

Chapter 7: Real-Time Analytics and Operational Decision Intelligence

7.1. Introduction

Real-time analytics is the processing of data as it becomes available and usually with minimal latency. Such analytics finds applications in diverse domains, including manufacturing, healthcare, transportation, finance, news media, sports, and telecommunications. A common requirement across these scenarios is the need for complex decision-making integrated into the real-time analysis process. Such analytics thus extends beyond traditional monitoring, alerting, and dashboarding toward automated and autonomous Operational Decision Intelligence (ODI).

Five trends help propel operational decision intelligence into the mainstream of real-time analytics: (1) data stream requirements in response to Service Level Agreements (SLAs) from business process applications; (2) automatic or semi-automatic decision models to process the data streams; (3) causal inference to better understand the processes being monitored; (4) explainability techniques such as counterfactuals to foster some level of trust in the results of automated decisions; and (5) integration of external operational constraints such as available inventory, and external signals such as market pricing or actions by competitors. Note that SLAs condition response times for the analytics, not for input latency.

7.1.1. Manufacturing and Supply Chain

Emerging and fully instrumented production environments and supply chains create new opportunities for real-time analytics as applied to transfer order fulfilment and customer delivery. A fashion retailer provides an example even more ambitious than that already stated, in aiming to deliver the latest fashion styles onto shop floors within a fortnight of design completion. Data from customer requests generated through social media, mobile

applications, and call centres are used, and supply chain optimisation aims to allow 24-hour orders from stores and subsequent fulfilment from distribution centres worldwide.

An early application to improve the throughput of a complex assembly line that paints car parts used an enhanced digital model and an interactive decision support tool to provide plant managers with an operational dashboard and support for the rapid calculation of rescheduling alternatives. The approach combined real-time data acquisition and visualisation with the development of a robust production model enabling easy integration into the decision support tool. Sensors capable of capturing production parameters were used, together with operational data provided by the manufacturing execution system. A real-time production dashboard was developed to visually represent the status, enhance communication, and support rescheduling.



Fig 7.1: Real-Time Analytics and Operational Decision Intelligence

7.1.2. Healthcare and Emergency Response

Real-time analytics also play a key role in healthcare, where timely interventions can have life-or-death consequences. Critical patients in intensive care and trauma units are monitored with a host of connected devices that generate streams of vital parameters. Delays in detecting abnormalities in these streams can lead to patient deterioration or even death. Several methods leverage AI techniques to automatically determine thresholds for these signals and warn medical practitioners when a signal goes beyond these thresholds. Such proactive treatments during the "golden hour" of trauma care can significantly improve patient survival.

Real-time analytics appears prominently also in the domain of emergency management for natural disasters such as floods, earthquakes, and tsunamis. These disasters trigger time-consuming chains of alerts, decisions, and natural-resource allocation. For example, during a tsunami, warning signals must be sent, people need to evacuate, protection measures in ports need to be taken, and rescue teams need to be put on standby—all as quickly as possible. However, people can only transmit information about earthquakes with an uncertain location and low altitude of the tremors. Consequently, tsunami networks combine traditional sensors with multiple other data sources to determine whether it is necessary to take action for the affected area.

7.1.3. Transportation and Smart Cities

Traffic complexities are oftentimes a window to city performance. Comprehensively providing and processing the related stream data indirectly represents city conditions. Actually, big data technologies are widely used to process traffic information and make transportation systems smarter. Smart transportation is one of the application fields of smart city constructions. These technologies are designed to ease traffic flow, become environmentally friendly and economically efficient. The focus is usually put on vehicular networks and transportation based on wired/wireless sensor networks (WSNs).

Connected vehicles enhance traffic management. Traffic can be controlled and monitored in real-time via data fusion and dissemination between the vehicles' infrastructure. Such technologies create new services for constant information use, improving road safety, increasing fuel economy, decreasing waiting times and even optimizing the allocation of police resources and emergency services. Intent-based Traffic Engineering focuses on users' requests and monitors SLAs, embedding the feedback from the entire communication path to optimize routing of in-network resources. Traffic estimation and prediction control data streams, aggregating forecasts, and safety management services. Road accidents are detected from image data and news. An accident and traffic time estimator controls SLAs. Traffic analysis and optimization, considering incoming flows both from external sources and vehicular ad-hoc networks, ensure vehicle fuel economy. Such coalitions of connected vehicles incrementally evaluate the aforementioned SLAs.

7.2. Foundations of Real-Time Analytics

The constituents of real-time analytics span diverse phases of operations, often employing a set of techniques deemed operational decision intelligence. Decision support can adopt real-time and near-real-time characteristics, enabling short-term workload and resource adjustments while guaranteeing least disruption of overall

objectives, during instance event-based actions. Real-time analytics is especially relevant within domains characterized by dynamic conditions and environments, including but far from limited to manufacturing, supply chain, healthcare, emergency response, the transportation and logistics sector, smart cities, and the financial services industry.

Much of ODA and ODA techniques of relevance to real-time analytics are directed toward aspects of decision-making associated with inference and response to operational disturbance, that is, reacting to the unexpected: a failure in one process unit or an unexpected demand on a product line, for example. Within these application areas, two types of technological capabilities are therefore required to achieve the stated goals. First, technology capable of accurately and rapidly assessing conditions and operation, through monitoring of process behaviour, status of control or supervisory layers, or telemetry data. Such assessment is usually performed with some form of automatic decision intelligence. Second, technology that brings high-fidelity data to the right place at the right time, supporting the speedy and reliable resolution of the unexpected.

7.2.1. Data Streams and Ingestion

In manufacturing, healthcare, transportation, smart cities, and many other domains, real-time analytics is playing a key role in optimizing mission-critical operations and decision support. Information systems make use of an expanding variety of source data that arrives in online streaming fashion (data streams) and from batch-based sources. Data pipelines ingest, transform, and redistribute the information in near real time with stream processing technology. A major concern in stream processing, which focuses on data basis and throughput characteristics, is achieving low latency while processing this information.

Data streams and batch-based data sources provide the information for operational decision intelligence (ODI), a decision-oriented extension of business intelligence, which also plays a key role in operational domains. Building on developments in business intelligence and online analytic processing, the aim of ODI is to better support the automated and human decision making of an organization to optimize the execution of primary processes. Automated decisions are made by software systems, but require underlying decision models. They are increasingly being made in real time or close to real time so that resources such as products and services are allocated, directed, and consumed in an efficient manner. Automated decisions play a central role in manufacturing operations, in service-oriented companies, and in organizations supporting mission-critical processes such as public safety.

7.2.2. Event Processing and Temporal Models

Temporal models for real-time analytics must capture events and processes within their own or an embedded time dimension. The time dimension of event streams differs from the time dimension in traditional databases. Event time denotes the timestamp of the originating event and is usually supplied by the data sources. Processing time indicates when an event arrives at the event streaming platform, which is usually “now.” The arrival of an event may be delayed because of the changing landscape or a transient failure in the network. The age of an event is the difference between time now and event time, and the result is relevant as long as the age of the event is below a defined threshold.

The notion of windowing is inherent in defining a temporal aggregate on an event stream during its time. An ad-hoc decision made for a shorter time interval may change with the passage of event stream time, and hence temporal aggregates must be updated. An update policy can capture the logic of when to update the decision.




Fig 7.2: Event Processing and Temporal Models

7.2.3. Metrics, SLAs, and Quality of Service

Operational Decision Intelligence applications and Cloud Computing workloads impose requirements on Quality of Service (QoS), and the corresponding metrics incorporate latency, throughput, and Head-of-Line (HoL) blocking. A cloud-based platform must provide a Service Level Agreement (SLA) for application developers. An SLA defines quantitative expectations with regard to system performance.

SLAs allow customers to determine anticipated costs and available resources versus performance of real-time processes. A misconception is that SLAs must be signed between the cloud provider and cloud user. In reality, a cloud user also requires metrics to size a workload and can use historical data to inform the definition of the SLA offered to ODI application users.

Given that real-time analytics applications run in a shared environment, a well-implemented SLAs takes into consideration the characteristics of the application and defines a set of non-functional requirements. Latency, throughput, and HoL blocking are the most commonly agreed non-functional requirements. Latency indicates the time taken to compute values from event arrival to the computed event being available to the next process. Total throughput of all the real-time processes in the ROI application is measured. The blocking prevents the first event to exit a process of a real-time application before X events are delivered.

7.3. Operational Decision Intelligence: Concepts and Frameworks

Operational decision intelligence relies on a combination of appropriate models and a robust orchestration layer directing operational decisions. The operational perspective takes the view that decisions act on data streams and require a continual breadth of decision horizons. Different decision models apply to different types of operations — users may want to make inferences in one context while performing planning in others. Consider the problems of recommending the next action versus inferring the effect of the latest action.

N-in-N analyses attempt to normalize these differences. Yet the complexities of effecting a recommendation in real time may not be compatible with guaranteeing a prescribed quality of service. Further, such guarantees depend on statistical, historical properties of the decision-making paragon. The concept of causal inference treats supply chains as dynamic systems and projections of state transition (effect propagation) models as understandable, trustworthy justifications for single-agent decisions.

7.3.1. Decision Models and Orchestration

The concept of Operational Decision Intelligence implies the existence of decision models used to describe, manage and possibly automate the decisions made within an enterprise. Decision models are a formal representation of business decisions. They leverage the conceptualization properties, flexibility and expressiveness of decision trees while providing a rigorous mathematical foundation for decision and risk analysis supported in decision science. The term "decision model" is used intentionally:

describing decisions using decision models requires an enterprise to analyze and archive both their predefined operational decisions as well as non-explicit diagnostic decisions – considered as special cases of selection decisions – so that they can be orchestrated in real time.

Decision Models were proposed about a decade ago to describe, document, govern and orchestrate business decisions in the scope of Business Process Management. As such, they form a bridge between Business Process Workflows and Business Logic. BPMN allows documenting the interactions between business processes, while Decision Model Notation (DMN) enables detail-level documentation of business decisions. The Decision Model and Notation (DMN) standard provides an expressive graphical language to model these operational decisions, an underlying theoretical framework to analyze and assess decisions, a semantic foundation that enables automated model execution, and tools that support both model authoring and execution. In the context of Operational Decision Intelligence, a formal specification of decision model orchestration is necessary to describe how decisions can be triggered when BPMN processes and process instances are executed.




Fig 7.3: Decision Models and Orchestration

7.3.2. Causal Inference in Operations

Causality determines connections among variables – for instance, how the delivery of raw materials affects the completion of an end product. Such information is crucial for operational excellence, where changes in a business driver are likely to lead to a related

change in a target variable. For example, when the supplier is asked to ensure delivery on time, one should reasonably expect that product delivery would also improve.

Discovering and quantifying causal relationships and exploring their temporal structure and directionality are three fundamental tasks in causal inference. A discussion of these tasks and examples from operations and supply chain management illustrate the practical relevance of causal inference. The approach separates correlation from causation, quantifies cause-and-effect relationships in operational systems, and guides decision-making, thereby improving the reliability of operational decisions in terms of the expected impact on the business.

Causality naturally plays a central role in operational decisions. A decision model – that is, a model encompassing several causes and their intertwined relations – is often stylized graphically as a directed acyclic graph (DAG), where nodes represent events, actions, or conditions, and directed edges indicate causal dependencies. Even when decisions are made sequentially, it can be useful to move beyond stepwise dependent choices and outline an overall decision model incorporating all influential variables.

7.3.3. Explainability and Trust in Automated Decisions

Rigorous testing is typically done to ensure that the results of predictive and prescriptive analytical models are trustworthy. The question of trusting the results of analytics in operational decision-making, however, is much more challenging. The analytics are generally much simpler models, sometimes the results of test runs rather than operational runs, or even based on operational experience. They could even involve the use of data from entirely different operational contexts. When operations are fully or partially automated, the use of such analytics raises the possibility of decision-making that is not explainable or interpretable by human users. Such methods are generally regarded as black boxes and trust is based more on the accuracy of the results than on the correctness of the reasoning process.

Explainable predictive models are critical in domains such as healthcare, finance, and operations, where risk management requires understanding which decision parameters most influence outcomes and why, to what degree of confidence, and in what contexts. For operational contexts involving a complex set of machines, teams of personnel, and intricate manuals, the use of AI/ML methods to automate decisions is an obvious approach. Causal inference also becomes important in such contexts, especially in the transportation industry (e.g., autonomous cars). Decision support systems based on white-box algorithms are thus likely to be the future in these domains, rather than matured AI/ML decision methods.

7.4. Architectural Considerations

Before discussing the technologies and methodologies of real-time analytics, it is necessary to consider supporting architectural aspects. Efficient data-management pipelines are essential for sustainable data ingestion, preparation, and integration, whether or not the data are processed in real time. Latency, throughput, and scalability are common concerns for operation and communications systems. Real-time processing systems must also include observability, monitoring, and testing capabilities that are either not needed or less prominently required in traditional batch-based systems.

Data pipelines aggregate enterprise and sensor sources into data lakes, where they can be refined for storage in enterprise data warehouses or mining in data science. Data-pipeline and platform architecture must consider the economics of cloud-processing costs against the cost of on-premise processing. Security, governance, compliance, and risk-management must also play a part in planning. Ingestion must scale to the volume, velocity, and variety of data in real time, proposing load-shedding and alternative sampling or derivation methods when scanning. for completeness needs checking and inferring for trustworthiness serving various real-time analytics tasks. Data quality is important for the data stored or processed and first-order sources.




Fig 7.4: Architectural Considerations

7.4.1. Data Pipelines and Platform Architecture

Analyzing ever-growing numbers of events at similar or smaller time scales for decision-making in every operational domain is a dream for domain experts. However, the shallow data pipelines being built in organizations today are far from the innovative

paths down which they would like to move. More than a decade's accumulation of work on real-time analytics, in conjunction with explorations of operation-strength decision intelligence, points towards the preferred architecture of data-union pipelines, built on five principles: 1. Perhaps a pivotal architectural aspect for operational domains, the data-union pipeline connects the data ingested from disparate sources to the data consumed by multiple destination applications, clubs, and displays. Building this pipeline enables precomputing or prefetching data for secondary processing, thus decreasing responsiveness at time of invocation; 2. Organizations are maturing towards hosting an operational decision-intelligence platform with an in-memory real-time analytics store-a purpose-built, distillation-like data repository, functioning as a single, trusted source of truth for lower-order analysis; 3. Reliability of data accesses at the destination is paramount, both for higher-order insights through machine learning and for consuming intelligent recommendations; 4. A scale-into-full-ability approach is mandated, so that scaling resource requirements up or down is a matter of operational-choice, not design; 5. Data feeds have to be observably healthy for dependable knowledge generation.

Figure 1 illustrates this preferred path of data pipelines and platform architecture.

7.4.2. Latency, Throughput, and Scalability

Applications of real-time analytics cover a wide range of domains, serving diverse business functions. Manufacturers monitor their operations via internet-connected devices, machines send alerts to engineers, and business analysis provides real-time insights into the supply chain, sales, marketing, and customer behavior. In the healthcare domain, resource managers allocate staff using real-time information from patient monitoring systems, hospital bed availability, and community infection rates. First responders receive information on earthquakes, fires, incubation periods, population density, and traffic data that can be leveraged to manage a coordinated response. Transportation networks use real-time information from vehicle sensors, passengers, and traffic lights to provide better services, allocate resources, and optimize operations. Cities use diverse sources of real-time information to reduce congestion, optimize bus and taxi services, and safely control crime.

Architectural considerations in real-time analytics focus on latency, throughput, and scalability. Latency is the end-to-end time for an event to pass through the system. In applications like fraud detection, tracking a streaming event is key, whereas for applications like stock price monitoring, a deviation from a predicted value may matter more than the temporal relation. Degree of parallelism in processing is used to increase throughput. The system needs to be continuously monitored, and any concerns should be detected early so that administrators can take corrective action before problems arise.

7.4.3. Observability, Monitoring, and Reliability

Architectural decisions for Real-Time Analytics, especially around latency, throughput, and scaling, have direct dependencies on observability and monitoring aspects. Early identification of problems enables repairing sites or data before they affect user quality assessment. Near real-time observability is crucial for meeting SLAs and maintaining customer trust. The data sources themselves must be monitored to ensure that analytics processes can function correctly without upstream failures going untested.

Failure to detect problems early may call for a risk mitigation process that can allow continuing operations safely, albeit not in the automatically optimal manner or quality of service originally desired. The warning signals of risk mitigation must be accurate, not just frequent. A reliable status classification for every operational decision made in an organization, instead of the specific case of risk mitigation, is an architecturally elegant solution for providing a safety net for a full-fledged Operational Decision Intelligence system.

Reliability of a system also extends down to the building blocks themselves. For components of the system present in-house—especially those capable of close observation for optimal latency—redundancy is often applied to provide graceful degradation when failures occur. The goal is to minimize the chance that an operational decision cannot be made for a specific concern due to factors that are internal to the organization. Stream processing and CEP engines are, however, built to run in clustered modes, offering redundancy out of the box and automatically providing better scaling than deploying independent instances.

7.5. Technologies and Methodologies

Real-time data processing technologies enable both real-time analytics on constantly flowing sensors, machines, and users, and operational decision intelligence that can respond automatically to local problems while higher-level processes remain in the background. Stream processing systems like Apache Flink, Apache Kafka, and Apache Samza move data very rapidly without downtime, supporting high throughput and low latency. Data are often captured within a distributed storage layer as a compact set of updates to an often huge collection of records about business activity, kept for a limited time and refreshed constantly. Special database products optimize end-user and business-operations analysis. Real-time textual streams from users or machines may be ingested and analyzed using complex-event-processing methods. Real-time-supporting libraries such as Apache Spark provide additional methods.

Production-ready complex-event-processing systems are evolving methods for triggering immediate alerts or responses to important situations and for deciding the right

action in potentially risky or dangerous contexts. Machine-learning systems are similarly transitioning from batch mode to streams. Recent developments make it possible to integrate such real-time processes with longer-term operational-prescriptive systems, including supplier logistics, demand fulfillment, inventory planning, and strategic marketing decision making, as well as with longer-term support-logistic processes such as scheduling and diagnosis. Nevertheless, business-operations planners should work with open systems to leave room for optimization.

7.5.1. Stream Processing and Complex Event Processing

Stream processing models and systems support the ingestion and processing of continuous data streams in real time. Processing is performed at low latency while data is in-flight, often using a micro-batch model. Stream processing systems, tailored toward the requirements of real-time analytics operations, discover, track, and respond to events of interest.

Complex event processing systems offer a higher level of abstraction than base stream processing systems. Intelligent stream processing patterns are supported beyond standard query models, often using SQL-like languages, incorporating built-in operators for joins across streams and windows, multi-pattern detection, and machine-learning-based detection. Automated rule management lets domain experts easily define, validate, and maintain the detection of complex events of interest.

7.5.2. In-Memory Computing and Real-Time Analytics Stores

In-memory computing refers to data processing techniques that rely on main memory rather than secondary storage devices. This approach avoids the bottlenecks of traditional data management by removing the input/output operations on disks. The response time thus becomes a function of the bandwidth of a single memory instead of alternating sequential disk accesses. In-memory computing has been adopted in big data processing using a distributed architecture.

Real-time analytics stores group memory computing with persistent data structures for real-time big data processing. Real-time analytics concentrates on providing the most up-to-date information possible, by ingesting updates from sources as they happen and delivering programmatic access to the latest information. The architecture of a real-time analytics store is similar to the analogy of a barbershop, where cameras placed in locations focus on the detailed events happening and allowing access to the most updated view rather than every last detail. The cameras may serve to monitor the ambient information, in contrast to the critical events that happen on main memory.

However, organizations have much larger volumes of information that require consideration. Although most of the information may be of historical value and less critical ingestion delays may be acceptable, a complete, up-to-date picture is required by access services that inevitably demand a number of concurrent queries. Make or break decision processes rely on the occurrence of critical events.

7.5.3. Machine Learning in Real-Time Contexts

The predictive capabilities of machine learning models are often employed within decision systems to take actions in real time, but they can also be applied in a more explicit manner as part of the decision-making process itself. For example, a system that detects an abnormal situation may use a change detection model in a machine learning framework to determine how to respond to the situation. Anomaly detection is, therefore, an established area of active research in machine learning.

Most machine-learning models operate on stand-alone data sets and are typically trained offline on historical snapshots. In a real-time context, however, snapshots of the variables of a predictive model are updated through incoming data streams—the features of the model—and, once the conditions for a new prediction are met, the model is invoked to generate an updated prediction. A similar approach can be followed for clustering models to trigger alerts when data are assigned to clusters associated with previous alerts. Thus, the concept of events or change detection can also be recognized when real-time monitoring considers systems in a much broader sense. Real-time monitoring can also include the detection of clusters, pseudo-clusters, and persistence, so that multiple machine-learning models can be integrated into event monitoring and detection systems.

Machine-learning methods experimentation that offers guidance regarding their suitability for a range of tasks in real-time contexts and use cases increasingly relies on transfer learning for analysis and performance. In particular, connecting multiple sensors and monitoring their information in real time opens up a wide variety of applications. A wide array of approaches can also be found in the literature for their application in network management, including those that manage performance, detection, and localization of events.

7.6. Conclusion

Real-Time Analytics and Operational Decision Intelligence support a spectrum of innovative applications across domains such as Manufacturing, Supply Chain, Healthcare, Emergency Response, Transportation, and Smart Cities. Foundations of

Real-Time Analytics cover the continuous nature of Data Streams, Event Processing mechanisms and appropriate Temporal Models, along with definition and monitoring of quality indicators such as Metrics, Service-Level Objectives, and Quality-of-Service definitions. Concepts and frameworks for Operational Decision Intelligence encompass Decision Models and orchestration of Decision Workflows, capture and exploitation of causal information for operational decisions, and considerations for Explainability and Trust in Automated Decisions.

Architectural Considerations highlight the role of Data Pipelines and Platform Architecture in supporting non-functional requirements related to Latency, Throughput, Scalability, and Reliability, as well as mechanisms for Observability and Monitoring. Technologies and Methodological Aspects cover the use of Stream Processing and Complex Event Processing Services for Event Processing, In-Memory Computing for operational Decision Support Systems and Decision Intelligence applications, and the contribution of Machine Learning to Real-Time Analytics and Operational Decision Intelligence.

References

- Akidau, T., Bradshaw, R., Chambers, C., Chernyak, S., Fernández-Moctezuma, R. J., Lax, R., McVeety, S., Mills, D., Perry, F., Schmidt, E., & Whittle, S. (2015). The dataflow model: A practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *Proceedings of the VLDB Endowment*, 8(12), 1792–1803.
- Hiremath, N. B., Kolla, S. K., Sunkara, R., G, S. Lal., & Sireesha, K. (2026). Adaptive and Intelligent Secure Multimedia Transmission for Dynamic Communication Systems. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1–8). IEEE. 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN). <https://doi.org/10.1109/iciscn67954.2026.11566143>
- Babcock, B., Babu, S., Datar, M., Motwani, R., & Widom, J. (2002). Models and issues in data stream systems. *Proceedings of the Twenty-First ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, 1–16.
- Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., & Tzoumas, K. (2015). Apache Flink: Stream and batch processing in a single engine. *IEEE Data Engineering Bulletin*, 38(4), 28–38.
- Chandramouli, B., Goldstein, J., Barnett, M., DeLine, R., Fisher, D., Platt, J. C., Terwilliger, J. F., & Wernsing, J. (2014). Trill: A high-performance incremental query processor for diverse analytics. *Proceedings of the VLDB Endowment*, 8(4), 401–412.
- Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 36(4), 1165–1188.
- Cugola, G., & Margara, A. (2012). Processing flows of information: From data stream to complex event processing. *ACM Computing Surveys*, 44(3), Article 15.
- Malhotra, A., Niharika, D. S., Mangalampalli, B. M., Nema, S., Gupta, Z., & Kriplani, K. (2026, June). Distributed and Edge AI for Continuous and Context-Aware Healthcare Intelligence.

- In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
- Mangala, N. (2025). Agentic Data Pipelines: Autonomous ELT Orchestration Using AI Agents on Microsoft Fabric and Databricks. *International Journal of Computer Technology and Electronics Communication*, 8(6), 11891-11907.
- Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods, and analytics. *International Journal of Information Management*, 35(2), 137–144.
- Kambatla, K., Kollias, G., Kumar, V., & Grama, A. (2014). Trends in big data analytics. *Journal of Parallel and Distributed Computing*, 74(7), 2561–2573.
- Jyoshna, K., Peddi, R. K., S, P. Nayak., M, Dhanamalar., & Punitha, S. (2026). Spatio-Temporal Graph Neural Network with Global Spatio-Temporal Network for the Traffic Flow Prediction. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1–6). IEEE. 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN). <https://doi.org/10.1109/iciscn67954.2026.11566462>
- Labrinidis, A., & Jagadish, H. V. (2012). Challenges and opportunities with big data. *Proceedings of the VLDB Endowment*, 5(12), 2032–2033.
- Marz, N., & Warren, J. (2015). *Big data: Principles and best practices of scalable realtime data systems*. Manning Publications.
- Paramasivam, R., Reddy, S. S., Reddy, V. A. R., Sandra, K., & Narayana, M. V. S. (2026). JellyFish Search Optimization with Arctic Puffin Optimization Algorithm for Efficient Routing Based on the Software Defined Communication Network. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1–6). IEEE. 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN). <https://doi.org/10.1109/iciscn67954.2026.11566144>
- Stonebraker, M., Çetintemel, U., & Zdonik, S. (2005). The 8 requirements of real-time stream processing. *ACM SIGMOD Record*, 34(4), 42–47.
- Piyush Kumar Pareek. (2026). Self-Evolving Analytics Pipelines for Reliable AI-Augmented Software Systems. *Journal of Intelligent Decision Making and Information Science*, 3(5s), 1558–1578. <https://doi.org/10.59543/jidmis.v3.1286>
- van der Aalst, W. M. P. (2016). *Process mining: Data science in action* (2nd ed.). Springer.
- Kolla, S. K., & Bandi, V. D. V. K. (2025). Autonomous Clinical Monitoring Platforms Using Reinforcement Learning and Deep Neural Networks. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 8(6), 13300-13313.
- Vakkalagadda, T., & Rajamanickam, V. (2026). Agentic AI Architectures for Next-Generation Investment Advisory and Portfolio Intelligence. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(9s), 1206-1219.