Review Article

Business Process Innovation with Artificial Intelligence: Levering Benefits and Controlling Operational Risks

Jana Koehler

School of Information Technology, Lucerne University of Applied Sciences and Arts, Rotkreuz, Switzerland

Email address:jana.koehler@hslu.ch**To cite this article:**Jana Koehler. Business Process Innovation with Artificial Intelligence: Levering Benefits and Controlling Operational Risks. *European Business & Management*. Vol. 4, No. 2, 2018, pp. 55-66. doi: 10.11648/j.ebm.20180402.12**Received:** February 1, 2018; **Accepted:** March 10, 2018; **Published:** April 10, 2018

Abstract: Artificial Intelligence (AI) is gaining a strong momentum in business leading to novel business models and triggering business process innovation. This article reviews key AI technologies such as machine learning, decision theory, and intelligent search and discusses their role in business process innovation. Besides discussing potential benefits, it also identifies sources of potential risks and discusses a blueprint for the quantification and control of AI-related operational risk.

Keywords: Artificial Intelligence, Operational Risk, Technology Benefits and Risks, Machine Learning, Decision Theory, Search Algorithms

1. Introduction

Recent years have seen many successful applications of artificial intelligence (AI) technology, accompanied by a strong interest of the public in their impact. Many of these applications improve the flexibility and/or efficiency of business processes, if not even implementing entire new business models. In this article, we review selected key AI technologies and their potential benefits and risks in business process management (BPM). A successful application of AI should carefully consider both aspects, i.e., take advantage of opportunities to achieve benefits, but also study, quantify, and control the operational risk of AI technologies. As operational risk of AI is a new, yet unseen challenge for many business process experts, we discuss the technology sources of operational risk in AI and identify selected attention points for the business.

A business process is a set of interrelated activities designed to achieve a specific business output or objective and provide value to a customer [8]. Any human activity, be it creative thinking, teaching, trading, manufacturing can be modeled and analyzed as a business process.

A business process is usually triggered by some event, e.g., a customer asking for a mortgage loan, and uses data, e.g., the credit history of the customer. Data drives the business

process towards its goal, which is for example to approve or deny the loan.

Today's business processes take more and more data sources into account, which has been made possible by modern IT technology. The data is mapped to a model, aggregated, condensed, and analyzed in order to arrive at an interpretation and assessment of the data with the goal to predict some aspect of the future. For example, the prediction in the mortgage approval process tries to anticipate the likelihood that a customer can pay down a loan. Based on the prediction, a decision is made and one or several actions are taken. For example, if the creditworthiness of the customer is predicted as solid, then the loan is approved, an account opened, the money transferred etc. Historically, humans have been responsible for the activities that drive the process from data gathering via prediction and decision making to selection of the action(s) to be executed.

Figure 1 presents an abstract model for the main activities performed in a business process that we will use to structure our discussion of AI technologies. ¹Three main transitions lead from data via prediction to decision and finally to action.

¹ We believe this model is well reflecting typical business process designs and widely accepted and applied by practitioners. We were not able to trace it to a specific source or reference, it rather seems to be common wisdom within the business process management community.

The model illustrates how these transitions build on each other. A process takes the given data to *forecast* the future in the form of some prediction. The prediction is the basis to

conclude on some decision, which in turn determines how a human or artificial agent will *behave* by executing some action.

Figure 1. A Model of Common Business Process Designs.

The model also anticipates the change that occurs while the business process and its underlying decision making proceed. An action influences the future and may render previous predictions obsolete, but also a prediction can influence the future, which is captured by the well-known phenomenon of the self-fulfilled prophecy. The future is looping back into the ‘now’ as the process repeats and new, changed data is gathered.

Modern business process management is focusing on the improvement of business processes. This involves making processes more efficient, but also more flexible and responsive to the changing needs of business users. Two current trends in the design of business processes support process optimization and flexibility:

First, business processes become more case-oriented [61, 58, 38] and rule-driven [15, 33, 55, 28, 27] to support the flexible work styles required in a modern economy of service networks. Cases provide an information-oriented basis for dynamic business processes that progress from milestone to milestone while putting the required activities together ‘on the fly’. Business rules describe business policies as guidelines and boundaries for flexible work styles, they make business decisions more transparent and manageable, and they are used to automate business decisions. Cases and rules are supported by recent OMG standards such as CMMN [39], DMN [40], and SBVR [41].

Second, business processes produce and consume more and more data. “Big Data” with its four *Vs* standing for *volume*, *variety*, *velocity*, and *veracity* best characterizes this trend [9].

² Big data is often reduced to these 4 Vs, however, Gartner’s original definition consists of three parts:

Big data is high-volume, -velocity and -variety information assets that demand cost-effective, innovative forms of information processing for enhanced insight and decision making.

It is part 2 and 3, the “cost-effective, innovative forms of information processing” for “enhanced insight and decision making” that require, but also enable further innovation and optimization in business process management. As it is said in [53], “elevator logs help to predict vacated real estate,

shoplifters tweet about stolen goods right next to the store, emails contain communication patterns of successful projects.... Business value is in the insights, which were not available before. Acting upon the insights is imperative”.

Both trends, case- and rule-orientation as well as big-data adoption, provide many opportunities, but also challenges for practitioners as well as researchers.

Artificial Intelligence (AI) is a family of powerful technologies that is particularly well suited to provide innovative forms of business process re-engineering. AI can be embedded into business processes to support humans by intelligent agents or to drive humans out of the process and replace them by fully automated solutions. The latter aspect of automation recently caused a wide discussion in the press about AI being a job killer, see for example [63]. In the following, we will take a closer look on how AI can be applied to support business processes. The model presented in Figure 1 is a good abstraction to study the most commonly deployed AI technologies within business processes. AI technologies can be applied along the entire cycle of *forecast-conclude-behave* to improve decision making and activity execution. AI can deal with huge amounts of data, it is widely applicable and replicable, and it can achieve full automation. This makes it often a technology of choice as it is a cost-effective way to improve the efficiency of business processes. Furthermore, it is innovative, providing insights that have not been possible before. Finally, AI accelerates automated decision making through its various technologies such as intelligent agents, planning, and others.

Our subsequent discussion is settled at a certain abstraction level to provide the ‘big picture’ on how AI technology is applicable in business processes and can thus neither be complete nor go deep into details. However, we hope that this article provides a good starting point for practitioners to further explore and understand the potential and risks of AI technology within business process management. In particular, we believe that addressing operational risks is very important when embedding AI within business processes as AI is a powerful technology.

The paper is organized as follows: We proceed with Section 2, which briefly introduces three key AI technologies: machine learning, decision/utility theory, and search

² Gartner originally contributed the first 3 Vs to the definition, the 4th V for veracity was added by IBM [22] for a good reason.

algorithms. We discuss how these AI technologies support the *forecast-conclude-behave* cycle and discuss opportunities and benefits. Machine learning can be used to improve/automate the forecasting, decision/utility theory can provide more insightful or automatic conclusions, and search algorithms can optimize desired behaviors. In Section 3, we will investigate the operational risks of the three AI technologies under consideration. We will discuss where operational risks origin within the technologies and give some practical hints what practitioners can do to quantify and control these risks. In Section 4, we extend our perspective beyond the technology-related risks and explore the surrounding contexts of business processes. We also propose a blueprint how to control risks across contexts and briefly investigate how control theory can provide control loops for each technology in the cycle as well as over larger business contexts. A key element seems to provide references of desirable behavior and to develop a systematic notion of stability for AI-based business process designs. Section 5 concludes the paper with a short summary of our discussion.

2. Applying AI in Business Processes: The Benefits

Artificial Intelligence is a wide field of research. A definition of artificial intelligence is not straightforward at all and seems to presuppose a good understanding of human intelligence. It is thus not surprising that the major textbook on AI [50] avoids a definition of the term and rather discusses it in the broader context of understanding its various forms of human and rational thinking and acting. A key metaphor underlying AI research for at least the last 25 years has been the notion of a *rational agent*:

A rational agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators. For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidences provided by the percept sequences and whatever built-in knowledge the agent has. [50, p. 35, 38]

Most AI research considers rational thinking and acting as a prerequisite for intelligent behavior. Similarly to the *forecast-conclude-behave* cycle that we use as an abstract model to discuss typical activities within business processes, AI has studied rational thinking and acting within a loop of “prediction, decision, action”. Of course, the separation of making forecasts/predictions from drawing specific conclusions and taking a decision, which finally leads to the choice of some action resulting in a specific behavior is a simplification. In reality, these activities are intertwined in intricate ways and several of these loops overlap each other in human behavior. However, separating these activities is useful to deeper understand them and to develop technologies that support or automate them. AI has done this across its various subfields over many years and developed key technologies such as machine learning, decision theory, and search

algorithms, respectively that we will explore more closely in the following.³ Again, in AI systems these technologies are often combined and embedded into each other. For example, search algorithms can be used within machine learning tools to optimize parameters and minimize classification errors. Machine learning can enhance search algorithms by learning and improving their underlying heuristics. However, to understand the benefits and risks of each AI technology, considering them in isolation is a good starting point and a venue that we follow in this paper. Combined technologies and overlapping loops can boost the benefits of the resulting hybrid solution, but also multiply their risks.

In the following, we can only summarize the key concepts and benefits of each technology, but neither provide a thorough introduction nor completely list successful applications. Instead, we provide selected pointers to further readings.

2.1. From Data to Prediction: Machine Learning

Machine learning methods [2, 11, 21] can be applied to the data available for a business process to compute a prediction. AI has developed many different machine learning algorithms such as for example neural networks or decision trees that are available today in ready-to-use libraries, for example [64, 25]. Learning allows an agent to improve its performance based on its perceptions made in the world. Machine learning research distinguishes among others between supervised and unsupervised learning methods. In unsupervised learning, an agent learns patterns from the input without receiving explicit feedback. In supervised learning, an agent learns from labeled examples of input-output pairs, the training data. The result of its learning process is evaluated on test data on which the agent makes a prediction of the output based on the given input. Learning algorithms differ in the representations of the inputs and outputs, the types of models (input-output functions) they can learn, and how they do the learning. The machine learning community has developed various measures such as accuracy, precision, recall, AUC and others to describe the quality of a learning algorithm. In addition, statistical methods provide information such as confidence intervals and standard deviation. When evaluating several machine-learning algorithms for a business application, these measures are applied in order to decide which learning method should be favored over another. If the percent values of the measures and tests are considered as being “good” or “adequate” (often expressed by values above a defined threshold for an application domain), the learned model is put into practice.

Machine learning is very interesting for business process management as it can detect patterns, i.e., functions that relate input with output, which have remained unnoticed by humans. The methods are also useful to learn patterns when humans have difficulties in describing properties of the input data that determine the output. For example, playing the

³ Other fields, such as knowledge representation for example, are important as well, however, their discussion would go beyond the scope of this paper.

game of Go has been recently mastered by a hybrid deep learning and search algorithm, despite it is very challenging for humans to evaluate a situation in a GO game and choose a particular move that is especially valuable in the light of a specific strategy [54]. In business processes, the ability to detect patterns is successfully used to spot deviating behavior, e.g., credit-card frauds, or to classify data, e.g., segment customers for marketing purposes. Machine learning methods can also be flexibly applied to changing inputs and thereby increase the flexibility of business processes, e.g., adjust the assessment of a case under changing circumstances. They can also be used to mine and detect business rules when monitoring business processes. Many other potential application scenarios exist or can be imagined.

2.2. From Prediction to Decision: Decision and Utility Theory

AI investigates decision and utility theory [26, 32, 6, 47, 43] to represent and compute preferences of rational agents and determine the actions than an agent should perform when maximizing its utility function. Decision theory has its roots in economics and investigates situations where agents have to deal with uncertainty about the current state of the world and the future (because the outcome of actions is uncertain) and where agents have to deal with conflicting goals, for example when they want to maximize the profitability while minimizing risks, see [23, 20] for two pointers to seminal early publications. Decision theory allows agents to evaluate uncertain and conflicting situations under finite or infinite decision horizons and a given set of preferences. These preferences are represented in the form of utility or reward functions that assign numeric values to possible states of the world in order to express the desirability of a world state for the agent. Rational agents will select actions based on the principle of maximum expected utility, i.e., an agent will prefer actions that maximize the expected utility or reward that the agent tries to achieve. The utility function determines whether an agent is risk-averse, risk-neutral, or risk-seeking. Utilities are modeled under the 'closed world' assumption, i.e., they are captured for the known part of the world and need to ignore the unknown. Furthermore, decision theory is normative and describes how a rational agent should behave. It does not describe how humans behave. In fact, it has been observed several times that human decision making often deviates from the mechanisms applied by AI-based theories and that humans quite often do not take rational decisions, but rather behave in an altruist manner. The decisions taken by AI systems can be evaluated by comparing them to existing examples of best practices (the so-called gold standard), an approach which is somewhat similar to the validation of machine learning algorithms on test data. An evaluation should also systematically vary the parameters of the decision model to explore how sensitive the decision outcome is to small changes of the assigned probabilities and utilities.

Decision-theoretic models can underlie any decision that is taken within a business process, e.g., to buy or sell stocks or to

give a recommendation to a customer. AI systems benefit from decision-theoretic models to better deal with uncertain information and to find out which questions to ask to improve their knowledge about the world. Current commercial trends such as personal digital assistants or cognitive computing [62, 56] provide interesting linkages to decision theory.

2.3. From Decision to Action: Search Algorithms

Search algorithms are studied in AI and mathematics (notably operations research) [42, 19, 36, 5, 59, 7, 60] to find optimal solutions to planning and scheduling problems, e.g., which activities to perform in a successful marketing campaign, how to schedule a given number of jobs on a set of machines to optimally produce a product, or how to plan a tour to optimally deliver goods across a set of scheduled sites. AI has developed many different search techniques, which are incorporated into powerful optimization and problem solving techniques today. These techniques often rely on a discrete model of the world that is described by a set of states and actions, where actions are applicable in certain states and their effects determine which state(s) an agent can reach next when executing the action. The states and actions of a problem define the search space where the states are the nodes of the underlying graph and the actions describe the possible transitions between the states represented by the edges of the graph. State-based search algorithms are widely applicable due to the many powerful discrete modeling techniques that are available. They have been scaled to very large search spaces by developing efficient encoding techniques for the models and by applying different strategies to explore large search spaces. In particular, we can distinguish between uninformed and informed search strategies. Uninformed search explores a search space using specific exploration strategies based on the structure of the underlying search graph, whereas informed search uses heuristics to measure the 'goodness' of a state and then decides which state to visit next. Properties such as completeness and optimality characterize a search algorithm, i.e., does the algorithm find a solution in a search space if one exists (for example by visiting all states) and does it return a solution, which is optimal based on some given optimization function.

The value of search algorithms to optimally solve planning, scheduling, and other problems in business process management is of no doubt and many successful applications in logistics, manufacturing, and other domains exist.

3. Inherent Technology Limits: The Sources of Operational Risk

Following [10], we understand operational risk as the risk of loss resulting from inadequate or failed internal processes, people, and systems or from external events.⁴

⁴ Whereas operational risk in banking mostly focuses on financial loss, we adopt a wider understanding where loss is not restricted to money, but can also comprise

The model in Figure 1 provides some basis for first observations of operational risk in business processes. First, data selection may be biased or data may be inaccurate or very different than expected. Second, decisions and actions carry the inherent risk that they are obsolete as the predictions have already changed the world and with that the data. Third, when automating the transitions, we accelerate decision making, but also the change of business conditions, which we can already widely observe in today's economy. Acceleration can lead to time-pressure and shorter decision-making horizons. Accelerated change affects among others the data on which processes are based making it more challenging to anticipate which change will occur and how the change affects any technology that is within a business process. Furthermore, we have to take into account that more and more businesses apply AI and they may impact each other in unforeseeable ways. Let us briefly discuss a specific example of an AI-based innovation of a business process to illustrate these challenges further.

3.1. Operational Risk in an Underwriting Process using AI

Figure 2. Growth of the loan volume, graphics as published in AI Magazine.

The Spring Edition of AI Magazine in 2008 featured an article [31] about an automatic underwriting system called *Custom DU* developed by a mortgage lender Fannie Mae.⁵ Fannie Mae purchased residential home loans from over 2500 lending institutions in the secondary market such as mortgage companies, thrifts, banks, and credit unions, and pooled them together as mortgage-backed securities for sale to investors with a “guarantee” of timely payment of principal and interest. The automated underwriting system allowed mortgage lenders to determine whether a loan complied with Fannie Mae’s underwriting guidelines, but also enabled lenders to define their own lending rules. Lenders could put their loan approval rules into the

easy-to-use system to fully automate the approval of loans in real-time and without the need to do any software development. The published article discusses technical aspects without revealing too many details on how new rules were defined and who was authorized to do so. At the core of *Custom DU* lied the IBM JRules Business Rules Engine, one of the leading business rule technology systems in the market, which uses well established and proven AI technology such as RETE pattern matching [17] and constraint reasoning [45]. *Custom DU* had a huge business impact, which was also demonstrated with the following illustration in the AI Magazine article, see Figure 2.

The illustration indicates a very strong growth of Fannie Mae’s loan volume, although no precise information on the y-axis was published. Improving the efficiency of the process and growing loan volume was a goal for business process innovation in this case, but one is wondering whether such a strong growth should not have alerted business concerns. Only six months later, on September 23, 2008 the Wall-Street Journal published an article where it detailed how it was possible that the vast accumulation of toxic mortgage debt that was driven by the aggressive buying of subprime mortgages and mortgage-backed securities by Fannie Mae and Freddie Mac poisoned the global financial system [13].

It is difficult to analyze details of the system without having access to the details of the model by which the system produced its decisions. However, the text gives some interesting information that is worth to look at. It is not relevant for our discussion to identify any specific mistake as a source of risk accumulation, for example, if the parameters of a decision model are not well designed or if a method used to produce decisions is inadequate in the context of an application. We are interested in principle sources of operational risk that will occur within an AI technology even if the technology is well designed and no mistakes are made.

The AI Magazine article emphasizes in several places how important it was to allow lenders to build their own customized rule sets based on their risk factors and operational challenges, but it never discusses risks and important practices of business rule governance. As it was pointed out in [57], business rules are a technology of great potential, but also of great risk. However, risk considerations received no attention in the article, although the integration of over 2500 lenders in one AI-based system led to economy-and technology-of-scale effects that made it easy to accumulate enormous risks. This accumulation of risks finally contributed to the financial crisis of 2008/09.

What went wrong? On the one hand, we had a too risk-tolerant regulatory regime that did not take action on known systemic and operational risks in the mortgage market. On the other hand, a team of computer scientists and psychologists led by a business manager took well-understood technology and applied it at a massive and distributed scale. We do not know if the team thought about potential impacts of their applications, at least the article does not discuss this aspect of the development. The following factors likely contributed to the problematic side effects of *Custom DU*:

human lives, material goods, hazard to ecosystems, etc.

⁵ In 2007, *Custom DU* was one of the award winning applications honored at the Conference for Innovative Applications of AI “for taking intelligent underwriting to a new level with a Web-based system that enables mortgage lenders to build their own automated underwriting applications” [1].

1. The focus of the development was mostly set on cost and time reductions with the system “giving its clients a decision in minutes, not days” and “the ability to deliver loans with larger amounts, provided they were underwritten by Custom DU”. This focus consequently led to a high incentive to take humans out of the mortgage approval loop. As the article emphasizes “according to a recent mortgage benchmarking study, lenders that deploy automated underwriting at the point of sale recognize the largest per loan cost savings and achieve the greatest per person loan capacity”.
2. The system also enabled a more flexible business process that expanded conventional boundaries, i.e., lenders could use the technology to serve customers who would not conform to the conventional and more cautious credit guidelines and they could define business rules that varied according to channel, customer, product, region, line of business, and other factors. As a consequence, a high variability in using the AI technology was achieved where “the business rule set to be executed can be different for each transaction”.
3. A risk-tolerant culture drove the development, which did not seem to put risk management in the focus of application development. Whereas we read about the user interface to support “policy administrators to create and update rules and messages, test rules and messages, create and update activation rules that determine how rule sets are associated with products, and customize findings”, we never read about the four-eyes-principle or any governance mechanisms that would validate and approve business rules prior to publication. A mechanism of real-time modification of business rules and deployment was implemented that enabled authors of business rules to easily change rules and make them effective for the next transaction.
4. The user-friendly wrapping of the technology “allowed business users to build rules as if they were experienced rule writers” and was used by people who had “never managed a rule-based system” before using *Custom DU*, i.e., the system potentially exposed a complex technology to novice users.

In a nutshell, the AI technology was used to speed up and automate complex decisions. *Custom Du* supported two types of humans: rule authors and loan approval agents. Rule authors defined the boundary conditions for the decisions in the form of business rules, whereas loan approval agents applied the automated decision making to their cases. We do not know whether each player in the application only considered his/her own context or also investigated beyond context boundaries. Defining the decision conditions was made so easy that even non-experts were able to do so. In the decision process itself, humans only played a minor or no role at all anymore.

In the following, we take a closer look at some limits of the AI technologies under consideration and explore in more detail how to systematically derive and quantify the operational risk of AI-based process solutions.

3.2. Limits in Machine Learning

Heterogeneous and unstructured data sources are increasingly combined to build prediction models. Most data is uncertain and one cannot be sure about the veracity of it. Practitioners aggregate and abstract the data, but this does not make the data more reliable. The output of a prediction model is thus inherently uncertain. As was observed in [46], many machine learning techniques focus on fitting a model without providing any estimate on the accuracy of the model.

To arrive at a prediction, machine learning, in particular supervised learning methods can be applied to the data available for a business process. These methods are trained using a set of training data and subsequently validated on a set of test data. All learning methods critically depend on the quality of the data and require among others that the training and test samples must be representative of the domain (normally distributed) and be chosen randomly. Both criteria are not so easy to meet in practice where data is usually collected over a certain period of time, the number of instances is limited as an access to the entire data population is often impossible, data can have hidden dependencies, or there may be an unconscious bias when selecting data due to the nature of business processes. Furthermore, it has been shown that machine learning algorithms can be easily mistaken by hidden states and patterns in the data [37]. These problems become more critical if learning approaches are applied to data where no reliable test oracles exist or when unsupervised learning methods are put into practice. In a nutshell, all that the validation methods can provide so far is a statement that a learner produces a result with some confidence (which is usually below 100%) on a given set of specific test data. For example, an algorithm that is tested as being 99% correct on the test data, will produce an incorrect result in 1 out of 100 cases. What does this error mean for the business process? For example, will it incorrectly assess the creditworthiness of a customer and what will this imply? Will it incorrectly classify a credit card transaction as a fraud or miss a fraud? Can we quantify this impact?

The risk associated with the learner is hidden in the confidence percentages and its (often unknown) dependency on the data, independently on how well the learning algorithm was trained.

As we discussed earlier, unpredictable data changes are likely to result from predictions and actions applied across the universe of businesses processes, but we have in fact no information on how a learning algorithm will behave on changed data, i.e., if the confidence will be different and what data changes will have critical effects on the confidence. In addition, learning algorithms seem to perform well on ‘narrow’ domains where all information is explicit in the data representation, but ‘wider’ domains where data contains hidden states, which are typically occurring in big data applications, seem to pose problems. Recently, the term *out of range behavior* has been introduced to capture the impact of changing data, see e.g., [3] as its challenges go far beyond the previously studied problems of noise and concept drift [51].

Recent results for deep neural networks illustrate in impressive ways how these learners can be mistaken by slight changes in data [37, 35]. Business users should therefore pay attention to the following questions:

1. Is it possible to quantify the data context on which a learned model depends?
2. Can one anticipate and describe potential changes of data and systematically generate test and training data that exhibit the change?
3. What impact on the business process can a data change have given the learned model?
4. Are there data properties that ensure that the confidence remains within given boundaries under data change?
5. Can one translate the confidence value returned by an evaluation method into a quantifiable risk for an application domain?
6. On highly variable data, should one integrate online testing of learned models to monitor whether the confidence of a model is changing?

3.3. Limits in Decision Theory

AI has developed decision and utility theory to represent and compute preferences of rational agents and determine the actions than an agent should perform when maximizing its utility function. A core concept of these theories is the *reward* that an agent tries to achieve. The risk of decision and utility theory can be found in the delta between the reward that the agent thinks it will receive and the reward it actually receives (reward delta). The case of *Custom DU* illustrated that the reward in terms of growing the loan volume was successfully maximized over a certain horizon. However, in the longer run, the agent not only lost all of its reward, but the agent also ceased to exist when the bank went out of business.

The source of the risks lies in the decision horizons considered by agents and underestimated decision consequences which cause a delta between the expected and received reward.

Any utility and decision model can only capture preferences, possible actions and their effect outcomes in a limited way and not map the full complexity of the business context, in which an agent operates. It is thus critical what content is built into the decision and utility models used by autonomous intelligent agents. It seems that when agents take a narrow and 'selfish' approach when considering the utility of their actions, they also tend to ignore long-term effects that may occur outside their immediate context of activity. Furthermore, the utility functions used by agents might not well represent how an agent should deal with sparse events that occur rarely. In addition, the computation of decisions over long horizons is computationally very expensive and may be impossible under accelerated change.

The decision models used by commercial applications are often confidential and it is thus not possible that these models could be inspected for accuracy or risk by the research community or any other independent agency. Although risk assessment and management is a very active research area in many industries, in particular in finance [4], we are not aware

of established standards of risk assessment for AI-based decision making. In addition, it has been observed several times that human decision making often deviates from the mechanisms applied by AI-based theories. Here are some questions for further consideration:

1. Should a business open-source its decision and utility models for transparency and allow customers to opt in or out when having reviewed links between preferences and action outcomes and the resulting rewards?
2. How can a business process use methods that allow intelligent agents to monitor their reward delta and to change (learn new) utility models when the delta increases for long-term effects?
3. How to systematically apply a risk-based comparison of different utility- and decision models for a particular business process?
4. What are systematic stress tests to explore the boundary conditions for a given business process using AI?
5. How do mixed models of human and AI-based decision making play together in a business context?

3.4. Limits in Search Algorithms

To derive an action from a decision, search and optimization algorithms are used that for example solve planning and scheduling problems. Soundness of these algorithms is established by mathematical proofs and we can usually assume that algorithms are correctly implemented, although practical soundness may constitute a problem as the correct engineering of algorithmic implementations is a challenge by itself [30].

It is not too difficult for users to validate whether a solution returned by an algorithm is correct. Thus, theoretical soundness does not seem to be a major issue and adding efficient validation is possible unless a domain requires solutions of huge size with non-polynomial bounded solution length.

However as for optimality, the picture is completely different. When systematic search algorithms are used, optimality can again be established by mathematical proof.

Unfortunately, most practical problems incur search spaces that are prohibitive for systematic search, for example that involve 10^{100} or more states. In this case, non-systematic, local search algorithms using heuristics come into the picture. Proving optimality of these algorithms usually comes in the form of a conditional statement. For example, simulated annealing [24] converges to the optimum if the schedule lowers the temperature *slowly enough*, i.e., an optimal solution is returned with a certain probabilistic confidence. In practical applications, it is often impossible to validate optimality of a solution unless for very small instances and search spaces. If a suboptimal solution is returned, the user does not know how far this solution is from the optimum unless an approximation algorithm with a quality guarantee is used.

Uncertain solution quality can be a source of risks when solving very large and intractable optimization problems.

Similarly to the other AI technologies studied in this paper, a change in the data (search instances) can have an

unknown impact on the (confidence in the) solution quality. Although structural investigations of search spaces have revealed interesting insights such as phase transition regions [14], the structure of search spaces resulting from practical applications, including those that result from combining heterogeneous data, is usually unknown and only few automated methods exist for the structural analysis of state and solution spaces, e.g., [62]. Furthermore, it is still an art to model a problem such that it becomes solvable by a specific algorithm. Practitioners often spend a long time to select an algorithm and fine-tune their model. Small changes in the model often lead to dramatic and unpredictable differences in the performance of an algorithm and again, the model and algorithm can only be tested on some instances of a domain. In addition, commercial applications rarely publish insights into the modeling or reveal details of their algorithms such that little scientific exchange on algorithm engineering for commercial applications can take place. Here are some starting points for further questions when dealing with large search problems in case that non-optimality can induce risks:

1. Does a non-optimal solution induce risks? Which risks?
2. Should one explore the search-space structure to understand the distribution of solutions in the search space and how often such a risk can occur?
3. Do different instances of a search problem yield search spaces with dynamically changing structure?
4. Can the model and encoding of an application be further improved to scale up an algorithm?
5. Should one look at automated methods to learn/generate heuristics?
6. How can different application encodings and search algorithms be compared with each other, in particular with respect to the ability of finding (near)-optimal solutions?
7. Is it possible to apply an approximation algorithm that provides bounds for solution quality?

4. Controlling Risks Across Contexts

In the previous section, we identified three sources of operational risk in the AI technologies under consideration:

1. *Confidence* expressing the accuracy of a machine learning algorithm.
2. *Reward* expressing the expectations of a decision theory model.
3. *Suboptimality* expressing the unknown or reduced solution quality returned by a search algorithm when failing to find an optimal solution in a very large search space.

Each risk can potentially occur even when the application was very well engineered as the data used by a business process and its boundary conditions can change over time in sometimes unpredictable ways. We proposed some questions as starting points for understanding how change can influence the risk sensitivity of an AI technology and to arrive at a quantification of operational risk.

One possible way how businesses can address these risks is to investigate ideas and solutions from control theory. In Figure 3, we introduce a blueprint for operational risk management. It shows how AI technology-related risks occur in each transition and how these risks potentially accumulate in a business process. The risks finally materialize in the actions executed by agents that act in the business process under consideration.

As the figure illustrates, the goal is not to avoid, but to control the risks at each step as well as over the entire process. One key insight of the previous discussion is that there could be inherent bias when data is selected to validate AI technology and that this bias must be made explicit and controlled, see also [52]. Furthermore, each technology comes with an inherent error and (similarly to humans) can make mistakes. Consequently, control loops over each technology as well as the combined technology mix need to be established. An example implementation of this blueprint using communicating automata for the control loop is described in [29].

Figure 3. Operational Risk Management Blueprint for AI.

Practitioners who develop AI-based applications usually establish specific control solutions for their context of deployment, but these are individually tailored solutions, about which little gets published. General methods and

frameworks for quantifying and controlling risk have not been widely researched so far. Control theory can help in this endeavor as we will briefly discuss next. Furthermore, approximate computing [18] is a new research paradigm that

combines imprecise hardware and software to compute high-quality results at lower costs, which could also be promising to study.

Feedback control theory [16] provides the essential elements to embed AI systems into control loops and vice versa. Given some reference, a controller supervises the output of the system and provides input based on the measured error to bring the system back to the reference behavior. This means that providing references and being able to measure errors (deviations from the reference) for predictions, decisions, and actions is the key to risk management of AI technologies. However, this is not straightforward at all and may be even impossible in some applications due to the lack of references. In such cases, one can argue that human insight and responsibility must provide the reference and humans should not be taken out of critical loops, see also [49]. For such application scenarios, more discussion is needed to determine where the border should be drawn between full automation and control and human supervision in business processes.

At a minimum, a control loop must establish boundaries for certain state variables of the system and recognize when the system behavior begins to move outside these boundaries, which also demands that AI systems are able to explain their behavior, recall the explanation components of early expert systems. Control theory has introduced the concept of system stability in case that some of the state variables cannot be controlled. Stability is an important property of human societies and needs to be ensured for mixed human-artificial and purely artificial scenarios. Stability and quantifiable notions of operational risk could help to refine the concept of robustness as for example discussed in [48].

One way to achieve stability is to establish rules that conduct behavior and normative bodies that observe and control behavior. This means that in order to control pre-

dictions and decisions we need to model the actions that they imply, assess the potential impact (effect) of actions, and establish rules that govern action effects. Assessing the impact of actions is easy for agents that act in narrow (technical) domains, but becomes harder when agents operate in open and complex environments. In multi-agent scenarios, effects can be amplified as in the case of our introductory example, or they can be altered and completely modified. A key question seems to be what regulations and rules should be established and how these can be put effectively into practice to detect (and prevent) un-anticipated effects that are considered as harmful, see also [44]. An early example are Asimov's three Laws of Robotics, but the discussion about their usefulness has been rather controversial, see [34]. Machine ethics [12] investigates these questions further.

Controlling risks becomes even more challenging when we consider the wider business context. Modern service networks integrate many players and bring together different services. In such a network, one service provides the prediction, another takes a decision, finally a third one selects and executes an action. Each service itself can be a conglomeration of many agents contributing to the service. The consequences of the prediction can thus lie completely outside consideration when the prediction service is designed and similarly, for all other services. The horizontal and vertical integration of AI technology within service networks and across many players and industries takes the technology out of its narrow context, in which it was usually developed by the community, into an open and unknown space with its own dynamics. Businesses can thus impact each other in unforeseeable ways. Assessing, quantifying, and controlling this impact and its associated risks is key to a successful and beneficial application of AI technologies.

Figure 4 summarizes levels of risks one could distinguish for further and deeper investigations.

Figure 4. Layers of Context for AI Applications.

We started with technology-related risks and identified their responsible sources. These can accumulate to process-specific risks. Within a business, the risks of single processes can sum up to business risks. Business risks can also result when

several players within the same industry accumulate technology and process risks. For example in the financial sector, process acceleration of online trading systems can magnify and speed up unforeseeable impacts. For example,

the increased volatility of the stock market can lead to events such as for example Black Monday on August, 24 2015. Finally, systemic risk can result, which express the possibility that an event at the enterprise level could trigger severe instability or collapse an entire industry or economy. Consequently, control loops over many processes, i.e., entire business contexts or economic systems may be required. Today, we can find such control loops for example in the forms of regulations, e.g., in the financial industry. The effectiveness of such regulations is, however, not completely known.

These wider context risks can result even when each player is carefully addressing risks in its own processes or business contexts as they may completely lie outside the horizon of imagination and investigation. Massive applications could lead to systemic effects that cannot be dealt with at the level of single processes, businesses, or industries. Table 1 lists some potential systemic risks that seem to be worth a further study, requiring to engage in interdisciplinary research between BPM, AI, and other fields. When applying control loops and reference behavior we may still miss to identify behavioral deviations that occur, but that are not monitored. In some way, any solution is closed-loop, but there are always risks beyond that loop. What are the control mechanisms that we need in such situations?

Table 1. Potential Systemic Risks For Selected Businesses.

Application Area	Type of Risks
Finance	financial, economic, political
Medicine/Health care	biological, ethical, societal
Industry 4.0	environmental, social
Social Media	psychological, societal

To conclude, here are a few suggestions of questions one could address in wider business contexts:

1. How can we model risks, study the impact of mixed AI technology solutions, and investigate effects of scale in wider business contexts?
2. What are notions of stability, including critical state variables and their boundaries for specific types of business processes or business models?
3. Could powerful explanation components used by AI systems be helpful within a control loop?
4. How can we improve the systematic and unbiased testing of combined AI technology - are there effective and scalable validation methods that go beyond testing?
5. Which mechanisms for human, artificial, and mixed responsibility should be established for a given context?
6. How can we investigate and develop risk monitoring and escalation mechanisms for combined AI technology mixes?
7. How should the fundamental principles of machine ethics be formulated and translated into rules and boundary conditions for AI-based business process solutions?

To illustrate the first point in the above list, let us return to business rules as an application area of AI, which lied also at

the heart of the *Custom DU* process solution. The recent Decision Model and Notation (DMN) standard of the OMG [40] defines a notation and exchange format for decisions, captured by business rules and represented in decision tables. The OMG standard allows users to override the logical reasoning over decision tables by a *hit policy*, which is applied in the case of overlapping rules, i.e., when several business rules match the input data. Hit policies are based on assigned rule priorities or on the position of the business rule in the decision table and provide a procedural conflict resolution that is used in practice, although overlapping rules often hint towards inconsistent business policies. By using mechanisms like hit policies, the result of the decision computation depends on the physical order of the rules in tables, which is not in line with the formal logical reasoning applied by the underlying AI-based rule engines. The OMG standard is aware of the problematic side of this approach, because a table containing such procedural definitions “is hard to validate manually and therefore has to be used with care” [40, p. 68], but it still admits such an approach as it is also common practice by business users. Should one systematically investigate and demonstrate the potential risks of such approaches and provide better solutions? What are valid combinations of logical and procedural reasoning? How can we automatically validate and quantify the risks of such combinations?

5. Conclusion

The optimization and increased flexibility of business processes drives three key technologies of artificial intelligence research into numerous application domains where these technologies are combined and applied at a larger scale: machine learning, decision and utility theory, and search algorithms. Each technology is very powerful and can provide many benefits, which we briefly discuss. Furthermore, for each technology we also examine its limits and sources of operational risk, and we discuss how we can further evolve our understanding of operational risks of AI in BPM.

We introduce a blueprint for the quantification of operational risk for combined AI technology mixes based on a forecast-conclude-behave cycle and discuss how the notion of risk is exhibited by each technology: confidence level, delta in rewards, and sub-optimality of solutions. We discuss questions that practitioners can ask to further study the quantification and management of risk for each technology as well as AI-based technology mixes within their business context. In addition, we propose to link AI-based process solutions stronger to control theory and approximate computing to provide control loops for each technology in the cycle as well as over larger business contexts. A key element seems to provide references of desirable behavior and to develop a systematic notion of stability for AI systems.

Acknowledgements

Deep thanks to colleagues who discussed and helped

shaping the ideas in this article: Carlos Linares Lo'pez, Donna Dillenberger, Marc-Oliver Gewaltig, and Malte Helmert.

References

- [1] AAAI. 19th conference on innovative applications of artificial intelligence (IAAI) - sampling AI-at-work in today's and tomorrow's business, science and government, July 22, 2007. press release.
- [2] Ethem Alpaydin. *Introduction to Machine Learning (3rd Ed.)*. MIT Press, 2014.
- [3] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [4] Silvio Angius et al. *Risk modeling in a new paradigm: developing new insight and foresight on structural risk*. McKinsey Working Papers on Risk, Number 13, 2011.
- [5] Krzysztof Apt. *Principles of Constraint Programming*. Cambridge University Press, 2010.
- [6] David Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2013.
- [7] Roberto Battiti, Mauro Brunato, and Franco Mascia. *Reactive Search and Intelligent Optimization*. Springer, 2008.
- [8] Tony Benedict et al. *BPM CBOK Version 3.0: Guide to the Business Process Management Common Body Of Knowledge*. CreateSpace Independent Publishing Platform, 2013.
- [9] Mark A. Beyer and Douglas Laney. The importance of 'big data': A definition. Gartner.com article from June 21, 2011, DOI G00235055.
- [10] Operational risk - supporting document to the new Basel capital accord, 2001. <http://www.bis.org> - Bank for International Settlements (BIS).
- [11] Christopher Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [12] Nick Bostrom and Elizer Yudkowsky. The ethics of artificial intelligence. In *Cambridge Handbook of Artificial Intelligence*, pages 316–334. Cambridge University Press, 2014.
- [13] Charles Calomiris and Peter Wallison. Blame Fannie Mae and Congress for the credit mess. *The Wall Street Journal*, Sept. 23, 2008.
- [14] Peter Cheeseman, Bob Kanefsky, and William Taylor. Where the really hard problems are. In *12th Int. Joint Conference on Artificial Intelligence (IJCAI)*, volume 91, pages 331–340, 1991.
- [15] Thomas H. Davenport. *Thinking for a Living: How to Get Better Performance and Results from Knowledge Workers*. McGraw-Hill, 2005.
- [16] John Doyle, Bruce Francis, and Allen Tannenbaum. *Feedback Control Theory*. Macmillan Publishing Co., 1990.
- [17] Charles Forgy. Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence*, 19 (1):17–37, 1982.
- [18] Vaibhav Gupta, Debabrata Mohapatra, Sang Phill Park, Anand Raghunathan, and Kaushik Roy. IMPACT: Imprecise adders for low-power approximate computing. In *17th IEEE/ACM Int. Symposium on Low-power Electronics and Design (ISLPED '11)*, pages 409–414. IEEE, 2011.
- [19] Jeff Heaton. *Artificial Intelligence for Humans*. Create Space Independent Publishing Platform, 2013.
- [20] Ronald A. Howard and James E. Matheson. Risk-sensitive markov decision processes. *Management Science*, 18 (7), 1972.
- [21] Rob Hyndman and George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2013.
- [22] IBM, 2015. <http://www.ibmbigdatahub.com/infographic/four-vs-big-data>, accessed on 11/05/2015.
- [23] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47 (2), 1979.
- [24] Scott Kirkpatrick. Optimization by simulated annealing: Quantitative studies. *Journal of Statistical Physics*, 34 (5):975–986, 1984.
- [25] KNIME Analytics Platform. <https://www.knime.org/>.
- [26] Mykel J. Kochenderfer. *Decision Making Under Uncertainty: Theory and Application*. MIT Press, 2015.
- [27] Jana Koehler. The role of BPMN in a modeling methodology for dynamic process solutions. In *2nd Int. Workshop on BPMN*, volume 67 of *Lecture Notes in Business Information Processing*, pages 46–62. Springer, 2010.
- [28] Jana Koehler. The process-rule continuum - can BPMN and SBVR cope with the challenge? In *13th IEEE Conference on Commerce and Enterprise Computing (CEC-2011)*, pages 302–309. IEEE, 2011.
- [29] Jana Koehler, Chris Giblin, Dieter Gantenbein, and Rainer Hauser. On autonomic computing architectures. Research Report RZ 3487, IBM Research - Zurich, 2003.
- [30] Richard Korf. How do you know your search algorithm and code are correct? In *Proc. 7th Annual Symposium on Combinatorial Search, SOCS*, pages 200–201, 2014.
- [31] Srinivas Krovvidy, Robin Landsman, Steve Opdahl, Nancy Templeton, and Syd-nor Smalera. Custom DU - a web-based business user-driven automated under-writing system. *AI Magazine*, 29 (1):41–51, 2008.
- [32] Max Kuhn and Kjell Johnson. *Applied Predictive Modeling*. Springer, 2013.
- [33] Roy Levin and Marco Iansiti. *The Keystone Advantage: What the New Dynamics of Business Ecosystems Mean for Strategy, Innovation, and Sustainability*. McGraw-Hill, 2004.
- [34] Do we need Asimov's laws? MIT Technology Review, <http://www.technologyreview.com/view/527336/do-we-need-asimovs-laws/>, May 16, 2014.
- [35] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. *arXiv preprint arXiv:1610.08401*, 2016.
- [36] Dana Nau, Malik Ghallab, and Paolo Traverso. *Automated Planning: Theory & Practice*. Morgan Kaufmann, 2014.

- [37] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 15)*, pages 427–436. IEEE, 2015.
- [38] Alan Nigam and Nathan Caswell. Business artifacts: An approach to operational specification. *IBM Systems Journal*, 42 (3):428–445, 2003.
- [39] Object Management Group. Case Management Model and Notation 1.0 (CMMN), 2014. <http://www.omg.org/spec/CMMN/1.0/>.
- [40] Object Management Group. Decision Model and Notation Version 1.0 (DMN), 2015. <http://www.omg.org/spec/DMN/1.0/>.
- [41] Object Management Group. Semantics of Business Vocabulary and Business Rules 1.3 (SBVR), 2015. OMG Available Specification, OMG document number formal/2015-05-07.
- [42] Judea Pearl. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison Wesley, 1984.
- [43] Martin Peterson. *Decision Theory and Social Ethics: Issues in Social Choice*. Cambridge University Press, 2011
- [44] Thomas Piketty. *Capital in the Twenty-First Century*. The Belknap Press, 2014.
- [45] Jean-Francois Puget. Applications et e'volutions d'ILOG SOLVER. In *JF-PLC'95, IVe'mes Journe'es Francophones de Programmation en Logique & Journe'e d'e'tude Programmation par Contraintes et applications industrielles*, page 429, 1995.
- [46] Jean-Francois Puget. Optimization is ready for big data: Part 4, veracity. *IBM Developerworks*, 2015. https://www.ibm.com/developerworks/community/blogs/jfp/entry/optimization_is_ready_for_big_data_part_4_veracity.
- [47] Michael D. Resnik. *Choices: Introduction to Decision Theory*. University of Minnesota Press, 1987.
- [48] Stuart Russell, Daniel Dewey, and Max Tegmark. Research priorities for robust and beneficial artificial intelligence, 2015. http://futureoflife.org/data/documents/research_priorities.pdf.
- [49] Stuart Russell et al. Autonomous weapons: an open letter from AI and robotics researchers, 2015. <http://futureoflife.org/open-letter-autonomous-weapons/>.
- [50] Stuart Russell and Peter Norvig. *Artificial Intelligence - A Modern Approach (3rd Edition)*. Morgan Kaufmann, 2010.
- [51] Jeffrey C Schlimmer and Richard H Granger. Beyond incremental processing: Tracking concept drift. In *Proc. 5th Nat. Conf. on Artificial Intelligence (AAAI)*, pages 502–507, 1986.
- [52] Yoav Shoham. Why knowledge representation matters. *Commun. ACM*, 59 (1):47–49, 2015.
- [53] Svetlana Sicular. Gartner's big data definition consists of three parts, not to be confused with three 'V's, 2013. Forbes.com article from March 27, 2013.
- [54] David Silver et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529:484–489, 2016.
- [55] Jim Sinur. The art and science of rules vs. process flows. Research Report G00166408, Gartner, 2009.
- [56] Siri. <http://www.apple.com/ios/siri/>.
- [57] Jim Hagemann Snabe, Ann Rosenberg, Charles Moller, and Mark Scavillo. *Business Process Management - the SAP Roadmap*. SAP PRESS, 2008.
- [58] Keith Swenson. *Mastering The Unpredictable: How Adaptive Case Management Will Revolutionize The Way That Knowledge Workers Get Things Done*. Meghan-Kiffer Press, 2010.
- [59] Hamdy Taha. *Operations Research: An Introduction*. Prentice Hall, 2010.
- [60] Peter Todd, Thomas Hills, and Trevor Robins. *Cognitive Search: Evolution, Algorithms, and the Brain*. MIT Press, 2012.
- [61] Wil M. P. van der Aalst, Mathias Weske, and Dolf Grünbauer. Case handling: a new paradigm for business process support. *Data & Knowledge Engineering*, 53 (2):129–162, 2005.
- [62] Jean-Paul Watson. An introduction to fitness landscape analysis and cost models for local search. In *Handbook of Metaheuristics*, volume 146 of *International Series in Operations Research & Management Science*, pages 599–623. Springer, 2010.
- [63] Oscar Williams-Grut. Robots will steal your job: How AI could increase unemployment and inequality. *businessinsider.com*, Feb 15, 2016. <http://uk.businessinsider.com/robots-will-steal-your-job-citi-ai-increase-unemploy>.
- [64] Ian H. Witten, Eibe Frank, and Mark A. Hall. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2011.