

Bringing BPM to the Workers: Towards Worker-Centric Management of Business Processes

Vinod Muthusamy, Vatche Isahagian, and Yara Rizk

IBM Research AI, USA

Abstract. Workflow management systems, as originally conceived, tried to impose structure to work, in a way similar to the assembly line. While imposing constraints may be needed because of regulations and compliance, it is important to recognize that much of knowledge work is done in an ad-hoc fashion. Case management is an admission that ad-hoc work exists, and an attempt to adjust to the flexibility of work. Unfortunately, it still requires some explicit steps. With recent changes in the nature of knowledge work, enterprises have increased the adoption of online communication tools such as chat and collaborative documents. Ad-hoc work has become the new norm, and knowledge workers are increasingly utilizing these communication tools to resolve and complete cases. In this work, we argue that business process management solutions need to engage workers in the channels where they already work and outline the challenges of bringing the tools to these environments.

Keywords: business processes · case management · artificial intelligence

1 Introduction

Business process management (BPM) comprises a spectrum of modeling and management approaches and tools, including workflow management and case management. Workflow management systems take a control-flow centric view, based on pre-defined business processes and explicitly defined activities with ordering constraints, where activities may be automated or assigned to human knowledge workers [18]. Case management systems take a more data-centric view, recognizing that knowledge work often needs to be flexible and prescriptive workflows are too restrictive, and thereby allow knowledge workers to define ad-hoc activities and the freedom to decide when to execute the activities. [3]

While case management acknowledges the flexible and ad-hoc nature of knowledge work, implementations of case management solutions still require centralized tracking of the case. Workers are free to carry out the work how they see fit, but need to return to the case management tool to record the results of their work. For example, a nurse may call a patient to check their status the day after getting discharged, but then needs to explicitly login to the case management tool, find the correct case for the patient, select the appropriate

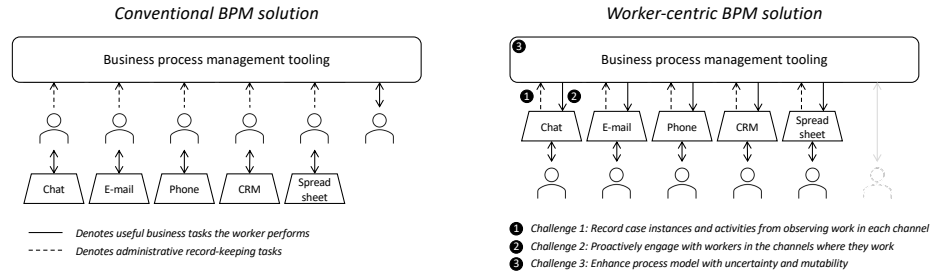


Fig. 1: A conventional BPM solution requires knowledge workers to perform the redundant administrative step of recording their work in the BPM tooling, and interact directly with the BPM tool when necessary. A worker-centric solution eliminates the manual record-keeping step, and offers more familiar channels to interact with the BPM tool when necessary.

activity, and record the patient’s status. As illustrated in Figure 1, the nurse is performing work twice here: once to carry out his task in the channel of his choice (a phone call in this case), and then an administrative step to record the status and results of this task in the BPM tool.

It is not enough to recognize that knowledge work is ad-hoc, unstructured, and takes place across different channels. We need BPM tooling to reach into the channels and modalities where people already work, observe what they do, and record the results in the BPM tool. We call this a *worker-centric BPM solution*. Having separate steps to carry out the actual business task and record-keeping is time consuming and risks the same problems with rigid workflows: work gets done “behind the system’s back” [3]. A worker-centric solution means that workers only concern themselves with performing their business tasks in the appropriate channels, and the BPM tooling takes care of the record-keeping. This is illustrated in Figure 1 where the record-keeping step, denoted by dotted arrows, is no longer a manual step in the worker-centric solution.

We recognize that there are situations where knowledge workers need to legitimately interact directly with the BPM tools. For example, they may need to examine the case history to find who worked on a particular activity so they can followup with questions, or they may want a list of pending overdue work items. A secondary argument in this paper is that these interactions with the BPM tools should also happen in the channels where workers work. For example, a conversational assistant would allow a knowledge worker to query for overdue work items through a “more natural” free-form chat interface. The multi-channel interfaces to the BPM tooling is illustrated in Figure 1 as solid arrows from the BPM tool to the workers’ channels. The figure also depicts our vision that direct interactions with the BPM tools should be deprecated in favor of more channel-specific ones.

A worker-centric BPM solution brings several benefits. First, knowledge workers no longer need to perform the redundant step of recording their work in the

BPM tool after already having performed it in the channel of their choice, such as a voice call. Second, workers that need to interact with the BPM artifacts have more choice in the modality. Third, process owners and workers get a more complete and accurate picture of the end-to-end “as-is” process [1] since directly instrumenting the channels where work is happening and automating the record-keeping means less work getting done “behind the system’s back”. This in turn provides a more accurate dataset for any process analysis to find opportunities for optimization or automation.

Realizing a worker-centric BPM solution will require a wide variety of research innovations in process modeling, tooling, and algorithms. To this end, we present three sets of broad research challenges. (1) Techniques that automate the record-keeping aspects of a business process across multiple channels and modalities are needed. (2) Solutions to allow multi-channel interactions to the BPM tools via a diverse set of more natural interfaces are required. (3) Business process models need to be extended to take into account uncertainty and errors in automated record-keeping. We discuss these challenges in more detail in Section 2.

2 Research challenges

In this section, we outline some of the research challenges, illustrated in Figure 1, that need to be solved to realize a worker-centric BPM solution. We also reference relevant existing techniques and technologies that offer initial solutions to these challenges that the community can build on.

2.1 Infer work by instrumenting across channels

Inferring the process instance. It is generally difficult to correlate work to a process instance [14], and it is particularly challenging in this case where the process instance needs to be derived from the context, which is often unstructured.

For example, when a nurse calls a patient, the context includes the phone number of the patient, the identity of the nurse, and the time the phone call was made. Cross-referencing this information with the patient profile, history of hospital visits, and the department where the nurse works can be used to infer that the call is in relation to a procedure that the patient had for which they were recently discharged. Increased adoption of personalized treatment pathways makes this task even more challenging.

Some existing work that address this problem for a simpler case of structured logs include an Expectation-Maximization approach to estimate a Markov model from an uncorrelated event log [8], and a Simulated Annealing approach that produces a correlated event log from a set of uncorrelated events and a process model [5].

Inferring the activity. In case management tools, workers log work by explicitly selecting from a pre-defined list of activities or by creating and naming a new

ad-hoc activity. There is no such explicit activity when workers perform their tasks in their native channels of work, such as a chat room.

For example, a radiologist may email her analysis of a patient’s x-ray scans. An AI model can process the email subject, body, and attachments to classify the email as completing a pre-defined “Analyze x-ray” activity, as well as extract metadata, such as the radiologist’s name and whether the diagnosis was negative, and record all of this as part of the process instance. AI models to classify documents and extract key-value pairs [10], and classify tasks from chat logs [6] are useful technologies to build on.

A more challenging problem occurs when workers perform tasks that don’t fall into the set of pre-defined activities. For example, the radiologist may forward the x-rays to a senior colleague for help interpreting the scans. It is not clear if this task should be treated as an unnamed sub-activity of the “Analyze x-ray” activity, or a new activity altogether. Perhaps the system should intervene and ask the radiologist to identify this task. Techniques from conversational systems to infer intent and resolve ambiguity are instructive here [17].

2.2 Proactively engage with workers across channels

Eliminating the direct interaction of users with the BPM systems will require these systems to proactively interact with users in their channels to close the loop. This translates to channel-specific alerting capabilities to inform users of any changes in the system, especially those that involve multiple actors that may change the state of the system (including automation bots that perform tasks in the process) [16]. Also, channel-specific query-result presentation allows users to understand the events in the process [13]. Finally, proactive engagement also allows the system to correct any mistakes made by the AI solutions used to address Challenge 1 in Section 2.1 [15]. This could entail resolving ambiguities in the identified activity or process, asking for missing information required by the process but not provided by the user, or simply asking the user to validate work done by automation solutions such as RPAs [9].

2.3 Enhance the process model semantics

Mutable case history. Due to the challenges in Section 2.1, the system may make mistakes in inferring the case instance or activity of some work. There needs to be a facility for these errors to be corrected manually at a later time by a knowledge worker or a case owner.

It is not clear how the provenance of these mis-classifications should be recorded. A process improvement specialist likely only wants to see the corrected case history, while a forensic auditor may be interested in the erroneous history. A knowledge worker may or may not want to see the errors depending on whether there were downstream activities that acted on those errors. This is somewhat related to compensations in workflow management systems [18] but these are more system-level errors.

Another subtlety is that in some cases the system may correct itself as more context becomes available. For example, consider an email thread that the system infers is related to the most recent procedure for a patient. However, a later email may provide additional clues that the thread is actually about a future procedure. In this case, the activity associated with the email thread needs to be removed from one process instance and added to another. Current BPM solutions do not accommodate mutating the history in this way.

The database community has developed solutions on how to update query results as streaming data arrives late or is corrected in the future [4], and some of the programming models and optimizations may guide researchers on how to solve this problem in the BPM context. In robotics, fault detection and correction are essential to successfully deploying robots in the environment, for which self-repairing algorithms have been investigated [11]. Frankly, we think this an unexplored area open for pioneering research in the BPM community.

Probabilistic process models. To address the reality that the system may make mistakes in recording work, we need to support a notion of activities that only probabilistically belong to a case. For example, an AI model may only have 60% confidence that an email thread belongs to a particular case instance. Both extremes of not including the activity associated with that email thread in the case or adding it with 100% certainty are less than ideal.

This will have fundamental changes to how processes are modeled, perhaps adopting probabilistic Petri nets [12], and rethinking all the process analysis methods built on this foundation [18]. As well, modeling notations [7] and the semantics of workflow management systems [2] need to be revised to reflect probabilistic assumptions.

3 A BPM solution that puts workers first

BPM systems have been developed with a top-down mindset, catering first to process owners and administrators that want to impose structure and order on the business processes and knowledge workers. This has resulted in solutions that force workers to perform administrative record-keeping tasks and learn to use unfamiliar tooling outside their comfort zone, and offers a view of the process devoid of any of the real-world uncertainty.

We think it is time for a more worker-centric BPM solution that brings in the latest in AI advancements, including computer vision, speech recognition, and natural language processing, in order to engage knowledge workers in the channels and modalities in which they are already most comfortable and productive. Doing this will also require fundamental extensions to how processes are modeled, including adding notions of uncertainty and mutability, and addressing the ripple effects of such core extensions.

While challenging, we believe that the intersection of BPM and AI can bring about a new form of process management that is more worker friendly, less redundant and more amenable to modernization.

References

1. van der Aalst, W.: "Mine your own business": using process mining to turn big data into real value (2013)
2. van der Aalst, W.M.: Three good reasons for using a petri-net-based workflow management system. In: *Information and Process Integration in Enterprises*, pp. 161–182. Springer (1998)
3. van der Aalst, W.M., Weske, M., Grünbauer, D.: Case handling: a new paradigm for business process support. *Data & Knowledge Engineering* **53**(2), 129–162 (2005)
4. Akidau, T., Bradshaw, R., Chambers, C., Chernyak, S., Fernández-Moctezuma, R.J., Lax, R., McVeety, S., Mills, D., Perry, F., Schmidt, E., et al.: The dataflow model: a practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *VLDB* (2015)
5. Bayomie, D., Di Ciccio, C., La Rosa, M., Mendling, J.: A probabilistic approach to event-case correlation for process mining. In: *International Conference on Conceptual Modeling*. pp. 136–152. Springer (2019)
6. Carenini, G., Murray, G., Ng, R.: Methods for mining and summarizing text conversations. *Synthesis Lectures on Data Management* **3**(3), 1–130 (2011)
7. Decker, G., Dijkman, R., Dumas, M., García-Bañuelos, L.: The business process modeling notation. In: *Modern Business Process Automation*, pp. 347–368. Springer (2010)
8. Ferreira, D.R., Gillblad, D.: Discovering process models from unlabelled event logs. In: *International Conference on Business Process Management*. Springer (2009)
9. Ghose, A., Mahala, G., Pulawski, S., Dam, H.: The future of robotic process automation (rpa). In: *International Conference on Service-Oriented Computing*. pp. 273–280. Springer (2020)
10. Li, J., Wang, H.J., Bai, X.: An intelligent approach to data extraction and task identification for process mining. *Information Systems Frontiers* **17**(6) (2015)
11. Liu, J., Harkin, J., McDaid, L., Halliday, D.M., Tyrrell, A.M., Timmis, J.: Self-repairing mobile robotic car using astrocyte-neuron networks. In: *2016 International Joint Conference on Neural Networks (IJCNN)*. pp. 1379–1386. IEEE (2016)
12. Liu, Y., Miao, H., Zeng, H., Li, Z.: Probabilistic petri net and its logical semantics. In: *2011 Ninth International Conference on Software Engineering Research, Management and Applications*. pp. 73–78. IEEE (2011)
13. Mittal, A., Saha, D., Jain, P., Sen, J., Jammi, M., Sankaranarayanan, K.: Bootstrapping chatbot interfaces to databases. In: *8th ACM IKDD CODS and 26th COMAD*, pp. 47–55 (2021)
14. Muthusamy, V., Slominski, A., Ishakian, V., Khalaf, R., Reason, J., Rozsnyai, S.: Lessons learned using a process mining approach to analyze events from distributed applications. In: *Proceedings of the ACM DEBS* (2016)
15. Rizk, Y., Isahagian, V., Boag, S., Khazaeni, Y., Unuvar, M., Muthusamy, V., Khalaf, R.: A conversational digital assistant for intelligent process automation. In: *International Conference on Business Process Management*. pp. 85–100 (2020)
16. Rizk, Y., Isahagian, V., Unuvar, M., Khazaeni, Y.: A snooze-less user-aware notification system for proactive conversational agents. *arXiv preprint arXiv:2003.02097* (2020)
17. Weld, H., Huang, X., Long, S., Poon, J., Han, S.C.: A survey of joint intent detection and slot-filling models in natural language understanding. *arXiv preprint arXiv:2101.08091* (2021)
18. Zur Muehlen, M., Shapiro, R.: Business process analytics. In: *Handbook on Business Process Management 2*, pp. 243–263. Springer (2015)