Deep Science Publishing, 2026
https://doi.org/10.70593/978-93-7185-114-5_7

Chapter 7: Intelligent Decision Support in Service and Production Environments

7.1. Introduction

Intelligent Decision Support Systems help manage dynamic systems in service and production environments. Evidence-Based arguments focus on Core Technologies, Applications in Service and Production, Architectural Considerations, and Evaluation Methodologies.

Evidence-based Intelligent Decision Support goes beyond Data-Driven Decision Making. The Decision Support dimension centers on Intelligent Systems in Service and Production Systems. In Service Contexts, Importance is placed on Forecasting Demand for Capacity Planning. In Production Contexts, Predictive Maintenance is critical for avoiding Downtime and Faulty Production. Thus, Intelligent Decision Support encompasses Evidence-Based Insights about Services in Dynamic Contexts.

7.1.1. Overview of Intelligent Decision Support Systems

Intelligent Decision Support Systems (IDSS) can be seen as knowledge generation devices that capitalize on the fact that decisions should be based on facts as much as possible and are more reliable if they can be backed by evidence. In other domains, such as the medical field, health professionals, doctors, and nurses want to base their assessments and recommendations on the best available information. They want skilled advisers to support them with their decisions, and these advisers ought to recommend actions with a sound and compelling scientific underpinning. In these specific domains, physicians and intelligent systems can work together and share knowledge; therefore, intelligent decision support systems can play an essential complementary role. In fact, any position of trust and authority within an organization would benefit from this kind of system usually reserved for medical activities. By compiling and analyzing all available data on that specialized unit, it can be shown to have refrained from expressing doubts about the action plan, aware that the recommended alternatives were based on hard data collected over the years or decades.

More broadly, systems can synthesize information from multiple data sources, including business records, public reports and news feeds, and external and national databases, such as bureau databases or the Census Bureau. Intelligent Decision Support Systems automatically retrieve these data sources, select the relevant type, and intelligently organize these data in a more effective structure that supports trial-and-error test capabilities of different decision alternatives. Evidently, intelligent decision support systems are not mandated to base their analysis or recommended actions on all possible data available.

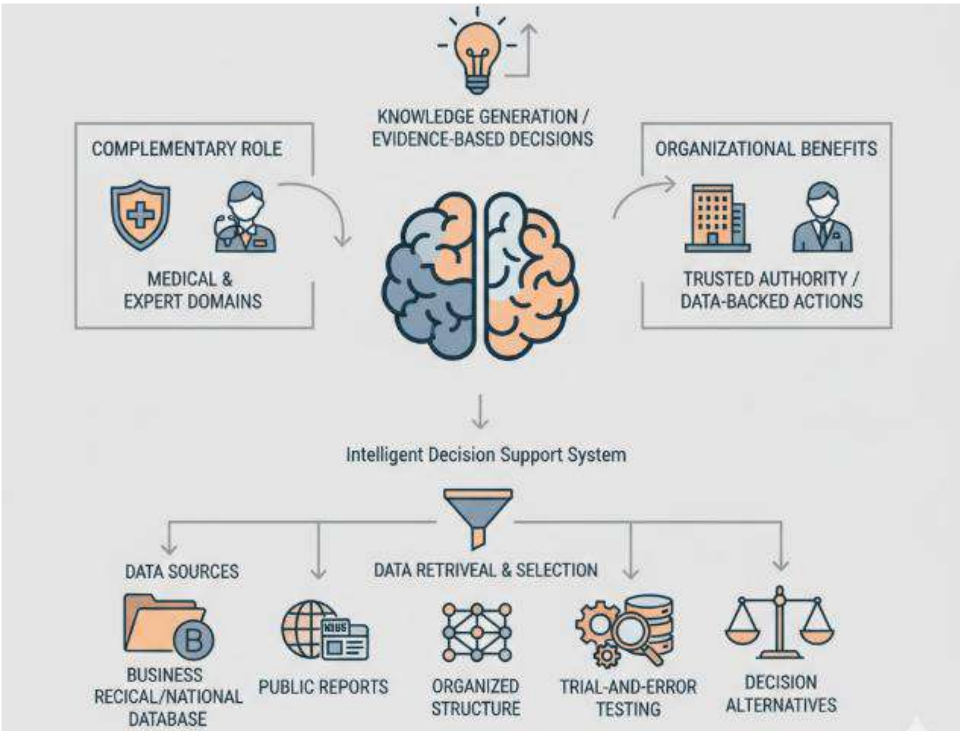

Fig 7.1: Epistemic Augmentation: A Framework for Evidence-Based Knowledge Generation in Intelligent Decision Support Systems

7.2. Foundations of Intelligent Decision Support

The foundations of Intelligent Decision Support can be viewed along two critical dimensions: data quality and governance, and data modeling and inference techniques. Evidence-based decision support requires a data foundation that is accurate, complete, up to date, consistent, relevant, and understandable. Integrating different data sets creates additional challenges around provenance, usage policies, semantic interoperability, and capture of additional context that enables correct interpretation. The use of multiple data models assists in capturing laws, relationships, and likelihoods among variables of interest, thereby expressing knowledge for use in identification, matching, and other kinds of inference. Data-driven, knowledge-driven, and model-driven inference techniques can help in making decisions when certain information is unavailable or uncertain.

Intelligent Decision Support in service environments is often focused on demand forecasting, capacity planning, resource allocation, or routing and scheduling of service delivery. Forecasting supply and demand is a critical task that underpins many of the decisions in these environments. Capacity often needs to be planned at a much lower frequency than the forecasting is done, depending on the cost and timeframe for acquiring new resources, yet does not always update to match fluctuations in demand forecast. Routing and scheduling problems must be solved rapidly using heuristics to meet dynamically changing demand. Common decision-support applications in production environments include predictive maintenance and failure detection, optimization of production processes, and allocation of resources to production jobs.

7.2.1. Data Quality and Governance

The utility and reliability of a decision support system depend fundamentally on data quality and governance. Data quality relies on a range of factors, including integrity, accuracy, consistency, completeness, time-consistency, and lineage. Integrity refers to satisfactory data structures and connections, whereas accuracy concerns the correspondence between data values and their real-world references. Consistency indicates a lack of contradictions within a data set and the normal relationship among different sets. Completeness indicates whether all required data are present at a given time. Time-consistency assesses whether data collected over the same time period fully mirror the actors being described. Finally, data lineage traces the origin of data and the set of processes that acted upon the data from origin to destination.

High data quality often requires a data governance framework that formally institutionalizes processes and policies regarding data control within an organization. The aim of data governance is to optimize data availability and ease of use while maintaining the required levels of privacy, confidentiality, and security. A mature data governance framework addresses issues related to people, processes, and data architecture, assigns individuals roles and responsibilities in relation with data, governs Master Data Management, establishes rules to ensure adequate data quality, and preserves data assets over time.

7.2.2. Modeling and Inference Techniques

Evidence-based arguments are central to intelligent decision support. Decision models formalize a decision situation and support deciding among available alternatives. Proponents of one alternative construct a dedicated model to demonstrate its superiority, without decision support if common knowledge suffices. Models also underpin behavioral explanations, generate partial utility functions, or help understand a complex setting. Static models verify performance for known inputs, while dynamic models rely on history. Simulation models estimate the outcomes of alternatives, removing common inputs, while deterministic models elucidate the choice for a specific state of nature.

Alternatives value an outcome by probabilistically weighting its utility. Probabilistic models for forecasting transitions or processes illuminate possible futures. Expert judgment augments limited sample data, commonly structured in a decision-analytic format. Bayesian networks generalize this idea, clustering the model in local probability tables. Sparse networks exploit context-specific independence or similar transitions. Theorem proving or model checking via reward machines controls concurrent systems with finite memory. Online learning adjusts decision-analytic weights and further leverages demonstration samples. Reinforcement learning establishes a policy from a known model and operates without knowledge or model.

7.3. Decision Support in Service Environments

Service environments create schedules and offers for service delivery to clients. Such decision support problems typically involve a sequence of decisions, starting with a long-term demand forecast, followed by a capacity planning, and then a detailed dynamic scheduling of operations. Since the context is that of decision support, the goal is not to precisely predict the future but rather to give estimates (forecasts) that are to be as accurate as possible for planning purposes.

The first set of decision problems corresponds to medium-term demand forecasting, capacity planning, and hiring, while the second involves a dynamic re-scheduling of the operations based on a sequence of up-to-date demand realizations. Demand forecasts based on state-space models with multivariate covariates indicate significant summer and winter peaks that need to be properly accounted for in the middle-term capacity planning. A real-world example of dynamic scheduling of school bus operations shows that the potential economic savings from a real-time re-routing of the buses during the day are indeed considerable.

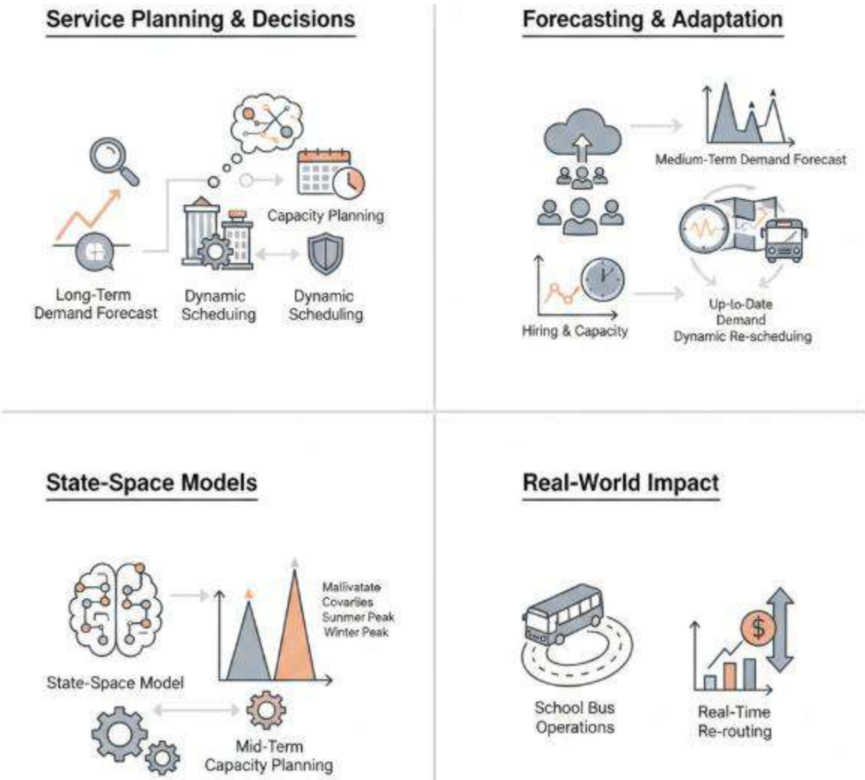

Fig 7.2: Dynamic Multi-Horizon Decision Support: Integrating State-Space Forecasting with Real-Time Operational Re-scheduling

7.3.1. Demand Forecasting and Capacity Planning

Followed by the forecasting step, the demand forecast can be defined as a time series of product/service requests or arrivals at the available or outsourceable resources. Based on these time-based arrival-distribution predictions, capacity should be planned for the next time interval or period. Such planning requires understanding whether all the incoming demand can be processed by the available/outsourceable resource assets at hand. Under-provisioning of the resources (not scaling for peak demand) leads to overstretched service times and customer dissatisfaction, while over-provisioning incurs additional costs and might be seen as an unpromarket move. Therefore, usually the capacity planning step decides whether additional resources, such as a scale-up (more resources for the same type of task), a scale-out (of the same resource type, e.g. Cloud solutions), or other resource type (for service, e.g. outsourcing) need to be provisioned.

Service-data-enabled automation now permits an even more dynamic approach to this task by feeding the forecasts directly. Such automated data-driven, IT (Information Technology)/OT (Operational Technology) integrated, and either self-tuning or pretuning solutions can now be termed DFO (Demand Forecasting and Order Book) automation. For highly variable demand patterns that proceed from ex-ante user-specified periods (weekdays working hours, weekends), hybrid fuels should be put into place with either fixed/elastic (bus services, e.g. BDC) or up-to consolidated capacity.

7.3.2. Dynamic Scheduling and Routing

Dynamic scheduling necessitates the concurrent assignment of multiple resources to time-bound tasks of variable duration in multi-constrained environments. Its main objective is to minimize total operational cost. Response time is a critical indicator of customer expectations in service systems involving multi-facility routing, such as airline flight scheduling. As most of the probable combinations of resource-task assignments are practically infeasible, complex problems in dynamic scheduling are more effectively solved using parallel lists organized by defined control structures. The resource-orientation approach achieves real-time stability, requiring only O(n log n) comparisons for n requests, and is applied in service environments subject to multiple constraints, i.e. travel time between facilities and available vehicle capacity.

Dynamic vehicle routing refers to the real-time determination of vehicle routes for a fleet of vehicles responding to requests for service coming from customers over time. Its goal is to minimize operational costs, including transportation cost and service delay, while satisfying constraints such as vehicle capacity and the precedence of service. At the center of the problem is the assignment of requests to vehicles and the determination of routing sequences that minimize the costs associated with the service. Dynamic vehicle routing can be considered as a multi-constrained problem scenario, known for its high complexity. Computational study has highlighted the effectiveness of Parallel Tabu Search in addressing complex problem scenarios.

7.4. Decision Support in Production Environments

Intelligent Decision Support finds broad applicability in production environments as well, where its primary aims are optimization and resource allocation. Predictive maintenance solutions create models for condition-based servicing, while anomaly detection models using time-series data serve to preview imminent problems. Process optimization yields optimal parameters for transforming input to output; resource allocation ensures both temporal feasibility and minimum costs.

Predictive maintenance systems serve to optimize servicing periods within the typical preventive maintenance strategy. Data describing component conditions may result from physical models or, more frequently, from sensor readings. Various data-fusion techniques create the final condition description, which typically still needs a-priori knowledge for interpretation. If an event-based condition description is available, a projection onto “normal” operation may assist in revealing the effect of deviation from normality on performance. Such a model can also serve to schedule the &% process and prepare for corrective measures. Neural networks with advanced architectures are well suited for representing functions in production processes, but their training requires high effort. Hybrid models—integrating physical knowledge with neural-networks learning capabilities—represent an interesting alternative.

Process optimization determines suitable values for parameters characterizing a process or service when transforming input into output to achieve some desired objective. Control parameters are defined during implementation, and control units are designed to apply adjustments during operation to keep the service within presumed limits. Resources—people, machines, and space— require scheduling to determine start times and assign resources to ensure temporal feasibility of the service plan with minimum costs.

7.4.1. Predictive Maintenance and Anomaly Detection

Intelligent Decision Support in service and production environments embraces using data and analysis for effective, holistic decision-making that improves efficiency and effectiveness. The services and operations research literature offers clear architectures for the Intelligent Decision Support activities in both service and production domains. The focus here is on predictive maintenance and anomaly detection in production environments.

Decision-support problems require predicting the future state and performance of a system for scenarios defined by the decision variables. Hence, maintenance schedules, potential fault alarms, and production service updates need to be anticipated in a timely manner. For service systems that operate under a threat of failure or which experience performance degradation, the scope of prediction is determined by factors such as service safety and customer goodwill. For production systems, supports such as physical-product condition, performance effectiveness, resource scarcity, and demand level shape the prediction task. These future conditions determine the appropriate assignment and execution of the task set of inspection, repair, replacement, and preventive maintenance.

7.4.2. Process Optimization and Resource Allocation

For production environments, Intelligent Decision Support encompasses medium- and short-term process and resource planning and management. These tasks, semantically close to the ones found in transportation and logistics environments, can be framed as optimization problems in manufacturing contexts. In particular, the planning and management of production processes are extremely resource-intensive, leading organizations to rely on Physical Internet-based transdisciplinary commons in order to offer service-based manufacturing models.

Process optimization and resource allocation problems are prevalent in both production and service environments. Indeed, in production, Intelligent Decision Support encompasses predictive modeling for process optimization and for the efficient allocation of human resources in short-term operations. Typical use cases involve the optimization of: (i) design processes, where sub-processes work during dedicated time intervals and require high expertise in specific areas; (ii) manufacturing processes, where the goal is to determine how many manufacturing and packaging lines should be activated over a time horizon in order to satisfy an uncertain demand; (iii) construction processes, with the aim of optimizing both the allocation of human resources within each activity and the sequence of activities that should be executed.

In close relationship with demand forecasting, Intelligent Decision Support contributes to process optimization by predicting the lead time of each activity and the general project lead time. Anomaly detection in process execution represents another relevant use case, focusing particularly on activities that require a high-level skill set. Smart contracts implemented in a Blockchain context can also be considered components for the Intelligent Decision Support of process optimization and resource allocation in construction, as they enable the automated control of material supply during the execution of the project.

7.5. Architectures for Intelligent Decision Support

Evidence-based arguments call for discussion of architectures that enable Intelligent Decision Support in both service and production environments. Particularly important are the Data Fabric and Integration Patterns that govern the management of Data Quality, Governance, and Protection together with Modeling and Inference techniques. Since the intelligence depends on the ability to learn from Data, enable Real-Time Analytics and Edge Computing, and improve Decision Support through practice and experience, the emphasis here is identifying new requirements for these underlying Data Platforms.

Data Fabric and Integration Patterns

A Data Fabric is defined as a unified architectural approach to simplify—via logical, policy-based implementation—the organization, management, and consumption of Data and Data Services across a growingly hybrid, multicloud landscape. It aims to serve all types of Data across all Data Management User Types—Development, Data Science, Analytical, and Business—of each of the Integration Categories—Application, Data, Process, and API Integration. As a layer for unifying and simplifying Data Management and Operations, it enhances the ingestion, storage, analysis, sharing, and accessibility of Data of all types. To fulfill these objectives, it combines principles for and capabilities of Logical Data Warehousing, multicloud Data Union services, Shared and Open Data Ecosystems, Data Governance, Data Catalogs, Application Programming Interface Management, Cloud Data Lakes, Federated Data Stores, web-scale Data Management, and DataOps, all documented separately.

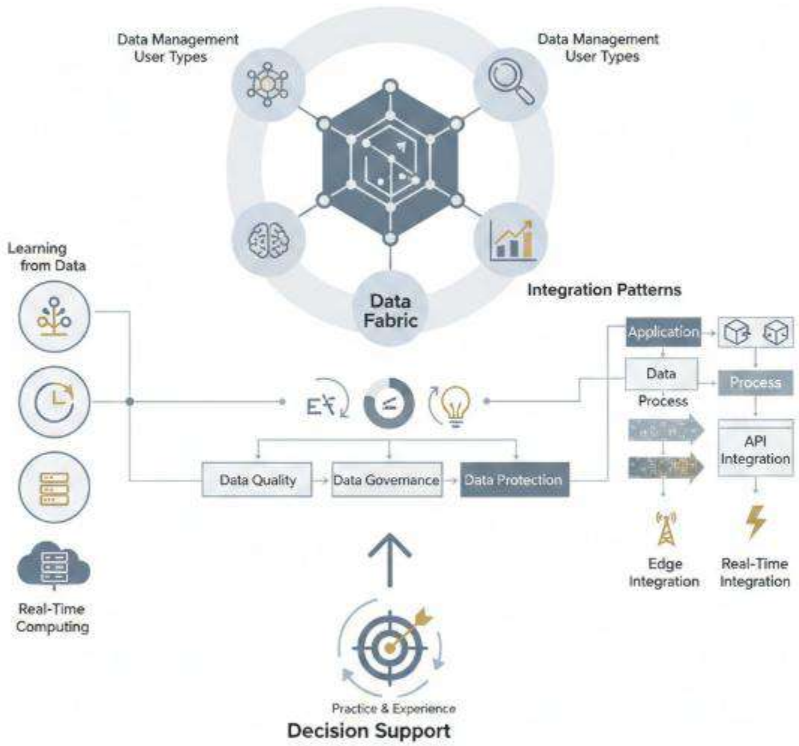

Fig 7.3: Integrative Data Fabrics and Federated Patterns: A Framework for Intelligent Decision Support in Hybrid Production Ecosystems

Production environments often mandate, due to safety or high performance imperatives, that Data Connections and Storage must be managed and optimized with Edge or Real-Time Integration Patterns. It is naïve to believe that a Data Integration pattern that achieves significant Data Sharing Requirements for Decision Processes in any particular environment (Web Hosting, Cloud Services, Traditional Enterprise, e-Commerce) can be seamlessly used in another environment without major reengineering costs. Thus, equally important as the Data Fabric approach to Data Infrastructure are Data Integration Patterns, portfoliated as families of similar Logical Integration Definitions, for a complete set of Data Sharing Capabilities for Decision Processes that are typical and critical in each Data Management Application Environment.

7.5.1. Data Fabric and Integration Patterns

A Data Fabric is a flexible architecture and set of data services combining end-to-end data management capabilities and supporting data discovery, governance, and consumption across a distributed environment. It enables an organization to easily share data across its entire ecosystem, from on-premises to multicloud environments. Data Fabric acts as an integrated layer (fabric) of data and connecting processes, regardless of location and format, consuming wherever it officially resides. With Data Fabric, data can be shared and moved securely, enterprise-wide or with customers and partners, through standardized services, thereby supporting business initiatives and providing the backbone for Data Mesh strategies.

Within the concept of the Data Fabric, various Data Integration patterns may be classified. The archetype of data integration pattern is data synchronization, in which data from one operational system must be transformed, cleansed, and loaded into another operational system. Such integration patterns should support consistency, freshness, and correctness of the data being moved. As load latency increases, it is commonly preferred to scale up a data warehouse or a data mart using data replication strategies or to use data federation techniques. Another frequent integration pattern is the need to produce a consolidated view of the enterprise's data in a data warehouse supporting Business Intelligence (BI) and analytics tools. The consolidated view must cover the entire set of data or at least a very large subset of them to enable cross data analysis. Although the data in the Data Warehouse are rarely updated, there is still a need to keep data load latency to a minimum.

7.5.2. Real-Time Analytics and Edge Computing

A key requirement of Intelligent Decision Support concerns processing speed and latency. In an increasing number of application domains, e.g., advanced manufacturing, printing service networks, hospitals, product distribution systems, and container vessels, intelligent decision support – sooner rather than later – depends on the capability to complete individual tasks with short turn-around times. This requires a decision support engine that operates at least partially in real-time and, when a service consists of temporally-related tasks, optimizes resource provision rapidly.

Support for these requirements involves integrating a business data fabric with an Edge Computing architecture. This enables real-time analysis and provides decision support capabilities in close geographical proximity to the individual service tasks. Such an architecture supports fast and resource-efficient response to partial tasks without imposing excessive cost and resource overhead when these tasks are linked together temporally. Such features contribute both to decision-action coherence and to the provision of time-critical services.

Data patterns may also suggest a different architecture. When there is demand for multiple similar services (e.g., distribution service including repair, maintenance, replacement, and replenishment), it is possible to cache the predicated requirement for these services. The support for a specific service relies on testing the cached information against its real service requirement; by developing and implementing predictive capabilities the cache enables both Intelli-gent Decision Support and Intelligent Service Control.

7.6. Evaluation and Validation Methodologies

Well-structured evidence is necessary to demonstrate that Intelligent Decision Support (IDS) improves the quality of support provided to decision makers or users compared to an alternative IDS or no IDS. Two approaches to evaluation and validation are considered. The first is based on quality assessment of prediction and decision support models and the second is based on end-to-end simulation within the service or production environment of interest.

Measuring both the effect of modelling and inference techniques on the quality of decisions supported and the actual effect of these decisions on the key performance indicators of the business goal of interest is a reasonable approach to support claims of effectiveness and efficiency, as the effect of quality on end objectives is established in prior literature.

7.6.1. Metrics for Effectiveness and Efficiency

Six categories of performance metrics and related measures help decide if decision support is “good enough” for a given purpose. These metrics help assess the effectiveness of decision support for complex, service, and production applications, and whether it runs efficiently enough to fit the application’s constraints.

A first group of effectiveness metrics quantifies how good, overall, and in relevant detail, the decisions are. A second group captures the amount of computational resources used to support the decisions, from data movement and processing in the supporting systems through user interaction to execution in real-world systems such as a transportation or production system. Subsequent groups of measures refine these high-level ideas.

These categories support decision support's many facets. No single metric suffices for the complete decision cycle, since each metric has strengths and weaknesses. Rather, composite scores drawn from a set of well-chosen, quantitative metrics combine the best of each area. Individual metrics or carefully crafted combinations of the six areas can also shed light on specific elements of the overall decision support problem, such as resilience for complex applications.

7.6.2. Experimental Design in Service and Production Contexts

When evaluating the impact of an Intelligent Decision Support System in either service or production conditions, replication of the practical environment forms an integral part of the experimental design. As introduction, frameworks for demand forecasting and capacity planning may utilize historic point-of-sale data from hotels, with a typical volume exceeding one million entries. Demand has been categorized according to business segments, produced origin, and geographical area, in an effort toward uncovering the underlying pattern of the demand series. Quantitative techniques such as regression analysis, econometric models, ARIMA models and neural networks have been assessed on one year's data for forecasting capacity requirements for six lead times, in addition to ongoing or near-future demand.

Two segments of the hotel have been modeled. One framework is strictly earmarked for business segment customers; the other serves tourism segment customers. Both frameworks incorporate the same business environment, System Dynamics technique, flex simulation, Monte Carlo simulation, operational research and simulation-based Adam forecasting heuristic algorithm, Stepwise Linear Regression, econometric models and Neural Networks. Factors affecting demand including governmental policy decisions and national holidays / festivals events are incorporated into the factory and along the supply chains to reflect the dynamic environment. بمعن futures are included to obtain the longer-term capacity plan to determine the budget for the long-term capacity-changing uncertainties. System Dynamics is applied to model the overall system and key parameters for the Monte Carlo simulation are generated. Integration with the hotel information systems allows for real-time alerting during next-day operations, while initiative parameters for near-future-period demand and supply adjusting are generated.

7.7. Conclusion

Intelligent Decision Support in service and production environments is examined through evidence-based arguments. The analysis encompasses data quality and governance, modeling techniques, and methods to evaluate effectiveness and efficiency, followed by decision-support processes for service and production systems. The selection, integration, and use of data and analytics are considered in the context of three Data-Intelligence Analytic Processing Factions (DIPFA).

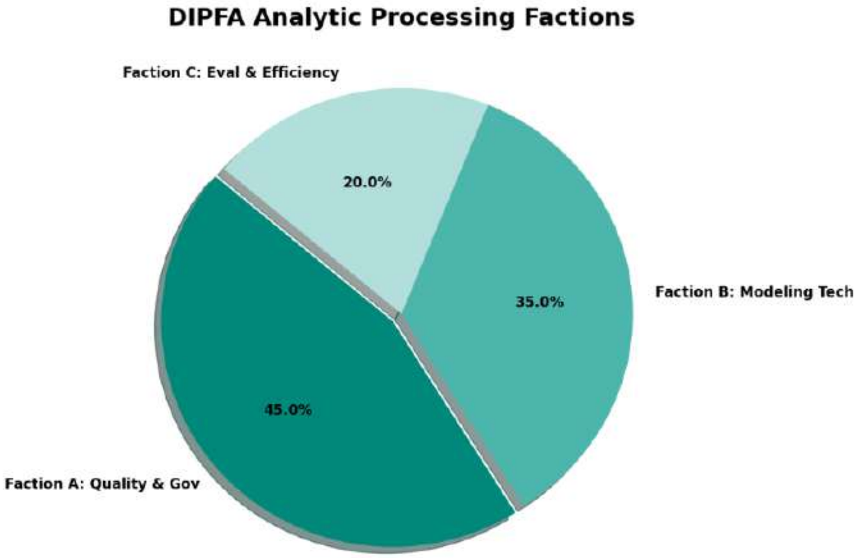

Fig 7.4: DIPFA Analytic Processing Factions

Intelligent Decision Support systems, covering demand forecasting and capacity planning, dynamic scheduling and routing in service contexts, as well as predictive maintenance and anomaly detection, process optimization, and resource allocation in production contexts. Although Artificial Intelligence is often equated with advanced statistical models, thoughtful preparation of input data is essential to support all models, including signal-based detection, prediction and prognostics. High-quality data services and services Data Fabric that address data flow to Data Sink are needed for all Intelligent Decision Support contexts.

7.7.1. Summary and Future Directions

The transformative influence of Information Technology, particularly the digitization of services, is opening exciting avenues toward increasing operational effectiveness and efficiency in service and production business environments. A key element of these developments is the emergence of Intelligent Decision Support Systems that support service and production organizations by autonomously making certain types of decisions or, at the very least, providing decision-makers with effective and efficient recommendations. Intelligent Decision Support Systems utilize data and analytics from the past, present, and future by leveraging the latest technologies in data sensing, collection, storage, and processing. Intelligent Decision Support Systems are characterized by two essential markers. First, Intelligent Decision Support Systems, directly or indirectly, deploy model-based or model-free methods of artificial-intelligence, machine-learning, and deep-learning inference to provide effective and efficient decision recommendations; and second, Intelligent Decision Support Systems can be validated against data-driven evidence-based measures of effectiveness and efficiency. A formal analysis of Intelligent Decision Support Systems provides a framework for organizing the literature and promising future research avenues in Intelligent Decision Support.

Intelligent decision support in service and production environments can be conceived as a collection of Decision Support Applications that fulfill two conditions. Decisions must be characterized as evidence-based predictions that do not change in real time but must still be made on a high-frequency basis (e.g. demand predictions, capacity-scheduling decisions for computers, computer clusters, and networks, Dynamic Routing Decisions in Transportation and Logistics, and Daily Crew Scheduling) and, in the case of Service-Enabling Applications, decisions that need to be made autonomously but are still subject to the demand-supply mismatch characteristic of services (e.g., Least-Cost Mode Selection at Airports). Decision Support Applications in Production-enabling Applications utilize all forms of failure, anomaly, and status data for the asset, process and resources to generate predictions, alerts, or advice whose purpose is to improve operational effectiveness defined as preventive reliability or preventive maintainability.

References

Dal Pozzolo, A., Bontempi, G., Snoeck, M., & Snoeck, M. (2020). Adversarial drift detection in credit card transactions. Pattern Recognition Letters, 136, 252–258.

Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. Information Sciences, 479, 448–455.

Nagubandi, A. R. (2023). Advanced Multi-Agent AI Systems for Autonomous Reconciliation Across Enterprise Multi-Counterparty Derivatives, Collateral, and Accounting Platforms. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 653-674.

Nagabhyru, K. C. (2023). From Data Silos to Knowledge Graphs: Architecting CrossEnterprise AI Solutions for Scalability and Trust. Available at SSRN 5697663.

Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. Expert Systems with Applications, 240, 122156.

Kaggle, M., Dal Pozzolo, A., Bontempi, G., & Snoeck, M. (2021). Calibrating probability with undersampling for unbalanced classification. Pattern Recognition Letters, 140, 224–231.

Guntupalli, R. (2025, June). AI-Powered Data Analytics in Cloud Computing. In International Conference on Data Analytics & Management (pp. 280-289). Cham: Springer Nature Switzerland.

Zhang, Z., & Zhou, Z.-H. (2021). Adversarial learning for generative models: A survey. IEEE Transactions on Knowledge and Data Engineering, 33(3), 1026–1040.

Varri, D. B. S. (2022). AI-Driven Risk Assessment And Compliance Automation In Multi-Cloud Environments. Available at SSRN 5774924.

Carcillo, F., Dal Pozzolo, A., Bontempi, G., & Snoeck, M. (2021). Scarff: A scalable framework for streaming credit card fraud detection with concept drift adaptation. Information Fusion, 71, 182–197.

Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. International Journal of Scientific Research and Modern Technology, 1(12), 227–237. https://doi.org/10.38124/ijsrmt.v1i12.1111

Juszczak, P., Adams, N. M., Hand, D. J., Whitrow, C., & Weston, D. J. (2008). Off-the-peg and bespoke classifiers for fraud detection. Computational Statistics & Data Analysis, 52(9), 4521–4532.

Rao, A. N., Garapati, R. S., Suganya, R. T., Kaliappan, A., & Kamaleshwar, T. (2025, August). Smart Solar Harvesting and Power Management in IoT Nodes Through Deep Learning Models. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.

Bahnsen, A. C., Aouada, D., & Ottersten, B. (2015). Cost-sensitive decision trees for fraud detection. Expert Systems with Applications, 42(3), 977–986.

Vadisetty, R., Polamarasetti, A., Goyal, M. K., Rongali, S. K., kumar Prajapati, S., & Butani, J. B. (2025, May). Cloud-Based Immersive Learning: The Role of Virtual Reality, Big Data, and Generative AI in Transformative Education Experiences. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-6). IEEE.

Wang, S., Cao, J., Yu, P. S., et al. (2022). Deep learning for anomaly detection: A survey. ACM Computing Surveys, 54(2), 1–38.

Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.

Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. ACM Computing Surveys, 52(1), 1–38.

Amistapuram, K. (2022). Fraud Detection and Risk Modeling in Insurance: Early Adoption of Machine Learning in Claims Processing. Available at SSRN 5741982.

Raghavan, S., & Manchanda, P. (2021). Behavioral analytics for fraud detection. Journal of Marketing Analytics, 9(2), 73–89.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

Siva Hemanth Kolla. (2022). Knowledge Retrieval Systems for Enterprise Service Environments. International Journal of Intelligent Systems and Applications in Engineering, 10(3s), 495–506. Retrieved from https://ijisae.org/index.php/IJISAE/article/view/8037

Goodfellow, I., Bengio, Y., & Courville, A. (2020). Deep learning. MIT Press.

Zhou, C., Paffenroth, R. C., et al. (2017). Anomaly detection with robust deep autoencoders. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 665–674.

Gottimukkala, V. R. R. (2021). Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows.

Chen, Y., Kou, G., Peng, Y., & Alsaadi, F. E. (2021). A multi-layer ensemble framework for credit risk prediction. Expert Systems with Applications, 174, 114762.

Eling, M., Nuessle, D., & Staubli, J. (2022). The impact of artificial intelligence along the insurance value chain and on the workforce. Journal of Risk and Insurance, 89(2), 1–38.

Xu, Y., Sun, J., & Liu, J. (2023). Fairness-aware machine learning for insurance risk prediction. IEEE Access, 11, 24501–24513.

Inala, R. Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective.

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). Statistical and machine learning forecasting methods: Concerns and ways forward. PLOS ONE, 17(3), e0265480.