Designing User Engagement for Cognitively-Enhanced Processes

Alexei Lapouchnian
University of Toronto
Canada
alexei@cs.toronto.edu

Zia Babar
University of Toronto
Canada
zia.babar@mail.utoronto.ca

Eric Yu
University of Toronto
Canada
eric.yu@utoronto.ca

ABSTRACT

A natural way to ease the introduction of cognitive computing capabilities into a user organization is through already well-established applications such as business process management (BPM) systems. Cognitive capabilities can enhance a business process by offering analytics-based recommendations on decisions and increasingly sophisticated automation through machine learning. Yet the organizational adoption of such advanced capabilities is not straightforward. Unlike conventional IT systems whose functionalities and correct operation are more transparent, user acceptance of advice and recommendations from an automated system requires development of trust over time. Additional supporting processes may emerge and evolve over a period of time to monitor, evaluate, adjust, or modify the cognitively-enhanced business process so as to enable personnel to adapt to the enhanced capabilities. In this paper, we propose that a systematic model-based approach can ease the transition to cognitive business operations. The use of suitable modeling techniques can facilitate the uncovering and analysis of obstacles to adoption, and guide the systematic search for viable modes of interaction and cooperation between human user and cognitive advisor.

KEYWORDS

Cognitive Computing, Cognitive Business Operations, Business Process Management, User Engagement, Requirements

ACM Reference format:

A. Lapouchnian, Z. Babar, and E. Yu. 2017. Designing User Engagement for Cognitively-Enhanced Processes. In *Center for Advanced Studies Conference, Markham, Ontario, Canada, November 2017 (CASCON'17)*, 7 pages.

1 INTRODUCTION

Introduction of cognitive systems into enterprises are resulting in greater operational efficiencies through better decision making and ongoing cycles of learning and improvement. However, the adoption of cognitive capabilities and their incorporation into well-established business processes can be challenging, requiring specialized skills (e.g., expertise in data science and machine learning) and extensive user training, while touching on multiple layers and aspects in the enterprise (systems, processes, organizational behaviour etc.), thus incurring significant costs and time investments.

Cognitive systems are deemed to have a number of critical characteristics, which help distinguish them from other enterprise information systems. These include being able to function with a degree of autonomy, demonstrate continuous learning behaviour, perceive and anticipate events in the surrounding environment, select among multiple alternatives, and take appropriate actions in response. These systems adjust over time based on the current context, past performance, etc. and continuously learn using sophisticated algorithms while producing recommendations for human decision makers. Recent years have seen the emergence of cognitive services platforms (for example, IBM Watson [1]) which provide the cognitive computing capabilities for enhancing decision making in business processes (BPs).

Introducing cognitive systems into existing enterprise business processes requires a rethink of their design, implementation and execution. A business process will evolve as it takes advantage of increasingly sophisticated cognitive capabilities, but could also evolve for other unrelated reasons [2]. When a cognitive advisor system is introduced, even if it appears to minimally alter the business process into which it is inserted, there can be many adjustments and adaptations in the surrounding processes, occurring at multiple levels and over different time frames. These processes may be strategic or operational in nature, have different execution frequencies, pertain to planning or designing activities, etc. These surrounding processes too would evolve and be executed differently in response to changes in cognitive capabilities, among other things. Thus, the direct and indirect effect on multiple processes by cognitive systems would need to be considered.

As human users continue to be responsible for some portion of the business process, issues of transparency, trust, autonomy, speed and cost of decisions, etc. would need to be considered when analyzing how users engage with cognitive systems [1][4]. Learning and adaptation would happen both on the human side and on the system side. Users' attitudes towards the system, as well as towards automation in general, will change according to the cumulative experiences with the system and the quality of the output it produces. These attitudes translate into different desired ways of interacting with cognitive systems. Users interacting with the system will need to make their own judgements about cognitive advisors' recommendations. Individuals who have a stake in the decision at various operational and managerial levels of responsibility and accountability may develop strate-

CASCON 2017, November 2017, Markham, Ontario, Canada

DOI: 10.1145/123_4

gies to take advantage of the strengths of the system and to cope with its limitations and weaknesses.

In this position paper, we consider new user engagement (UE) experiences in cognitively-enhanced business processes, including possible changes to and the creation of new surrounding processes, to facilitate user and system engagements while:

- ensuring that the adoption of such systems is done in a manner of least disruption and highest efficiency;
- complying with an organization's internal requirements and constraints as well as those from the business domain;
- factoring in predictable changes to allocation of work among human users and cognitive systems;
- focusing on system performance, users' trust in systems' recommendations, and other important characteristics.

2 A RANGE OF USER ENGAGEMENTS FOR DECISION PROCESSES

Given the possible nature and forms of user engagement, and how they can be assembled to suit a diverse range of enterprise requirements, a systematic approach is required to visually depict, evaluate and reason about the possible space of process configurations that covers *User Engagement Actions* (UEAs), which represent the elementary interactions that comprise user engagement with cognitive systems. Some are executed by human users, while others by (possibly automated) assistants. The set of these primitives that can potentially be selected depends on the nature of the task that humans and automated 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 assistant (A) can include the tasks listed below:

- **Communicate case data:** H communicates the details of the case/BP or decision instance to A (e.g., the relevant information about a loan application). Alternatively, A can obtain the relevant data themselves.
- **Communicate decision parameters:** H communicates the desired parameters for the decision to A, including the criteria for making it, the desired confidence levels, etc.
- **Present Recommendation:** A presents recommended decision option to H (e.g., about whether to approve or reject a particular loan). For decisions with more than two options, a ranked list of recommendations 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 justifications of the recommendation are presented by A to H. E.g., A can explain the decision to reject the loan by pointing to particular characteristics of the applicant's financial situation.
- **Present Decision:** A presents a previously made decision to H (or to a specially designated auditor) for auditing. Variants include the presentation of batches of decisions.
- **Audit Decision:** H (or a designated auditor) audits the decision previously made by A. 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 automated cognitive assistants with 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 human decision makers.

How a cognitive system is received at an enterprise can be very specific to the situation at that organization. Factors affecting adoption success may include the knowledge and skills of involved personnel, aptitude for understanding the capabilities of the cognitive system and its limitations, willingness to learn and adapt, attitude towards automation, labour relations, reward structures, etc. Various human roles in an organization need to discover how they can best work with the cognitive system as there may be many unknowns. Both sides (the user organization and the cognitive solution provider) will need to adjust and eventually converge to a workable state. But even this state will continue to evolve as the cognitive system gets better through machine learning, or gets new features, and on the user organization side, as the personnel gain experience or learns new skills. Changing business contexts and organizational requirements may influence UE as well.

Towards this end, we describe some general examples of user engagement modes (UEMs) – ways users can engage with cognitive systems – in the following sub-sections which allude to a trajectory of increasing cognitive sophistication in the enterprise. These cases are not meant to be prescriptive or definitive, but rather indicate the types of possible engagement configurations that can exist, and highlight some important differences among them. We use these examples to lay the foundation for a systematic modeling approach that helps with a search for workable UE arrangements and enables reasoning about why one engagement approach works but another does not.

Example 1: Decision Processes with no Cognitive Systems Support

As an initial starting point, we take the situation where no cognitive capability currently exists in the enterprise and where human decision making in a given BP is facilitated by non-cognitive enterprise systems. These systems allow humans to make informed business decisions as per corporate policies and rules by pre-processing the data and bringing it to a state where humans can use it to make appropriate decisions. Common examples of such enterprise systems are Decision Support Systems, Business Intelligence solutions, and ERP systems.

Here, business processes are relatively simple and non-changing, and are usually designed at systems implementation time. The processes can be gradually improved by automating key process activities however the nature, execution frequency, and order of the activities within the BPs themselves do not change. Incoming information and data would be processed based on programmatically-defined business rules to enable human decision-making capabilities. There is limited dependency on systems and it is mostly focused on resources and data preparation or on performing simple tasks while providing an output. The decisions within the business processes that are made by humans, while being informed, are based on experience, gut-feel, with the decision-making process taking into account various contextual and situational factors. The decisions themselves may not be uniform and consistent and may vary significantly amongst various users and decision makers. The systems do not perform decision making themselves nor do they learn over time.

Example 2: Decision Processes with Advisory Cognitive Systems Support

The previously described business process can be enhanced by including enterprise cognitive systems to help the human decision maker by providing recommendations and advice at key junctures in the process execution cycle.

Introducing these systems in the business process would result in two significant changes. Firstly, new UEAs (particularly pertaining to recommendations) involving the human user and the cognitive system would be needed: human users still make decisions, with system-produced recommendations presented to them as advices for approval or rejection. Secondly, additional and surrounding processes are present that pertain to the planning and designing of key process activities that influence the execution and usage side of the cognitive system itself. The way these surrounding processes are executed can be different from the primary business process. For example, the impact assessment and system fine-tuning activities can be done offline and not be a part of the process to produce a recommendation. Thus, they will have a lower execution frequency than the primary process, with their output being (re)used in multiple primary BP instances.

The introduction of these cognitive systems into the business process is also accompanied by changes in how non-functional objectives (represented as goals) are met, at both enterprise and process participant (actor) level. E.g., cognitive systems may provide recommendations that increase overall business process efficiency or accuracy, which are enterprise-level goals. However, actor goals, such as trust in cognitive systems, would also need to be considered since their high satisfaction levels drive the adoption of such cognitive systems.

Example 3: Decision Processes with Autonomous Cognitive Systems Support

A further transition to the greater adoption of enterprise cognitive systems happens when recommendations give way to decisions being made by the system itself without the assistance or involvement of the human users.

The shift towards minimizing human involvement in making decisions within the primary BP drives the overall design and configuration of business process architecture. Processes that produce specifications for how the various software tools and artifacts are to be utilized in the BP emerge, along with processes that design these tools for future use. Greater use is made of the available context and various means and forms of context capture (e.g., sight, voice, and data inputs). There is an emergence of additional supporting processes (such as monitoring, evaluation and auditing, adjusting, redesigning) for the primary cognitively-enhanced BP to work. Most process activities run at machine speed as the involvement of human users in BP execution is limited. Feedback and feed-forward paths are now visibly present in the BPA with the results from previous recommendations fed back to improve future ones.

Decisions are only executed if there is a high degree of confidence in their correctness. They are expected to be (in terms of accuracy) equal to, or better than, the decisions made by humans. This results in a substantial change in relationships between humans and cognitive systems. Involvement of users in operational and transactional processes is significantly reduced. These processes are directly managed and executed by systems. Humans are in more governance-related roles now, tasked with reviewing the overall decision making and the surrounding BPs. Such a change manifests itself in the form of changed actor goals and dependency relationships on the cognitive systems.

Table 1: Example Modes of User Engagement with Cognitive Systems in an Enterprise

Example 1: No Cognitive Systems Capabilities (Unassisted)	Example 2: Advisory Cognitive Systems Capabilities (Assisted)	Example 3: Autonomous Cognitive Systems Capabilities
Decision making is limited to programmatically defined information systems with no true autonomous behaviour.	Decision making enabled by cognitive systems that are not fully autonomous but assistive in nature.	Decision making enabled by Cognitive systems that are fully autonomous and able to make independent decisions and solve problems on their own.
Relatively straight-forward processes with limited number of process design choices and no cognitive UEAs.	Introduction of UEAs and processes to the primary business process for enable cognitive-based decision making for users.	Introduction of UEAs and processes for autonomous decision making at machine-scale requiring changes in process execution frequency.
Systems are programmatically defined to execute certain instructions with minimum to no continuous learning present.	Systems deemed to be reliable to inform decision making, but still require human oversight to provide additional context.	Systems are deemed to be sufficiently reliable to guide and inform decision making.
Simplistic nature of systems ensures robust design with predictable behaviour while conforming to enterprise regulations.	Low trust and confidence in the decision-making capabilities of cognitive systems. Reliability of decision-making not deemed to be at par with humans.	High degree of trust and confidence in cognitive systems. Decisions made are at minimum considered to be as good as those of humans.
No concept of trust, reliability or confidence as no cognitive-based decision making by the systems.	Human users still involved in operational decision making but have dependencies introduced on cognitive systems.	Humans mostly involved in governance related activities rather than operational decision making.

Organizations and individual decision makers will strive for user engagements that reflect their changing enterprise requirements, business domain constrains, the level of trust (of both human decision makers and organizations) in analytics in general and in (cognitive) assistants they're employing in particular, and the quality and availability of the relevant data and contexts. This leads to different *combinations and configurations of UEAs* selected at different times. These configurations 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, for many UEAs, there exist options about when, how frequently, etc. to execute them. E.g., decision parameters can be selected once and for all, 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 the previous ones, etc. It is apparent that the navigation of the possible space of user engagement configurations, while adhering to various enterprise functional and non-functional requirements, and balancing them against actor motives and objectives is hard.

In the remainder of the paper, we propose a systematic approach based on a modeling framework that provides notations and accompanying methods to represent, analyze, and ultimately select the appropriate means of user engagement.

3 TOWARDS A FRAMEWORK FOR DESIGNING COGNITIVELY-ENHANCED PROCESSES

3.1 Modeling Cognitively-Enhanced Processes

Overall, we need to be able to characterize the space of alternative user engagement configurations reflecting the whole spectrum of UEAs, their potential combinations, frequency and scope of their execution, and context, among other things. E.g., the Explain Recommendation UEA (mentioned in Section 2) can be omitted in case of a high existing level of trust in the advisor while audits can be added due to industry regulations. Standard process modeling notations are not well equipped to represent these options. To analyze this 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.

To address the above challenges, here we extend and apply a previously introduced conceptual modeling framework *hiBPM* (for *higher-order* BPM) [5][6][7][8] for modeling user engagement processes together with organizational and automated processes that surround them, as well as their interrelationships. 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 *business process architectures* (BPAs) of organizations. Fig. 1 presents a simplified BPA for a loan approval scenario, focusing on the domain-specific BPs (Loan Approval and Repayment) in its centre, cognitive systems-specific processes (e.g., Analytical Model Creation) at the bottom, and BPs managing user engagement (e.g., UEM Selection) at the top of the figure.

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 cognitive advisors. For a loan approval scenario (Figs. 1 and 2), 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. To represent this, in hiBPM, PEs can be placed in various *process stages*, which are characterized by their execution frequencies. E.g., a decision maker (a loan officer) can supply the desired decision parameters to a cognitive advisor 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 change with 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 for some) loan approval instances (note the 1:N annotation on the outgoing flow, which indicates the *recurrence* relationship among stages – that is, for each execution of the first stage, multiple executions of the second one are possible). This is appropriate if those parameters do not depend on the characteristics of each loan approval case. Also note that stages are annotated with actors responsible for their execution and feedback paths are also indicated to capture, e.g., that users' approval/rejection of recommendations feeds back into 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 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 the UEA to be executed for the loan approval process at any given time. Design/Use models a stage that produces artifacts/tools to be (re)used by one or more subsequent stages. E.g., in Fig. 1, Analytical Model is being used in the loan approval BP to help with decision making.

hiBPM uses goal models [9][10] to systematically represent and analyze BPA configuration options and trade-offs among competing quality objectives, such as flexibility, cost, performance, etc. to help design BPAs for particular organizations.

3.2 Designing and Managing User Engagement

In this section, we illustrate the proposed modeling framework using examples from the loan approval domain and outline a methodology for identifying and managing user engagements.

Given a particular organization, a cognitive system, and a problem (i.e., a decision) at hand, how do we design and manage user engagements? 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 manually approve or audit all recommendations made by cognitive systems (as illustrated by all 3 user engagement configurations shown in Figs. 1 and 2).

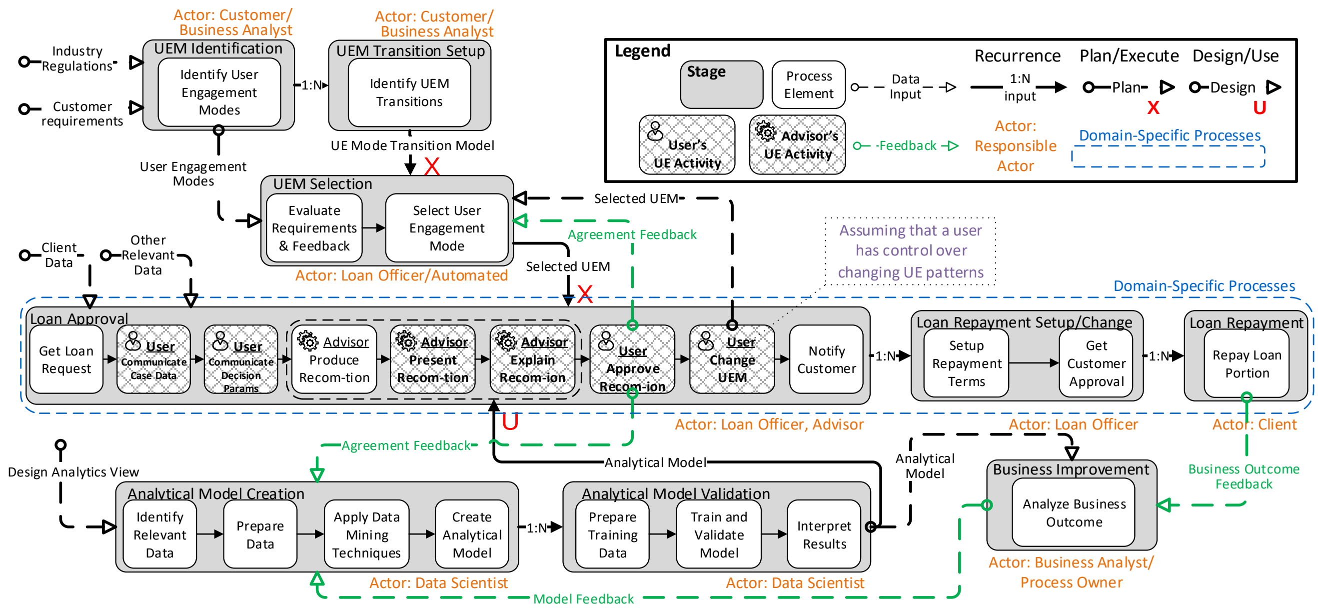

Figure 1*:* **A hiBPM model showing a BP architecture for a cognitively-enhanced loan approval process.**

As already stressed, in most real-life situations, UE needs to be evolving to accommodate changing user/system capabilities and user attitudes (with trust being the prominent one), requirements, and contexts. Thus, UE cannot be static and needs to be managed concurrently with domain-specific BPs as well as cognitive systems and UE management processes, which necessitates focusing on process architectures rather than individual BPs.

When designing UE, we need to explore the large space of options defined by various combinations of UEAs and influenced by requirements/constraints as well as business context. Designing and evolving UE at the level of UEAs results in a fine-grained and flexible approach that supports changing user engagement by adding/removing particular activities as well as repositioning them within the BPA. A goal-driven approach inspired by [9] is envisioned, where goal models are used to capture functional goals, focusing on alternative ways of attaining them, with quality requirements playing the role of criteria for selecting among the options. These goal models are linked to process actions, each of which achieves a certain functional goal. Various configurations of these actions are selected using goal reasoning algorithms based on functional and non-functional requirements (and possibly on context). The benefits of this approach is that goal models are amenable to human analysis and thus offer transparency and predictability.

Alternative approaches based on declarative process modeling [11] or AI-type planning can be employed as well. These powerful approaches would enable dynamic identification of new UE configurations based on changing goals/contexts, which matches cognitive systems' abilities to be adaptive and self-learning and to support (with the appropriate feedback) the handling of previously unidentified conditions, unpredictable shifts in data patterns, etc. On the other hand, such dynamism makes it difficult to predict the evolution of UE and introduces additional complexity.

In many situations and for many organizations, the UEA-based approach's flexibility may not be needed. Coupled with its complexity and lack of transparency/predictability, this necessitates a simpler approach for well-known types of decisions and for organizations that require more transparency and predictability when dealing with cognitive systems. 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. This set of UE patterns captures the whole space of UE configurations as finer-grained adjustments (at the level of UEAs) are not possible. This reduces the space of options for UE (simplifying its management at runtime) and provides more transparency, thus favouring organizations that may be cautious when it comes to automation and cognitive technology.

Here, UE patterns are designed with the help of goal-driven analysis techniques available in the hiBPM approach. Given high-level quality objectives, some of which have already been identified in Table 1, appropriate UE patterns are identified for particular types of decisions (e.g., an approve/reject type vs. that with the higher number of possible choices). These patterns are available upfront and can be analyzed. They can be reused by many organizations for handling similar problems. To further customize user engagement, the set of UE patterns can be pruned based on requirements and constraints. E.g., in some domains and/or organizations, audit of decisions may be mandatory and only patterns containing the audit-related UEAs will be included. These steps of the process are represented by the stage UEM Identification in Fig. 1.

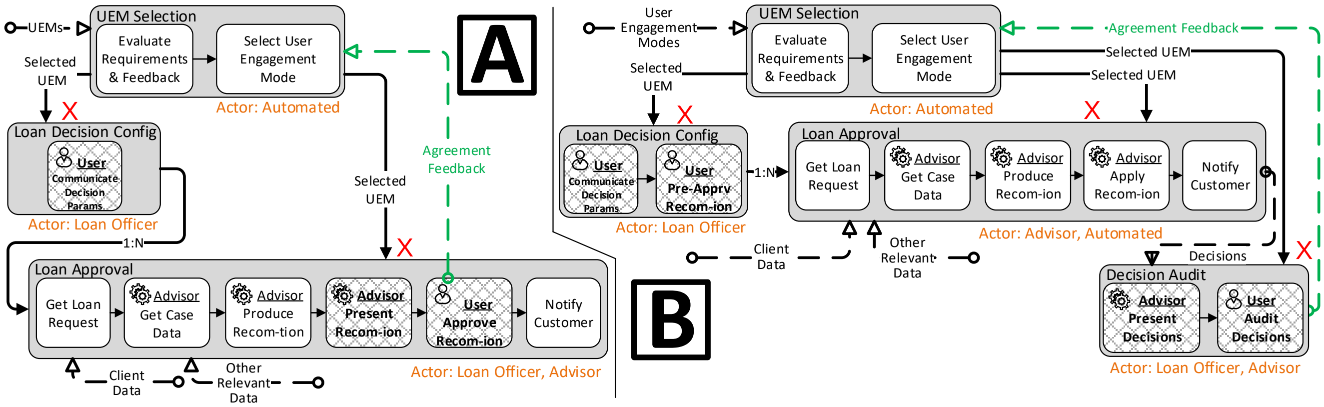

Figure 2: **hiBPM models illustrating alternative user engagement patterns for the loan approval BP: A) preset decision parameters and no explanations of advisor's recommendations; B) automated decision making with manual audits.**

Possible UE patterns for decisions can include the following:

- **P1: Supervised learning**. Decisions are made by a human expert while a cognitive advisor (CA) monitors his work. CA is using case data, context and the decision outcome as the input to a supervised learning algorithm.
- **P2: CA as an Advisor**. Decisions are made by a human expert. CA's recommendations are presented as advice.
- **P3: Human approval of CA's decisions**. Human user approves or rejects CA's recommendations (per decision instance).
- **P4: Human is informed of decisions by CA**. CA makes decisions, while a human is informed of them (per instance).
- **P5: CA's decisions with (batch) human audits**. Human users audit automatically produced decisions once per N number of decision instances, once in a time interval, etc.
- **P6: CA's decisions with human audits on request**. Humans have the opportunity to audit CA's recommendations whenever they wish.
- **P7: CA's decisions with automated self-audits**. No humans are involved by default. Humans can review the self-audits.

The presented patterns can have a number of variations. Fig. 1 shows a fragment of the BPA for handling loans (including approval/rejection and repayment) with the help of a cognitive system. The selected UE pattern is evident from the UEAs (shown as shaded activities) that are present in the Loan Approval process. One can see that the decision maker (here, the loan officer) communicates case data and decision parameters in the Loan Approval stage, which means that this is done for every instance of the loan approval decision, indicating that different decision parameters may be needed for every loan request. After presenting its recommendation, a cognitive system explains it to the user, pointing to the fact that the loan officer is not yet fully trusting the system or is learning on the job. It is a variation of the UE pattern P3.

Next, we focus on the dynamic aspect of user engagement and specify its possible evolution paths by defining transitions between UE patterns (see UEM Transition Setup in Fig. 1) triggered by conditions, such as reaching a certain level of trust in the cognitive advisor (e.g., the UE pattern in Fig. 2A lacks the task Explain Recommendation, possibly due to a higher level of trust), quality of recommendations (lower quality of recent recommendations may warrant a decrease in automation and the shift, say, from autonomous to human-assisted decision making). A change in the context requiring a different UE pattern may be recognized as well – e.g., a significant change in characteristics of decision instances compared to the dataset on which the analytical model was trained. Until enough new data has been gathered and the algorithms have been trained, a more conservative UE pattern may be selected. The transitions will also take into consideration quality requirements and characteristics of organizations and problem domains. Also, as trust plays a crucial role in UE, manual transitions (or the ability to override transitions) will likely be needed. The explicit transitions, which are specified using statecharts, support transparency and predictability of the UE infrastructure. While this approach can only handle predefined conditions and evolution trajectories since they are defined upfront, all steps of the methodology can be re-executed under some conditions. Thus, to update the set of UE patterns and/or to change the transitions among them, the corresponding UEM Identification and UEM Transition Setup will be re-run, producing potentially new output, which will be reused in the subsequent stages in the hiBPM model.

Having described the main aspects of the model in Fig. 1, we briefly outline the different UE patterns shown in Fig. 2A and 2B. They intend to showcase the range of options for specifying UE, both in terms of the set of UEAs and their positioning in the architecture. Fig. 2A is a BPA fragment illustrating a UE pattern where decision parameters are preconfigured (see the Loan Decision Config stage executed once for possibly many loan approval instances), recommendations are no longer explained, and manual UEM change is no longer available. It is a variant of the P2 UE pattern. Fig. 2B illustrates a fully automated loan approval process – note that Pre-Approve Recommendation is executed in the stage prior to Loan Approval, pointing to the fact that CA's recommendations in the loan approval BP are pre-approved, thus highlighting a high level of human trust in the system.

4 DISCUSSION & CONCLUSIONS

In this position paper, we focused on the issues of designing and evolving user engagement with cognitive systems, which needs not only to take into consideration trust, organizational and decision makers' requirements and constraints, contexts, etc., but also evolve with changes in these parameters as well as in the user and system capabilities. We identified UEAs as the building blocks of user engagement and outlined a systematic model-driven approach based on the hiBPM process architecture modeling framework that can represent a space of options for user engagement through multiple levels of process modeling, higher-level relationships among BPs, and other advanced features and analyze them using requirements-driven techniques focusing on reasoning about objectives, alternative solutions, and trade-offs.

This approach is a starting point for developing a comprehensive method aimed at simplifying enterprise adoption of advanced cognitive systems by enabling the fine-tuning of user engagement with such systems based on organizational, social, and technical needs and constraints and supporting the ongoing evolution of such engagements driven by changes in all of these dimensions. There are many research and practical challenges. This challenging domain requires complex, multi-level modeling with powerful, yet tractable analysis methods. While we are building on the foundation of goal and process modeling, where many such analysis methods are available, many features of the proposed approach are novel and need to be further studied and validated. Practicality and usefulness of the proposed approach is being evaluated with an industrial partner. At present, we are focusing on a limited, but diverse set of decisions types that are typical in the enterprise context and that are increasingly being handled with the aid of advanced analytics. They provide the platform for further development and validation of our approach.

REFERENCES

[1] J. Kelly III, and S. Hamm. 2013. Smart Machines: IBM's Watson and the Era of Cognitive Computing. Columbia University Press.

[2] R. Hull and H. Nezhad. 2016. Rethinking BPM in a cognitive world: Transforming how we learn and perform business processes. In International Conference on Business Process Management, BPM'16.

[3] B. M. Muir. 1987. Trust between humans and machines, and the design of decision aids. International Journal of Man-Machine Studies, 27(5-6), 527-539.

[4] J. D. Lee, and K. A. See, 2004. Trust in Automation: Designing for Appropriate Reliance. Human Factors 46(1), 50–80.

[5] A. Lapouchnian, E. Yu, A. Sturm. 2015. Re-Designing Process Architectures: Towards a Framework of Design Dimensions. In RCIS 2015.

[6] A. Lapouchnian, E. Yu, A. Sturm. 2015. Design Dimensions for Business Process Architecture. In Proc. 34th International Conference on Conceptual Modeling (ER 2015).S

[7] A. Lapouchnian, and E. Yu. 2014. Exploiting emergent technologies to create systems that meet shifting expectations. In Proceedings of 24th Annual International Conference on Computer Science and Software Engineering, pp. 371-374, IBM Corp.

[8] Z. Babar, A. Lapouchnian, and E. Yu. 2015. Modeling DevOps Deployment Choices using Process Architecture Design Dimensions. In The Practice of Enterprise Modeling, Springer Berlin Heidelberg.

[9] A. Lapouchnian, Y. Yu, and J. Mylopoulos. 2007. Requirements-driven design and configuration management of business processes. In International Conference on Business Process Management BPM'07, pp. 246–261. Springer, Berlin Heidelberg.

[10] A. Lapouchnian, Y. Yu, S. Liaskos, and J. Mylopoulos. 2006. Requirements-Driven Design of Autonomic Application Software. CASCON'06.

[11] G. De Giacomo, M. Dumas, F. Maggi, M. Montali. 2015. Declarative Process Modeling in BPMN. In Proceedings of International Conference on Advanced Information Systems Engineering CAiSE'15,