

Rethinking BPM in a Cognitive World: Transforming How We Learn and Perform Business Processes

Richard Hull¹(✉) and Hamid R. Motahari Nezhad²

¹ IBM T.J. Watson Research Center, New York, USA
hull@us.ibm.com

² IBM Almaden Research Center, San Jose, USA
motahari@us.ibm.com

Abstract. If we are to believe the technology hype cycle, we are entering a new era of Cognitive Computing, enabled by advances in natural language processing, machine learning, and more broadly artificial intelligence. These advances, combined with evolutionary progress in areas such as knowledge representation, automated planning, user experience technologies, software-as-a-service and crowdsourcing, have the potential to transform many industries. In this paper, we discuss transformations of BPM that advances in the Cognitive Computing will bring. We focus on three of the most significant aspects of this transformation, namely: (a) Cognitive Computing will enable “knowledge acquisition at scale”, which will lead to a transformation in Knowledge-intensive Processes (KiP’s); (b) We envision a new process meta-model will emerge that is centered around a “Plan-Act-Learn” cycle; and (c) Cognitive Computing can enable learning about processes from implicit descriptions (at both design- and run-time), opening opportunities for new levels of automation and business process support, for both traditional business processes and KiP’s. We use the term *cognitive BPM* to refer to a new BPM paradigm encompassing all aspects of BPM that are impacted and enabled by Cognitive Computing. We argue that a fundamental understanding of cognitive BPM requires a new research framing of the business process ecosystem. The paper presents a conceptual framework for cognitive BPM, a brief survey of state of the art in emerging areas of Cognitive BPM, and discussion of key directions for further research.

1 Introduction

Business Process Management (BPM) remains a central, foundational element of running organizations today. This paper explores how BPM will be impacted by advances in *Cognitive Computing* [3, 8, 10, 12], an emerging family of technologies that include natural language processing (NLP), machine learning, and the ability of systems to improve through experiential learning. We believe Cognitive Computing, combined with evolutionary advances in other fields including knowledge representation, automated planning, software-as-a-service, user experience technologies, and crowdsourcing, will transform the BPM ecosystem in

fundamental ways. Cognitive Computing will enable “knowledge acquisition at scale” with the help of emerging methods for natural language understanding and machine learning at scale. It will change the nature of Knowledge-intensive Processes (KiP’s) [4], including a shift in their underlying process meta-model. Advances in Cognitive Computing will enable new ways of learning and enacting processes at both design- and run-time. It will open opportunities for new levels of automation and business process support for all types of processes including KiP’s.

Over the last decade BPM research has been expanded in many directions, to support more flexible business process [18], process mining [22], case management applications [13, 19, 23], and social BPM [1]. There has been some work around bringing AI planning into business process space [14, 15] and processing textual information related to processes [21]. However, most of this research has remained close to the traditional BPM framework, with fairly clear separation between process models (or “schemas”) and process instances. Less attention has been given to laying a BPM foundation for bringing the benefits of process automation over unstructured and semi-structured information, and contexts where the separation between process model and process instance is blurred or essentially non-existent. A key theme of this paper is to advance this discussion towards *Cognitive Process Enablement*, that will enable a whole new level of flexibility in business processing while nevertheless enabling traditional levels of auditing, monitoring, reporting, provenance, and also learning from experience. Important work in this direction is provided by recent exploration of KiP’s [4, 6]. A pioneering work is also [16], that applies cognitive techniques to learn processes as they progress, and help to guide and facilitate them along the way.

Another important, but largely unexplored, aspect of BPM is that of learning the business processes that are *implicitly described* in text or other forms rather than explicitly modeled. In [17] these have been termed “descriptive processes”, as opposed to “prescriptive processes”, to highlight the fact that there is no formal process model specification. Instead these processes are described by the process instances themselves, by “digital exhaust” such as communications between parties (e.g., emails, forums), and by purpose-built natural language documents (e.g., processing guidelines, best practices, regulations, and corporate policies). In the context of transaction-intensive processes that are not automated, the emerging capability for *Cognitive Process Learning* can enable a cost-effective approach for mapping the processes to formal process models and then deploying them. Furthermore, we argue that Cognitive Process Learning, together with a Plan-Act-Learn based meta-model, has the potential of enabling the formal specification and automation of many semi-structured and unstructured business processes that today are not formalized in process models. This will rely on advances in NLP and Machine Learning that hold the promise of extracting goals, best practices, actor intentions, commitments, promises, process fragments and the like from implicit descriptions, including those produced during run time.

We use the term *Cognitive BPM*¹, to refer to a new paradigm in BPM which encompass all BPM contexts and aspects of the BPM ecosystem that are impacted and enabled by the application of Cognitive Computing technologies. This paper provides a conceptual framework for Cognitive BPM in general, introduces key modeling abstractions that will be used in Cognitive BPM, discusses cognitive learning of business processes, and how Cognitive Computing will enable a new style of KiP’s. We present the basis for a new process meta-model, called *Plan-Act-Learn*, which can support the full range from structured to unstructured processes in a seamless and systematic manner. The paper describes recent research advances, and identifies key research challenges going forward.

Organizationally, Sect. 2 provides the context for our discussion. Section 3 lays out a framework for Cognitive Computing and BPM, and briefly surveys how Cognitive Computing is already surfacing in the BPM marketplace. Sections 4, 5, 6 discuss, respectively, cognitive Process Model abstractions, Process Learning, and Process Enablement. Section 7 provides a summary of the discussion.

2 The Context

This section describes an overall context in which we explore the impact of Cognitive Computing on BPM. We identify three broad classes of Business Processes, and briefly discuss the emerging trends in Cognitive Computing that are relevant to BPM.

2.1 Three Classes of Business Process

Figure 1 identifies three broad types of Business Process, as follows.

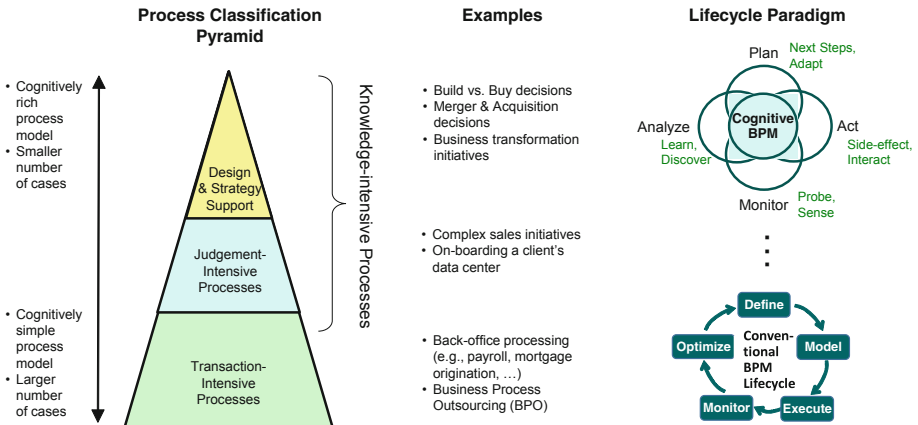


Fig. 1. Workforce pyramid and alternative BPM lifecycles

¹ The term was coined in our earlier work [17].

Transaction-intensive Processing, i.e., processes that are well-defined and are executed many times. Typical examples include week-to-week payroll processing, supply chain management, accounts receivables, etc., within the enterprise, and typical on-line purchasing and self-help in the retail and service industries.

Judgement-intensive Processes, i.e., human-driven operational work requiring many judgements involving complex information, organizations and systems. Examples include processes such as managing new sales relationships, performing project management for large scale IT or other on-boarding engagements, or investigating fraudulent activities. Adaptive Case Management has emerged to help support these kinds of operations, but in practice many of these processes are still *ad hoc*, manual procedures managed using spreadsheets.

Design and Strategy Support Processes, that involve open-ended creative collaborative work, including many decisions based on broad areas of knowledge and analysis. Examples here include the early stages of merger and acquisition explorations, of build vs. buy decisions, and also business model transformation explorations. While these processes may follow a family of best practices, they often have an unstructured and *ad hoc* nature, because of the many possible directions that may need to be explored.

Processes in the second and third categories are often referred to as “Knowledge-intensive Processes (KiP’s)” [4] because of the amount and complexity of knowledge that is used, acquired, and manipulated as they progress.

Actually, most Transaction-Intensive Processes have some characteristics related to KiP’s. In particular, there is a substantial amount of contextual knowledge that is relevant to the effective operation of transactional processes (e.g., the business motivations for the processes, how data flows into and out of the system, and regulations and corporate). More concretely, there are typically *ancillary processes* that are needed to ensure that the routine processes have appropriate data to work with (e.g., in most applications there is exception handling, and in Payroll processing there are ancillary processes for aspects such as incorporating new hires, processing terminated employees, etc.) While these ancillary processes should be routine, a non-trivial percentage of the process instances end up requiring judgements based on an experiential knowledge of the underlying business context, organization policies and overall processing environment.

As suggested in Fig. 1, the three levels of Business Process range from “cognitively simple” processes to “cognitively rich” processes. In today’s world most of the cognitive aspects of these processes are performed by humans, but over time we expect more and more of the cognitive functions to be performed by machines – often in the form of “cognitive agents” – with varying degrees of human oversight. As suggested on the right side of Fig. 1, with increased reliance on Cognitive Computing capabilities and automation we anticipate a shift in the BPM lifecycle paradigm. Transaction-intensive processes will still rely on a formal process model, and the now classical Define-Model-Execute-Monitor-Optimize cycle. In contrast, as suggested in the Introduction, automation of cognitively-rich KiP’s will rely on a new kind of BPM lifecycle, where the separation between process model and process instance is largely blurred or non-

existant. As described in Sect. 4 below, relatively static process models will be replaced by iterative planning; monitoring will occur continuously both within individual instances and across larger families of them; and analysis will be used for learning in a range of areas, including process refinements and also about the application domain and the particular instance at hand.

2.2 Overview of Cognitive Computing

The area of Cognitive Computing is still emerging, and there is no widely accepted definition. Many companies, including IBM [10], Hewlett Packard [8], Deloitte [3], and KPMG [12], are offering visions for what Cognitive Computing is and how it will impact industry and our world. Reference [10] states that “[c]ognitive computing refers to systems that learn at scale, reason with purpose and interact with humans naturally. Rather than being explicitly programmed, they learn and reason from their interactions with us and from their experiences with their environment.” This vision of Cognitive Computing is still coming into reality, but Cognitive Computing techniques are already being considered in the BPM context (see Subsect. 3.2).

For this paper, we are most interested in how key technologies considered within the Cognitive Computing umbrella will enable extending automation of traditional processes, and to support automation of less structured processes, including Judgement-Intensive and Design and Strategy Support ones, that are generally not supported by formal process models today. The technologies we focus on include: natural language understanding (both rules-based and statistical), machine learning (especially in connection with text) knowledge representation and reasoning about knowledge, planning, and experiential learning. These capabilities are becoming available both as traditional in-house functionalities and as cloud-hosted Software-as-a-Service (SaaS). We anticipate that these services will become widely available and relatively inexpensive in the coming years, enabling “always-on Cognitive Computing”.

3 Towards a framework for cognitive BPM

This section provides a conceptual framework for understanding the key ways that Cognitive Computing will transform BPM in the coming years. Also included is a brief survey of emerging cognitive capabilities in the BPM industry.

3.1 Four Pillars

As shown in Fig. 2, the framework rests in part on the new kinds of information that Cognitive Computing can make sense of, including unstructured data, Internet of Things (IoT) data, new kinds of “smart” devices, and etc. Cognitive Computing will leverage and improve on human-to-human collaboration because of new capabilities to ingest and reason about natural language communication,

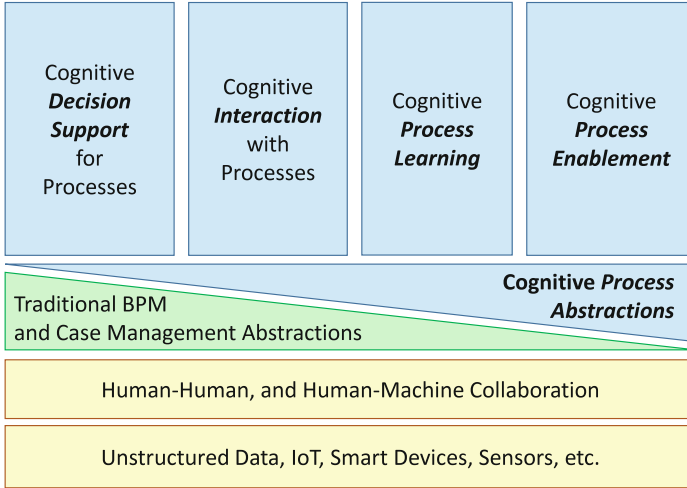


Fig. 2. A conceptual framework for cognitive BPM

and will improve on human-machine collaboration through better communication, and through machines being immensely better at understanding, reasoning about, and carrying out human goals and intentions.

The next layer of Fig. 2 highlights the fact that Cognitive Computing will be applicable in traditional BPM and Case Management contexts, and will also call for and enable new classes of Business Process not supported by process automation today. Cognitive Computing can accelerate the arrival of the next generation in BPM, by enabling the development of a fundamentally new family of process abstractions that will support much richer, more adaptive, more proactive, and more user-friendly styles of process coordination (see Sect. 4).

The four pillars across the top of Fig. 2 correspond to the primary ways that businesses and users will experience the impacts of Cognitive Computing on BPM.

Cognitive Decision Support: Many processes today, from structured to unstructured rely on human effort to make decisions based on deep experience and with reference to large volumes of unstructured data. Cognitive Computing will enable a mammoth increase in the quantity and breadth of such decisions.

Cognitive Interaction: Most human interaction with Business Process systems today is confined to screens, and often relies on constrained sequencing of steps. Advances in multi-modal human-computer interaction and in Cognitive Computing offer a rich opportunity for dramatically improving these interactions by supporting new interaction channels and devices. Importantly, these can enable new styles of collaborative work, e.g., to support collaborative goal formulation and collaborative decision making (with active participation from cognitive agents).

Cognitive Process Learning: Many processes are described only implicitly, in purpose-built documents, digital exhaust, and system logs. Across the full spectrum from structured to unstructured processes, Cognitive Computing can help to capture and codify process specifications, to enable much more automation while still retaining the requisite flexibility.

Cognitive Process Enablement: The separation of process model and process instance as found in classical BPM and Case Management is too confining for cognitively rich KiP's. The vision of Cognitive Process Enablement is to enable a vastly different style of business process support that puts the users back in charge. The underlying process model is highly event-driven, and focused on on-going goal formation, learning of relevant knowledge including constraints, and planning and decision making.

3.2 Cognitive in Today's BPM Marketplace

KPMG's report [11] provides insights into how the industry is incorporating automation into business operations. They use the overall term *Robaotic Process Automation* to refer to three classes of automation. First is *Basic Process Automation* which focuses on the automation of manual tasks through "screen scraping" and application of rules engines on structured data. These capabilities have been available to the industry for several years. Second is *Enhanced Process Automation*, is essentially the Cognitive Decision Support pillar of Fig. 2. These applications are becoming available and are still maturing. The third stage is termed *Autonomic/Cognitive*; this is essentially the Process Learning and Process Enablement pillars of Fig. 2. The report suggests that common application of the Class 3 automation in industry is at least three to five years in the future. The report also suggests that different industries will adopt Robotic Process Automation in different time frames, with IT as the earliest adopter; Sourcing/Procurement, Finance and Accounting, Human Resources, and Supply chain/Logistics in a next wave of adoption, and Real Estate Financial Modeling and Legal after that.

There are early-stage products and offerings in Class 3, the Cognitive Process Learning and Enablement space. For example, the Amelia offering from IPsoft [9] and the Ignio product from Digitate (and offshoot of Tata Consulting) [5] apply machine learning and other Cognitive capabilities to increase automation and optimization of IT services delivery. Also, the Holmes Cognitive System from Wipro is applying cognitive computing capabilities in variety of enterprise and business process management scenarios. The Wipro Holmes web site [24] describes solutions that apply Cognitive Computing in areas such as IT Help Desk, prescription fulfillment, retail purchasing assistance, and compliance.

4 Abstractions for Cognitively-Enabled BPM

As suggested in Sect. 2, a new BPM paradigm is needed to take full advantage of Cognitive Computing. At the same time, Cognitive Computing will help

to enable this new paradigm. This section identifies some of the key building blocks that are anticipated in the new paradigm. The abstractions described here are most relevant to the Judgement-Intensive and Design and Strategy Support processes discussed in Sect. 2. They may also become relevant to the more knowledge-intensive portions of Transaction-Intensive Processes.

The abstractions needed for cognitively-enabled BPM have significant overlap with those discussed in the emerging field of Knowledge-intensive Processes (KiP's). Indeed, several of the key abstractions that we highlight below are present in some form in the extensive survey of KiP requirements [4], in the KiP ontology [6], and the discussions reported in [2]. There are two points of divergence, however. The first is that Cognitive Computing brings the possibility of "knowledge at scale", because cognitive techniques can be used to automatically sift through vast amounts of unstructured data and harvest correspondingly large amounts of knowledge relevant to a process instance. The second is defining abstractions that enable systematic process support for the full spectrum from structured to unstructured processes.

4.1 Key Building Blocks

The key building blocks for cognitively-enabled business processes include the following.

Knowledge, Including Constraints: Knowledge at scale is the fundamentally new element that Cognitive Computing brings to BPM. The possibility of knowledge coming from virtually unlimited sources, and being applied to many different aspects of an on-going process instance, dramatically increases the need for highly flexible processing, that can react to unexpected new information quickly and appropriately. Constraints, rules and policies on the process form an important aspect of the overall knowledge base. These may relate to costs, availability of resources, timing, allowed limits and behavior, and many other factors. The constraints may change over time, and impact decisions and planning.

Goals/Subgoals: A key concept in cognitively-enabled business processes is the notion of *goals*. Initial top-level goals may be specified in advance, and additional goals and sub-goals can be formulated dynamically, based on events in the environment, the current context, new learnings, best practices, and a myriad of other factors.

Agents (Human and Machine): Cognitively-enabled processes will be centered around both human and (automated) cognitive agents. These agents will have varying intentions, roles, and specialties. Collaboration between these agents will be rich and on-going. Communications between the agents may be captured, analyzed, and used in future aspects of a process.

Decisions: Agents will make decisions based on information and knowledge acquired through the process. These decisions may lead to new goals or the achieving of Goals, to Actions (see below), or to Plans (see below).

Actions: An action is an atomic unit of work performed by the agent. The actions in a cognitively-enabled process may side-effect the external environment and/or may lead to new learning.

Plans: A plan (or an action plan) consists of one or more related actions and is used to achieve goals, and may introduce sub-goals. The plans in cognitively-enabled processes may be revised as new information comes in and/or new decisions are made. Plans can be viewed as process model fragments, but their usage is quite different: Plans may be created frequently, modified frequently, and will generally be updated after taking action(s), within the plan, as a result of decisions made.

Events: Most cognitively-enabled processes will be highly event-driven. This is feasible in part because of the highly flexible process model (see below). It is also feasible because automated cognitive agents will be able to rapidly analyze the significance of most incoming events, thereby enabling (human or automated) decisions about whether and how to respond to them. The events may come from the external environment, from the results of information analysis or knowledge acquisition, or from decisions made by agents.

4.2 Plan-Act-Learn Cycle for Cognitively-Enabled Processes

Cognitively-enabled BPM will require a highly flexible process model, that can nevertheless support systematic monitoring and reporting, audits, provenance and explanation, and repeatability. Similar to KiP's [4], processing might range from structured to unstructured, and all points in between.

We conjecture that an appropriate process meta-model for cognitively-enabled BPM will be based on a *Plan-Act-Learn* cycle, as illustrated in Fig. 3. One part of this triad is focused on the planning and deciding portions of the process, which may be carried out by humans, by machines, or by a combination. The plans and decisions may lead to world side-effecting actions (e.g., using resources, transferring funds or goods). The plans and decisions may also lead to learning activities, e.g., the ingestion and analysis of relevant data. This will feed into an ever-expanding knowledge base. Events from the environment, and also environmental reactions to process actions, may contribute to the knowledge base. Finally, the knowledge base will lead to further decisions, goals, and plans.

In terms of more conventional BPM systems, we anticipate that the Plan-Act-Learn cycle can be supported by a kind of “universal” cognitive case management




Fig. 3. The Plan-Act-Learn cycle for next-generation cognitively-enabled BPs. (Dashed lines indicate optional pathways.)

system that is integrated with a knowledge management system. In such a system, case instances might be used to manage plans and plan fragments, and also goals and sub-goals. These case instances could record progress of plans and towards goals, and also information relevant for provenance and audits. Newly arriving events could be processed by multiple of the case instances in parallel, and might also lead to the creation of new case instances. Note that a traditional BPM or Case Management process model can be supported in this framework, in essence by including the process model in the knowledge base, and having the plan-act steps repeatedly refer to that process model when deciding what to do next.

The high variability of Plan-Act-Learn-based process instances will call for a re-thinking of how to support traditional BPM capabilities such as monitoring, auditing, and improvements through analytics on history. What remains fairly constant across the highly variable ways that Plan-Act-Learn-based instances might play out? A possible answer is to shift the focus of monitoring, etc., towards the higher-level goals and also “control points”, that is, business-relevant observable actions (side-effects) that need to occur in all or a substantial percentage of the process instances.

4.3 Recent Initiatives Embodying Key Abstractions

This subsection surveys three recent research systems, each of which embodies a subset of the key abstractions just described.

Citation [20] describes a system that supports the Plan-Act-Learn cycle by using a form of “universal” case management system as advocated above. The system is illustrated by using an example in city government, where advisors to a mayor collaborate to make recommendations to the mayor about various proposals. For any given proposal there may be multiple processing steps and multiple activities to gather information using a variety of (primarily manual) mechanisms. The system is implemented on top of a Guard-Stage-Milestone (GSM) business artifact system (a foundational element for modern Case Management [13]); this provides the advantage of modeling both processing steps and data as first-class citizens.

Reference [7] provides an important early step towards the rich goal-driven style of process management that will be needed for cognitively-enabled BPM. The focus of this paper is to dramatically simplify the job of business analysts when designing a business process that is intended to achieve certain goals. This research models goals as a collection of boolean conditions (in Conjunctive Normal Form), models a library of tasks that include their effects modeled as boolean conditions, and also models a set of constraints on how tasks can be sequenced. The paper describes an algorithm based on successive refinements of the goals into subgoals that yield an ontological match with capabilities of available tasks.

The SmartPM model and prototype system [14] provides on-going contextual awareness and automated planning capabilities. The system uses BPMN to provide a process model, but enables flexible adaptations to a process instance as required by incoming events or new information. The adaptation may involve

the creation of a new plan, which can thereby resolve exceptions that were not designed into the original process. The family of tasks used by the system is defined at design-time, and includes pre- and post-conditions expressed in terms of data objects and attributes, also defined at design time. SmartPM provides a very important demonstration of how on-going planning can be incorporated into a business process system; extensions will be needed to support the incorporation of new data types and tasks during run-time, and to permit richer kinds of data and knowledge in the planning.

5 Towards Cognitive Learning of Business Process

A significant impediment to the automation of business processes today stems from the challenges of learning processes that are described only implicitly, i.e. not explicitly specified and modeled. Cognitive Computing holds the promise of automatically learning many aspects of such processes, thereby substantially reducing the cost of automation. Furthermore, as discussed in Sect. 6, in the case of Judgement-Intensive Processes and Design and Strategy Support Processes, the learning can be interleaved with process enablement to provide recommendations and guidance along the way.

Three dimensions of learning about process are considered: from structured data, from purpose-built documents, and from (unstructured) “digital exhaust”. The structured and unstructured cases are considered here separately, but in practice they will be used in combination.

A variety of structured data sources may be available in connection with an implicitly described process. For the ancillary processes around Transaction-Intensive processes, system logs of the core processes can provide a wealth of information. In particular, techniques from Process Mining [22] may be applied to learn the process models that underly both the core processes, and also some of the ancillary processes. For example, the steps that were taken to insert a new hire into a payroll system might be identified by looking at log entries of tasks that involve data about the new employee at times around when that employee started work. While most process mining work is focused on process, it will be important in cognitive learning to gain a holistic understanding of the overall process, including data manipulation, and constraints on data and processing.

Similar techniques might be applied for Transaction-Intensive and Design and Strategy Support Processes, although there may be less log data available, logs may contain less- or semi-structured information and the available log data may be harder to find and extract.

We turn now to learning from unstructured data. One kind of unstructured data consists of *purpose-built documents*, that is documents that were created specifically to describe aspects of a process. These include actual natural language descriptions of (parts of) a process. These also include documents that give high-level guidelines and/or constraints about a process, including best practice descriptions, corporate policies and government regulations. The other kind of unstructured data is called here *digital exhaust*, and consists of documents and

other digitally available records that are created during execution of process instances. This can include emails between process participants (and hopefully, live conversations between them), entries into process-relevant wikis or forums, the contents of trouble tickets, calendar entries, and also informally structured documents such as spreadsheets, powerpoints, etc.

Techniques are emerging for learning processes from both kinds of unstructured data. For example, [21] uses information extraction techniques to identify tasks and their sequencing from textual process descriptions, to enable comparison the text description with a formal specification (e.g., in BPMN) of the process. The text analytics is performed primarily using the Stanford Parser, which provides a rich family of rules-based capabilities for text analytics. While the approach of this paper assumes availability of a formal specification of a process, it appears that the techniques could be expanded to learn a fair amount about a process model from the text description alone.

An emerging sub-area that is gaining attention is to apply text analytics to government regulations to extract rules and constraints on processes to ensure compliance. A representative work in this area is [25]. This combines both statistical and rules-based approaches to NLP: first, statistically-based techniques are used to classify sentences that are deemed to hold regulatory information; second, these are processed using a rules-based approach, Rules relating to both industry-specific terminologies and an industry-specific ontology are also used.

Techniques are also emerging for learning process from digital exhaust, primarily email. The use of NLP techniques on email is a well-traveled field, with several tools now available that analyze email to provide personal assistance (e.g., Google Now, Microsoft Cortana, Amazon Echo). In contrast, there are only a handful of papers focused on extracting process-relevant aspects from emails, such as tasks and actions. We highlight here the eAssistant system [16], which combines both statistical and linguistic, rule-based techniques key process building blocks. The focus there is on *actionable statements*, which include both promises and requests, and *actions*, which include adding to a “to-do” list, adding to a “follow-up” list, responding to a question, scheduling a meeting, etc., and action lists (i.e., process fragments). In addition to finding these, eAssistant includes an adaptive component, that enables extensibility of feature sets being looked for, and supports online, continuous trainability. (eAssistant can also help to guide processes at runtime; see Sect. 6).

What about the accuracy of the information learned from the above techniques? The use of NLP techniques to learn BPM-relevant information is in its infancy, and so improvements will be on-going. Current techniques are essentially classification-based, and it is typical to measure accuracy in terms of precision (of the objects classified as target what percentage are actually target objects) and recall (what percentage of target objects are classified as being target). These measures generally have an inverse relationship. Speaking broadly, automated NLP techniques typically have precision and recall in the 70% to 95% range. This highlights the importance of enabling humans to understand the outputs of automated learning.

6 Towards Cognitive Process Enablement

Cognitive Process Enablement refers to the ways that Cognitive Computing, taken broadly, can enhance the actual processes that carry out business operations, considered at the level of process modeling. We are focused here on how the processes themselves will be impacted by Cognitive Computing.

Cognitive Computing will impact both classical BPM (and Case Management) processes, and also processes that follow the Plan-Act-Learn meta-model. A central impact, relevant to both settings, is that Cognitive Computing will lead to processing constructs that are at a semantic level higher than those of conventional BPM – including goals, plans, policies/rules, and constraints. These constructs, called here *cognitive BPM constructs* will be both human-understandable and (directly or indirectly) machine-executable. A second key impact, relevant mainly to the Plan-Act-Learn cycle, is that the Cognitive Computing capabilities will be applied repeatedly and in near real time to provide input into the Plan/Decide step.

Although not addressed here, knowledge of implicitly described processes and the perspective provided by the cognitive BPM constructs can be used in other ways, e.g., to verify compliance of a process with regulations, or to streamline modifications of structured processes.

6.1 Classical BPM Setting

Suppose that automated learning as described in Sect. 5 is used on implicit process information in order to build a deployable process model. Because NLP techniques are generally not 100 % accurate, the learned process model will have to be vetted and revised by humans, and will also need to be tested. There is also the question of what kinds of job roles will be needed to vet, adjust, and test the learned process models (and model fragments). Effort should be made to enable Business Analysts to perform all or most of the process modeling adjustments, so that the added cost and delay of bringing in IT specialists and software engineers is minimized. It will be beneficial to present the model using both standard process constructs and cognitive BPM constructs.

These requirements help to envision an overall framework and system for cognitive enablement of the learning and deployment of classical BPM and Case Management process models. A main component is for the learning, including identification and ingestion of implicit process descriptions, logs, and etc. This component will be akin to many Data Science application frameworks, with a rich on-going combination of people and programs to learn and refine the process model, and also to help with evolution over time. Key outputs from the learning, in addition to actual process model constructs, will be *explanations* of the constructs, including, e.g., how they relate to the implicit descriptions. Another main component of the framework will be for testing and refinement. This should be aimed at Business Analysts, and should include Cognitive Computing capabilities to aid with identifying appropriate tests and process model improvements.

6.2 Plan-Act-Learn Setting

In the grand vision of Cognitive Process Enablement, a family of (automated) cognitive agents are used as smart, creative, and pro-active helpers that assist the human in the enactment of processes, and learn the human users' goals for each initiative, and learn context, preferences, and best practices over time. Cognitive agents should understand the capabilities of all resources available, including the agents (both human and automated). Agents are supported in launching new sub-activities, hypothetical explorations, trials, and conventional processes in a free form way. The cognitive agent serves as a proactive project manager, proactively suggesting ideas and approaches, providing expert advice and decision support, analyzing many what-if?scenarios, proactively performing investigations across structured and unstructured data on its own, identifying resources (including personnel), keeping track of schedules, managing and guiding collaborative activities, and recording decisions.

While this vision is some years off, there is a broad base of research to draw from, including in knowledge representation, planning, and multi-agent systems. As one illustration, the eAssistant system [16] mentioned above brings together learning about inflight processes with knowledge representation to provide runtime guidance and support for Judgement-Intensive processes. More broadly, we anticipate that cognitively-enabled processes will be founded on a Plan-Act-Learn cycle, so that they can quickly respond to new events and newly learned knowledge.

We briefly mention several of the key near- and medium-term challenges in the process management space raised by this grand vision. Advances in *goal identification* and *planning* are clearly needed. Massive amounts of application-specific knowledge acquisition creates a challenge in *knowledge representation, prioritization, and explanation*, that is, enabling agents to take advantage of knowledge that is relevant to a decision or task at hand, and ignore knowledge that is irrelevant.

Advances are needed in *process-specific knowledge acquisition*. A specific challenge relates to *event monitoring and triage*, and in particular, tools that enable appropriate response to incoming events, be they from the environment, from agents, or from newly acquired knowledge.

Cognitive Computing holds the potential of automating large swathes of the *Project Management* function, i.e., keeping track of the overall process, deadlines, shifting requirements, etc., and to alert relevant stakeholders to new events, trends, requirements, and delays.

Finally there is the challenge of *trust*. Mechanisms to encourage trust will need to be built into all levels of cognitively-enabled processes; this includes the services that can explain and support testing of essentially all of the automated decisions and plans that are made.

7 Summary and Key Steps for Cognitive BPM Research

This paper has laid out a framework for understanding how Cognitive Computing will impact the practice of BPM over the next several years, and focused primarily on emerging perspectives for cognitive process abstractions, cognitive process learning, and cognitive process enablement. Our findings are relevant to the full spectrum of business processes, from Transaction-Intensive, to Judgement-Intensive, to Design and Strategy Support.

This paper describes many of the research challenges that Cognitive Computing brings to BPM. We conclude by reiterating the most important of the research themes that are most central to the BPM community.

Automatic Learning about Business Process. This learning will be at “design time” (e.g., from purpose-built documents, historical digital exhaust, and system logs), and at “run time” (e.g., from asserted requirements and goals, fresh digital exhaust including human collaborations, and the process instance history). The learning needs to be geared towards process automation and enhancement, including semi-automated Project Management, pro-active knowledge acquisition, and guiding of human activities.

Embracing Flexibility: The Plan-Act-Learn cycle. A new kind of process meta-model is needed for KiP’s in the context of “knowledge acquisition at scale”. We have proposed to base this on the Plan-Act-Learn cycle. But there is a huge distance between this high-level proposal and a robust framework and technology base that can support benefits such as monitoring, provenance, auditability, and ability to refine based on previous performance.

Trust: Explanation, Testing and Manual Adjustment. Trust is essential for automation to be successful. Tools and techniques developed for cognitively-enabled BPM must include confidence-building components at many levels.

Acknowledgements. The authors wish to thank several IBM colleagues for numerous inspirational discussions on the topics presented in this paper, including Currie Boyle, Robert Farrell, Janet Hunter, Matthias Kloppmann Rong Liu, Mike Marin, Manoj Mishra, Nirmal Mukhi, Jae-eun Park, Karthikeyan Ponnalagu, Michael Oland, Eniko Rozsa, Stuart Strolin, and John Vergo. The authors also thank members of the working group [2] on Knowledge-intensive Processes (KiP’s) at the Dagstuhl workshop on “Fresh Approaches to Business Process Modeling” held in April, 2016, where the discussions were very stimulating and informative.

References

1. Brambilla, M., Fraternali, P., Vaca, C.: BPMN and design patterns for engineering social BPM solutions. In: Business Process Management Workshops, pp. 219–230 (2011)
2. Brucker, A., Gal, A., Herwix, A., Hull, R., Mecella, M., Nezhad, H.R.M., Santoro, F.M., Slaats, T., Wong, W.: Knowledge-intensive Processes. Unpublished manuscript created by a working group during the Dagstuhl workshop on “Fresh Approaches to Business Process Modeling”, 8 May 2016

3. Deloitte. Artificial intelligence, real results. <http://www2.deloitte.com/content/dam/Deloitte/global/Documents/About-Deloitte/gx-gr15-artificial-intelligence-computing-capabilities.pdf>. Accessed 01 July 2016
4. Di Ciccio, C., Marrella, A., Russo, A.: Knowledge-intensive processes: characteristics, requirements and analysis of contemporary approaches. *J. Data Semant.* **4**(1), 29–57 (2015)
5. Ignio: Neural automation system for enterprises. <https://www.digitate.com>. Accessed 01 July 2016
6. dos Santos França, J.B., Netto, J.M., do E.S. Carvalho, J., Santoro, F.M., Baião, F.A., Pimentel, M.G.: KIPO: the knowledge-intensive process ontology. *Softw. Syst. Model.* **14**(3), 1127–1157 (2015)
7. Ghose, A.K., Narendra, N.C., Ponnalagu, K., Panda, A., Gohad, A.: Goal-driven business process derivation. In: Kappel, G., Maamar, Z., Motahari-Nezhad, H.R. (eds.) *Service Oriented Computing*. LNCS, vol. 7084, pp. 467–476. Springer, Heidelberg (2011)
8. Hewlett Packard. Augmented intelligence: Helping humans make smarter decisions. <http://www8.hp.com/tw/zh/software-solutions/asset/software-asset-viewer.html?asset=2195447&module=1970414>. Accessed 01 July 2016
9. IPSoft home page. <http://www.ipsoft.com/>. Accessed 01 July 2016
10. Kelley III, J.E.: Computing, cognition, and the future of knowing: how humans and machines are forging a new age of understanding. http://www.research.ibm.com/software/IBMResearch/multimedia/Computing_Cognition_WhitePaper.pdf. Accessed 01 July 2016
11. Robotic Revolution – separating hype from reality, 5 October 2015. <https://home.kpmg.com/xx/en/home/insights/2015/09/separating-hype-from-reality.html>. Accessed 01 July 2016
12. KPMG. Embracing the cognitive era, February 2016. <https://assets.kpmg.com/content/dam/kpmg/pdf/2016/03/embracing-the-cognitive-era.pdf>. Accessed 01 July 2016
13. Marin, M., Hull, R., Vaculín, R.: Data-centric BPM the emerging Case management standard: a short survey. In: *Business Process Management Workshops*, pp. 24–30 (2012)
14. Marrella, A., Mecella, M., Sardiña, S.: SmartPM: an adaptive process management system through situation calculus, indigolog, and classical planning. In: *Proceedings of Conference on Principles of Knowledge Representation and Reasoning KR (2014)*
15. Marrella, A., Mecella, M., Sardiña, S.: An adaptive process management system based on situation calculus, indigolog, and classical planning. In: *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI) (2016, to appear)*
16. Nezhad, H.R.M.: Cognitive assistance at work. In: *AAAI Fall Symposium Series*. AAAI Publications, November 2015
17. Nezhad, H.R.M., Akkiraju, R.: Towards cognitive BPM as the next generation BPM platform for analytics-driven business processes. In: *Business Process Management Workshops*, pp. 158–164 (2014)
18. Reichert, M., Weber, B.: *Enabling Flexibility in Process-Aware Information Systems - Challenges, Methods, Technologies*. Springer, Berlin (2012)
19. Swenson, K. (ed.): *Mastering the Unpredictable: How Adaptive Case Management Will Revolutionize The Way That Knowledge Workers Get Things Done*. Meghan-Kiffer Press, Tampa (2010)
20. Vaculín, R., Hull, R., Vukovic, M., Heath, T., Mills, N., Sun, Y.: Supporting collaborative decision processes. In: *2013 IEEE International Conference on Services Computing, Santa Clara, CA, USA, June 28 - July 3, 2013*, pp. 651–658 (2013)

21. van der Aa, H., Leopold, H., Reijers, H.A.: Detecting inconsistencies between process models and textual descriptions. In: Proceedings of International Conference on Business Process Management (BPM), pp. 90–105 (2015)
22. van der Aalst, W.M.P., et al.: Process mining manifesto. In: Business Process Management Workshops, pp. 169–194 (2011)
23. van der Aalst, W.M.P., Weske, M., Grünbauer, D.: Case handling: a new paradigm for business process support. *Data Knowl. Eng.* **53**(2), 129–162 (2005)
24. Wipro Holmes web page. <http://www.wipro.com/holmes/>. Accessed 01 July 2016
25. Zhou, P., El-Gohary, N.: Ontology-based information extraction from environmental regulations for supporting environmental compliance checking. In: Computing in Civil Engineering, pp. 190–198 (2015)