Integrating Artificial Intelligence into Airline Operational Control Centers: Opportunities, Challenges, and Safety Implications for Flight Dispatch Operations

Jamil Akhtar

Flight Operations Supervisor | Licensed Flight Operations Officer (FOO) | Licensed Flight Dispatcher

Qatar Aviation Services, Doha, Qatar

Abstract

Airline operational control centers (OCCs) operate at the intersection of meteorology, aircraft performance engineering, regulatory compliance, and real-time risk management. Dispatchers and Flight Operations Officers (FOOs) are required to synthesize large volumes of heterogeneous data weather products, Notices to Airmen (NOTAMs), aircraft performance limitations, fuel burn projections, and network constraints under significant time pressure, often for dozens of flights simultaneously. This paper examines the integration of artificial intelligence (AI) and data-driven decision-support tools into airline OCC workflows, with particular attention to flight dispatch operations. Drawing on operational experience and a critical synthesis of the scholarly literature on decision support systems, predictive analytics, naturalistic decision-making, and human-AI teaming in aviation, the paper develops a conceptual framework comparing traditional and data-driven dispatch processes across eight performance dimensions: dispatch processing time, route efficiency, fuel consumption, delay reduction, dispatch decision accuracy, situational awareness, risk mitigation effectiveness, and flight completion efficiency. The analysis indicates that AI-enabled dispatch tools can meaningfully reduce flight plan preparation time, improve the consistency of fuel and route recommendations, and enhance the dispatcher's ability to detect emerging operational risk — provided that such tools are positioned as decision aids rather than autonomous decision-makers. The paper also identifies persistent challenges, including data integration complexity, algorithmic opacity, automation complacency, regulatory acceptance, and the risk of eroding the recognition-primed expertise that experienced dispatchers bring to non-routine events. The discussion concludes with industry implications for OCC design, training, and certification, and outlines a research agenda for validating AI-assisted dispatch tools under operational conditions.

Keywords

Flight dispatch; airline operational control center; decision support systems; predictive analytics; fuel optimization; weather-informed routing; aviation safety management; human-AI teaming

1. Introduction

The airline operational control center is the nerve center of daily flight operations. Within the OCC, Flight Operations Officers and licensed dispatchers hold joint responsibility with the pilot-in-command for the safety of each flight, a regulatory arrangement that places the dispatcher's judgment at the center of pre-flight planning and in-flight monitoring. In practice, this means a single dispatcher may be responsible for the flight following of ten to twenty active flights at once, each requiring continuous reassessment of weather along the route, fuel sufficiency against changing winds, alternate aerodrome suitability, NOTAM-driven restrictions, and the operational status of the destination airport.

The volume and velocity of information available to dispatchers has increased substantially over the past decade. Satellite-derived weather products, ADS-B-based real-time aircraft tracking, machine-readable NOTAM feeds, and airline-specific performance databases now generate far more data than a single individual can manually reconcile within the time available for flight release. This growing information burden has motivated airlines, vendors, and researchers to explore artificial intelligence and machine learning as a means of supporting — rather than replacing — the dispatcher's decision-making process.

This paper examines how AI-enabled, data-driven dispatch systems are reshaping operational control center practice. The analysis is grounded in the day-to-day realities of flight dispatch: the construction of operational flight plans, the application of fuel policy (including contingency, alternate, and discretionary fuel), the interpretation of significant meteorological information (SIGMETs), and the execution of flight watch during irregular operations. Rather than treating AI as a generic efficiency tool, the paper situates its discussion within the specific constraints of Part 121-equivalent dispatch regulation, OCC shift structures, and the collaborative relationship between dispatchers, flight crews, and air traffic management.

The remainder of the paper is organized as follows. Section 2 critically reviews the literature on flight dispatch decision support, OCC decision-making, predictive analytics, and human-AI teaming, identifying research gaps relevant to operational dispatch. Section 3 describes the conceptual research methodology, including the comparative framework and evaluation criteria used in the analysis. Section 4 presents a data-driven flight dispatch framework, structured around four architecture figures. Section 5 presents results and discussion, organized around the eight performance dimensions identified in the methodology, supported by four data tables. Section 6 discusses industry implications for OCC design, training, and regulatory acceptance. Section 7 outlines limitations and future research directions, and Section 8 concludes.

2. Literature Review

2.1 Decision Support Systems in Air Transport

The application of decision support systems (DSS) to air transport has a long history, predating the current wave of machine learning by several decades. Early work on airline scheduling and rescheduling established the conceptual foundation for treating dispatch and recovery problems as structured decision processes amenable to heuristic and optimization-based support — contributions that trace back to interactive, network-based scheduling tools developed for major carriers in the early 1990s and real-time recovery support systems explored through the mid-1990s. A more recent strand of this literature frames DSS broadly as any system that provides information necessary to support decision-making in environments characterized by unstructured or semi-structured problems, a description that fits the operational control environment closely because the air transport system is a safety-critical, human-in-the-loop, distributed and complex environment that is fertile ground for decision support applications.

A second important line of work addresses the integration of decision support across the boundary between maintenance control and operational control. In this framing, the operations control center monitors and manages daily flight operations and provides support when a dispatch decision may lead to a flight delay or operational constraint arising from a deferred maintenance item, while the captain retains the authority to refuse a proposed temporary fix if it is judged to impair safe continuation of the flight. This literature emphasizes that effective decision support in dispatch is not solely a technical integration problem; it is also a coordination problem involving multiple stakeholders — dispatch, maintenance control, and the flight crew — each with distinct authority and information needs. A system that shares consistent information across these stakeholders has been shown to reduce miscommunication independent of how the maintenance organization is structured relative to the airline.

Earlier foundational work on decision support for airline system operations control demonstrated the integration of computer science and operations research techniques into a single platform supporting real-time flight following, aircraft routing, and irregular operations recovery. Such systems were designed to integrate real-time flight following, aircraft routing, maintenance, crew management, gate assignment, and flight planning with dynamic aircraft re-scheduling and fleet re-routing algorithms for irregular operations, with algorithms intended to help flight controllers optimally re-route aircraft, crews, and passengers when operational problems disrupt the schedule. While the underlying computing architectures described in this line of work are now dated, the functional decomposition — flight following, routing, recovery, and crew/passenger impact — remains a useful organizing structure for evaluating modern AI-enabled dispatch tools.

2.2 Operational Control Center Decision-Making: A Naturalistic Perspective

A distinct but complementary body of research applies naturalistic decision-making (NDM) theory to understand how OCC personnel actually make decisions under time pressure, rather than how normative models prescribe they should decide. A mini-ethnographic case study conducted within a major South American airline's OCC found that operations controllers constantly and actively look for signs of disruption while monitoring normal operations and rebalancing resources, with decision-making distributed and decentralized across multiple functions, each relying on a repertoire of strategies to deploy solutions to dynamic scenarios. This finding is significant for AI integration because it suggests that dispatch decision-making is not a single centralized optimization problem but a distributed cognitive process in which multiple actors —

dispatchers, network controllers, crew schedulers — each hold partial situational pictures that must be reconciled.

This NDM perspective draws on the recognition-primed decision (RPD) model, in which experienced practitioners recognize a situation as typical and rapidly retrieve an appropriate course of action rather than exhaustively comparing alternatives. For AI-enabled dispatch tools, this raises an important design question: should the system attempt to replicate RPD-style pattern recognition (for example, flagging that a given combination of weather, NOTAM, and aircraft status resembles a previously encountered disruption pattern), or should it instead supplement RPD by surfacing options the dispatcher might not otherwise consider? The literature on training for aviation decision-making suggests that decision aids are most effective when they support — rather than substitute for — the perceptual and pattern-recognition processes that underlie expert performance, since the development of robust decision skills depends on exposure to realistic operational conditions and the evaluation of resource requirements against available capabilities.

2.3 Predictive Analytics for Flight Delay and Disruption

Flight delay prediction has become one of the most heavily studied applications of machine learning in aviation operations, reflecting both the economic significance of delay and the relative availability of historical operational data. A recent evaluation of multiple machine learning algorithms for arrival delay prediction found that Random Forest combined with random oversampling and Decision Tree combined with SMOTE both achieved the highest predictive performance, with accuracy and F1-scores around 0.90, while simpler models such as Naive Bayes remained competitive under balanced conditions. This result is methodologically important for dispatch applications because it underscores that resampling strategy and validation design — not merely model architecture — substantially affect predictive performance on the imbalanced delay datasets typical of airline operations.

More architecturally sophisticated approaches have combined recurrent, graph-based, and convolutional neural network components to capture spatial and temporal dependencies across an airport network. One such hybrid architecture, combining LSTM, graph convolutional, and 3D convolutional components, was developed specifically to incorporate an airport-level congestion measure into spatio-temporal learning, with the explicit aim of improving both predictive accuracy and the interpretability of delay forecasts through an attention mechanism. A separate IEEE review of recent flight delay prediction studies organized the literature around a six-step comparison framework spanning data collection and processing, feature and model selection, and evaluation, reflecting the maturation of this research area into a recognizable methodological pipeline.

Beyond arrival delay, machine learning has also been applied to the prediction of holding patterns — airborne delay maneuvers driven by congestion, weather, or airspace constraints. Graph-based machine learning approaches applied to large trajectory datasets from major European airports have been used to detect and predict holding behavior, with implications described in this literature for fuel efficiency, delay reduction, and passenger experience. For dispatch operations, the ability to anticipate holding probability before departure has direct fuel-planning implications, since anticipated holding time can be factored into contingency fuel decisions rather than discovered only in flight.

2.4 AI and Predictive Analytics for Fuel Efficiency and Route Optimization

A growing literature addresses the application of AI to fuel consumption prediction and route optimization. Recent work on a predict-then-optimize approach to minimizing aircraft fuel consumption argues that while traditional methods rely on fixed formulas or empirical data considering parameters such as air distance, aircraft weight, and altitude, AI models can instead learn from high-dimensional historical data to uncover complex relationships that traditional statistical methods might overlook, including real-time weather data, aircraft-specific performance metrics, and historical flight information. This distinction is operationally meaningful: traditional flight planning systems compute fuel burn from performance tables and forecast winds, whereas AI-based approaches can, in principle, learn systematic biases between forecast and actual conditions for specific routes, aircraft types, or seasons, and adjust fuel recommendations accordingly.

Industry-facing literature on fuel efficiency applications of AI emphasizes real-time operational adjustment as a key value proposition, with predictive models assessing weather, traffic, and other environmental factors to support fuel-efficient procedures such as continuous descent approaches. While such sources are not peer-reviewed and should be read with appropriate caution regarding vendor claims, they are useful in identifying the operational use cases that airlines are actively pursuing: pre-flight fuel load recommendations, in-flight route and altitude advisories, and post-flight performance feedback loops that refine future predictions.

A complementary stream of work addresses contrail avoidance as an environmental dimension of route optimization. Trajectory optimization studies have developed algorithms to compute wind-optimal cruise trajectories while avoiding airspace regions favorable to persistent contrail formation, formulating contrail-prone regions as soft constraints within an optimal control problem. For dispatch operations, contrail-aware routing represents an emerging category of weather-informed optimization that extends beyond the traditional focus on turbulence, icing, and convective weather avoidance, and is likely to become more operationally relevant as airlines respond to non-CO_2 climate impact regulation.

2.5 Weather Data Integration and NOTAM Processing

The integration of meteorological data into trajectory planning has been addressed in the patent literature describing dynamic aircraft trajectory management systems, which envision the use of convective weather products such as automated probabilistic thunderstorm height forecasts, alongside System Wide Information Management (SWIM) data sources, to produce trajectory plans in seconds rather than through the longer manual flight planning process. Such systems are explicitly designed to account for disruptions including weather, traffic, passenger, pilot, maintenance, airspace procedure, and infrastructure factors within a unified fleet optimization framework. While patent disclosures do not constitute peer-reviewed evidence of operational performance, they indicate the technical direction of commercial development in dispatch decision support.

On the NOTAM side, industry commentary has highlighted the application of data analytics and AI to improve the relevance and timeliness of NOTAM information, with AI-powered analysis of large datasets proposed as a means of identifying patterns and predicting potential hazards to enable proactive risk mitigation. This is operationally significant because NOTAM overload — the well-documented problem of dispatchers and flight crews receiving large volumes of low-relevance notices alongside safety-critical ones — has long been identified as a barrier to effective pre-flight risk assessment. AI-based filtering and prioritization of NOTAM content represents one of the more immediately actionable applications of machine learning to dispatch workflow, though the literature in this area remains largely industry-facing rather than peer-reviewed.

2.6 Human-AI Teaming and Automation Trust in Aviation Operations

As AI systems move from advisory tools to more active participants in operational decision-making, the human factors literature on automation trust and human-AI teaming becomes directly relevant to dispatch. Recent work applying a HABA-MABA-AABA (Humans Are Better At / Machines Are Better At / AI Are Better At) framework to flight deck operations argues that effective integration depends on recognizing the distinct capabilities and limitations of humans, automation, and AI, with task allocation, trust, and role clarity identified as the central design challenges. The same body of work cautions that allocating responsibilities to an AI teammate could degrade an operator's ability to quickly and accurately handle tasks that the AI is assigned, and that concerns about skill degradation in highly automated environments are well documented and reinforced by current research framing how AI-based task allocation may change operator responsibilities.

Broader human factors guidance on human-AI teaming in aviation emphasizes the need for realistic, human-plus-AI simulation environments in which operators can observe AI performance under both normal and degraded conditions. This literature notes that contemporary AI has well-known weaknesses, including data biases, edge or corner-case effects, and outright hallucinations, such that in the mid-term AI will almost certainly be partnered with human expertise whose judgment monitors and tempers AI outputs — a principle already reflected in regulatory requirements for human agency and oversight in safety-critical domains such as aviation. The same source warns that a low-situational-awareness condition combined with a sudden workload spike following loss of AI support, compounded by an inability of the AI to explain its recommendations, could constitute a significant safety hazard.

Research on adaptive human-AI teaming in air traffic control identifies a closely related tension: operational personnel often express high general trust in automation while simultaneously expressing low trust in specific novel forms of automated decision support, particularly where short-term tactical control is concerned. Controllers in one study were found to prefer retaining short-term tactical actions under their own control while being more willing to delegate longer-term strategic planning and routine tasks to automation — a distinction that maps closely onto the dispatch environment, where dispatchers may be more receptive to AI support for strategic fuel and route planning than for real-time, safety-critical flight-watch decisions during an active disruption.

2.7 Research Gaps and Synthesis

Several gaps emerge from this synthesis that are directly relevant to flight dispatch operations. First, the predictive analytics literature on delay and fuel prediction is methodologically mature but has been validated predominantly using historical schedule and weather datasets rather than within live OCC workflows; there is limited published evidence on how dispatchers actually use these predictions during flight release decisions, or how prediction outputs are reconciled with regulatory fuel policy. Second, the NDM literature on OCC decision-making provides a rich account of how experienced dispatchers reason about disruptions, but has not yet been extended to examine how the introduction of AI-generated recommendations changes the distributed cognitive process described in that literature — for example, whether AI recommendations become an additional information source that controllers must reconcile, or whether they begin to anchor the dispatcher's initial situation assessment in ways that could introduce confirmation bias.

Third, the human-AI teaming literature has developed largely around the flight deck and air traffic control, with comparatively little attention to the OCC and dispatch function specifically, despite the dispatcher's unique regulatory role as a co-equal decision authority with the pilot-in-command. Fourth, weather and NOTAM integration research is dominated by patent disclosures and industry commentary rather than peer-reviewed evaluation, leaving open questions about false-positive rates in AI-prioritized NOTAM feeds and the operational consequences of over- or under-alerting. These gaps motivate the conceptual framework developed in the following sections, which is intended to organize future empirical work around the specific decision points encountered in day-to-day dispatch.

3. Research Methodology

This paper adopts a conceptual research methodology rather than an empirical or experimental design. The objective is to construct a structured framework for comparing traditional dispatch operations with data-driven, AI-enabled dispatch operations across a set of performance dimensions that are meaningful to operational practice and that connect to the research streams identified in Section 2. This approach is appropriate at the current stage of AI adoption in OCCs, where deployments are heterogeneous across airlines, proprietary, and not yet subject to the kind of large-scale comparative evaluation that would support a quantitative meta-analysis.

3.1 Comparative Framework Construction

The conceptual framework was developed by mapping each stage of the dispatch workflow — pre-flight planning, flight release, in-flight monitoring, and post-flight review — onto the corresponding traditional and data-driven practices. Traditional practice descriptions reflect standard operational control procedures: manual review of weather charts and SIGMETs, fuel calculation using flight planning software with dispatcher-applied policy overlays (contingency fuel, alternate fuel, discretionary fuel), NOTAM review by individual reading, and flight watch conducted primarily through periodic position reports and ETA monitoring. Data-driven practice descriptions reflect the categories of AI-enabled tools identified in the literature reviewed in Section 2: predictive delay and holding models, AI-assisted fuel burn prediction, automated NOTAM prioritization, and continuous flight-monitoring platforms with anomaly detection.

3.2 Evaluation Criteria

Eight evaluation criteria were selected to span efficiency, safety, and decision-quality dimensions of dispatch performance:

1. **Dispatch Processing Time:** The elapsed time required for a dispatcher to construct and validate an operational flight plan, from initial weather review to flight release.
2. **Route Efficiency:** The degree to which the selected route minimizes flight time and track miles subject to operational, regulatory, and airspace constraints.
3. **Fuel Consumption:** Actual fuel burn relative to planned fuel, and the accuracy of fuel load recommendations relative to operational requirements.
4. **Delay Reduction:** The frequency and duration of departure and arrival delays attributable to dispatch-controllable factors.
5. **Dispatch Decision Accuracy:** The degree to which dispatch decisions (route, fuel, alternate selection) prove appropriate in light of conditions actually encountered.
6. **Situational Awareness:** The dispatcher's ability to maintain an accurate, current picture of all flights under their control, including emerging hazards.
7. **Risk Mitigation Effectiveness:** The degree to which operational risks (weather, NOTAM-driven, mechanical, crew) are identified and addressed prior to becoming safety events.
8. **Flight Completion Efficiency:** The proportion of flights completed as planned without diversion, return, or significant in-flight replanning.

For each criterion, the paper presents an illustrative comparison between traditional and data-driven dispatch, drawing on operational ranges that are representative of medium-to-large airline OCC environments rather than figures drawn from any single carrier's internal performance data. These ranges are presented as illustrative reference points to support discussion and hypothesis generation, not as validated empirical findings; the Limitations section addresses this distinction explicitly.

3.3 Assumptions and Analytical Logic

The framework assumes a medium-to-large network carrier operating a mixed narrow-body and wide-body fleet on a hub-and-spoke or point-to-point network, with an OCC staffed on a 24-hour rotating shift basis and dispatchers responsible for flight following across multiple simultaneous flights. It assumes that data-driven tools are deployed as decision-support layers integrated with, rather than replacing, the existing flight planning system, consistent with current regulatory frameworks that retain the dispatcher as the accountable decision-maker for flight release. The analytical logic proceeds by identifying, for each evaluation criterion, the specific information-processing bottleneck in traditional practice, and then assessing how the categories of AI-enabled tools identified in the literature review address — or fail to fully address — that bottleneck.

4. Data-Driven Flight Dispatch Framework

This section presents the conceptual architecture underlying a data-driven flight dispatch environment, organized around four figures. Each figure represents a distinct layer of the overall framework: the integrated data architecture (Figure 1), the OCC decision-support relationships (Figure 2), the weather-informed route optimization process (Figure 3), and the broader operational efficiency enhancement framework within which dispatch sits (Figure 4).

4.1 Integrated Data-Driven Flight Dispatch Architecture

Figure 1 depicts the overall data architecture supporting AI-enabled dispatch. Four categories of input data — meteorological systems, NOTAM databases, aircraft performance systems, and airport operations data — feed into a dispatch decision engine responsible for integration, normalization, and rule-based pre-processing. From the decision engine, processed information flows into three functional modules: a risk assessment module that applies hazard scoring and threshold checks; an optimization layer that performs route and fuel optimization; and a flight monitoring platform that provides real-time tracking and alerting. The outputs of these modules converge on the dispatcher decision output — the flight plan, fuel load, route, and risk briefing that the dispatcher reviews, modifies as necessary, and uses as the basis for flight release. A feedback loop returns post-flight performance data (actual fuel burn, actual routing, realized delays) to the input layer, enabling the underlying predictive models to be retrained on operational outcomes.

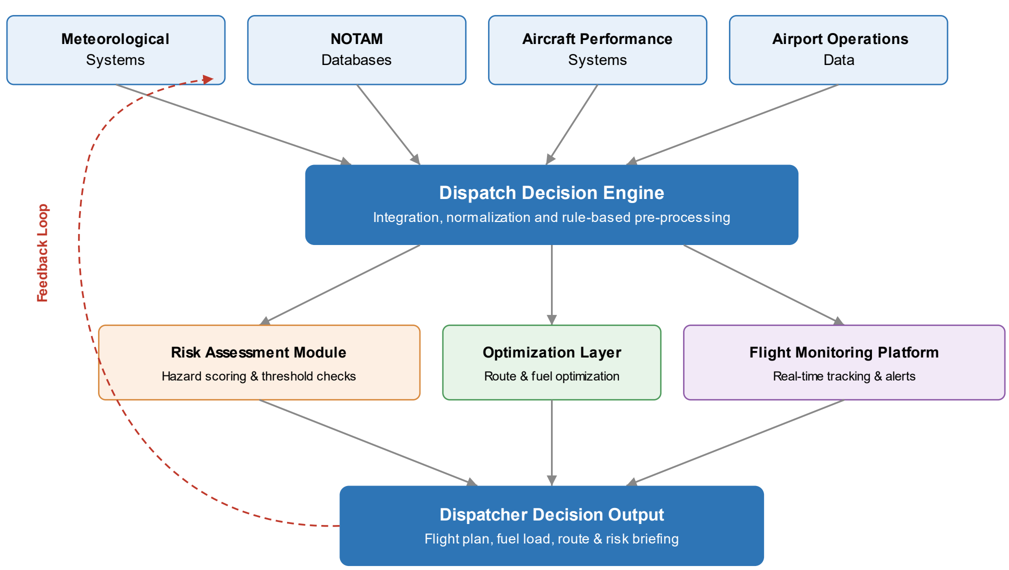

Figure 1. Integrated Data-Driven Flight Dispatch Architecture

Figure 1. Integrated Data-Driven Flight Dispatch Architecture

A critical feature of this architecture, from an operational standpoint, is that the dispatcher decision output remains a distinct stage rather than a terminal automated action. This reflects the regulatory requirement that a licensed dispatcher or FOO retains accountability for flight release; the architecture is therefore deliberately structured to position AI-generated outputs as inputs to dispatcher judgment rather than as substitutes for it.

4.2 Operational Control Center Decision-Support Framework

Figure 2 focuses on the relationships within the OCC itself. The operational control center functions as the central coordination hub, connected to the dispatcher and to flight crew as the two primary human actors in the decision loop. Four information domains feed into this hub: weather intelligence (SIGMET, METAR, and TAF feeds), fuel planning (burn forecasts and tankering analysis), risk assessment (flight risk assessment tool, or FRAT, scoring and hazard flags), and operational monitoring (ADS-B-based tracking and estimated time of arrival alerts). These domains converge on a decision outputs layer encompassing flight release, route revisions, fuel orders, and hold or diversion advisories.

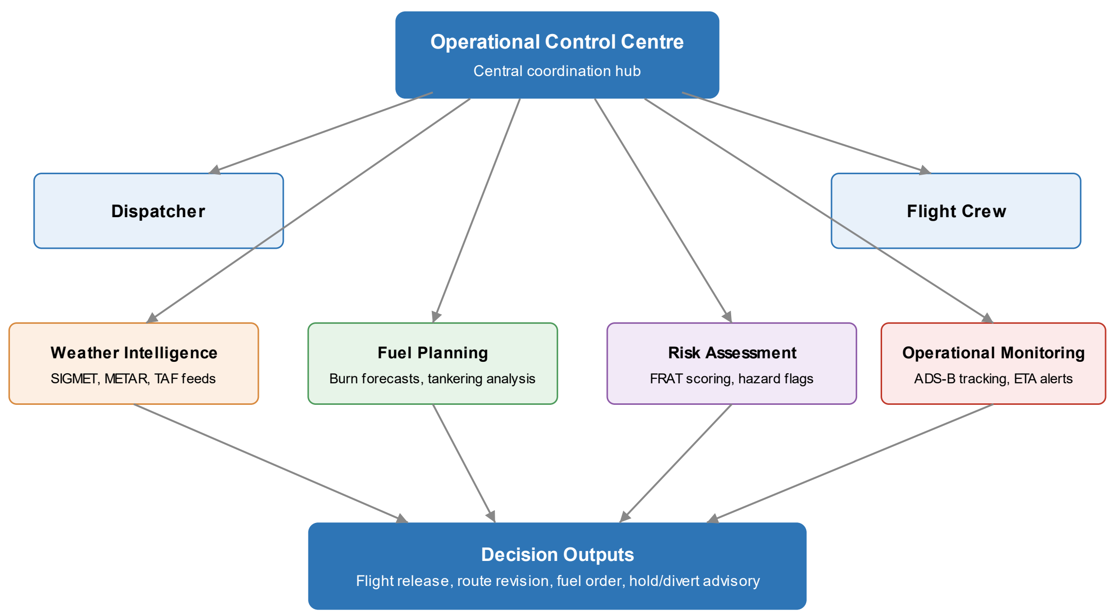

Figure 2. Operational Control Center Decision-Support Framework

Figure 2. Operational Control Center Decision-Support Framework

This framework highlights a point emphasized in the NDM literature reviewed in Section 2.2: the OCC is not a single decision-maker but a network of roles, each of which may interact with AI-generated information differently. A network controller monitoring fleet-wide connectivity impacts, for example, may weight an AI-generated delay prediction differently than a dispatcher focused on a single flight's fuel sufficiency. Effective integration of AI tools into the OCC therefore requires attention to how outputs are presented to each role, not merely whether the underlying prediction is accurate.

4.3 Weather-Informed Route Optimization Model

Figure 3 details the weather-informed route optimization process, which represents one of the more mature applications of data-driven dispatch tools. Weather inputs are processed through a hazard assessment stage that classifies convective activity, turbulence, icing, and other hazards along candidate routes. Route generation produces a set of candidate trajectories, which are evaluated through fuel impact analysis and checked against operational constraints — including airspace restrictions, aircraft performance limitations, and NOTAM-driven closures. These streams converge on an optimization engine implementing a multi-objective cost-index solver, which balances fuel consumption, flight time, and operational cost index targets to produce a final route selection.

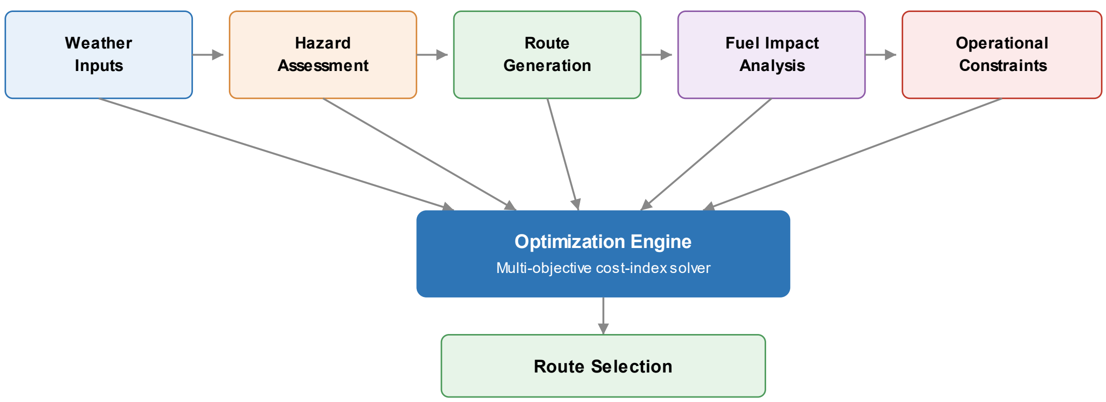

Figure 3. Weather-Informed Route Optimization Model

In operational practice, this process does not eliminate the need for dispatcher review of the selected route. Rather, it reduces the number of candidate routes that the dispatcher must manually evaluate from a potentially large set of viable alternatives to a small set of optimization-ranked options, each accompanied by the fuel and hazard rationale underlying its ranking. This preserves the dispatcher's role as the final arbiter while substantially reducing the manual workload associated with route comparison, particularly on long-haul sectors where multiple track structures and step-climb profiles may be viable.

4.4 Operational Efficiency Enhancement Framework

Figure 4 situates dispatch-specific tools within the broader operational efficiency enhancement framework of the airline. Predictive analytics sits at the center of this framework, with bidirectional relationships to dispatch intelligence, safety management, fuel optimization, delay mitigation, operational control, and performance monitoring. A continuous improvement loop connects fuel optimization and delay mitigation outcomes back into the predictive analytics core, reflecting the feedback relationship already described in Figure 1 at the architecture level.

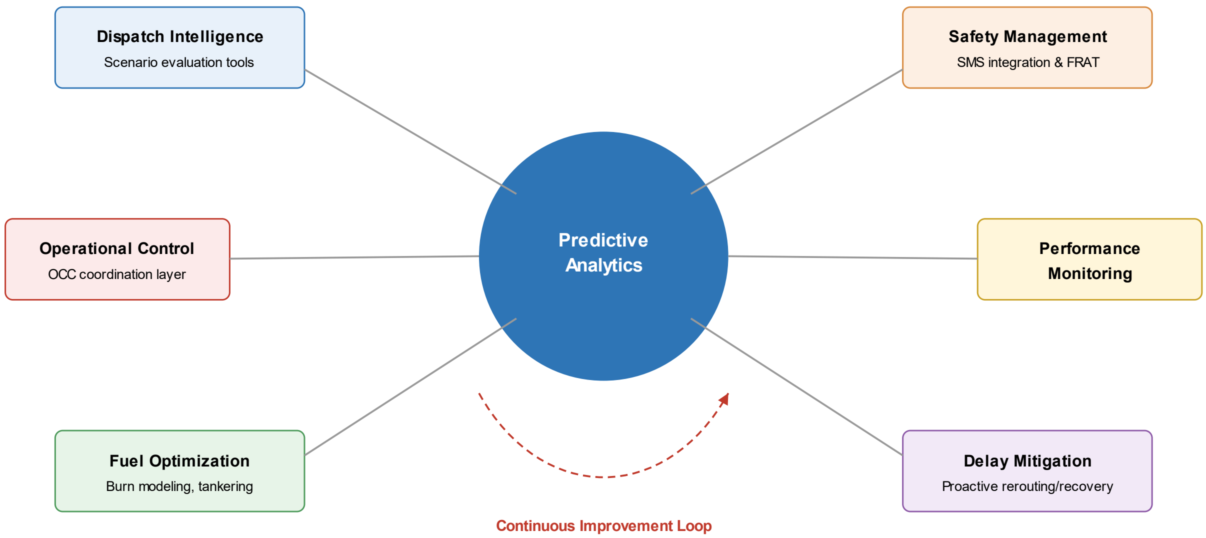

Figure 4. Operational Efficiency Enhancement Framework

This broader framing matters because dispatch does not operate in isolation. A route optimization recommendation that improves fuel efficiency for a single flight may have network-level implications — for example, an earlier or later arrival time that affects downstream crew legality, gate availability, or connecting passenger itineraries. The operational efficiency enhancement framework is intended to make explicit that dispatch-level AI tools must be evaluated not only on flight-level metrics but on their interaction with the broader operational control function.

5. Results and Discussion

This section presents the comparative analysis across the eight evaluation criteria introduced in Section 3.2, supported by four tables. The discussion that follows each table interprets the comparison in light of operational practice and the literature reviewed in Section 2, and critically considers both benefits and limitations.

5.1 Operational Challenges in Modern Flight Dispatch Operations

Table 1 summarizes the principal operational challenges currently facing flight dispatch, organized by category, with an indication of how each challenge is addressed (or not addressed) by current data-driven tools.

Challenge Category	Description	Operational Impact	Current Mitigation Approach
Weather data volume	Dispatchers must reconcile multiple, sometimes conflicting, weather products (model output, SIGMET, PIREP) across long-haul routes.	Increased flight plan preparation time; risk of overlooking localized hazards.	AI-assisted weather summarization and hazard flagging integrated into flight planning displays.
NOTAM overload	High volume of low-relevance NOTAMs alongside safety-critical notices for departure, en-route, and destination/alternate aerodromes.	Risk of missing operationally significant NOTAMs amid routine administrative notices.	AI-based NOTAM prioritization and relevance scoring, still maturing industry-wide.
Fuel policy complexity	Contingency, alternate, discretionary, and tankering fuel decisions must reconcile company policy, regulatory minimums, and route-specific risk.	Risk of over- or under-fueling; over-fueling imposes recurring fuel cost penalties.	AI-assisted fuel burn prediction informed by historical route and aircraft-specific performance data.
Multi-flight workload	A single dispatcher may simultaneously monitor 10-20+ active flights during peak operations.	Reduced time available per flight for situational assessment, particularly during irregular operations.	Automated flight monitoring platforms with exception-based alerting to focus attention on flights requiring intervention.
Irregular operations coordination	Disruption events (diversions, mechanical issues, crew legality) require rapid coordination across dispatch, maintenance control, and crew scheduling.	Time pressure increases risk of miscommunication between OCC functions.	Shared decision-support platforms integrating maintenance, crew, and dispatch data, as identified in the OCC literature.

Algorithmic opacity	AI-generated recommendations (route rankings, fuel predictions) may not be fully explainable to the dispatcher.	Risk of automation complacency or, conversely, dispatcher distrust and disuse of valid recommendations.	Explainability layers presenting underlying rationale (weather inputs, historical deviation) alongside recommendations; an area requiring further development.
Regulatory acceptance	Dispatch decision authority is regulatorily assigned to the licensed dispatcher/FOO; AI tools must operate within this framework.	Constrains the degree of autonomy that can be delegated to AI systems regardless of technical capability.	AI positioned strictly as advisory input to dispatcher-of-record decision, preserving human accountability.

The challenges summarized in Table 1 share a common thread: most are not purely technical problems but problems of information presentation, workload distribution, and accountability allocation. This is consistent with the NDM literature's emphasis on the distributed, multi-actor nature of OCC decision-making. A technically accurate AI prediction that is poorly integrated into the dispatcher's workflow — for example, presented in a separate system requiring manual cross-referencing — may produce little operational benefit and could even increase workload during high-tempo operations.

5.2 Traditional Dispatch versus Data-Driven Dispatch Systems

Table 2 compares traditional and data-driven dispatch practice across the eight evaluation criteria defined in Section 3.2. The figures presented are illustrative operational ranges intended to support discussion of the direction and approximate magnitude of expected differences, not validated measurements from a specific carrier.

Evaluation Criterion	Traditional Dispatch	Data-Driven Dispatch	Illustrative Direction of Change
Dispatch Processing Time	Manual weather review, NOTAM reading, and fuel calculation typically require 20-40 minutes per long-haul flight plan.	Pre-processed weather summaries, prioritized NOTAMs, and AI fuel predictions reduce manual review time, with dispatcher focus shifted to exception review.	Reduction in active review time per flight, with time reallocated toward flights flagged as higher-risk.

Route Efficiency	Route selection from a limited set of pre-filed or commonly used tracks, manually compared against weather charts.	Optimization engine evaluates a wider set of candidate routes against weather, fuel, and constraint data, presenting a ranked subset.	Modest improvement in track-mile and fuel efficiency on routes with multiple viable track options (e.g., North Atlantic, Pacific).
Fuel Consumption	Fuel loads based on flight planning system output plus dispatcher-applied policy buffers, calibrated primarily on aggregate fleet experience.	AI-assisted fuel prediction incorporates route-, aircraft-, and season-specific historical deviation between planned and actual burn.	Reduction in average over-fueling on routes with consistent historical bias, without compromising contingency fuel policy compliance.
Delay Reduction	Delay causes identified largely after the fact through standard delay coding; proactive identification depends on individual dispatcher experience.	Predictive delay and holding models flag elevated delay probability before departure, enabling proactive fuel and routing adjustments.	Earlier identification of delay-prone flights, enabling targeted mitigation rather than uniform buffer application.
Dispatch Decision Accuracy	Decision accuracy depends on the individual dispatcher's experience and the currency of information manually reviewed.	AI-generated recommendations provide a consistent baseline informed by historical outcomes, reviewed and adjusted by the dispatcher.	More consistent baseline decision quality across dispatchers of varying experience levels, with experienced dispatchers retaining override authority.
Situational Awareness	Awareness maintained through periodic position reports, ETA monitoring, and dispatcher-initiated checks.	Continuous flight monitoring platform with exception-based alerting surfaces deviations from	Earlier detection of in-flight deviations (route, fuel burn, ETA) requiring dispatcher attention.

		plan in near-real time.	
Risk Mitigation Effectiveness	Risk identification relies on dispatcher review of weather, NOTAMs, and aircraft status against experience-based heuristics.	Risk assessment module applies consistent hazard scoring across all flights, surfacing combinations of factors that may not trigger individual heuristics.	More consistent identification of compound risk scenarios (e.g., marginal weather plus a relevant MEL item) across the full flight schedule.
Flight Completion Efficiency	Diversions, returns, and significant in-flight replanning addressed reactively as they occur.	Pre-flight risk and weather assessment reduces the proportion of flights released into conditions likely to require in-flight replanning.	Modest reduction in weather-related diversions and holding-pattern fuel burn on routes with reliable predictive coverage.

The comparison in Table 2 reflects a consistent pattern: data-driven dispatch tools primarily shift the dispatcher's effort from broad manual review toward exception handling and judgment application on flagged cases. This is operationally desirable in principle, since it concentrates expert attention where it adds the most value. However, this shift also introduces a new risk identified in the human-AI teaming literature — if the AI system fails to flag a genuine risk (a false negative), the dispatcher's reduced engagement with the routine review process may mean that risk goes undetected, whereas under traditional practice the dispatcher's manual review, however time-consuming, provided a more uniform level of scrutiny across all flights.

5.3 Impact of Predictive Analytics on Flight Dispatch Performance

Table 3 examines the impact of specific predictive analytics applications — delay prediction, fuel burn prediction, NOTAM prioritization, and holding pattern prediction — on dispatch performance, drawing on the categories of models identified in Section 2.3 and 2.4.

Predictive Analytics Application	Underlying Model Approach (per literature)	Dispatch Decision Affected	Operational Benefit	Operational Limitation
Arrival delay prediction	Ensemble methods (e.g., Random Forest, Decision Tree) with resampling	Contingency fuel decision; passenger connection risk assessment.	Earlier identification of flights at elevated delay risk,	Predictive accuracy reported in the literature (e.g., ~0.90

	for imbalanced data; hybrid LSTM/GCN/3D-CNN architectures for spatio-temporal patterns.		supporting proactive fuel and connection planning.	accuracy/F1 in controlled evaluations) reflects historical datasets; live operational accuracy may vary by route and season.
Holding pattern prediction	Graph-based machine learning models applied to large trajectory datasets from major airports.	Contingency fuel allocation for destination-area holding.	Fuel planning that accounts for route- and time-specific holding probability rather than a uniform buffer.	Model performance demonstrated primarily on European airport trajectory data; transferability to other regions requires validation.
Fuel burn prediction	Predict-then-optimize approaches using historical flight data, aircraft performance, and weather inputs.	Fuel load recommendation (trip, contingency, and discretionary fuel).	Reduction in systematic over- or under-fueling relative to flight-planning-system defaults on routes with stable historical deviation patterns.	Requires sufficient route- and aircraft-specific historical data; less reliable for new routes, new aircraft types, or atypical operating conditions.
NOTAM prioritization	AI-based pattern analysis of NOTAM content combined with automation of dissemination.	Pre-flight risk assessment; identification of operationally significant restrictions.	Reduced risk of overlooking safety-critical NOTAMs amid high-volume administrative notices.	Industry-facing descriptions of this capability remain largely vendor-

				driven; peer-reviewed validation of prioritization accuracy is limited.
Contrail-aware route optimization	Optimal control formulations treating contrail-favorable regions as soft constraints alongside wind-optimal trajectory solutions.	Route and altitude selection on applicable long-haul sectors.	Potential reduction in non-CO_2 climate impact alongside conventional fuel-efficiency objectives.	Adds an additional optimization objective that may, in some cases, trade off against fuel-minimal routing, requiring explicit policy on relative weighting.

A recurring theme in Table 3 is the gap between model performance reported under controlled evaluation conditions and the conditions under which dispatch decisions are actually made. Predictive models for delay and holding are generally trained and validated on historical datasets from specific airports or regions; their performance when applied to a different airline's network, fleet mix, or operating environment is not guaranteed to match published results. This is not a reason to avoid such tools, but it does suggest that initial deployment should be accompanied by airline-specific validation — comparing model predictions against actual outcomes on the airline's own routes before the predictions are used to materially adjust fuel or routing decisions.

The contrail-aware optimization case in Table 3 also illustrates a broader point about multi-objective optimization in dispatch: as additional objectives (environmental impact, noise abatement, crew duty time) are incorporated into route and altitude optimization, the dispatcher's role increasingly involves understanding and, where necessary, adjusting the relative weighting applied to these objectives — a responsibility that requires the dispatcher to understand the optimization logic at a conceptual level, even if not at the level of the underlying algorithm.

5.4 Operational Efficiency, Safety, and Fuel Performance Improvements

Table 4 synthesizes the preceding analysis into an overall assessment of efficiency, safety, and fuel performance improvements that may be achievable through data-driven dispatch, alongside the conditions under which each improvement is most likely to be realized.

Performance Dimension	Improvement Mechanism	Illustrative Magnitude (operational range)	Conditions for Realization
Flight plan preparation efficiency	Pre-processed weather, NOTAM, and fuel data reduce manual review time per flight plan.	Reduction in active dispatcher review time, reallocated toward flagged-exception flights.	Requires integration of AI outputs into the primary flight planning display rather than a separate system.
Fuel cost reduction	Route- and aircraft-specific fuel burn prediction reduces systematic over-fueling on routes with consistent historical bias.	Small but recurring per-sector fuel savings on high-frequency routes with stable historical patterns; magnitude varies substantially by fleet and network.	Sufficient historical data volume per route/aircraft combination; dispatcher confidence in predictions sufficient to adjust discretionary fuel.
Delay-related cost reduction	Earlier identification of delay-prone flights enables proactive crew, gate, and connection management.	Reduction in downstream disruption costs (missed connections, crew duty extensions) for flights correctly flagged in advance.	Sufficiently low false-positive rate to avoid unnecessary proactive interventions on flights that would not have been delayed.
Situational awareness during irregular operations	Continuous monitoring platform with exception-based alerting surfaces deviations across the full flight schedule simultaneously.	Earlier dispatcher awareness of in-flight deviations, particularly during high-tempo multi-flight disruption events.	Alert thresholds calibrated to avoid alert fatigue; alerts presented with sufficient context for rapid assessment.
Risk identification consistency	Risk assessment module applies uniform hazard	More consistent identification of compound risk	Risk scoring logic validated against actual safety

	scoring across all flights regardless of individual dispatcher workload.	factors across the full operation, independent of any individual dispatcher's available time.	event data; scoring thresholds reviewed periodically by safety management system processes.
Weather-related diversion/holding reduction	Pre-flight weather-informed routing and holding prediction reduce the proportion of flights encountering unanticipated weather-driven replanning.	Modest reduction on routes and seasons with reliable predictive coverage; limited effect on genuinely novel weather patterns.	Predictive models retrained regularly on recent weather and operational outcome data; dispatcher retains authority to override on emerging information.

Three observations follow from Table 4. First, the improvements identified are consistently described in conditional terms — dependent on data availability, alert calibration, and dispatcher trust — rather than as guaranteed outcomes of deploying AI tools. This reflects both the conceptual nature of the framework and a genuine feature of the underlying literature, which similarly frames AI benefits in aviation as contingent on integration quality rather than automatic. Second, the magnitude of fuel and cost-related improvements is likely to be route- and network-dependent: airlines operating dense, high-frequency networks with substantial historical data per route are better positioned to realize fuel prediction benefits than airlines operating predominantly seasonal or low-frequency routes. Third, the safety-related improvements — situational awareness and risk identification consistency — are arguably the most operationally significant, since they address the multi-flight workload challenge identified in Table 1 directly, but they are also the hardest to quantify in advance of deployment, since their value is realized primarily in the relatively infrequent cases where a risk would otherwise have gone unnoticed.

5.5 Illustrative Operational Scenario

To ground the preceding analysis, consider an illustrative scenario representative of operations in the Middle East/South Asia traffic flow during the summer convective season. A dispatcher releasing a flight from Doha to a South Asian destination during the afternoon monsoon period must account for the likelihood of convective activity near the destination at the planned arrival time, the fuel implications of potential holding or diversion to an alternate, and any NOTAMs affecting the destination's approach procedures. Under traditional practice, the dispatcher reviews available SIGMETs, satellite imagery, and terminal area forecasts manually, applies a contingency fuel buffer based on general policy and personal experience with the destination's seasonal weather behavior, and reviews the NOTAM list for the destination and alternates.

Under a data-driven dispatch framework consistent with Figures 1-3, the dispatcher would instead receive a pre-processed weather hazard summary specific to the route and arrival window, a fuel recommendation incorporating the historical relationship between forecast and actual conditions at that destination during the monsoon season, a prioritized NOTAM list highlighting any approach-relevant restrictions, and — if a holding pattern prediction model is deployed — an estimate of holding probability and expected holding duration based on similar historical days. The dispatcher's task shifts from assembling this picture from multiple sources to reviewing, validating, and where necessary adjusting a pre-assembled picture, with the final fuel and routing decision remaining the dispatcher's responsibility.

This scenario illustrates both the potential benefit — a more complete and consistent risk picture assembled in less time — and the corresponding risk: if the underlying predictive models have not been validated specifically for this destination's monsoon-season behavior, a confidently presented but locally inaccurate prediction could lead the dispatcher to apply less scrutiny than the traditional manual process would have prompted. This underscores the importance, raised in Section 5.3, of route- and season-specific validation before predictive outputs are relied upon for fuel and routing decisions in operationally significant scenarios.

6. Industry Implications

6.1 OCC Design and Workflow Integration

The analysis in Section 5 suggests that the operational value of AI-enabled dispatch tools depends heavily on how they are integrated into existing OCC workflows and display environments. Tools that require dispatchers to consult a separate system, outside the primary flight planning and monitoring displays, are likely to see limited adoption regardless of underlying model accuracy, consistent with the broader human factors finding that decision support is most effective when embedded within the operator's existing task structure rather than added as a parallel process. Airlines considering AI dispatch tools should therefore prioritize integration with existing flight planning, weather, and NOTAM systems over standalone analytics platforms, even where the latter may offer more sophisticated modeling.

6.2 Training and Competency Development

The introduction of AI-generated recommendations into dispatch workflow has implications for dispatcher training and competency maintenance that extend beyond familiarization with new software. Training programs should explicitly address the conditions under which AI recommendations are likely to be less reliable — new routes, atypical weather patterns, periods following fleet or network changes — so that dispatchers develop calibrated trust rather than uniform trust or uniform skepticism. This is consistent with the human-AI teaming literature's emphasis on operators understanding both the capabilities and the failure modes of AI systems they work alongside.

A related training implication concerns the maintenance of manual dispatch skills. If AI tools substantially reduce the routine manual review burden, airlines should consider how dispatchers maintain proficiency in manual weather interpretation, NOTAM review, and fuel calculation — skills that remain necessary during AI system outages and that underpin the dispatcher's ability to critically evaluate AI outputs. This mirrors long-standing concerns in the flight deck automation literature regarding manual flying skill retention in highly automated cockpits.

6.3 Safety Management System Integration

The risk assessment module described in Figure 1 and the risk mitigation effectiveness criterion in Table 2 should be integrated with the airline's broader safety management system (SMS) rather than operated as a standalone dispatch tool. Hazard scoring logic embedded in AI dispatch tools constitutes a safety-relevant algorithm whose performance should be monitored through the same data-driven safety assurance processes applied to other SMS components — including periodic review of cases where the tool's risk scoring diverged from actual outcomes, whether through false negatives (missed risks) or false positives (unnecessary alerts that may contribute to alert fatigue).

6.4 Regulatory and Certification Considerations

Because dispatch decision authority is regulatorily assigned to the licensed dispatcher or FOO, the introduction of AI tools into dispatch does not, under current frameworks, change the locus of accountability for flight release decisions. However, regulators and airlines will need to develop clear documentation standards for how AI-generated recommendations are presented, how dispatchers are expected to review and document their assessment of those recommendations (particularly where a recommendation is overridden), and how the performance of AI tools is monitored over time. The absence of such documentation standards could create ambiguity in post-incident investigations regarding the role that an AI recommendation played in a dispatch decision.

7. Limitations and Future Research

This paper has several limitations that should be addressed by future research. First, the comparative framework presented in Section 5 is conceptual and illustrative; the operational ranges presented in Tables 2 through 4 are intended to support discussion of direction and approximate magnitude, not to serve as validated benchmarks. Future research should pursue empirical validation through controlled comparisons — for example, comparing fuel load accuracy, dispatch processing time, and risk identification rates before and after deployment of specific AI tools within a single airline's OCC, using a consistent measurement methodology.

Second, the literature reviewed in Section 2 is unevenly distributed across the topics relevant to this paper. Predictive analytics for delay and fuel prediction is supported by a substantial and growing peer-reviewed literature, while weather and NOTAM integration is addressed primarily through patent disclosures and industry commentary. Human-AI teaming research relevant to dispatch specifically — as distinct from the flight deck or air traffic control — remains particularly sparse. Future research should prioritize dispatch-specific studies of human-AI interaction, ideally using the naturalistic decision-making methodologies that have proven productive in studying OCC decision-making more generally.

Third, this paper has not addressed the cybersecurity implications of increased data integration across meteorological, NOTAM, aircraft performance, and airport operations systems, which represents an additional risk dimension as OCC architectures become more interconnected. Fourth, the paper's framing assumes a medium-to-large network carrier; the applicability of the framework to smaller carriers, charter operators, or cargo operators — which may have different data volumes, route structures, and OCC staffing models — has not been examined and would benefit from separate study.

Finally, future research should examine the longitudinal effects of AI integration on dispatcher expertise development. If AI tools handle an increasing share of routine cases, newly licensed dispatchers may have fewer opportunities to develop the recognition-primed expertise that the NDM literature identifies as central to effective performance during non-routine events. Understanding whether, and how, this expertise can be developed in an AI-supported environment — potentially through deliberate exposure to historical case scenarios during training — represents an important area for future work.

8. Conclusion

This paper has examined the integration of artificial intelligence into airline operational control centers, with a focus on flight dispatch operations. Drawing on a critical synthesis of the literature on decision support systems, naturalistic decision-making in OCCs, predictive analytics for delay and fuel prediction, weather and NOTAM data integration, and human-AI teaming in aviation, the paper developed a conceptual framework comparing traditional and data-driven dispatch practice across eight performance dimensions. The analysis suggests that AI-enabled dispatch tools hold genuine potential to reduce flight plan preparation time, improve the consistency of fuel and route recommendations, and enhance situational awareness across the multi-flight workload that characterizes modern OCC operations.

At the same time, the analysis identifies that these benefits are conditional on careful workflow integration, route- and season-specific validation of predictive models, and training that develops calibrated rather than uniform trust in AI outputs. The dispatcher's regulatory role as a co-equal decision authority with the pilot-in-command should remain the organizing principle for AI integration: AI tools are most valuable, and most consistent with current safety frameworks, when positioned as decision aids that expand the dispatcher's effective situational picture rather than as autonomous decision-makers. Future research grounded in operational OCC environments, using the naturalistic decision-making and human-AI teaming methodologies identified in this review, will be necessary to validate these conceptual findings and to guide the next generation of dispatch decision-support systems.

REFERENCES

Bratu, S., & Barnhart, C. (2006). Flight operations recovery: New approaches considering passenger recovery. Journal of Scheduling, 9(3), 279-298.

Castro, A. J. M., & Oliveira, E. (2008). The rationale behind the development of an airline operations control centre using gaia methodology. International Journal of Agent-Oriented Software Engineering, 2(3), 350-377.

Fogaça, L. B., Henriqson, E., Carim Junior, G., & Lando, F. (2022). Airline disruption management: A naturalistic decision-making perspective in an operational control centre. Journal of Cognitive Engineering and Decision Making.

Hassan, M. A., Larsson, R., Tongur, V., & Mårtensson, J. (2021). Fuel-optimal cruise control for diesel-electric class 8 trucks. Control Engineering Practice.

Klein, G. (2008). Naturalistic decision making. Human Factors, 50(3), 456-460.

Klein, G., Orasanu, J., Calderwood, R., & Zsambok, C. E. (Eds.). (1993). Decision making in action: Models and methods. Ablex Publishing.

Koornneef, H., et al. (2020). A decision support framework and prototype for aircraft dispatch assessment. Reliability Engineering & System Safety, 200, 106932.

Lan, S. (2003). Planning for robust airline operations: Optimizing aircraft routings and flight departure times to achieve minimum passenger disruptions (Doctoral dissertation, Massachusetts Institute of Technology).

Lettovsky, L. (1997). Airline operations recovery: An optimization approach (Doctoral dissertation, Georgia Institute of Technology).

Patriarca, R., Di Gravio, G., Woltjer, R., Costantino, F., Praetorius, G., Ferreira, P., & Hollnagel, E. (2020). Framing the FRAM: A literature review on the functional resonance analysis method. Safety Science, 129, 104827.

Rosenberger, J. M., Johnson, E. L., & Nemhauser, G. L. (2003). Rerouting aircraft for airline recovery. Transportation Science, 37(4), 408-421.

Richters, F., Schraagen, J. M., & Heerkens, H. (2016). Assessing the structure of non-routine decision processes in airline operations control. Ergonomics, 59(3), 380-392.

Teodorovic, D., & Stojkovic, G. (1995). Model to reduce airline schedule disturbances. Journal of Transportation Engineering, 121(4), 324-331.

Yu, G., Argüello, M., Song, G., McCowan, S. M., & White, A. (2003). A new era for crew recovery at Continental Airlines. Interfaces, 33(1), 5-22.

AlBassam, S. A. A., & AlShahrani, D. N. (2025). Flight delay prediction: Evaluating machine learning algorithms for enhanced accuracy. PLOS ONE, 20(12), e0335141.

Belcastro, L., Marozzo, F., Talia, D., & Trunfio, P. (2017). Using scalable data mining for predicting flight delays. ACM Transactions on Intelligent Systems and Technology, 8(1), 1-20.

Okwir, S., Amouzgar, K., & Ng, A. H. (2025). Exploring prediction accuracy for optimal taxi times in airport operations using various machine learning models. Journal of Air Transport Management, 122, 102684.

Song, C., Ma, X., Ardizzone, C., & Zhuang, J. (2024). The adverse impact of flight delays on passenger satisfaction: An innovative prediction model utilizing wide & deep learning. Journal of Air Transport Management, 114, 102511.

Wong, J.-T., & Tsai, S.-C. (2012). A survival model for flight delay propagation. Journal of Air Transport Management, 23, 5-11.

Springer Nature. (2025). Prediction and optimization of civil aviation flight delays based on machine learning algorithms. International Journal of Computational Intelligence Systems.

Graph machine learning for flight delay prediction due to holding manoeuvre. (2025). arXiv preprint arXiv:2502.04233.

A reliable predict-then-optimize approach for minimizing aircraft fuel consumption. (2025). Transportation Research Part C: Emerging Technologies.

Khan, W. A., Ma, H. L., Chung, S. H., & Wen, X. (2021). Hierarchical integrated machine learning model for predicting flight departure delays and duration in series. Transportation Research Part C: Emerging Technologies, 129, 103225.

Xue, D., Ng, K. K. H., & Hsu, L. T. (2020). Multi-objective flight altitude decision considering contrails, fuel consumption and flight time. Sustainability, 12(15), 6253.

Sustainable aviation practices and decision support systems for eco-efficiency assessment: A small-scale literature review. (2014). Air transport research literature synthesis.

Korentsides, J., et al. (2026). The use of artificial intelligence (AI) in the flight deck: Enhancing human-AI teamwork in aviation. Journal of the Air Transport Research Society, 6, 100099.

Human factors requirements for human-AI teaming in aviation. (2025). Aerospace, 5(2), 42.

A framework for supporting adaptive human-AI teaming in air traffic control. (2023). Eurocontrol/SESAR research publication.

Flight-based control allocation: Towards human-autonomy teaming in air traffic control. (2024). Aerospace, 11(11), 919.

Endsley, M. R. (2017). From here to autonomy: Lessons learned from human-automation research. Human Factors, 59(1), 5-27.

Bainbridge, L. (1983). Ironies of automation. Automatica, 19(6), 775-780.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230-253.

Dekker, S. W., & Woods, D. D. (1999). To intervene or not to intervene: The dilemma of management by exception. Cognition, Technology & Work, 1(2), 86-96.

Plant, K. L., & Stanton, N. A. (2017). The development of the Schema-Action-World (SAW) taxonomy for understanding decision making in aeronautical critical incidents. Safety Science, 99, 23-35.

Praetorius, G., Hollnagel, E., & Dahlman, J. (2015). Modelling vessel traffic service to understand resilience in everyday operations. Reliability Engineering & System Safety, 141, 10-21.