Deep Science Publishing
https://doi.org/10.70593/978-93-49910-41-6_8

DeepScience
Open Access Books

Chapter 8: Real-time risk monitoring: From batch processes to live analytics

8.1. Introduction

The term 'real-time analytics' has recently seen a large increase in the frequency of its use in both the academic and popular press. The availability of data from a large number of disparate sources is now spawning hundreds of articles that predict future trends or needs, or that analyze either the past or the present in the attempt to beat competitors and obtain maximum profits. It's amongst the present that our emphasis lies: we are passionate about devising models and systems that can help to recognize problems on a real-time basis. As healthcare professionals, we are especially interested in potentially using huge datasets in the effort to create 'early warning' systems that can allow monitoring in very high-risk patients and avoid preventable deteriorations. These might then potentially enable us to develop interventions aimed at enhancing the patient's life, thus avoiding extended hospitalization and containing hospital costs (Adnan et al., 2019; Ghosh & Scott, 2018; Hassani et al., 2018).

The key to real-time monitoring is to have a perception of what is happening every minute, as well as knowing how important that minute is. That is not a trivial task, especially for processes that can have a large number of decisions external to the control system. Even when the control system is itself the focus of analysis, most statistical process control and quality-related literature is concerned with finding suitable models that can describe the historic pattern of variability. The availability of advanced modeling methods, together with the increase of computing power, has given rise to a new class of models that can handle far larger numbers of predictors. These methods also permit automatic fitting in a wide variety of settings. They are referred to as models and look for linear and nonlinear relationships between input and target data over very short or fractional time intervals. By their very nature, they allow real-time estimates of where an automated system is 'lost in the woods', or whether an activity is going in the right direction. A less transparent use of conventional models can be the real-time monitoring of business data using model-based quality function-related time series. With a forecast-driven approach, we show how these models can be selective in the search for starting points of process malfunctions. These time series models have no intention at all to capture full information at a detailed level, but are constructed by building in the biggest disturbances as a tailwind, thus generating a huge potential for 'error-catching' situations, while at the same time aiming at a measure that gives enough warning of shifting dynamics (Tiwari et al., 2020; Venkatramanan et al., 2018).

8.1.1. Overview of the Study

The main objective of this study is to validate a real-time model that estimates the consequences of incidents in an industrial plant. In today's chemical industry, model predictive control techniques are widely used as an advanced control strategy that operates at the edge of the existing process limits. Many industrial fields, including automotive, finance, and telecommunications, make heavy use of their historical data or stock up incoming data for further analysis. The possible number of sensors and computing power are no longer the only limiting factors in a state-of-the-art data analysis model. However, the continuously increasing amount of data can hardly be used in the existing process monitoring techniques in the chemical process industry.

A general task of monitoring large processes is the detection of incidents, the identification of the root cause for this event, and the timely estimation of the consequence. In batch production, where incidents normally result in off-quality products that can no longer be sold, those timely estimation models are highly regarded as they have proven to be effective in several industrial applications. This creates a short cycle of incident, consequence, and model maintenance. For continuous production industries, using the same idea for predictive models, we have several further problems: the batch processes learn the relation between the operational data, the incidents, and the consequences from a very complex data structure. In a continuous process, online quality assays are relatively rare; thus, the operational data is much simpler, and we can expect quickly time-varying phenomena of interest.

8.2. Understanding Risk Monitoring

The role of real-time monitoring has had a very specific origin in the world of process control. In the industries represented by chemicals, oil, gas, and later power generation, the monitoring of critical control systems in near real time, or sometimes completely in real time, is the basic foundation to standards of speed, quality, performance, and most importantly, safety. In this section, we would like to attempt to explain in simple language what the principal concept of real-time risk monitoring is and how it has evolved into important analytics performed within minutes of an event and using vast amounts of information. Consider an example from power generation: a supercritical Type-III turbine fault could potentially cost utilities up to 5% of their annual revenue and result in at least one unplanned maintenance event a year. Similarly, in the chemical industry, faults such as easy handling can result in damage to reactors, which could lead to a cost as high as the annual revenue of the plant. Hence, estimation of risks of such events as quickly as possible and taking immediate preventive steps are critical.

How does monitoring occur in the chemical industry? One of the very important operators in the industry is the outside area operator, whose responsibility is ensuring that the reactors are running smoothly outside of the operational area. Monitoring is done at fixed hourly intervals and enormously relies on the training and experience of the operator. For example, the flow of HCl in the reactors should not be greater than the feed rate of N2. The reactor effluent is expected to be clear, and lack of it might indicate improper phase separation. Discussions with operators revealed that such day-to-day monitoring is almost like driving on a superhighway, and more rarely, the operators are asked to take a diversion when it is observed that certain process parameters monitored are lying around thresholds or outside the mandated operating envelope. These observations are also further compared by experienced engineers in the control room, who advise further action if necessary. The foremost skills are to understand the process and know exactly at which points critical parameters are monitored. Given the vast amount of detailed data, large capital investment in the reactors, and the expected revenues from production across all reactors, we were next interested in understanding why day-to-day the operators are responsible for detecting deviations rather than using tools and computational procedures to utilize this vast amount of real-time information.

8.2.1. Definition of Risk Monitoring

Risk monitoring might be defined as the periodic review of industrial and environmental risks and hazards, with a primary goal to ensure that additional controls are applied to minimize and mitigate their impacts on the process plant (both inside and beyond the perimeter). This definition is based on the typical tasks performed in risk monitoring, which are the reviewing and analysis of hazards in the plant, and corresponding incidents that have occurred on-site or off-site of the installation. These tasks are performed through the backward flow risk analysis. It also fulfills an important role in keeping the quantitative risk indicators and event frequencies up to date in the Risk Analysis Model when the models are to be used first in the pre-start-up safety analysis and later during the life cycle of the installation.

Users performing risk monitoring and study for a particular project might include land use planners, facility fire and safety officers, corporate risk managers, external corporate auditors, facility insurance risk underwriters, and internal health, safety, and environmental managers. All of these roles are represented in most operating companies. Planning and coordination of risk monitoring for a particular project are important sub-processes for risk monitoring in the company chosen and selected to monitor and study the risk for a certain chemical installation. These two sub-processes are activated at the beginning of the life cycle of a project. The role of the facility risk advisor to propose and obtain the planning and coordination should thus take place in time.

8.2.2. Importance of Real-Time Analytics

The day-to-day business of most large corporations involves an evolving mix of routine operations, predictable events, and unexpected experiments. As the world becomes more interconnected, the distinction among these three parts of the business can blur. In some cases, business promises to become more predictable, as more transaction flows are standardized and more customer choices are restricted by the preferences of the groups to which they belong. Yet new differences have emerged that also confound the existing means of monitoring business and detecting emerging problems. In particular, our digital world unfolds much more quickly than the physical one. Consequently, many critical business decisions shift almost imperceptibly from batch mode to real time.

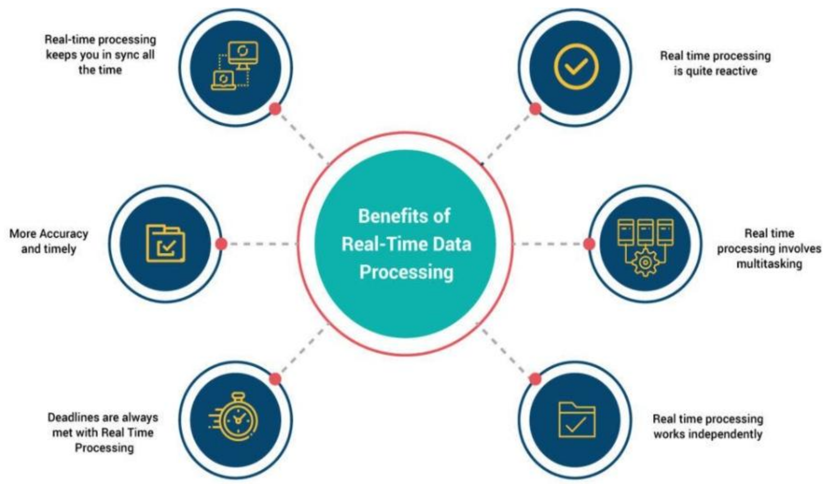

Fig 8 . 1 : Real-Time Processing

One way that businesses operate more quickly is through large-scale, low-latency analytical techniques. Digital decision-making has the capacity for mashing up data to create new customer offerings, for reacting in the night to optimize transient revenue opportunities, for routing around bottlenecks and malfunctions before they create disasters, and for many other things. However, a poorly thought-out decision can translate these capacities into a fast track to unprofitability, and the faster they go, the less time businesses have to react to mistakes. Consequently, fast decision-making also requires a fast, fault-tolerant system for monitoring decisions before and after they are made. And it requires a bank of subject-matter experts who can fiddle with offerings, invent options that can be fulfilled if contracted for, and diagnose and correct emerging problems—all without adding too much friction to those automated fast data streams.

8.3. Historical Context

A discussion about real-time risk monitoring entails the historical context of risk and uncertainty. In practice, most observations of interest are backward-looking. Information is costly to acquire and often incomplete or inaccurate. Technically, estimation and forecasting are based at least on the data that a decision-maker observed at some point in the past. Learning from experience presupposes some measure of performance, and this includes recognizing and diagnosing failure. However, organizations must move beyond the reactive model common in batch processing artifacts such as corporate financial accounts. The challenge is to develop and maintain an anticipatory model that provides insight in advance of performance data return.

The expectation is that high-frequency information can promote timely decision-making with both improved local performance and better learning about the structure and the current status of the process generating the observations. The resulting analytics are real-time, signaling the need for decision and action or inaction. We expect that modeling with rapidly acquired observations of diverse structure and form will constitute a powerful diagnostic tool that supports the continuous appraisal of a system and of the likely performance of a decision. The operationalization of real-time monitoring extends traditional functions, extending but improving the performance of management and governance, compliance and audit, and strategic oversight tasks throughout an organization, shortening the feedback loops that focus information flow and learning.

8.3.1. Evolution of Risk Monitoring Techniques

Over the past decade, progress in customer requirements, data processing, and predictive analytics techniques has shifted corporate governance risk management from primarily a static, batch process to something that is both data-driven and in near real-time. Certainly, identifying potential instances of governance, compliance, or risk issues seems endemic. However, the ability to quantify and track individual behaviors along with possible instances of bad governance in near real-time is a significant step forward. Historically, most risk management exercises have been post-event: collecting, aggregating, and scoring incidents. This satisfied regulatory requirements, be they in the financial, healthcare, access to education and research, or audited governance domains. However, as the number of incidents has continued to increase, this after-the-fact approach has been shown to be administratively burdensome and incomplete. Furthermore, providing rapid responses to red flags requires that data be processed concurrently with the actual events being observed. Improvements in mining, statistical, big data, and business intelligence techniques have led to tailoring detection algorithms to the governance attributes of client data, in a manner that combines the qualitative expertise of the risk manager in spotting anomalies with the quantitative power of machine algorithms that are able to monitor ever larger sets of attributes.

8.3.2. Batch Processes in Risk Management

In risk management, batch processes have existed for a long time and are often executed overnight or on weekends under very heavy resource conditions. The basic goal of operational risk capital allocation or risk management is to quantify downside risk, i.e., the likelihood of an operational loss of a certain magnitude or a certain frequency, often to three decimal points. In contrast to market risks, there is only a small amount of loss data available. In addition, it is not sufficient to analyze and quantify risk along conventional dimensions. The allocation should provide incentives to act proactively. In most allocation methodologies, three dynamically linked steps are to be performed. First, a loss data process has to be defined and quantified. As a second step, a statistical model for the loss frequency distribution and the loss severity distribution should be selected. In the third step, a capital allocation and, potentially, a risk management instrument have to be derived. All capital allocation methodologies require that the features of the allocation are dynamically linked to the loss distribution process.

Based on a loss data process, a statistical model suitable for capital allocation should show certain features. In addition to well-suited statistical properties, the model needs to be relatively easy to understand and comprehend, particularly if the methodology is used or to be used for risk management decisions. Traditionally, capital allocation was performed by allocating various risk types blockwise. The pool approach still underlies most capital models. For market risk, centralized formulas have been developed using the loss distribution technique that allows the different sources of market risk to be aggregated. This approach allows for a scenario-blocking approach without losing any information. However, this approach is not feasible for operational risk. In consequence, a quite large number of distinct loss distribution processes exist.

8.4. Technological Advances

In recent years, industrial automation and IT solutions have undergone a revolution. Specifically, the deployment and adoption of process optimization have recently driven investments in cornerstone solutions for operations management, i.e., the deployment of operational SQC and predictive analytics on full-process automation solutions. This has been made possible first by the availability of powerful parameter-estimation tools and fast processing power; it has been made possible second, essentially, by the deployment of rule-based solutions often involving naive statistical models. Naive is here taken to mean "instant response," not "Poor Man's Model."

The development of technological solutions is clearly a significant driving force. A noted abuse of technology actually lies in the over-marketing of wireless sensors to measure, by possible solutions providing surveillance background solutions. One primary abuse is due to a trend. Generally speaking, predictive analytics and real-time solutions are not common tools for chemical engineers and marketers alike. One roadblock is a strict, unfortunately often blinkered, school education. The sophisticated tools proposed by the most advanced research often include highly abstract tools for highly specialized purposes. However, "school" real-time tools based on classical monitoring (or rule-based monitoring, which is a digital duplication of actual simple expert knowledge to generate instant decisions) are generally proposed with continuous updating of another statistical tool, making the general acceptance of these tools complex; the result can easily be centuries-long industrial proposals.

8.4.1. Big Data and Analytics

The increase in content and prediction accuracy in big data feeds the analytics process. With faster, better, more data, we have the ability to process it more quickly, leading to better results and further automation. This process can create a circle that pushes the bounds of what we think can be done in a specific domain. Advances in technology, such as the speed at which more powerful GPUs are being developed, help to shore up the capability of the analytical engine. They are perpetuating this feedback loop, and prodigious amounts of data enable machine-learning algorithms to make sense of increasingly sophisticated models.

On the predictive model side, actual model operation speeds impact the operation decision cycle. Faster insights lead to the ability to create models faster, but this speedup should not be in isolation. It must also be met by operational speed. This point is underscored by analytics operating in the live world based on big data analytics. The risk plane of the most unfortunate domain, equity markets, has evolved from humans shouting at each other from the pits to traders in the late 20th century to machines quietly talking to each other in the 21st. This talk is all about extracting a profit from the difference in prices; this talk is based on algorithms that originate from a predictive engine based on machine learning. Such systems are in live operation mode, where academic or design time is gone; out-of-sample super forecasting is critical.

8.4.2. Machine Learning Applications

The extraction of useful knowledge from data has long been recognized as an essential part of decision-making. By means of machine learning methods, one can discover interesting patterns in data, personalize models for predictions and rankings, detect anomalies, or even optimize system parameters, among others. Modern ML algorithms have founded their success on the iterative approach of updating model parameters based on the minimization of a predefined fitness function over a given dataset. This has enabled the application of ML techniques to a broad spectrum of real-world prediction problems, including applications ranging from computer vision to healthcare, among others. Over the last few years, a growing interest in ML-based tools has appeared within the process systems engineering and industrial analytics communities. The application of any of these ML-based techniques has the potential for improving various facets of the current real-time monitoring strategies.

8.4.3. Cloud Computing in Risk Monitoring

Traditionally, businesses and public sector units have invested in large server farms, which take up substantial real estate in multiple layers of office buildings. These units create a substantial carbon footprint and require data centers. Such legacy systems call for regular upgrades and maintenance costs that could easily match the benefits of the systems in a few years and cause an explosion in the CAPEX and OPEX costs. Cloud computing has transformed these legacy systems by not only lowering the costs for computing, storage, and network, but also drastically reducing the carbon footprint. Computing has analogous economics that are similar to the power industry, and cloud computing has solved the illiquidity premium associated with capital-intensive costs. Enabled by rapidly decreasing cost per unit of computational power and density, cloud computing offers economies of scale along with the real-time flexibility to initiate and dismantle loosely coupled parallel processes.

8.5. Real-Time Data Processing

1. Creating Real-Time Data Pipelines

Now let's move a plant into a modern era, when we can easily read its sensors, store, and process the data. A common platform for doing so is a distributed stream processing system that makes it easy to publish and subscribe to data streams, integrate large amounts of data from different sources, and allow real-time analysis. Due to its fault-tolerant and durable nature, it enables replaying historical data, which is required for rebuilding models or making data available for batch processing.

At any given time in a stored data set, there may be only one value for any particular data item that is valid at that time. If one wants to refer back to the state of the world at some particular time, the data set must therefore be written in such a way that it records all the changes to the state of the world rather than merely reflecting the current state of the world. There are established ways of solving that problem, ranging in complexity from using database functionality for certain applications to using specialized tools, and with it, it is a bit exotic, but it is quite possible.

2. Data Consistency

There are some limitations to watch for, mostly surrounding data consistency. The fact that features come from potentially stale data means we cannot join data from the future. More generally put, these considerations are what make stream processing useful - every piece of stream data is joined with the latest snapshot. The alternative to check for is "what have I learned lately?" and make sure that we don't build a model that depends on a particular look-back window we no longer have.

8.5.1. Stream Processing Frameworks

Given the importance of stream processing in the real world and the potential of its use as a substrate for distributed systems, it's not a surprise then that in recent years different general-purpose solutions have been proposed. MapReduce++ is a MapReduce extension proposed to support more iterative data processing patterns. Moreover, Cloud Dataflow generalizes MapReduce to act as a management and execution tool for simple parallel data processing and differs from other systems due to the fact that it covers the full process from data ingress and aggregation to analysis and visualization. Included, we find projects such as Apache S4, aimed at stream processing, and Apache Spark and Apache Flink, striving to provide fast in-memory execution for workloads different from simple MapReduce. Other well-known systems are: Twitter Heron, a real-time continuous stream processing engine; Apple's Loom, designed to ingest, analyze, and react to millions of messages in real time; MOA and MOA-Anemos, providing stream clustering algorithms across different tasks; or MOA-SelectStream, dedicated to the visualization and selection of items from a stream.

8.5.2. Data Ingestion Techniques

One of the important factors that should be considered when it comes to building real-time applications is effective data ingestion. Data that is processed asynchronously should arrive at the time that allows the turn-out plans to be quickly deployed into the cluster to meet the needs of real-time allocation, even in systems that require absolute minimal latency data updates. In addition, rapid recovery from failures and maintenance of the application's flexibility in choosing scale is also important. Finally, the design must also consider the topics of backpressing data, slowly modifying database data, and any other interesting stopped scenarios.

We will explore the common data ingestion techniques traditionally used for batch processes, and then through careful consideration of implementation, we will demonstrate how this idea can be applied to the problem of real-time data allocation. Batch processes can typically pull data from their source and bring it back to the system's storage at a slow rate. This is advantageous when the data is being received across time zones or the load can be condensed from in-memory storage and released. Furthermore, because a batch process works with data in a static delivery, it remains failure-tolerant inherently. However, the most important advantage of the function is that it is easy to know how long each parameter needs to be processed, allowing the system to optimize simplified resource allocation.

8.5.3. Latency Considerations

Risk scenario validations tend to be performed in an offline manner, using batches of risk details to detect failures in some risk calculation or compliance rule. The cycle might range from a day to a month, given the batch calculations that might be required to really get a customized assessment. Yet the urge to analyze events with live data usually marches on, so the latency of real-time infrastructure design has become a new area of study. Basically, many fast processing tools need to be concentrated around some very crucial real-time interface gears. The typical approach starts with a complex event processing system that aims to provide algorithms that can handle high-speed processing of large streams of event data. However, the number of events that might need further scrutiny can become a real resource eater. The sorting and detailing of that day’s trades is paving the way for the next layer of fast analytics, where exceptional events on secondary markets get their notch in the real-time infrastructure. Finally, the possible issues need to be analyzed in more detail; for example, using some more complex or informative algorithms that are about to crunch the historical data. Still, the calculations and their output need to finish within a trading day.

8.6. Risk Indicators and Metrics

Every large firm is able to pull financial and legal information together. It is no surprise, therefore, that these are two of the most commonly used types of data in risk modeling. The entire sequence of published financial statements, together with the annual report, provides a detailed picture of the balance sheet, the income statement, and the statements of cash flow. In the case of a bank, notes are available on many aspects of the balance sheet and the income statement, and these are supplemented by other information that banks are required to produce and make publicly available. There is also a wealth of data on publicly traded firms, from the market value of equity to daily returns on stock and stock indices. More often than not, what is vital to a firm is revealed in the actions it takes, i.e., the complex sequences of financial and trading transactions underlying the exchange of cash between banks. The money trails revealed through internal monitoring offer unique insights into the risks being run by the firm's business lines and do so with no delay in reporting.

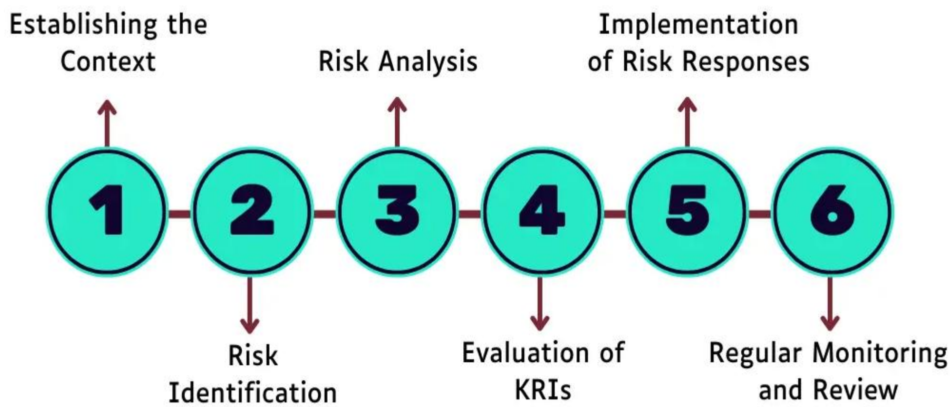

Fig 8 . 2 : Risk Monitoring

Although financial data is readily available, it needs to be supplemented before it can be used as the foundation for building, refining, and applying risk models. For example, a traditional statistical model is nearly always built on past data with the aim of being used on future observations. Time series data is easy to put together for many financial series, but the concepts that apply to time series data need adjusting when time series data is used for model estimates. Prior to model building, the original data needs to be turned into exogenous variables, such as macroeconomic series or other external market data, which may be correlated with the future variable of interest. In a similar manner, the temporal ordering of data is of far less concern when it is used in dynamic risk models, whether multivariate or univariate. However, the behavior of financial and trading transactions must be considered to determine which particular combined statistics might be efficient, given that certain data is contemporaneously available.

8.6.1. Key Risk Indicators (KRIs)

A large number of banking and insurance institutions define, calculate, and monitor these key risk indicators. The concept of the definition and monitoring of these KRIs has advanced, and insurance companies have traditionally implemented systems to control these KRIs. Ideally, the definition, calculation, and presentation of these KRIs should be computed step by step, bottom-up, involving not only enterprise-wide risk but also all defined risks including credit, market, operational, insurance, liquidity, interest rate, and exchange rate. These are defined using definitions and measures for business, supporting, and other transformation processes.

Risk is defined as a measure of the consequences of threats from the environment that oppose the objectives and the associated means of the enterprise. Risk can have a positive outcome (opportunity) or a negative consequence (threat). Opportunity is just the other side of the coin, risk. Opportunities and threats should reflect all defined objectives, the utilized means to achieve these goals which depend upon the strengths and vulnerabilities, and the responses or strategy of the organization. Ideally, the definition and calculation of these KRIs should be bottom-up, not only including enterprise-wide risk but also all defined risks derived from a transformation process. Enterprise-wide risk management is supported by an approach that combines both top-down and bottom-up processes.

8.6.2. Performance Metrics for Real-Time Monitoring

In Sect. 5.2, we presented a performance evaluation approach that specifically focused on online streaming monitoring. We compared it with common techniques that could certainly be used in a streaming setting but were not tailored to it. In this section, we describe the performance evaluation of RTWeaver, accomplished through a recorded industry case, against the performance metrics discussed. Due to data sensitivity and the client’s policy, some of the details of the addressed scenario are omitted. The aim of this case is not therefore to validate the general applicability of RTWeaver, but rather to provide a practical example of use and to highlight how the specific continuous evaluation metrics introduced are useful for assessing real-time analytics tools.

In this chapter, we reported the ongoing work for assessing the performance of a continuous evaluation approach for analytics in batch and streaming settings. We considered three concepts of accuracy from the literature and inspected their properties. We observed that there may be unexpected trade-offs in the performance of continuous and very granular evaluation of the models, in terms of data load and processing pressure on the systems being monitored. Since the research instrument creates a control loop, we foresee the ability of RTWeaver to minimize the impact of this newly detected problem. However, we will only assess if this is indeed happening once the case is made public, after the specifics are sufficiently masked.

8.7. Case Studies

The case studies presented here aim to argue that the kind of challenges presented in previous chapters are being faced today. This can be done by illustrating that these and related problems have been recently addressed in concrete, data-driven settings. A secondary goal of these descriptions is to contextualize the tool components described in the previous chapter. A final goal of this section is to describe several progresses in this direction being made by other researchers.

These will include both a level set as to where the tool stands as well as some of the applications that it has found relevant to this type of risk monitoring. We expect that by gaining some grounding in concrete examples, we will better appreciate the embedded background knowledge. It is our hope that these case studies will enlarge our mutual sense of the possible.

We provide a selection of practical case studies demonstrating the deployment and use of the tools of real-time risk monitoring in the process industry and risk management in general. Note that by the choice of these representative vignettes, some otherwise very important areas of application of these tools – in human processes, in signal processing and communications problems, and in the performance of real-time risk monitoring tools; however, these topics are treated by other works.

8.7.1. Financial Sector Applications

Financial Sector: Large commercial banks maintain extensive data mining operations. The data includes personal information about account usage such as customer concurrent activity, money transfer patterns, cash card use, and maintenance of cash card balances. Each activity is assigned to a customer model that has corresponding transaction-based thresholds, and transaction patterns are then compared in real time to the thresholds for a set of customer models. This real-time risk monitoring system identifies suspicious transactions. For example, identified transactions have the following possible content: they are unusual in the context of structured rules predicated on customer activities within a given time span; send money to a known money launderer; accumulate large chip card balances rapidly; visit other customer addresses, often at times that differ from those of the main customer; or use a personal credit card for chargeable activities followed by cash, bill pay, or electronic withdrawal activities. Risk of potential money laundering is inferred by the latter rule.

Mortgage defaults and delinquencies can be predicted by using the live scheme FRAUD DETECT LAG, which implements the fraud detection rules derived from nationwide studies involving default and delinquency rates. The rules identify behavior that significantly deviates from the behavior observed among those loans that remain current. Two sets of rules are considered within this project: mortgage holdout at issuance and mortgage performance after issuance. Within these two sets, behavioral patterns are investigated that could raise suspicions of fraudulent activity during the loan process or the renegotiation process. After loan issuance, seven patterns could suggest that renegotiation plans differ considerably from those of typical homeowners. The forecast is then incorporated into a multiple classifier system that issues fraud warnings once certain thresholds are crossed.

8.7.2. Healthcare Risk Monitoring

The availability of massive datasets from nationwide electronic health record system programs or disease screening surveys allows researchers to extensively study complex medical issues that were difficult to examine in the past. Monitoring datasets generated continuously in healthcare industries demonstrate more informative variability within the contexts. This capability provides risk models with the opportunity to enhance their descriptive powers given the contextual heterogeneity. As populations grow older and more patients have multiple chronic diseases with specific comorbidities, the development of healthcare big data will be the most important application of the century of information sciences.

Healthcare risk monitoring consists of a systematic mechanism for the purpose of early identification of possible future failures registered in biomedical practice, where the risk is concerned about certain conditions that may become true.

Healthcare professionals focus on patients who have received diagnostic procedures with abnormal status for predictive or preventive consideration. Target variables of risk monitoring models are usually collected through not only expensive diagnostic channels but also invasive monitoring devices. Additionally, the sample size from the target population with abnormal diagnostic results can be very small compared to the general population. This limitation from the distributional nature leads to an unbalanced class problem in which the number of observations from the overrepresented class substantially exceeds that in the small minority class. Therefore, the evaluations of classifier performances aiming at correctly classifying this positive class become extremely important. In this chapter, we will discuss related research on the development of healthcare risk monitoring models from the perspectives of risk identification with target class imbalance. After that, we consider the temporal aspects of the monitoring system, which involves functional hardware, real-time alerting and interventions, online analytics of practice control, and offline diagnosis with documented learning experience.

8.7.3. Manufacturing and Supply Chain Risks

Threats should be mapped with preventive measures aiming at reducing the likelihood of an event occurring, and damage control actions intended to prevent or mitigate harmful effects in the case of failure of the preventive measures. These measures include technical preventive actions consisting of redundancy design, analytic resource optimization, robust operations, transformation of the detection thresholds of traditional process monitoring, advanced fault detection and fault isolation of distributed systems, near real-time training of decision resolution, advanced fault detection and isolation, and advanced process monitoring. It is important to quantify the desired risk reduction versus the impacts of the measure on performance and operational costs. Further, there must be an understanding of the dynamic interactions between the manufacturing supply chain and the threats, to minimize unwanted interactions and to maximize wanted ones, to increase proactiveness and response; and there must be the driving of a strategic programmatic and cultural change to improve security by design in the manufacturing supply chain and among stakeholders in the manufacturing supply chain. It is also important to measure and observe both the output areas, driving both data or model-driven and feedback control prospective actions, while optimizing control strategies that consider the potential change in physical contraction that arises during out-of-normal conditions. People involved in analytical optimization tools are very much interested in the improvement of tools and the development of industry needs, and are willing to share technology with the industry. Public funding institutions will accelerate the improvements in the emerging nanomanufacturing community.

8.8. Challenges in Real-Time Risk Monitoring

The increasing complexity of banking and the financial system in an age of deregulation, the global nature of the banking system, and most importantly, the new technologies and instruments developed, made the notion of systemic risk in the financial system a practical issue by the turn of the century. It is a matter of great concern for central banks, regulators, and public officials. In many leading countries, supervisory authorities have published their opinions and research on systemic risk. According to the Basel Committee, "Banking supervisors must be satisfied that the internal risk measurement systems that banks use to determine regulatory capital requirements for market risk are as reliable as possible."

The major problem with transaction data is not the source itself but the promptness and the continuous accessibility of information. For deposit accounts, transaction data can be quite scattered for most of the month, except for a few days at the end of the month when most depositors settle their account charges and credits. As stated earlier, loans and funds accounts could be updated daily, but they would be collected and processed weekly on the addition of the Departmental Statement on Form 1 of the Field Account Statement and sometimes on a quarterly basis based on the summary commercials of Form 2B. These cycles add so much "noise" to the rationale of real-time computation that, thus far, the calculations and the surveillance processes have been performed on a weekly basis, at best.

8.8.1. Data Quality Issues

Data quality in the real-time data warehouse should be assessed often. Such an automatically recurring process should read error and exception tables that are created by the real-time ETL process itself. In the case of manual data entry, each pertinent dataset's last update time could be used to improve the efficiency of quality checks by identifying the last updated data irrespective of whether there might have been an error.

Data quality issues are of great importance. After the data warehousing process begins, errors are introduced - most of which are nobody’s fault. The errors can be divided into a number of categories: 1. Missing Data 2. Elided Records 3. Elided Columns Only 4. Misspelled Data 5. Added Data 6. Wrong Facts 7. Bad Relationships 8. Ambiguous Records 9. Duplicate Records 10. Incorrect Formats 11. Computational Results 12. Classified Results

8.8.2. Integration with Legacy Systems

A frequent problem in creating any sort of enterprise-scale system is seamless integration with the existing legacy system. This is particularly true with respect to legacy alarm systems. Many companies have complex alarm management systems in place to deal with their existing manufacturing operations. These systems are often highly developed. This is particularly true in the refining and petrochemical sectors of the process industries. They provide many layers of alarm inhibiting, suppression, and hierarchical displays to help present a clearer picture of operations to the engineer.

A proposed real-time risk monitoring system, by its very nature, sieves vast quantities of data in inventive ways to look for signs of potential accident scenarios. This is not how traditional alarm systems work—they were not designed to deal with vast statistical quantities; they tend to look for hard rule-based alarm states. Moreover, operators are used to this methodology. They respond quickly and efficiently. Implementing automated risk monitoring often cannot use the same methods. Analysts may instead have to give additional warnings of a possible problem based on their own analysis of monitored data. This only creates more alarms for the operator to deal with. There is also the problem of not just how to display the information to the operator but also what limits to set upon the analysis in the first place. The considerable work carried out so far in the development of the risk signaling system shows clearly the major problems that occupy the great majority of the time of the researchers in this particular research team. Extremely tight validation and comparison studies with real plant current alarm handling systems are planned at key stages of risk signal and risk threshold factor development. Although there will be an impact on the operator from these extra risk-aware alarms, it is hoped that these will frequently correspond with intervention checkpoints for the operator to take action on a possible hazardous uncontrolled transition in the process.

8.8.3. Scalability Concerns

Real-time risk management can quickly become impractically expensive for investment banks as more data needs to be kept within business hours, and additional cycles of rapid computation and decision-making occur. Buying approximately 10 times as much expensive technology to perform 29-second calculations can be a tough sell, especially if those calculations are performed with 100 times as many building blocks that create "data tying problems," in which both the inputs and outputs of one building block are inputs to another, in a tightly serial sequence—and with some blocks that reach outside of business hours.

Fig 8 . 3 : Real-Time Analytics

Another reason for not doing risk management live: There are often "risk containers" with small actively managed financial position pools, as well as larger long-term risk portfolios. By separating these two, more radical risk management technologies can be used for those portfolios that can benefit the most. However, jurisdictions continue to expand and pressure trading limits, which can both force trading into smaller time windows and make scheduling trading costs more consistent with post-close profitability. Banks have spent the last several years very publicly getting ready for even further mandated risk overhaul, including a more complex and volatile mix of entity limits, where they need some way to figure out at what cost entity risk can be hedged together in a trade. Concurrence: The energy desk will also want pre-trade analytics on its much shorter tenure risk positions before placing a trade.

8.9. Regulatory Considerations

Regulatory requirements and disclosure to investors are important. Investors are increasingly interested in the real-time monitoring of risk, as not doing so may lead to bad surprises in the future. Banks need to develop models to estimate their risks in various asset classes quarterly. Anecdotal evidence suggests that banks often use longer data histories to build such models. In practice, stress tests have identified that the banks' estimates of risk can be quite far apart from other assessments. Why might this happen? It could be that the bank's model is wrong or that the bank did not follow the model. In either event, the bank is penalized. A dynamic tool that the bank could consult regularly in order to compare its results with the said tool is going to be beneficial in helping the bank focus on the right areas.

We believe that the discussion of regulatory requirements is likely to intensify in the current environment where banks and financial institutions are going to be subject to more capital requirements. Our firm belief is that banks and financial institutions need to do whatever they can in order to demonstrate to the supervisor that they are being careful. Chief risk officers or those responsible for risk management within banks or financial institutions are going to be scrutinized in order to determine if they are up to the job. Banks need to be capable of offering as much useful information to staff members as possible. In our view, we believe that banks need to do this on a real-time basis. The purpose of the capital ratios is, after all, to ensure that the bank can withstand distress happening quicker than in the past. If something is going to happen at a pace faster than historically, live analytics is going to be helpful in observing, identifying, and ameliorating the risk.

8.9.1. Compliance Requirements

A number of industry-specific and global compliance requirements demand that companies adopt good enterprise risk management processes and operations. This publication has listed three such regulations, one from the global financial services and the other two from the United States. The Strategy SPIR—the Special Purpose Industry Report for this perspective—has a list of more than 30 pieces of regulation in this area. These regulations are typically of the form, "you need good financial discipline and controls, and if you have ever done business in any of the 35 member countries, you are subject to this law." Much has already been written on the details, processes, and operations surrounding the legislation, but the focus here is the provision of connectivity services, offered as umpire, in the regulatory infrastructure debate and typically centered on the governance and accountability aspects of these and the strategies of business models of those providing real-time and near real-time risk connectivity in compliance.

The motivation for taking on this regulatory debate is from the more than 50 years now of services delivered by financial exchanges spanning five continents of the world. On these venues, the world’s savers have sorted out for more than 10,000 companies, the winners and losers. All exchanges state these missions differently, but the work clearly demonstrates that not one of the functions of these institutions can afford to fail on a large scale; these functions are not failsafe, hence masked in compliance regulation, which is open to interpretation of existing preparedness readiness and avoiding the problems. But has everyone said again implies medium or high severity if anticipation and adaptive characteristics are not inherent in the support services offered by technology response and recovery. Additional reporting arguing that financial intermediaries do not inherently avoid speculation and have not stepped up considerably in the efforts to investigate and mitigate system-wide risk only shifts different systemic weaknesses elsewhere, either in time or to other venues.

8.9.2. Impact of Regulations on Risk Monitoring

Investment banks are under significant pressure to increase the robustness of their analytics and push more 'live' processing onto unfiltered data. More powerful trading strategies and systems are also coming on stream. This movement to ever-faster systems, however, comes alongside the usual need to ensure that trading risks and other operational concerns are well controlled. Cloud computing solutions offer the intriguing possibility of allowing risk controls to match the evolving, live strength of investment banks' trading systems, wherever those systems are residing. Financial regulations have been heading in the same direction. Most post-crisis financial regulations have aimed at achieving greater transparency and stability for financial services. Critically, the development of real-time risk monitoring systems that are scalable, live, and global should not be held back. It is also increasingly used to mitigate fraud, both within the bank and from external attackers who could gain massively from attacking a tightly interwoven financial infrastructure.

Credit risk models, which are used, for example, to calculate capital against loans, are themselves essentially probability of default models. Banks must own models that can predict trading and counterparty risks. These must not be run because a bank's position may have changed the moment it is asked for the information. Therefore, just like many of the bank's live look-up systems, these models too are increasingly moving to real-time systems. Such models produce a larger ecosystem of live analytics because it is nearly as critical to predict not just operational data like interruptions and fraud, but also how factors in the wider world might affect a bank's positions. Banks that can analyze externally produced data in real time have a better chance of making, and holding on to, their profits than those who are still running behind their data. These more advanced live analytics focus investment bank attention.

8.10. Future Trends

In time, live analytics will revolutionize how batch processes are controlled to ensure better safety. Instead of static risk-level information that activates safety control and action, operators will use dynamic risk-value transformation to make real-time mitigation decisions. In the long term, principles of superior safety based on disproportional risk reduction and cognitive risk management will become the norm. Through this, chemicals will be manufactured more safely, and the products developed for end users will be substantially safer. Due to the improved culture, safety monitoring will extend beyond high-risk control of batch processes to improve safety across all chemical manufacturing activities.

Real-time risk-state estimation, encompassing rapid risk-state computation and real-time risk-level visualization, is made possible through integrated risk model development, advances in computational software and hardware capabilities, and the evolution of wireless communication. As real-time monitoring matures, higher accuracy models can be incorporated for a more complete view of plant safety. Because of the growing importance of safety and the rapid technological evolution of real-time data gathering, mature live analytics-based state estimators will be available far earlier than the time spent by legacy risk models in industrial operations. While risk-state computation is expected to continue gaining enhanced visibility through visualization, increasingly automated safety control and improved sustainability will make scenario-based real-time safety decision-making the natural course of evolution.

8.10.1. AI and Predictive Analytics

AI, and in particular machine learning with predictive analytics, can surface hidden patterns and score the likelihood of future events during live processing. Examples include predictive maintenance and quality in manufacturing; service load forecasting in real-time scheduling; fraud detection in financial transactions; and recommendation engines used in e-commerce, digital marketing, and social networks. Predictive analytics can also be used by risk managers who are responsible for anticipating bad events and mitigating the severity of their consequences. These front-line risk managers use predictive models as real-time decision support to help them continuously monitor risk and change strategy. In these applications, batch training and prediction deployment processes are often an unacceptable latency.

There are several reasons why predictive analytics, with sophisticated machine learning, must be used in real time. A major challenge is that most operational decisions in the big data era must be machine-assisted because the demands of real-time digital business exceed the cognitive abilities of their human decision makers. Even where there is human governance, there are simply too many decisions influenced by big data analysis for risk to be managed in a systematic way. Even supervising these decisions using live analytics is not realistic, and real-time risk management is increasingly described as live analytics.

8.10.2. Decentralized Risk Monitoring Systems

The first issue, operational phase decentralization, results in diversification of the decision centers and real-time execution of indicative operations. From the risk monitoring standpoint, as long as consistency problems are resolved, decentralized local managers, with pertinent context, should be better able to react to local issues than a central service. Each manager observes micro patterns and general system trends, but typical situations are simple. This stands in contrast to the risk monitoring service, which is quality-oriented and processes global events and system disturbances that often must be rated as 'possibly urgent,' with significant delays and quality problems, without real-time reactions. On the operational side, due to the makers setting their rules and the much superior gap in the quality of the undertaken decisions, urgent issues require status transition 'owing to out of control/breach of limits' detected by the monitoring service acting either as an external or an executive service. The ex post role of properly chosen monitoring commands as 'reactive power' is widely recognized in engineering control and information systems.

This generic recipe of operational phase decentralization may be applied when practical. It may be effective and equally simple when monitoring is considered a service function—no alterations to the device interface, no alterations to data collection on the device side, but immediate generation of indicative if-then notifications on the monitoring service side, where design decision limitations are scattered. Only observation reuse through a single measure of observation quality is advisable. If quality is observed in addition to the observation, supervised machine learning serves as a generic warrant of operational quality. If the level of control is observed next to quality, a lurking proactive monitoring may be constructed; this requires more serious data analysis for predictive-analytical support of the phase transition. Proper anticipation of process malfunctions is repeatedly confirmed as a major benefit of predictive analytics and IoT monitoring.

8.11. Conclusion

The frequency of cyber-physical attacks is on the rise faster than ever, thanks to the increasing adoption of remote access to factories and critical infrastructure. Therefore, factory operators need to be aware of these threats in order to respond proactively and effectively. This work addresses this by taking the first steps in creating a real-time knowledge engine for a factory operational risk. This situational awareness technology allows the factory operator to understand what is going on in online processes, a foundation for understanding the anomalies that may be due to attacks. The technology employs the technique of change-point detection using a one-class classifier, an unsupervised technique that requires no labeled training data. We highlight operator usability requirements and demonstrate the technology and its user interface with a case study for a real-world, batch-wise pharmaceutical blending process. This live, risk-metric system uses this algorithm to recognize abnormal batches continuously.

In this study, we extend the concept of risk from financial risk, in which anomaly detection is labeled with the ground truth, to a more general type of operational risk in which unlabeled anomaly detection can happen. The risk philosophy is extended to the real-time risk-metric system for the purpose of operational risk management in the industrial domain. According to the pharmaceutical production scenarios, the process contains ingredients weighing, temperature setting, blender starting, end-point value detection, and discharge. By evaluating the different process operational decisions and the surveillance tasks, we create the corresponding operational risk prediction system. In future work, it would be interesting to research risk mitigation through proactive process planning and schedule modification in the real-time risk-metric system.

8.11.1. Summary of Key Insights and Recommendations

Real-time risk monitoring through live analytical systems can unleash enormous value by significantly reducing the time required to generate risk insights and thus enabling near-real-time enterprise decision making. Increasing threat landscapes, sophisticated threat actors, increasing frequency of attacks, and the increasing blast radius of security incidents demand speedy responses and have led to the build-out of live analytical systems within security information and event management offerings. For an enterprise undergoing digital transformation, transforming security information and event management architectures and regulations can advance security into becoming an enabler of future IT-enabled enterprise initiatives, rather than just a protector against established threats. Recommendations to accelerate the journey to real-time risk monitoring include embracing data, technology, and process technology. The insights and recommendations discussed focus on demystifying the approach to real-time risk monitoring; they can be used as the initial checkpoint to assess an enterprise's aspiration and readiness to implement the approach.

References

Hassani, H., Huang, X., & Silva, E. S. (2018). Big Data and Climate Change. Big Data and Cognitive Computing, 2(4), 27. https://doi.org/10.3390/bdcc2040027

Tiwari, A., Weeber, S., & Pape, S. (2020). Real-Time Analytics for Event Data Streams: A Survey. ACM Computing Surveys, 53(6), 1–44. https://doi.org/10.1145/3417980

Venkatramanan, S., Lewis, B., Chen, J., Higdon, D., Vullikanti, A., & Marathe, M. (2018). Optimizing Real-Time Decision Making Using Predictive Analytics: Challenges and

Opportunities. IEEE Intelligent Systems, 33(4), 52–59. https://doi.org/10.1109/MIS.2018.043141657

Adnan, M., Maqbool, O., & Lee, Y.-K. (2019). Big Data in Financial Services: Real-Time Analytics and Challenges. Journal of Intelligent & Fuzzy Systems, 36(5), 4863–4874. https://doi.org/10.3233/JIFS-179169

Ghosh, R., & Scott, J. E. (2018). Enabling Real-Time Risk Analytics in Financial Markets: A Data Stream Processing Approach. Journal of Risk and Financial Management, 11(2), 25. https://doi.org/10.3390/jrfm11020025