Designing Process Architectures for User Engagement with Enterprise Cognitive Systems

Alexei Lapouchnian, Zia Babar, Eric Yu, Allen Chan, Sebastian Carbajales

► To cite this version:

Alexei Lapouchnian, Zia Babar, Eric Yu, Allen Chan, Sebastian Carbajales. Designing Process Architectures for User Engagement with Enterprise Cognitive Systems. 10th IFIP Working Conference on The Practice of Enterprise Modeling (PoEM), Nov 2017, Leuven, Belgium. pp.141-155, 10.1007/978-3-319-70241-4_10 . hal-01765247

HAL Id: hal-01765247

<https://hal.inria.fr/hal-01765247>

Submitted on 12 Apr 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

Designing Process Architectures for User Engagement with Enterprise Cognitive Systems

Alexei Lapouchnian¹, Zia Babar², Eric Yu^{2,1}, Allen Chan³, Sebastian Carbajales³

¹Department of Computer Science, University of Toronto

²Faculty of Information, University of Toronto

³IBM Canada Ltd.

alexei@cs.toronto.edu, zia.babar@mail.utoronto.ca,
eric.yu@utoronto.ca, avchan@ca.ibm.com, sebastia@ca.ibm.com

Abstract. Cognitive capabilities can enhance a business process by offering automated analytics-based recommendations on key decisions by applying machine learning techniques. Yet the organizational adoption of such advanced capabilities is difficult as user acceptance of advice and recommendations from an automated system requires the development of trust over time. Business processes and the processes responsible for user engagement with enterprise cognitive systems need not only to be designed, but also have to change together with the supporting processes that can emerge and evolve over a period of time to monitor, evaluate, adjust, or modify the cognitively-enhanced business processes to enable employees to adapt to the enhanced capabilities of cognitive systems. In this paper, we propose a systematic modeling approach to study and guide the design and configuration of process architectures that help enterprises implement and evolve systems with cognitive capabilities. The use of suitable enterprise modeling techniques can facilitate the exploration and analysis of obstacles to adoption and guide the systematic search for viable modes of interaction and cooperation between a human user and a cognitive system.

Keywords: Cognitive Computing, Business Process Management, Enterprise Modeling, Enterprise Cognitive Systems, Business Process Architecture

1 Introduction

Organizations are increasingly looking to adopt and incorporate cognitive capabilities to help with decision-making tasks within the enterprise. Introducing enterprise cognitive systems (ECS) into any enterprise involves adoption challenges across multiple stakeholder and organizational perspectives. Users must learn to adapt to such systems as part of executing their business process (BP) activities, while the systems themselves have to be designed and configured as per each enterprise's unique requirements [1]. Upon reaching the desired levels of trust, confidence, accuracy, and reliability, these systems can increasingly become entrenched into the enterprise, which results in changes to the design and execution of business processes and the responsibilities of

users in charge of these BPs [2][3]. Enterprise cognitive systems (ECS) refer to software applications that utilize cognitive services, artificial intelligence, and machine learning techniques (among others technologies) to assist with and provide decision-making recommendations for human process participants [4]. Compared to traditional non-cognitive information systems, they do not have to be programmatically instructed to handle all future decisions and can be self-learning and self-adjusting while gradually taking on more decision-making responsibility in business processes [1].

Shifts in work allocation between human users and cognitive systems can be envisioned. They are triggered by well-defined conditions that focus on system performance, users' trust in systems' recommendations, and other important characteristics of human-system engagement. This engagement is not static. It varies over time based on not just evolving capabilities of enterprise systems, but also on user requirements and enterprise context and objectives [5]. In this paper, our view of user engagement (UE) is different from the notion of user engagement (sometimes called *customer engagement*) in marketing, which usually refers to how frequently, how long, etc. users interact with a (software) product [6]. We, on the other hand, use user engagement to describe the ways in which human users interact with enterprise cognitive systems, including the assignment of responsibilities for communication/collaboration, issues of trust, and the ability to override the system. Users' attitude towards automation in general, and the system they are interacting with in particular, can change based on the accumulated history of UE and the quality of the output the system produces – e.g., from feeling threatened, to cautiously optimistic, to wildly enthusiastic about it, and then possibly back to being cautious. These attitudes translate into different desired ways of interacting with the cognitive systems. The nature and impact of these changes to UE may not be limited to the BPs containing the decisions that cognitive systems are assisting humans decision makers with, but can also cause the introduction and modification of the additional supporting processes that monitor, evaluate, adjust, or modify the cognitively-enhanced business processes to enable humans and cognitive systems to adapt to their changing capabilities as well as business contexts and requirements. These business processes need to be studied together as a business process architecture (BPA) in conjunction with the drivers that cause shifts in user engagement and the many possible UE configurations resulting from those shifts.

2 Considering User Engagement During Business Processes Design

Recent years have seen a significant advancement in the sophistication of cognitive computing and its practical utilization in any enterprise. Maturity of cognitive computing and cognitive science is allowing for its increased penetration in the enterprise, with nearly every industry expected to be disrupted by the introduction of cognitive capabilities [8]. We use domain examples below to motivate for the approach discussed in the subsequent sections of the paper:

- Employees are responsible for repetitive execution of enterprise business processes, such as inbound and outbound customer service, processing customer-initiated loan applications, etc. Key user decisions made during process execution, while being informed by relevant data and policies, are based on experience and gut-feel, with the decision-making process considering various contextual and situational factors. As a result, the decisions may not be uniform and consistent and can vary significantly among different decision makers. Increasingly sophisticated cognitive systems can initially help their human users with recommendations for key decisions and later take complete responsibility for decision making, without the assistance or involvement of the users [9][10]. This would, additionally, alleviate the problem of inconsistent or inaccurate decisions.
- Enterprise IT and cloud infrastructure is ever increasing in size and has stringent requirements for uptime, security, and performance while adhering to tight technology budgets. Consequently, IT departments increasingly rely on automation of key technology activities (like the overall management and monitoring of IT assets) to aid their engineering teams [11]. Cognitive systems are another tool in their overall arsenal that can take advantage of useful application and systems data to streamline deployment, orchestration, and monitoring of various technology assets [12]. By allowing such systems to take over menial and routine tasks, various technology and engineering teams can focus on higher-valued governance-related activities, resulting in the overall shift in user responsibilities and functions.
- Enterprises are typically in the business of providing services to their customers. To reduce operational costs and improve customer services, cognitive systems are increasingly being used as a front-end channel for user communication where they provide a human-like interaction experience to the end-user by utilizing a slew of cognitive capabilities. Financial and wealth management advisory services [13] or personal assistants [14] are some common applications of such systems. The concerns of end-users interacting with these cognitive systems would be different from internal enterprise users and should be considered separately when designing enterprise solutions for cognitive systems integration.

The above examples provide some interesting points to consider. From an *enterprise* and *user* perspective, users are engaging with the new types of cognitive systems being introduced while dealing with real-world business situations. Both sides (i.e., the users and the systems) need to adapt and adjust to each other and eventually converge to a workable state: the users learning to execute their assigned business processes while the systems undergoing cycles of iterative improvements to make them significantly more efficient and intelligent. Factors affecting the adoption success (of the systems by the users) may include the knowledge/skills of involved personnel, their aptitude for understanding the cognitive systems' capabilities and limitations as well as their trust in such systems, willingness to learn and adapt, attitude towards and trust in automation in general, labour relations, reward structures, business domain regulations, etc.

From a *systems* perspective, how a cognitive systems solution is accepted in an enterprise can be very specific to the situation in that organization. But even this situation will continue to evolve as the cognitive system gets better through machine learning or

acquires new features, and on the side of user organization, as employees gain experience or learn new skills. Thus, enterprise cognitive systems should be capable of supporting a variety of enterprise business process configurations, with their own roles ranging from assistive, to advisory, to complete responsibility for decision making. Designers of cognitive systems cannot be expected to predict or prescribe exactly how the human side is going to use these systems.

From a *process* perspective, the impact of process-level user engagement is not limited to just direct system interactions, but includes the related processes as well. By related (or surrounding) processes we mean upstream processes that contribute in some way to the primary cognitively-enabled BP, or downstream processes that benefit from the output of the cognitively-enabled BP. These surrounding BPs too evolve in response to changes in cognitive systems' capabilities. Thus, multiple processes need to be considered for analyzing and designing user engagement. Some of these processes may operate at the transactional level while others may execute infrequently. Navigating the space of such possible process configurations that support multiple modes of decision making (from fully manual to partially autonomous to fully autonomous decision making) while considering functional and non-functional objectives of organizations and their users is difficult and may result in trial-and-error practices being employed without convergence to an acceptable solution.

Hence a large space of possible options for user engagement with cognitive systems needs to be considered, with a variety of factors and the complexity of the domain to be taken into consideration. Enterprise architects need to be prepared to help the user organization explore this space by employing techniques from their repertoire to point out the design decisions, possible alternate configurations, while considering aspects such as trust and confidence in the systems' decision-making ability, effort required to help with decision making, compliance with industry regulations and company rules, the cost of deployment, etc. Towards this end, we describe some general types of modes of user engagement suitable for different levels of cognitive sophistication in the enterprise. These lay the foundation for a systematic modeling approach (presented in the later sections) that helps model cognitively-enhanced business processes and considers the evolving needs, capabilities, and expectations in enterprises when designing and operating cognitive systems while also helping to search for workable user engagement arrangements and enabling the reasoning about why certain engagement approaches work and others do not.

3 Towards a Framework for Designing Cognitively-Enhanced Processes

In the previous sections, we highlighted the large space of possibilities for user engagement with cognitive systems that needs to be analyzed based on enterprise and user requirements, the level of trust in the system, the quality of system output, and system and user capabilities. Moreover, we stressed the fact not just the process containing the decision in question, but many additional processes can be impacted when user engagement evolves. This points to the need to model and analyze user engagement not (just)

at the level of BPs, but also at the level of a BP Architecture. In this section, we outline an approach for modeling user engagement configurations at both the process and process architecture levels using an advanced BPA modeling approach.

3.1 User Engagement Actions

Let us first look at how UE can be designed from the ground up by focusing on *User Engagement Actions* (UEAs), which represent the elementary interactions that comprise UE with cognitive systems. Some are executed by human users, while others by cognitive systems. The set of these primitives that can potentially be selected for some UE configuration depends on the nature of the task that humans and cognitive systems are solving. For example, if a task is a decision to be made, then the relevant UEAs involving a human decision maker (H) and a cognitive system (C) can include the following tasks:

- *Communicate Case Data*: H communicates the details of the particular decision instance to C. Alternatively, C can obtain the relevant data itself.
- *Communicate Decision Parameters*: H communicates the decision parameters to C, including the criteria for making it, the desired confidence levels, etc.
- *Present Recommendation*: C presents the recommended decision to H. For decisions with more than two potential options (those that are non-binary), a ranked list of options may be given.
- *Approve Recommendation*: H either approves or rejects the previously presented recommendation. Variants for decisions with more than two options may include the ability for the user to pick an alternative option.
- *Explain Recommendation*: Explanation and justification of the recommendation are presented by C to H.
- *Present Decision*: C presents a previously made decision to H (or to a specially designated auditor). Variants include the presentation of batches of decisions for subsequent audits.
- *Audit Decision*: H (or a designated auditor) audits the decision previously made by C. Variants include the audit of batches of decisions.

The above is just a sample of possible actions that may exist to support interactions between humans and cognitive systems. Note that these UEAs can take many forms and employ various media. E.g., recommendations from advisors can be presented as text, voice, video, etc. The same applies to the UEAs that are executed by decision makers. In this paper, we abstract from these details.

3.2 Modeling Cognitively-Enhanced Processes

Organizations and individual decision makers will strive for UEs that reflect their changing enterprise requirements, business domain constraints, the level of trust (of both human decision makers and organizations) in analytics in general and in cognitive systems they are employing in particular, the quality and availability of relevant data, and

contexts. This leads to different combinations and configurations of UEAs selected at different times. E.g., the Explain Recommendation UEAs can be omitted in case of a high existing level of trust in the system while audits of system-made decisions can be added due to industry regulations. We refer to these UE configurations as User Engagement Modes (UEMs). UEMs need to evolve together with the changes in the above parameters as well as due to the feedback reflecting how they meet their objectives. Also note that for many UEAs, there exist many options about when, how frequently, etc. to execute them. E.g., decision parameters can be selected once for all instances of the decision, changed periodically, or provided for every decision instance. Recommendations can be approved per decision instance or “pre-approved” for a group or for all instances depending on the level of trust in the advisor, the similarity of the current decision instances to previous ones, etc.

Fig. 1. BPA for a loan approval scenario with manual approval of automated recommendations.

Overall, we need to be able to characterize the space of alternative UE configurations reflecting the whole spectrum of UEAs for a given decision, their potential combinations, frequency and scope of their execution, and context, among other things. Standard BP modeling notations are not well equipped to represent these options. To analyze such as space, we need to take into consideration several levels (industry, organization, individual) of requirements and constraints. Further, there would be transitions across sets of UEAs due to changing enterprise requirements, contexts, etc. Also, as previously noted, changes in UE frequently affect related processes and give rise to new processes (e.g., for auditing system-produced decisions).

To address the above challenges, here we apply a previously introduced conceptual modeling framework hiBPM (for *higher-order* BPM) [15][16][17][18] for modeling user engagement processes together with organizational and automated processes that surround them and how they relate to each other. This is possible since the framework, which has BPMN as its basis, is being developed for modeling not just single BPs, but multiple BPs and their relationships, thus focusing on BPAs of organizations. Fig. 1 presents a simplified BPA for a loan approval scenario, with the loan approval/rejection decision being the focus of our analysis. Fig. 1 illustrates the domain-specific processes (e.g., Loan Approval and Loan Repayment), cognitive system-specific processes (e.g., Analytical Model Creation), and BPs that are part of the proposed approach for managing UE, such as UEM Selection (see Section 4 for further details on the method).

Fig. 2. HiBPM fragments showing alternative UE patterns for loan approval: A) Preset decision parameters and no recommendation explanation. B) Automated decision making with manual audits.

As part of applying hiBPM, we map UEAs into primitive process activities (referred to as *Process Elements* or PEs) in the BPA. Then, different configurations of PEs will correspond to different UEMs. The chosen set of UEAs needs to be injected into processes executed by both human decision makers and the cognitive systems. In a loan approval scenario (Fig. 1), the UEAs are modeled as shaded activities. It is not just the set of UEAs that can change, but also how these UEAs are executed, how frequently, by whom, etc. In hiBPM, PEs can be placed in various processes called *stages*, which are characterized by their execution frequencies. E.g., a decision maker (a loan officer) can supply desired decision parameters to cognitive system for each instance of a loan approval process (Fig. 1) – note the activity *Communicate Decision Params* in the *Loan Approval* process stage. This makes sense if those parameters are likely to be different for every decision instance. Alternatively, the user can preset decision parameters, which is modeled in Fig. 2A by placing the corresponding PE into the *Loan Decision Config* stage that is executed once for all (or some number of) loan approval instances (note the 1:N annotation on the outgoing flow, which indicates the recurrence relationship among stages: for each execution of the first stage, many executions of the second one are possible). This is appropriate if those parameters do not depend on the characteristics of each loan approval case. Additionally, note that stages are annotated with actors responsible for their execution and that feedback paths are also indicated to

model, for instance, that users' approval/rejection of recommendations feeds back into the stages responsible for analytical model design and for UE management.

hiBPM supports two higher-order relationships among stages: *plan/execute* and *design/use* (indicated by X and U in the hiBPM models, respectively). Plan/execute models that one process stage produces a specification (plan) to be executed by one or more stages. E.g., in Fig. 1, the UEM Selection stage specifies which UEAs are to be currently executed for the loan approval process. In Design/Use, a stage can produce artifacts/tools to be (re)used by one or more subsequent stages. E.g., in Fig. 1, Analytical Model is being used by Loan Approval BP to help with decision making.

hiBPM uses goal models [19][20] to systematically represent and analyze BPA configuration options as well as the trade-offs among competing quality objectives, such as flexibility, cost, performance, etc. for helping design BPAs that suit particular organizations.

3.3 A Fine-Grained Way for Handling User Engagement

To design UE, a large space of options defined by numerous UEA combinations and impacted by requirements, domain constraints, and business context needs to be explored and analyzed. Design and evolution of UE at the UEA level produces a granular and flexible approach capable of changing UE design by adding/removing certain UEAs and positioning them in the BPA. A goal-driven approach inspired by [19] is envisioned, with goal models being used to represent functional goals and alternatives for achieving them and non-functional requirements (NFRs) used as criteria for selecting the appropriate options. Goals are then linked to BP actions responsible for achieving them. Appropriate configurations of these actions are chosen with the help of goal reasoning algorithms based on functional and non-functional requirements and context. Using this approach improves transparency and predictability since goal models are human-readable.

Similarly, methods based on declarative process modeling [21] or AI-type planning can be used. These approaches would support the dynamic identification of new UE configurations based on changing goals and contexts. These capabilities match cognitive systems' self-adaptation and self-learning capabilities and their support (with the appropriate feedback) for handling previously unidentified conditions, unexpected changes in data patterns, etc. Conversely, such power and dynamism make the evolution of UE quite unpredictable and introduces additional complexity. Development of such approaches is part of our future work.

4 A Methodology for Designing Pattern-Based User Engagement

This section presents a methodology for identification, selection, and management of user engagement at the UE pattern level that can be used as an alternative to the UEA-based approach outlined above.

For many organizations in many business domains, the fine-grained UEA-based approach described in Sec. 3.3 might be overly complex while at the same time lacking in transparency and predictability and providing more flexibility than required. This points to the need for a simpler approach for well-known decision types and for organizations requiring more transparency and predictability in their users' engagement with cognitive systems. One way to address this issue is to propose an coarser-grained approach focusing on UEMs rather than UEAs. In such an approach, given a type of a decision problem, a set of UEMs, each representing a typical UE pattern – a tried-and-tested selection of UEAs and their allocation to the appropriate process stages – is identified. Since finer-grained adjustments (at the UEA level) are not possible, the identified set of UEMs captures the space of UE configurations completely. Such an abstraction reduces the number of options for UE, thus simplifying its management at runtime while also providing more transparency and predictability. Moreover, important types of decision problems can be analyzed ahead of time and the relevant set of UE patterns can be identified for them, to be reused by many enterprises. This supports organizations that may be cautious in relation to cognitive technology and automation in general and simplifies the adoption of cognitive systems in general.

Given some organization, a cognitive system, and a problem at hand (i.e., a decision to be made), how do we design and manage UE in this context? The solution should be a method that takes into consideration organizational goals as well as personal goals of the users involved in decision making, constraints originating within the enterprise and coming from the business domain, and various contexts. For instance, for loan approval, some of these constraints might be the need to either manually approve or audit all recommendations made by automated systems (as illustrated by all the user engagement configurations shown in Figs. 1 and 2). Moreover, a method for managing user engagement needs to support its evolution.

Below we present the steps of the method for designing and evolving UE at the UEM level. Note that with the help of hiBPM, UE management processes can be represented in the same model as runtime domain-specific BPs, such as those involved in loan approvals (see Fig. 1).

Step 1: UE Pattern Identification (See UEM Identification stage in Fig. 1). This step is typically done more rarely than the following steps. This is captured by the recurrence relationship (1:N) the stage has with the subsequent stages. Here, UE patterns are designed using goal-driven analysis techniques available in the hiBPM approach [15][16] and then empirically validated. Given high-level quality objectives, appropriate UE patterns are identified for particular types of decisions (e.g., approve/reject vs. that with a higher number of possible choices). For instance, for an approval type decision (i.e., with only *yes/no* options), some of the UE patterns that can be identified include:

- *P1: Supervised learning.* Decisions are made by a human decision maker while an enterprise cognitive system (ECS) monitors his work. ECS uses case data, context, and the decision outcome as the input to a supervised learning algorithm.
- *P2: ECS as an Advisor.* Human decision maker makes the decisions while ECS's recommendations are presented as advice.

- *P3: Human approval of ECS's decisions.* ECS's recommendations are approved or rejected by human decision maker (per decision instance) – e.g., see Fig. 1.
- *P4: Human is informed of decisions by ECS.* A human decision maker is informed of ECS's decisions (per instance).
- *P5: ECS's decisions with (batch) human audits.* Humans audit ECS's decisions periodically (once per N number of decision instances, once in a certain time interval, etc.) – see Fig. 2B as an example.
- *P6: ECS's decisions with human audits on request.* Humans can audit ECS's recommendations whenever they wish.
- *P7: ECS's decisions with automated self-audits.* Humans are not involved by default, but can review automated self-audits.

In Figs. 1 and 2, selected UE patterns are evident from the UEAs (shown as shaded activities) that are present in the Loan Approval process. Variations of the above patterns with slightly different UEA arrangements are possible. In Fig. 1, one can see that the decision maker (a loan officer in this case) communicates case data and decision parameters within the Loan Approval stage (i.e., for every instance of that loan approval decision), indicating that volatility is expected in these areas. Upon presenting its recommendation, a cognitive system explains it to the user. This indicates that the loan officer does not yet fully trust the system. It is a variant of the UE pattern P3.

Step 2: UE Pattern Pruning (See UEM Pruning stage in Fig. 1). To further customize user engagement, the set of UE patterns identified in Step 1 *can* be pruned based on particular organization's and/or user's requirements and constraints. E.g., in some domains and/or organizations, audit of decisions may be mandatory and only the patterns containing the audit UEAs will be selected. This step of the approach is normally executed for every decision problem to select the set of UE patterns available for each decision that requires cognitive system assistance.

Step 3: UE Pattern Transition Setup (See UEM Transition Setup in Fig. 1). Here, we switch our attention to the dynamic aspect of UE and model its possible evolution trajectories by specifying transitions between UE patterns triggered based on a certain level of trust in the system (e.g., Fig. 2A shows the UE pattern missing the task Explain Recommendation, perhaps due to a high level of user trust in a cognitive system), a certain quality of recommendations (recent recommendations of lower quality may warrant a reduction of automation level and the shift, say, from automated decision making to cognitive systems only providing advice to humans). This step is usually run once per decision problem instance – e.g., for a particular loan approval BP.

Context change may also trigger a transition to a different UE pattern – e.g., a significant shift in the properties of decision instances compared to the dataset on which the analytical model was trained can cause a selection of a more conservative (less-automated) UE pattern until the model has been retrained. Quality requirements and organizational or problem domain constraints will affect the transitions as well. Similarly, the ability to manually trigger transitions or override automatically triggered transitions, which puts humans in control of the automation, is important as trust plays a crucial role in UE. All of this helps with transparency and predictability in UE evolu-

tion. While this approach can only handle predefined conditions and evolution trajectories since they are explicitly defined upfront, all methodology steps can be re-run when needed. Thus, UEM Identification and UEM Transition Setup can be re-executed to update the set of UE patterns and/or to change the transitions among them. This produces a potentially new output to be reused in the subsequent stages in the hiBPM model.

Statecharts are used to define possible UE transitions, with states representing UE patterns and transitions referring to the possible UE evolution paths. Fig. 3 shows a statechart capturing one possible UE evolution specification using the patterns discussed in Step 1 and modeled as states in the model. Transition labels specify conditions that trigger changes in the active UE configurations. E.g., once the analytical model is trained in P1, the UE state changes to P2 where a cognitive system is acting as an advisor to humans. Then, once its recommendations are deemed to be of high quality, the system starts making decisions and applying them with human approval (P3). If a change in business context is detected in the states P2-P4, the supervised learning UEM (P1) is activated. Greyed out states (P5) are those removed in Step 2 of the approach.

Fig. 3. An abstract statechart illustrating UEM transitions.

Step 4: UE Pattern Selection (See UEM Selection in Fig. 1). In this step, a process running concurrently with the BP containing the decision (Loan Approval in our examples) monitors the various parameters through feedback (focusing on decision quality, etc.), evaluates triggers, and drives the appropriate changes in user engagement. It uses a model like the one in Fig. 3 to enact UE changes. Looking at the hiBPM model in Fig. 1, we see that UEM Selection *executes* the transitions *planned* by UEM Transition Setup (thus, the plan/execute relationships between these BPA stages). Moreover, UEM Selection itself implements the chosen UE pattern by instantiating the appropriate stages and instructing which UEAs to execute (also modeled by plan/execute stages). Different UE patterns create different BPA variations, as evident from the comparison of Figs. 1, 2A, and 2B.

Fig. 2A and 2B give examples of the range of options for specifying UE, showing both the set of UEAs and their positioning in the BPA. Fig. 2A is a fragment BPA illustrating a UE pattern where decision parameters are preconfigured (see the separate Loan Decision Config stage), no recommendation explanation is provided, and manual UE pattern change is no longer available (compared to Fig. 1). Fig. 2B shows a fully automated loan approval process. There, Approve Recommendation is executed in the stage prior to Loan Approval, indicating that cognitive system's recommendations are approved for a certain time frame or number of decision instances, thus stressing the

high level of trust in the system. Note that this “pre-approval” is about system recommendations and is different from the domain-specific action of pre-approving loans.

5 Related Work

Decision making and decision theory are much-studied areas in psychology and having an understanding of the deliberation process and the outcome can help in designing artificial cognitive systems that emulate a human decision-making process. Rational and descriptive approaches to decision-making and judgements are discussed in [22], including how decisions can be reached by deliberating among complex alternatives while considering emotions. Actors may not always invoke an optimum logic to select the best alternative within a set of possibilities, but rather have preferential solutions to a problem based on prior experience; this selection of a solution as part of routinized decision making is a component of an ongoing learning and evolving process [23].

Within any enterprise, there exist multiple business processes that can be collectively considered as part of a business process architecture, where the BPA “is a collection of business processes and their interdependencies with each other” [24]. In the enterprise architecture area, the notion of BPA is referred to as business process cooperation. For example, in ArchiMate, such cooperation includes the following aspects: *causal* relationships between business processes, *mapping* of business processes onto business functions, *realization* of services by business processes, and the *use of shared data* [7], and can also imply the type of relationships among business processes. Compared to ArchiMate, our hiBPM approach supports the meta-level relationships among BPs, such as plan/execute, which helps integrate multiple levels of processes into BPAs. For instance, it supports the modeling of UE evolution.

Cognition is the process by which “an autonomous system perceives its environment, learns from experience, anticipates the outcome of events, acts to pursue goals, and adapts to changing circumstances” [25]. The impact of cognitive computing on business process management (BPM) is covered in [26] where multiple types/levels of business processes are discussed – transaction-intensive, judgement-intensive, and design & strategy support processes – resulting from the incorporation of cognitive capabilities within an enterprise and how cognitive processes enablement can be attained.

Human-Computer Interaction (HCI) involves the study of the communication of information and data between a human user and a computer through an interface, using various modes and mechanisms such as input/output peripheral devices, voice, text, sound, gesture control etc. [27]. HCI aims to greatly improve the experience of the user interaction by leveraging the principles of information design, interaction design and sensorial design [28]. We approach the problem of user engagement with machines differently in that we consider the shifts in user approaches for decision making (as part of a routine business process execution) that are caused by changes in the design and implementation of cognitive systems.

Agent-Oriented approaches to software engineering provide design and implementation guide for software agents that need to demonstrate autonomous behaviour and

social ability while responding (both proactively and reactively) to changing environmental stimuli [29]. Software agents need to have intentions, beliefs and mental states based on which they'd plan, negotiate and perform actions which are deduced rather than programmatically executes. In this paper, we considered cognitive systems with much less social ability compared to a typical software agent. The body of research on agents and multiagent systems can provide a wealth of techniques and ideas for handling user engagement with cognitive systems.

6 Conclusions & Future Work

This paper focuses on the issues of designing and configuring BPAs where user engagement with enterprise cognitive systems is considered. Such user engagement needs to take into consideration trust in the systems' decision-making abilities, organizational and user requirements and constraints, and evolving business and technological context. A UEM – a way a human interacts with a cognitive system – and the BPA implementing it are not static and must evolve since the above parameters are also dynamically changing. To support the design and evolution of UE, user engagement actions were proposed as the building blocks of UE. A systematic model-driven approach based on the hiBPM BPA modeling framework is outlined. This framework is utilized for modeling the space of UE options through multiple levels of process modeling, higher-level relationships between business processes, and other advanced features. Although not illustrated in this paper, to analyze the space of options, hiBPM's requirements-driven techniques focusing on reasoning about objectives, alternative solutions, and trade-offs can be utilized.

This approach is a starting point for the development of a comprehensive method aimed at simplifying enterprise adoption of advanced cognitive systems by enabling the systematic design of UE with such systems based on organizational, social, and technical needs and constraints and supporting the ongoing evolution of such engagements driven by changes in these parameters. Overall, this method is aimed at supporting a systematic and controlled introduction of automation into organizations. Among other things, future research in this area will focus on the following:

- As stressed in the paper, the issue of trust is of paramount concern when it comes to user engagement. We plan to investigate how to incorporate the analysis and evaluation of trust into our approach together with the identification of various UEAs aimed specifically at establishing, maintaining, and increasing trust.
- Introduction of cognitive systems in enterprises usually results in increasing automation and thus causes changes in responsibility assignments among humans and automated systems. We plan to study this dimension of change in more detail while focusing on the social and organizational impact of such changes.
- We plan to focus on more complex types of decisions/recommendations (compared to the approve/reject type discussed in this paper) and identify new sets of UEAs and UE patterns as well as the typical transitions among these patterns for those decision types. Moreover, we plan to develop the corresponding BPAs for those UE configurations to simplify the adoption of cognitive systems by enterprises,

- We would also like to concentrate on a more thorough modeling and analysis of feedback in the context of user engagement and on identifying the typical metrics and triggers to be used in specifying UE evolution.
- While we are building on the foundation of goal and process modeling many features of the proposed approach are novel and need to be further studied, implemented, and validated.

At present, we are working with a large enterprise software provider and aim to introduce the proposed approach into its business process management offering and subsequently validate it in real-world engagements.

Acknowledgement: This work was partially funded by IBM Canada Ltd. through the Centre for Advanced Studies (CAS) Canada (Project #1030).

7 References

1. Ogiela, L., Ogiela, M. R.: Advances in cognitive information systems, Vol. 17, Springer Science & Business Media (2012)
2. Muir, B. M.: Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies*, 27(5-6), pp. 527-539 (1987)
3. Lee, J. D., See, K. A.: Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, 46(1), pp. 50-80 (2004)
4. Schatsky, D., Gurumurthy, R.: Cognitive technologies – Applying machine intelligence to traditional business problems, Deloitte University Press (2015). Retrieved from <https://dupress.deloitte.com/dup-us-en/focus/signals-for-strategists/trends-cognitive-technology-business-issues.html>
5. Hull, R., & Nezhad, H. R. M.: Rethinking BPM in a cognitive world: Transforming how we learn and perform business processes. In *International Conference on Business Process Management*, pp. 3-19, Springer (2016)
6. O'Brien, H. L., & Toms, E. G.: What is user engagement? A conceptual framework for defining user engagement with technology. *Journal of the Association for Information Science and Technology*, 59(6), pp. 938-955 (2008)
7. ArchiMate 3.0 Specification. Retrieved from <http://pubs.opengroup.org/architecture/archimate3-doc/>
8. Chui, M.: Artificial intelligence the next digital frontier?, McKinsey and Company Global Institute, (2017). Retrieved from https://www.mckinsey.de/files/170620_studie_ai.pdf
9. Bond, A. H., & Gasser, L. (Eds.): *Readings in distributed artificial intelligence*, Morgan Kaufmann (2014)
10. Kushniruk, Andre W.: Analysis of complex decision-making processes in health care: cognitive approaches to health informatics, *Journal of biomedical informatics*, 34(5), pp. 365-376 (2001)
11. Cherubini, G., Jelitto, J., Venkatesan, V.: Cognitive Storage for Big Data, *Computer Magazine*, 49(4), pp. 43-51 (2016)
12. Talia, D.: Towards Internet Intelligent Services Based on Cloud Computing and Multi-Agents. In: Gaglio S., Lo Re G. (eds) *Advances onto the Internet of Things. Advances in Intelligent Systems and Computing*, Vol 260, Springer (2014)

13. Benjamin, B., Burrowes, K., Gemes, A., Hogan, A., Lenzhofer, A., Ong, J., Spellacy, M.: Sink or swim: Why wealth management can't afford to miss the digital wave, Strategy&, PwC (2016). Retrieved from <https://www.strategyand.pwc.com/media/file/Sink-or-swim.pdf>
14. Myers, K., Berry, P., Blythe, J., Conley, K., Gervasio, M., McGuinness, D. L., Morley, D., Pfeffer, A., Pollack, M., Tambe, M.: An intelligent personal assistant for task and time management, *AI Magazine*, 28(2), pp. 47-61 (2007)
15. Lapouchnian, A., Yu, E., Sturm, A.: Re-designing process architectures towards a framework of design dimensions. In *Research Challenges in Information Science (RCIS 2015)*, IEEE 9th International Conference on, pp. 205–210. IEEE (2015)
16. Lapouchnian, A., Yu, E., Sturm, A.: Design Dimensions for Business Process Architecture. In *34th International Conference on Conceptual Modeling (ER 2015)*, pp. 276-284. Springer (2015)
17. Lapouchnian, A., Yu, E.: Exploiting emergent technologies to create systems that meet shifting expectations. In *24th Annual International Conference on Computer Science and Software Engineering*, pp. 371-374, IBM Corporation (2014)
18. Babar, Z., Lapouchnian, A., Yu, E.: Modeling DevOps Deployment Choices using Process Architecture Design Dimensions. In *The Practice of Enterprise Modeling*, pp. 322-337, Springer (2015)
19. Lapouchnian, A., Yu, Y., Mylopoulos, J.: Requirements-driven design and configuration management of business processes. In *Business Process Management*, pp. 246–261. Springer (2007)
20. Lapouchnian, A., Yu, Y., Liaskos, S., Mylopoulos, J.: Requirements-Driven Design of Autonomic Application Software. *IBM CAS Conference (CASCON)* (2006)
21. De Giacomo, G., Dumas, M., Maggi, F., Montali, M.: Declarative Process Modeling in BPMN. In *International Conference on Advanced Information Systems Engineering*, pp. 84-100, Springer (2015)
22. Hastie, R., Dawes, R. M.: *Rational choice in an uncertain world: The psychology of judgment and decision making*, 2nd Edition, Sage Publishing (2010)
23. Betsch, T., Haberstroh, S., Hohle, C.: Explaining routinized decision making: A review of theories and models, *Theory & Psychology*, 12(4), pp. 453-488 (2002)
24. Eid-Sabbagh, R., Dijkman, R., Weske, M.: Business process architecture: use and correctness. In *Proc. 10th International Conference on Business Process Management (BPM'12)*, pp. 65-81, Springer (2012)
25. Vernon, D.: *Artificial cognitive systems: A primer*, MIT Press (2014)
26. Hull, R., & Nezhad, H. M.: Rethinking BPM in a cognitive world: Transforming how we learn and perform business processes, In *International Conference on Business Process Management*, pp. 3-19, Springer (2016)
27. Dix, A. *Human-computer interaction*. In *Encyclopedia of database systems*, pp. 1327-1331, Springer (2009)
28. Shedroff, N.: *Information interaction design: A unified field theory of design*. *Information design*, pp. 267-292 (1999)
29. Bresciani, P., Perini, A., Giorgini, P., Giunchiglia, F., & Mylopoulos, J.: Tropos: An agent-oriented software development methodology. *Autonomous Agents and Multi-Agent Systems*, 8(3), pp. 203-236 (2004)