the Fisher Information Matrix, treating each model parameter independently.

The loss term

$$\mathcal{L}_{EWC} = \mathcal{L}_{new}(\theta) + \frac{\lambda}{2} \sum_i F_i (\theta_i - \theta_i^*)^2$$

is central to preserving knowledge while supporting adaptation.

This code is illustrative and may require modification depending on the specific architecture and task formulation. It provides a foundational strategy for deploying continual learning in real-world aviation diagnostics.

Appendix B. Computational Framework for Continual Learning in Adaptive Flight Routing

This appendix provides a conceptual outline and pseudocode for implementing a continual learning system for adaptive flight routing across diverse airspaces, as described in Case Study 2. The system must retain routing knowledge from a previously trained airspace domain (e.g., North America) while adapting to new regional rules and conditions (e.g., Eurocontrol airspace).

B.1 Model Architecture and Inputs

The routing model may be implemented as a reinforcement learning (RL) agent, a supervised deep neural network, or a hybrid policy-based system. The model takes the following inputs:

- **Aircraft state:** current location, heading, speed, altitude
- **Environmental features:** weather patterns, jet stream data, turbulence zones
- **Airspace regulations:** regional altitude rules, NOTAMs, restricted zones
- **Traffic density or slot availability**
- **Flight intent:** origin, destination, route constraints

The output is a sequence of waypoints or control decisions that optimize for criteria such as fuel efficiency, time, or safety under evolving constraints.

B.2 Continual Learning Integration

To ensure adaptability without forgetting previously learned airspace behavior, continual learning mechanisms are introduced:

- **Replay-based methods:** The model maintains a buffer of experiences (flight episodes) from previous environments. These are replayed alongside new data to prevent forgetting.
- **Regularization-based methods:** Elastic Weight Consolidation (EWC) or similar approaches constrain updates to critical model parameters identified from prior training.
- **Modular approaches:** Progressive Neural Networks (PNNs) create task-specific modules for each airspace, preserving old knowledge while enabling new learning through lateral connections.

B.3 Pseudocode for Replay-Based Continual Learning

The following pseudocode outlines a simplified training loop using experience replay to retain performance on prior airspace domains:

```
CopyEdit
# Initialize model and memory
model = initialize_model()
replay_buffer = initialize_replay_buffer()

# Load past experiences from North American operations
replay_buffer.load_old_domain_data("NorthAmerica")

# Training on new domain: Eurocontrol
for episode in training_episodes_Europe:
    state = environment.reset()
    done = False

    while not done:
        action = model.select_action(state)
        next_state, reward, done = environment.step(action)

        # Store new experience
        replay_buffer.add(state, action, reward, next_state)

        # Sample batch containing both old and new data
        batch = replay_buffer.sample_mixed_batch()

        # Compute and apply loss
```



```
loss = compute_loss(batch)
model.optimize(loss)
```

```
state = next_state
```

This architecture ensures that while the model adapts to new routing behaviors, such as altitude transition rules or regional separation standards, it maintains proficiency in older environments, such as North Atlantic track systems or FAA-based separation logic.

B.4 Practical Considerations

- **Data simulation tools** such as BlueSky or OpenSky APIs can be used to generate realistic traffic and weather inputs for training.
- **Evaluation** should include transfer tests, where the model is assessed on previously learned environments after learning the new one, to quantify retention.
- **Modularity** is recommended when regulatory shifts are stark, to isolate and manage region-specific knowledge.