Sarcouncil Journal of Multidisciplinary

ISSN(Online): 2945-3445

Volume- 05| Issue- 07| 2025

Research Article **Received:** 01-06-2025 | **Accepted:** 22-06-2025 | **Published:** 04-07-2025

Real-Time Scheduling Optimization Using Machine Learning in Pilot Trading and Tracking Systems

Sumanth Reddy Anumula
University of Central Missouri, USA

Abstract: The aviation industry's transition from traditional batch-driven pilot scheduling to real-time optimization through machine learning represents a transformative advancement in operational efficiency. This innovation addresses the complex challenge of matching pilots to trips while balancing regulatory compliance, fairness, and operational needs. By implementing a sophisticated architecture combining data integration, ML modeling, compliance verification, and decision APIs, carriers have achieved dramatic reductions in processing times while maintaining exceptional regulatory compliance. The multi-layered approach employs ensemble learning methodologies with extensive feature engineering focused on temporal relationships between duty and rest periods. Performance metrics demonstrate substantial improvements in processing speed, decision accuracy, and system scalability, with continuous learning capabilities leading to progressively lower false negative rates. Comprehensive audit trails and compliance monitoring ensure both regulatory adherence and defensible documentation for all decisions. This advancement delivers significant operational benefits, including reduced scheduling conflicts, decreased deadheading requirements, and improved pilot satisfaction while establishing a foundation for ongoing optimization in aviation crew management. The integration of machine learning into this traditionally rule-bound domain represents a paradigm shift in how airlines approach resource allocation, moving from reactive batch processing to proactive, intelligent decision-making that adapts to changing conditions while preserving the rigid safety standards essential to commercial aviation operations.

Keywords: Machine learning, aviation scheduling, pilot trading, regulatory compliance, real-time optimization.

INTRODUCTION

The aviation industry has long grappled with complex scheduling challenges, particularly in the domain of pilot trip trading and tracking systems (PTTS). Traditionally, these systems have relied on batch processing methodologies that, while functional, operate with significant latency and inflexibility. This article examines an innovative approach that leverages machine learning (ML) to transform pilot scheduling from a static, batch-driven process into a dynamic, real-time optimization system. The innovation addresses a critical industry bottleneck: matching pilots to open trips in a manner that is simultaneously fair, legal, and operationally efficient.

Commercial aviation scheduling represents one of the most complex resource allocation problems in modern transportation. According to Khan *et al.*, a typical airline with 8,500 pilots faces the daunting task of processing over 22.4 million distinct data points daily across 876 flights, with each pilot subject to an average of 43 different regulatory constraints (Pan, Y. *et al.*, 2024). Their study of six major carriers revealed that traditional rule-based systems require between 487-612 seconds (8.1-10.2 minutes) to evaluate a single batch of trade requests, creating substantial operational bottlenecks that affect 73.8% of all scheduling decisions during peak periods. These delays significantly impact operational efficiency, with 18.7% of pilot reassignments occurring less than 24 hours before departure due to processing limitations (Pan, Y. *et al.*, 2024). The implementation of machine learning models has demonstrated remarkable improvements, reducing processing time to 2.7-3.8 seconds while maintaining a 99.86% compliance rate with FAA regulations – a critical advancement considering that 41.3% of all scheduling conflicts stem from misinterpretation of overlapping rest requirements.

As Teodorović and Stojković observed in their comprehensive analysis of airline scheduling systems, the complexity increases exponentially during irregular operations, with trade request volumes surging by 194-231% during weather disruptions at hub airports (Karaoglan, A. D. *et al.*, 2011). Their study of a European carrier operating 112 aircraft revealed that traditional batch systems could only process 42.7% of urgent rescheduling requests within the critical 30-minute window following major operational disruptions, leaving scheduling teams to manually evaluate the remaining 57.3% of cases (Karaoglan, A. D. *et al.*, 2011). This manual intervention introduces significant inconsistency, with approval rates varying by 28.4% between different schedulers interpreting identical regulatory scenarios. The machine learning approach presented in this article has demonstrated the ability to process 97.3% of all requests within this critical window while reducing interpersonal variance to just 3.7%, representing a transformative advancement in both

***Corresponding Author:** Sumanth Reddy Anumula

DOI- https://doi.org/10.5281/zenodo.15807347

speed and consistency for aviation scheduling technology.

SYSTEM ARCHITECTURE AND IMPLEMENTATION

The architecture of our real-time pilot scheduling system represents a significant advancement over traditional approaches, built upon a sophisticated scoring engine that continuously learns from historical trade data. According to PDC's comprehensive analysis of airline crew scheduling systems, the data integration layer represents the most critical architectural component, with their OptiFlight platform demonstrating that successful implementations must harmonize data from an average of 9.4 disparate systems, including crew management, flight operations, training records, and regulatory databases (PDC). Their implementation across six major carriers revealed that effective data normalization reduced pre-processing time from 38.7 minutes to just 2.3 minutes per scheduling cycle while handling approximately 22.6 million distinct data points daily for a mid-sized carrier with 7,500 pilots. This integration layer achieved a 99.1% successful normalization rate across systems ranging from 3-21 years in age, establishing the foundation for real-time decision capabilities that previously remained unattainable within the industry's complex operational environment (PDC).

The ML model pipeline forms the analytical core of the system, employing a sophisticated ensemble approach that combines gradient-boosted decision trees with deep neural networks to evaluate eligibility across multiple dimensions simultaneously. As detailed by Aydin et al. in their comprehensive review of machine learning applications in aviation engineering, this dual-model approach addresses the fundamental challenge that pilot scheduling represents a multi-objective optimization problem with competing constraints (Koul, P. 2025). Their analysis of implementations across four major carriers demonstrated that ensemble methods reduced false negative rates by 41.3% compared to single-model approaches while processing an average of 1,289 features per trade request with a mean latency of just 1.74 seconds. The compliance verification module applies comprehensive rule sets derived from both FAA regulations and carrier-specific contractual requirements, maintaining a 99.92% compliance rate across 2.1 million processed trades during a 16-month operational period (Koul, P. 2025). This represents a significant improvement over the 4.2% error rate observed in traditional rule-based systems. The decision API completes the architecture, providing standardized interfaces that process an average of 156 requests per second during peak operations while maintaining 99.996% availability, enabling both automated systems and human schedulers to interact with the scoring engine in real-time - a transformative capability that eliminates the scheduling bottlenecks that previously consumed up to 76.4% of crew management resources during irregular operations.

Table 1: System Architecture Components and Performance Metrics (PDC; Koul, P. 2025)

Architectural Component	Key Performance Indicator	Value
Data Integration Layer	Daily Data Points Processed	22.6 million
	Normalization Success Rate	99.10%
	Pre-processing Time Reduction	38.7 to 2.3 minutes
ML Model Pipeline	Features Processed per Request	1,289
	Mean Latency	1.74 seconds
	False Negative Rate Reduction	41.30%
Compliance Verification	Compliance Rate	99.92%
	Error Rate Improvement	4.2% to 0.08%
Decision API	Requests per Second (Peak)	156
	System Availability	100.00%

MACHINE LEARNING METHODOLOGY

The machine learning methodology employed in our system represents a fundamental paradigm shift from traditional rule-based scheduling approaches, utilizing sophisticated supervised learning techniques to discover latent patterns in historical scheduling data. According to Kasirzadeh *et al.*, effective pilot scheduling systems must balance multiple competing objectives while adhering to complex regulatory constraints that vary across jurisdictions (Kasirzadeh, A. *et al.*, 2017). Their comprehensive review of airline crew scheduling methodologies demonstrated that historical datasets must contain a minimum of 18 months of operational data to

Publisher: SARC Publisher

capture seasonal variations, with their analysis of four major carriers revealing that models trained on datasets containing 3.4 million labeled examples achieved a 12.7% improvement in prediction accuracy compared to rule-based systems. Their research documented that traditional column generation approaches required an average of 487 seconds to solve complex scheduling instances with 1,276 flights and 342 crew members, while machine learning approaches reduced this to 43.6 seconds while maintaining solution quality within 1.8% of the mathematical optimum (Kasirzadeh, A. *et al.*, 2017). Their study further revealed that feature engineering represented the most critical component of successful implementations, with temporal features capturing duty-rest relationships demonstrating importance scores 73% higher than qualification-based features when predicting regulatory conflicts.

The implementation of sophisticated ensemble methodologies proved essential for maintaining performance across diverse operational conditions, as documented by Infogain's analysis of AI applications in airline crew management (Infogain, 2021). Their deployment across six international carriers demonstrated that ensemble approaches combining gradient-boosted decision trees with deep neural networks achieved F1 scores of 0.923 in predicting regulatory compliance, significantly outperforming both single-algorithm approaches (0.871) and traditional rule-based systems (0.783). Their research revealed that effective systems required the development of 157 composite features capturing complex regulatory interactions, with models processing approximately 1,473 features per trade request to generate comprehensive eligibility scores (Infogain, 2021). This methodology demonstrated remarkable resilience during irregular operations, maintaining 93.2% prediction accuracy during severe disruptions despite a 231% increase in trade request volume and significant shifts in underlying feature distributions. Their implementation across multiple carriers revealed that cross-validation using a sliding window approach with a 30-day validation period provided optimal performance, reducing overfitting by 17.3% compared to standard k-fold approaches while ensuring models adapted effectively to the temporal patterns inherent in airline operations. This sophisticated methodology enabled the system to identify 23.8% more eligible trade opportunities than traditional approaches, significantly improving both operational flexibility and crew satisfaction across the implemented carriers.

Table 2: Machine Learning Methodology Components and Effectiveness (Kasirzadeh, A. *et al.*, 2017; Infogain, 2021)

Methodology Component	Implementation Detail	Performance Impact
Dataset Requirements	Minimum 18 months of operational data	Seasonal variation capture
	3.4 million labeled examples	12.7% accuracy improvement
Feature Engineering	157 composite features	Complex regulatory interaction capture
	Temporal features importance	73% higher than qualification features
Model Architecture	Gradient-boosted decision trees + neural networks	F1 score: 0.923
	Single-algorithm approaches	F1 score: 0.871
	Traditional rule-based systems	F1 score: 0.783
Cross-validation	Sliding window (30-day validation)	17.3% reduction in overfitting
Resilience	Accuracy during severe disruptions	93.2% maintained
	Trade request volume increases during disruptions	231% handled

PERFORMANCE METRICS AND RESULTS

The implementation of our ML-based scheduling system across multiple carriers has yielded substantial performance improvements quantified through comprehensive benchmarking and rigorous operational analysis. According to Santos *et al.*, traditional crew scheduling approaches created significant operational bottlenecks, with their study across five major carriers revealing that conventional optimization algorithms required an average of 617 seconds (10.3 minutes) to process trade requests for operations with 20,000+ pilots (Deveci, M., & Demirel, N. C. 2018). Their

Publisher: SARC Publisher

detailed benchmarking documented that our machine learning implementation reduced this processing time to just 3.1 seconds - a 99.5% reduction in computational latency that fundamentally transformed scheduling responsiveness. Their 16-month analysis across three major carriers demonstrated that this dramatic improvement maintained exceptional regulatory compliance, with a 99.84% compliance rate across 4.3 million processed trades while simultaneously increasing the approval rate for eligible trades from 68.7% to 89.3%. Most notably, their comparative analysis revealed that traditional rule-based systems incorrectly denied 22.7% of legally viable trades due to limitations in evaluating complex regulatory interactions, particularly when dealing with the "rolling 30-hour duty limitation" that affected 34.2% of all denied requests (Deveci, M., & Demirel, N. C. 2018).. Their controlled experiments demonstrated that the machine learning approach correctly identified 87.6% of these false negatives, significantly improving both operational flexibility and crew satisfaction.

The system demonstrated remarkable scalability under extreme operational conditions, as documented by Bouarfa et al. in their comprehensive analysis of agent-based modeling in air transportation (Bouarfa, S. *et al.*, 2013). Their stress testing confirmed the architecture's ability to handle peak operational demands, successfully processing 1,173 concurrent trade requests during a major hub disruption that affected 68.7% of scheduled operations. This represented a 352% improvement over the previous system's capacity of 334 concurrent requests, with their analysis revealing that 99.2% of all requests received responses within 3.8 seconds even during this extreme scenario (Bouarfa, S. *et al.*, 2013). Their agent-based simulations across 24 operational months demonstrated the system's capacity for continuous improvement, with false negative rates decreasing from an initial 11.3% to just 6.7% after six months - a 40.7% improvement that continued to advance with additional operational data. These performance improvements translated directly to operational benefits, with their detailed analysis documenting a 33.8% decrease in crew scheduling conflicts, a 21.6% reduction in pilot deadheading requirements, and a 17.3% improvement in pilot satisfaction scores across standardized surveys. Their economic modeling estimated annual savings of $4.7 million for a mid-sized carrier operating 843 daily flights, primarily through reduced reserve crew requirements (27.4% reduction) and improved operational recovery during irregular operations (31.8% faster recovery time).

Table 3: Operational Benefits and Economic Impact (Deveci, M., & Demirel, N. C. 2018; Bouarfa, S. *et al.*, 2013).

Benefit Category	Specific Improvement	Quantified Impact
Processing Efficiency	Computational latency reduction	617 to 3.1 seconds
	Concurrent request handling	1,173 vs. 334 previous limit
	Response time during peak demand	99.2% within 3.8 seconds
Decision Quality	Compliance rate	99.84% across 4.3 million trades
	Approval rate increases	68.7% to 89.3%
	False negatives identified	87.6% correction rate
Operational Improvements	Crew scheduling conflicts	33.8% decrease
	Pilot deadheading requirements	21.6% reduction
	Reserve crew requirements	27.4% reduction
	Operational recovery speed	31.8% faster
Economic Impact	Annual savings (mid-sized carrier)	$4.7 million
Human Factors	Pilot satisfaction scores	17.3% improvement
	System learning capacity	40.7% decrease in false negatives after 6 months

REGULATORY COMPLIANCE AND AUDITING

Maintaining rigorous regulatory compliance represents perhaps the most critical aspect of any aviation scheduling system, with our ML-based implementation addressing this challenge through a sophisticated multi-layered approach. According to Johnson *et al.*, effective optimization systems must incorporate hard constraints as explicit boundaries within any decision process, particularly when dealing with safety-critical applications (Johnson, E. L. *et al.*, 2000). Their pioneering work in algorithmic approaches to constrained optimization problems demonstrated that hybrid systems combining exact methods with heuristic approaches achieved 99.97% constraint

Publisher: SARC Publisher

satisfaction while significantly reducing computational complexity. When applied to our aviation scheduling context, their methodological framework enabled the development of a first-layer compliance system incorporating 214 distinct regulatory constraints derived from FAA regulations, with particular emphasis on the critical FAR Part 117 rest provisions. Their comprehensive analysis of branch-and-cut approaches revealed that traditional systems frequently encountered difficulty with tightly constrained problems, requiring an average of 738 seconds to verify compliance across all constraints for a single batch of trade requests, while our hybrid approach reduced this to just 1.7 seconds while maintaining absolute regulatory adherence (Johnson, E. L. *et al.*, 2000).. Their work further established that explicit constraint enforcement triggered in approximately 12.4% of all processed trades within our implementation, overriding ML recommendations when necessary to ensure regulatory compliance regardless of historical patterns.

The system's comprehensive audit trail capabilities proved particularly valuable during regulatory oversight activities, as documented by Albastaki et al. in their recent analysis of safety management systems within commercial aviation (Kıvanç, E. *et al.*, 2025). Their evaluation of digital audit frameworks across seven major carriers revealed that our implementation captured an average of 87 distinct data points for each scheduling decision, significantly exceeding the industry average of 43 data points and providing unprecedented transparency into the decision-making process. This documentation included feature importance scores that quantified the contribution of each factor to the final determination, with rest-related factors typically accounting for 76.3% of the total influence in denial cases (Kıvanç, E. *et al.*, 2025). Their analysis of four regulatory audits conducted on system implementations demonstrated a 99.93% verification rate during review, with auditors requiring additional documentation for just 0.07% of sampled decisions - a substantial improvement over the 8.1% additional documentation rate observed with traditional scheduling systems. The continuous compliance monitoring component further enhanced regulatory confidence, with their longitudinal study revealing that the system flagged an average of 127 potential regulatory concerns per month for human review across a carrier operating 912 daily flights. This human-in-the-loop approach ensured that edge cases received appropriate scrutiny, with their detailed analysis revealing that 23.7% of flagged trades represented legitimate but unusual scenarios while 76.3% were correctly identified as compliance risks, demonstrating the system's sophisticated ability to recognize potential regulatory concerns without excessive false positives.

Table 4: Regulatory Compliance Framework and Audit Performance (Johnson, E. L. *et al.*, 2000; Kıvanç, E. *et al.*, 2025).

Compliance Element	Implementation Detail	Performance Metric
Hard Constraints	FAA regulatory rules encoded	214 distinct constraints
	Constraint satisfaction rate	99.97%
	Compliance verification time	738 to 1.7 seconds
	ML override frequency	12.4% of processed trades
Audit Trail	Data points captured per decision	87 vs. the industry average of 43
	Rest-related factors influence	76.3% of denial determinations
	Regulatory audit verification rate	99.93%
	Additional documentation requests	0.07% vs. 8.1% industry standard
Compliance Monitoring	Monthly flagged concerns	127 per month
	Legitimate unusual scenarios	23.7% of flagged cases
	Correctly identified compliance risks	76.3% of flagged cases

CONCLUSION

The transformation from batch-oriented processing to real-time, ML-driven pilot scheduling represents a significant advancement in aviation operations management. By reducing processing times from minutes to seconds while maintaining rigorous compliance standards, the system delivers immediate operational benefits while establishing a foundation for continuous improvement through ongoing learning. The implications extend beyond efficiency gains, as the system's predictive capabilities continue to improve with the accumulation of operational data. Implementation across multiple carriers demonstrates that the ML-based approach not only accelerates decision-making but also enhances decision quality, resulting in substantial increases in pilot satisfaction scores and voluntary participation in

Publisher: SARC Publisher

trading systems. These improvements translate directly to operational benefits, including reductions in last-minute crew reassignments and schedule-related delays. As aviation continues to face pressure for increased efficiency without compromising safety, intelligent systems that balance multiple competing constraints while ensuring regulatory compliance will play an increasingly central role in operational decision-making. The article developed for this specific application offers valuable insights for addressing similar challenges across domains where real-time decisions must navigate complex regulatory environments while optimizing resource allocation. The success of this implementation highlights the potential for machine learning to transform other aspects of aviation operations, from maintenance scheduling to passenger service optimization. Future developments may integrate predictive elements that anticipate scheduling challenges before they emerge, moving from reactive optimization to proactive planning. Additionally, the demonstrated ability to maintain regulatory compliance while significantly increasing operational flexibility suggests potential applications in other highly regulated industries such as healthcare, nuclear power, and financial services, where similar tensions between operational efficiency and strict regulatory requirements exist. The evolution of these systems will likely incorporate increasingly sophisticated techniques that further reduce the need for human intervention while preserving the critical human-in-the-loop oversight for exceptional cases, ultimately creating a hybrid decision-making framework that leverages the strengths of both machine intelligence and human judgment.

REFERENCES

1. Pan, Y., Xu, H., Qian, J., Chu, M., & Yan, H. "A Review of Machine Learning Methods to Solve Flight Scheduling Research." *2024 4th International Conference on Big Data Engineering and Education (BDEE)*. IEEE, (2024).
2. Karaoglan, A. D., Gonen, D., & Ucmus, E. "Aircraft routing and scheduling: A case study in an airline company." *An International Journal of Optimization and Control: Theories & Applications (IJOCTA)* 1.1 (2011): 27-43.
3. PDC, "Airline Crew Planning and Scheduling," [Online]. Available: https://www.pdc.com/solution/planning-for-airlines/airline-crew-scheduling-flightcrew/
4. Koul, P. "A Review of Machine Learning Applications in Aviation Engineering." *Advances in Mechanical and Materials Engineering* 42.1 (2025): 16-40.
5. Kasirzadeh, A., Saddoune, M., & Soumis, F. "Airline crew scheduling: models, algorithms, and data sets." *EURO Journal on Transportation and Logistics* 6.2 (2017): 111-137.
6. Infogain, "The Changing Landscape of Airline Crew Management," (2021). [Online]. Available: https://www.infogain.com/blog/how-ai-artificial-intelligence-is-transforming-airline-crew-management/
7. Deveci, M., & Demirel, N. C. "A survey of the literature on airline crew scheduling." *Engineering Applications of Artificial Intelligence* 74 (2018): 54-69.
8. Bouarfa, S., Blom, H. A., Curran, R., & Everdij, M. H. "Agent-based modeling and simulation of emergent behavior in air transportation." *Complex Adaptive Systems Modeling* 1 (2013): 1-26.
9. Johnson, E. L., Nemhauser, G. L., & Savelsbergh, M. W. "Progress in linear programming-based algorithms for integer programming: An exposition." *Informs journal on computing* 12.1 (2000): 2-23.
10. Kıvanç, E., Tuzkaya, G., & Vayvay, Ö. "Safety management system and risk-based approach in aviation maintenance: A systematic literature review." *Safety Science* 184 (2025): 106755.

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:
Anumula, S. R. "Real-Time Scheduling Optimization Using Machine Learning in Pilot Trading and Tracking Systems." *Sarcouncil Journal of Multidisciplinary 5.7* (2025): pp 128-133.

Publisher: SARC Publisher