

Cognitive Augmentation in Processes

Moshe Chai Barukh¹, Shayan Zamanirad¹, Marcos Baez²
Amin Beheshti³, Boualem Benatallah¹, Fabio Casati⁴
Lina Yao¹, Quan Z. Sheng³, and Francesco Schiliro^{3,5}

¹ University of New South Wales, Sydney, Australia
{mosheb,shayanz,boualem}@cse.unsw.edu.au, lina.yao@unsw.edu.au

² Université Claude Bernard Lyon 1, Lyon, France
marcos-antonio.baez-gonzalez@univ-lyon1.fr

³ Macquarie University, Sydney, Australia
{amin.beheshti,michael.sheng}@mq.edu.au

⁴ ServiceNow, Santa Clara CA, USA
fabio.casati@servicenow.com

⁵ Australian Federal Police, Australia
francesco.schiliro@hdr.mq.edu.au

Abstract. We present a vision for the next generation of process technology based on cognitive augmentation. Starting from current process technology, we show how by augmenting layers of cognitive intelligence to combine advances in machine-automation, crowdsourcing and more importantly adaptation and reasoning, we can advance support for emerging requirements of highly changing environments. We believe the challenges lie in the synergy between human and machine, in understanding how to orchestrate and combine their contributions. This vision paper sets forth a roadmap for future research by introducing a framework for cognitive augmentation, identifying the relevant research and technologies, and discussing its application amidst real-world use cases.

Keywords: Bots · Cognitive Computing · Processes

1 Introduction

Processes are an integral part of everyday life. Often, the most prevalent are those we are least mindful about, yet highly pervasive in everyday tasks (e.g., send an email, schedule a meeting, record notes and gather feedback). Colloquially referred to as “shadow processes”, these snippets of the overall process are typically being performed ad-hoc using a variety of cloud-based software tools, while the end process remains hidden [8]. These are highly unstructured processes. On the other extreme, there is an enormous body of work into structured processes. Formally referred to as Business Process Management (BPM), this technology proved monumental in allowing organizations to embrace workflow automation of tasks. However, the challenge in this approach stems for the inherent presupposition that processes are well-defined; they thus fail to cater for much needed agility in today’s dynamic environments [2].

At present, many have struggled to bridge the gap between highly structured and very unstructured processes – with many solutions ending up closer to either one extreme. We believe the challenges lie in the synergy between human and machine. For some tasks, humans are far more superior than machines, such as in judgement-oriented work. Whereas in other tasks that require consistent iterations, a machine would far outperform human capabilities.

For this reason, we envision the next generation of process capability resembling a humanoid. More broadly termed “cognitive computing”, it should be capable of assisting humans in human tasks, while augmenting machine-level capability. It should be capable of thinking, acting as well as learning autonomously akin to the human mind. Our vision is a world where everyday existing work platforms will converse with end-users via digital-assistance services – thus acting to mediate humans and work tools, and between different tools. Underlying all this, we envision a backend powered by several layers of cognitive intelligence, with data as the common factor connecting these disparate work tools.

In this paper, we present a vision that sets sail into this journey. We see this consisting of three main layers: (1) As the foundation, existing process systems, together with apps, tools and services must continue to be used. However we will rely strongly on the “everything-as-a-service” model, whereby such tools, even including sensors and physical monitors will be programmatically accessible. (2) On top of this will be several layers of “cognitive enablement”. These layers of intelligence will act in hierarchy where higher layers can be composed (and utilize) lower layers. More so this will include crowdsourcing and methods for continuous learning. (3) Finally, we have a layer of “cognitive delivery”, which means a seamless interface for human workers, in the form of bots that offer digital assistance through conversation. Putting it all together, we refer to this idea as “cognitive augmentation”

2 Process Technology Foundations

To project an accurate vision of the future, we must thoroughly understand the past. A fundamental view of a process is the coordination of tasks, data and the communication between tasks and data as well as stakeholders. Beyond this, the remaining technological landscape for processes can be abstracted simply as parameterizations of these three fundamental aspects.

For structured processes, classical business process management systems focused on the process-centric methodology – automating ‘tasks’ with secondary support for other aspects such as data and communication. Other structured processes systems shifted focus to data, known as artifact-centric systems, such as: structured data repositories, document engineering, artifact governance policing (e.g. IBM Governor [13]), and artifact lifecycle management (e.g. Gelee [4]). From this various synergies emerged. Such as between data and tasks, where the notion of “Business Artifacts (BA)” was introduced to assist in describing the data of business processes. Event-Driven BPM similarly shifted focus offering more powerful control of communication and its synergy with process

(sub-)tasks. With an event-driven approach, events produced by the process engine can in turn be prescribed to trigger or influence the execution of another task, and even cross-enterprise business processes.

Many process systems oscillated focus of support between *tasks*, *data* and *communication*, as well as synergies of these – with the goal towards increased flexibility. Ultimately however, it was hard for these approaches to separate between models (or “schemas”) and process instances. Even non process-oriented solutions struggled to agree upon accepted rules- or event- processing language.

On the other end of the spectrum, unstructured process support systems are typically present as Web-based SaaS tools, each targeting a specific type of task/s. Such as, communication and collaboration tools, project and task management tools, artifact management as well as visualization and direct-manipulation tools. This approach offered the much needed flexibility. However, multiple different tools are often needed to meet a typical end-to-end solution, resulting in “shadow processes” that are managed manually and difficult to track.

We should now better understand the goal of cognitive augmentation. Until now the mistaken mentality was an automate “everything” approach, and both structured and unstructured processed were incapable of this. The solution rather lies in a part-human part-machine approach - a “humanoid”; it’s then about the right type of automation being applied. For structured processes this means empowering human workers by automating the pre- and post- processing steps (e.g., translating natural language into low level commands and vice versa). On the other hand, for unstructured processes this means introducing automation by leveraging existing algorithms and APIs to automate both basic and complex micro-tasks.

In practice today, many enterprises have adopted case-management to draw closer to the reality that most processes are neither fully structured nor fully unstructured, and in fact requires both manual control as well as automation. These types of “semi-structured” processes are devised as a set of repeatable process patterns, yet each specific “case” can take upon its own variation. Case management offers interaction channels between people, services and data sources thus empowering open communication, and moreover Web services are being leveraged for enhanced automation opportunities (e.g., semantic tagging of artifacts to better work with the intensity of data). *ProcessBase*[5] is a unique framework that offers a *hybrid processes* approach to combine from structured to very unstructured processes. In the future, cognitive augmentation would enable autonomic process that ultimately thinks and learns like humans, and with this vision we can move into a reality of model-free processes that are self-descriptive rather than prescriptive.

3 Cognitive Process Augmentation

The next generation of tools are not just about integrating artificial intelligence (AI). It is about augmenting (not reinventing) existing tools, services and process systems (from structured to unstructured) with the rich and already mature

advances in data curation [7, 6], machine learning as well as crowdsourcing, and delivering this to end-users as natural and interactive digital assistance.

Figure 1 proposes a three-faceted framework to realize this vision. We start by leveraging *current process technology*, including structured, unstructured and case management. We analyze this rudimentary layer with respect to: data, tasks and communication capabilities. We then identify what enables *cognitive augmentation*, and this depends on utilizing advances in machine-driven automation, human workers in the crowd and most importantly reasoning and adapting. Cognitive processes must iteratively discover, learn and customize based on accumulated knowledge and experience. Finally, to the end-user cognitive processes means delivering a digital administrative assistant (a “humanoid”). It must support natural language interactions resembling the work practices of humans (providing guidance, advice, recommendation, contextualization and problem solving in decision making). The benefits of cognitive process will felt across the range of information systems, providing in-task assistance from email, groupware, workspaces to enterprise social platforms.

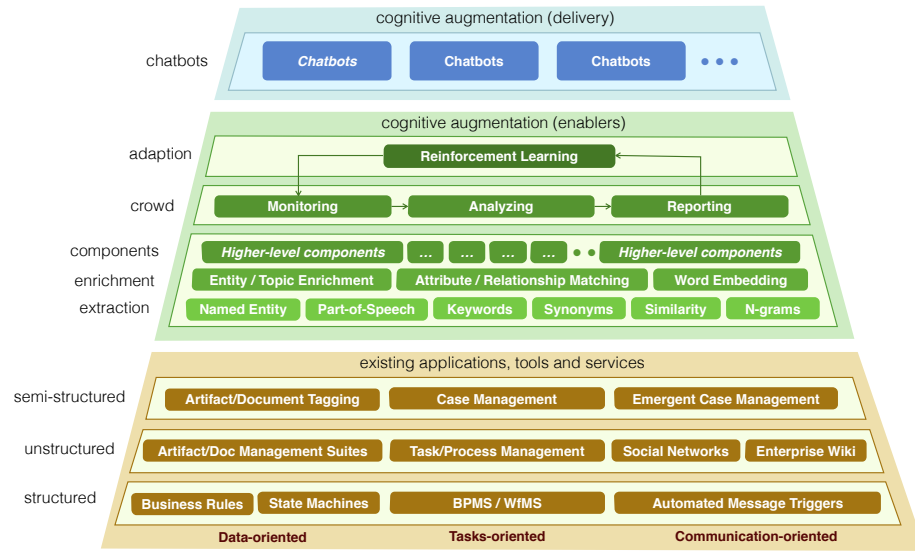


Fig. 1. Framework for Cognitive Augmentation in Processes.

4 Use Cases

Cognitive augmentation in real-world processes would significantly increase the productivity of processes, as well as the ability for enhanced insights and effective decision making. To illustrate this vision, we explore a typical use-case scenario, showing how cognitive capabilities can be enabled and delivered to the process worker. The same would apply to many other real-world scenarios, such as *investigative journalism*, *systematic literature review* or *activity recognition*.

4.1 Law Enforcement Investigations

Modern police investigations are complex projects that can span for years. As shown in Figure 2, investigators collect and manage information, as well as ensure evidence collected is relevant, admissible and sufficient to prove offenses at court beyond reasonable doubt. Evidence may be sourced from “witness statements”, “forensic reports” and “telephone intercepts”. Investigators must not only find content but apply their own cognitive efforts to extract meaning. For example, an investigator may retrieve the passenger manifest for all flights over a given time, and must then search for evidence that their person of interest, with a given passport number, traveled at the time of interest.

The overall process is highly cognitive both with respect to collecting and analyzing information, as well as to inferring interdependencies between data to eventually produce a storyline brief to present in court. Today an enormous amount of relevant data is available, from social media to tracking personal devices (e.g., monitoring a suspect’s location and social interaction can provide vital information to a case). Traditional tools are simply inadequate and thus most cognitive tasks are performed manually; this is no doubt tedious, error prone and highly insufficient. In the recent Bali attacks, investigators revealed several perpetrators were left unprosecuted only due to limited manual processing power.




Fig. 2. Law Enforcement Investigation Process.

Cognitive Enablement. With highly knowledge-intensive processes, it becomes paramount to prepare raw information into contextualized knowledge. Raw data is useless to both humans and machines unless processed in the correct order to derive valuable insight. For example, if we were to classify the topic of Tweets, we would first need to apply natural language extraction (e.g. to extract and identify nouns and verbs), before applying a classification algorithm.

As data accumulates during an investigation, it becomes vital to keep track of relevant events and detecting possible offenses from raw evidence logs. The analysis of such text-based logs involves a great deal of qualitative analysis that can be a lengthy process, and cases can even be cut short leaving criminals unprosecuted. Cognitive support can therefore significantly improve productivity:

1. **Offense Detection.** Typically, at the start of an investigation, an allegation statement is composed (e.g., the extract as shown below):

“Peterson was found to be in possession of 500 grams of Methamphetamine. It is alleged that Johnson may have sold Peterson the controlled substance discovered.”

Ordinarily, manually sifting through legislation, such as the Criminal Code Act 1995 which codifies thousands of criminal offenses, would be a very exhausting task. We would need to find the right offense (and all the offenses) that match a particular allegation, such as: Section 308.1 (“Possessing controlled drugs”) and Section 400.3 (“Dealing in proceeds of crime”).

2. **Event Recognition.** Next, the investigator records all types of evidence, and these logs are later used to prove the elements of an offense. Once again, cognitive support would not only help extract events, but also analyze and attach semantics. Moreover, it could also assist in reconstructing chains of events to simulate how the case developed, the identification of parties involved, understanding of its temporal dynamics, among other aspects. For example, for the sample evidence log show below, we can extract event types such as “phone call”, “bank transaction” or “travel movement”.

“On 23 Feb 2011, Peterson went to Lancaster to meet a person named Johnson in a pub, they watched a football match together until 8.30pm.”
(travel movement)

“Peterson, used his phone to transfer 6 thousands dollars to Johnson on 23/02/11:20:18.” (bank transaction)

Cognitive support can be applied at various layers of granularity. For example, Figure 3 illustrates a potential cognitive stack suitable for this scenario

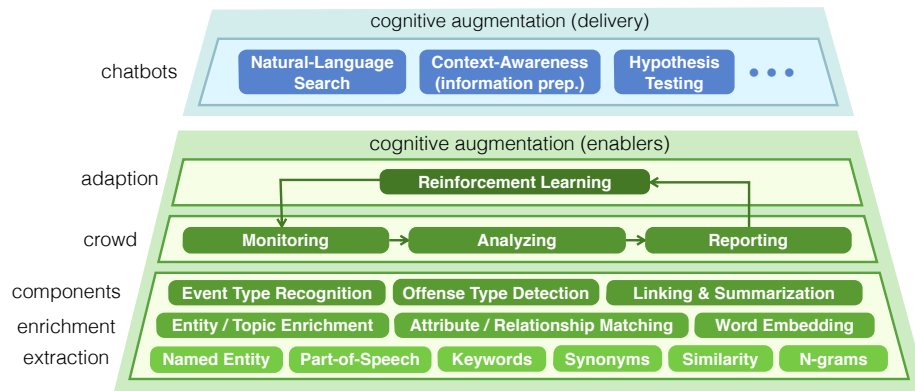


Fig. 3. Cognitive Augmentation Stack for Investigations Process.

– it shows the enablers needed for this type of cognitive support (and ultimately deliver them as end-user digital assistance components). At the fine-grained level, we have various information “extraction” components, such as: named-entity (using lexical analysis); part-of-speech (using synthesis of natural language to identify nouns and verbs); synonyms (using the urban dictionary); and timestamps (using parsers). In the case of event recognition, these rudimentary components assist to lexically deconstruct raw evidence logs.

3. **Linking & Summarization.** We should now appreciate that the underlying objective of the investigative process is to link evidence to elements of the the offense (i.e., we try to prove or disprove an offense based on the evidence). This is where an investigator would spend most of his time, sifting through, in many cases, thousands of pieces of evidence and linking them to possible offense violations. Once again cognitive support for this could be used to filter through key facts (such as the events recognized earlier), and for example, using event pattern/templates correlate such sets of events to indicate whether the elements of a particular offense have been committed. For example, to be charged under Section 400.3 (“Dealing with proceeds of a crime”), intent to carry out the crime must be established. In some cases, it can done by linking several pieces of evidence to reconstruct the picture.

<p>“Section 400.3(1)(b)(ii) the person intends that the money or property will become an instrument of crime”</p> <p>→ witness statement #1: phone call</p> <p>→ witness statement #2: meeting took place</p> <p>→ bank transaction #1: exchange of monies</p>
--

Moreover, the key information obtained here (and throughout the investigation) could be auto-summarized and chronologically compiled into a single evidence brief. Effective summarization are vital to present the case in a simple and organized manner in court.

4. **Action Generation.** During the process, investigators may also need to take certain action. For example, approving a search warrant to obtain missing information, or anything else to finalize the case. Once again, cognitive support in the form of summarization (e.g. over existing evidence) can be used to auto-generate tasks and remind/guide investigators about what actions are required.

Cognitive Delivery (Bots). The second part of cognitive augmentation is delivering to end-users a collaboration model that connects people, tools, processes, and automation into a transparent work environment. As mentioned earlier, the key is to balance between humans and machines. In fact, in most work processes, humans require machines as much as machines require humans. We envision conversational bots will achieve this, where end-users can express in a

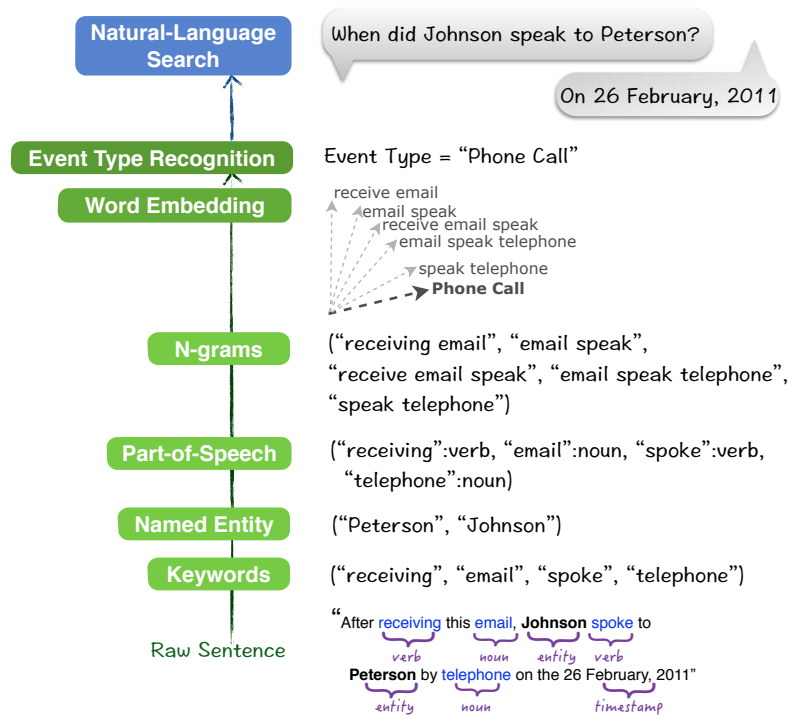


Fig. 4. Illustration of Cognitive depicting the "Natural Language Search" Bot in the Law Enforcement Investigations Use-case Computation.

controlled natural language the tasks they want to perform, or provide the requisite feedback, to interact with underlying cognitive services that drive the overall process towards its goal. The following describes two (due to space limitations) types of digital assistance for this scenario:

- **Natural Language Query.** Investigation data can be made available using controlled natural language queries (e.g., search person of interests, documents, artifacts, organization knowledge, people to ask questions, relationship and hypothesis-based based search, conversations to construct answerable queries). For example, Figure 4 shows how a simple question could be asked in natural language. This capability is powered by a number of techniques such as natural language processing, query intend discovery, entity mention discovery, knowledge graphs and deep learning algorithms to perform entity mentions and relationships based indexing over investigations as well external data.
- **Context Awareness and Proactive Information Preparation.** Proactively providing the right information at the right time is a proven technique to improve productivity and reduce information load. Cognitive services in

this category capture context (e.g., a task an investigator is working on like a line of inquiry, meeting information) and proactively surface relevant information (e.g., availability status, prepare and recommend information that is relevant to perform a task, advice to correct or complete missing information to increase information quality).

4.2 Systematic Literature Reviews

Systematic Literature Reviews (SLRs) aim at analyzing and synthesizing research evidence by following accepted community guidelines, in response to postulated research questions. They are one of the most important forms of publications in science, and are the basis for evidence-based practices and even government policies. As illustrated in Figure 5, the SLR process is typically carried out in phases, such as: (i) the definition of a goal and scope of the review (e.g., “studies on the effect of technology-supported interventions to reduce loneliness”); (ii) the output of which are the identification of relevant papers through a search strategy that stems from the research question; papers may also be annotated adding additional semantics and insights; (iii) the screening of these candidate papers – very often thousands of them – based on specified inclusion and exclusion criteria (e.g., “Filter out papers without loneliness as primary outcome”); and then (iii) the analysis of the selected literature and synthesizing summaries based on the findings, along with the discussion of potential biases.

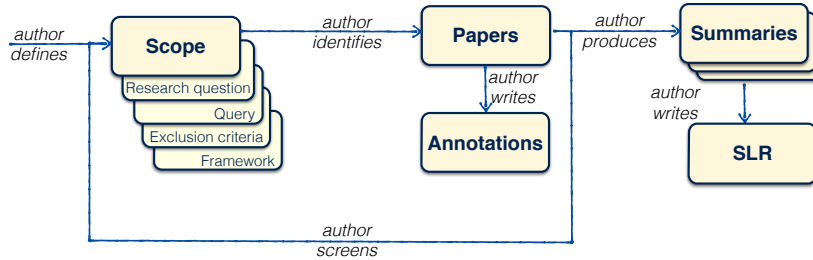


Fig. 5. Systematic Literature Review (SLR) Process.

While extremely valuable and of considerable impact, SLRs are very time-consuming, become rapidly outdated and are not easily maintained. The considerable effort required by researchers combined with the acceleration in the research production pace of the scientific community and the lack of adequate tools to support this process makes carrying out an SLR a very challenging endeavor [11]. Indeed, studies have shown that literature reviews might miss from 30% to 50% of relevant papers at the time of publication [9], either due to compromises to make the process manageable or new articles published.

Cognitive Enablement One of the most labor intensive yet critical phases is the identification of relevant scientific articles [11]. Focusing on this phase alone, we explore how cognitive support can be enabled to deliver a number of assistive components, working together to enable authors in performing a reviews that are unbiased, systematic, inclusive, yet tractable in terms of effort and latency.

This scenario, in particular, exemplifies how cognitive support is enabled using a mesh of automated (i.e., algorithms, services and AI) and crowd-driven techniques. The goal of the crowd in this scenario is both to conduct micro-tasks along the way that require human intuition, as well as providing feedback for ongoing adaptation, such as input to reinforcement algorithms. At the crowd layer, this may be abstracted into the following components that feed into each other but can also be iterated as the authors (and algorithms) gain more insights of the outcomes of each phase: i) search, referring to setting the scope and identifying the relevant papers; ii) annotate, the activity of labeling, filtering out and classifying scientific articles, and iii) synthesize, as the activity of extracting and deriving knowledge in relation to the research question and overall analytical framework of the review.

In the following, we identify some of the relevant cognitive support areas for this scenario:

1. **Query Definition.** Defining the initial query requires capturing the relevant properties of papers, typically by matching keywords found in title, abstract and description. Even prior to this, cognitive ability is required to translate the review scope into a viable set of query keywords. This phase proves challenging as it requires identifying all possible alternative keywords to a specific concept, a process that can take many iterations, involve trade-offs and be prone to error [1]. Automated support for this could be word similarity algorithms (by using word embeddings, either from general language knowledge or by specializing word vectors for a field of science). This word similarity can then be used for keyword expansion. These algorithms could also be enhanced by feeding stronger domain-knowledge, such as from scientific knowledge bases, and low-level data extraction components can be used to extract and curate this knowledge.
2. **Paper Screening.** Even upon refining our query and finding relevant papers, in many cases it is important to filter out papers that are out of scope. This requires a clear definition of the criteria for excluding the papers, namely the exclusion criteria. The selection of primary studies is one of the most difficult tasks in the SLR process, with a direct negative impact on the outcome of the review [11]. Once again automated support could be obtained, for example by using machine-learning classifiers to label papers [10]. An additional benefit of the screening process is that it helps obtain a global view on the body of work in a specific area of research – this also helps further refine the query (e.g., incorporating new keywords) as well as the inclusion and exclusion criteria, and feeds back into the refinement loop of the process.

3. **Recommending Papers.** Recommendation is a useful technique to attenuate the complexities of the overall process, which can complement query-based search strategies [1]. This may be accomplished by leveraging AI approaches based on word similarity, clustering and network analysis to recommend papers and encourage exploring related topics.

Cognitive Delivery At the end-user level, process workers should be provided with a unified work environment where they express in a controlled natural language the tasks they want to perform and interact with underlying cognitive services to refine their requests and perform desired tasks. In this specific scenario, examples of cognitive capabilities include:

1. **Query Expansion.** The identification of relevant literature can be facilitated by digital assistants that can support authors in scaling search strategies, otherwise unfeasible, using natural language. For example, authors could expand keywords, ask for additional papers from references, or iterate on the search and screening phases to receive query refinement suggestions. To make this possible, the chatbot should orchestrate the combination of crowd and AI support, while providing insights about the impact of each alternative in terms of cost, effort and information retrieval (IR) metrics such as precision and recall.
2. **Multi-Predicate Filtering.** Filtering out non relevant papers is an iterative and time-consuming phase that can greatly benefit from augmentation. Digital assistants offer an appealing interface for assisting authors in for example, understanding of the impact and quality of the different exclusion criteria, and making recommendations for the next iteration of the screening process.

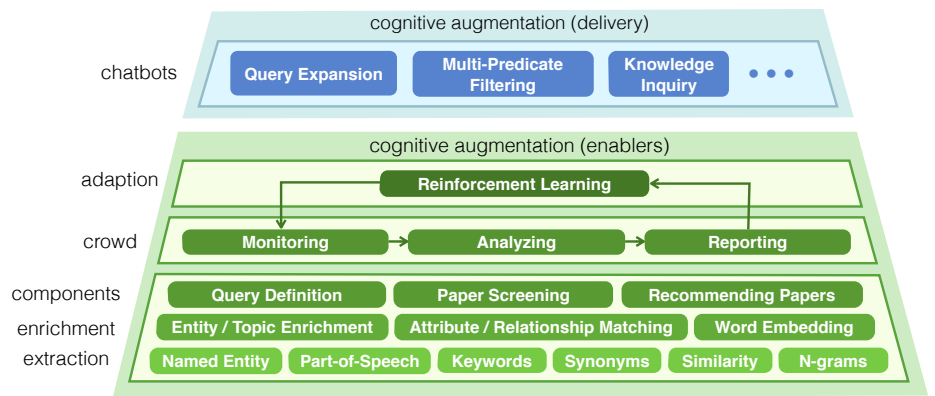


Fig. 6. Cognitive Augmentation Stack for SLR Process.

3. **Knowledge Inquiry.** Getting insights on the knowledge residing in the corpus of papers is another very relevant activity where digital assistants can significantly improve productivity and efficiency. For example, elaborating claims and supporting evidence from the literature usually requires authors to go back and forth from writing and preparing summaries to re-reading papers – an activity that requires significant attention and coordination among authors. Digital assistants could allow authors to elaborate queries in natural language to check claims, as well as prepare summaries and insights to inform authors (e.g., summary tables).

Figure 6 summarizes the above in a potential augmentation stack.

4.3 Augmentation of IoT-enabled Processes

Despite the early adoption, IoT based services are still only in their preliminary stages of development, with several unsolved technical challenges stemming from the lack of effective support for complex IoT services management and data analysis processes.

More specifically, a commonly overlooked limitation of current systems is that they do not make federated analytics over IoT services accessible to analysts and decision makers. There is an imperative need to integrate common user productivity services (e.g., spreadsheets applications and tools such as dashboards and collaboration tools) with underlying IoT data capture and management. Analysts often need to access, manipulate and analyze data from various federated IoT and other data services and should be empowered, like data scientists, to also benefit from the power of advanced analytics in analysis and decision making tasks.

The objective of work in this area is to usher increased productivity, effectiveness through greater simplicity, augmented intelligence and automation over IoT and data services. For instance, layering advanced data analysis and digital assistance capabilities on top of IoT, data, crowdsourcing, task management and collaboration services, may bring several advantages to IoT enabled processes (e.g., smart city, policing and health processes). Regarding knowledge-intensive law-enforcement processes, the implementation of information and communications technology has been a success factor for conducting data-driven and knowledge-intensive processes in law enforcement. The focus on making police work more efficient with new technologies is still valid and consists of many trends: extracting and analyzing large repositories of data gathered from various data sources such as open, private, social and IoT data islands. In this context, a knowledge-intensive process (a type of data-driven processes which comprises activities based on acquisition, sharing, storage, and reuse of knowledge) can benefit extensively from IoT. For instance, in law enforcement processes such as police investigation, knowledge workers (e.g., police investigators) can be augmented with smart entities (e.g., smartphones, smartwatches and smart police uniforms) to collect data (e.g., recording voice, taking photos/videos, using location-based services and leveraging sensory systems to detect explosives) in real-time and

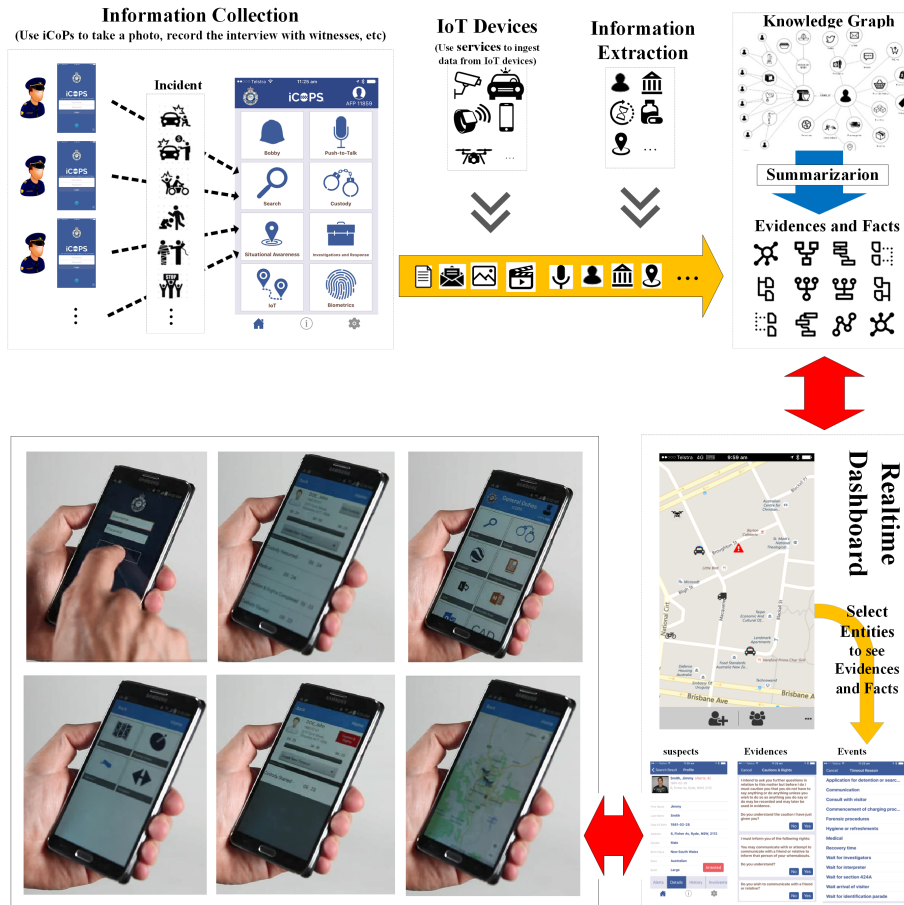


Fig. 7. The iCOP Architecture and screenshots [12].

relate this data to process analysis. This will accelerate the investigation process for cases such as Boston bombing (USA) where fast and accurate information collection and the analysis would be vital.

For example, in iCOP [12], an IoT-enabled framework was presented to explore how an evidence-based interface on a smart mobile device can be used in policing processes to provide a coherent and rigorous approach, to interrogate a “policing knowledge hub”: an IoT infrastructure that can collaborate with internet-enabled devices to collect data, understand the events and facts and assist law enforcement agencies in analyzing and understanding the situation to choose the best next step in their processes. Figure 7 illustrates the iCOP architecture along with some screenshots of the iCOP application.

5 Roadmap to the Future

The proposed vision provides an exciting opportunity to the entire community: from research scientists, to engineers and developers, as well as businesses people. This is because we rely on reuse across the spectrum and advocate against the “one-solution-fits-all” or “automate everything” mentality. Until now, many tools often arise strengthening one aspect of improvement (be it data management, control flow or communication) while neglecting the other. Nevertheless, each of these tools carry merit of their own (e.g., an algorithm designed by a research scientist or best practices developed by a business). To put it another way, cognitive augmentation will be about filling in the gaps between these disparate tools, algorithms or services.

Accordingly, we set forth the following roadmap (and identify some of the key challenges) towards the realization of this dream:

Cognitive Enablement should involve a new method of using AI, a “conversational AI” where end-users are able to iteratively and interactively tune the logic needed to achieve their goal.

- We see AI components packaged from highly defined low-level to less defined (blueprint) high-level functions.
- Nevertheless, the purpose of tuning or conversing with these AI components is two-fold: (i) to help define the logic for the process at hand; (ii) but also to train the system to learn the moves with less reliance on humans in future processing.
- Over time, using the above we project a probabilistic execution model rather than a deterministic model. A class of “modelless processes” not requiring to prescribe or implement the whole component beforehand. AI can use humans and observe how they work (along with continuous feedback) to derive the programming logic. We have begun early research and development into auto-mapping NL intent into API calls [14].

Cognitive Delivery should empower end-users to drive the section of the process that requires human intervention. This should be in the form of natural-language bots, that either proactively prompts the end-user to trigger some action; or reactively responds only if the end-user inquires. One of the major challenges is understanding user intent from an expression in order to translate into a executable command. We began to work on this direction with some initial research results achieved.

Cognitive augmentation has also the potential to empower populations, such as blind and visually impaired (BVIP) users. BVIP and other populations have been traditionally challenged by the current interaction paradigm for accessing information and services on the Web. We believe that cognitive augmentation can enable and deliver more natural experiences, and help close the digital inequality affecting the Web today. Thus, in the same way the problems we highlighted in this paper are amplified for vulnerable populations, so the benefits and potential social impact of cognitive augmentation. We have made our first steps in this direction [3] and call on the community to join us.

Putting all the above together, *Cognitive Augmentation* (i.e., AI + chatbots = conversational AI) should be packaged as first-class citizens to existing work tools, in a manner similar to what Service Oriented Architecture (SOA) achieved for Web services.

References

1. Al-Zubidy, A., Carver, J.C.: Identification and prioritization of slr search tool requirements: an slr and a survey. *Empirical Software Engineering* **24**(1) (2019)
2. Alotaibi, Y.: Business process modelling challenges and solutions: a literature review. *Journal of Intelligent Manufacturing* **27**(4), 701–723 (2016)
3. Baez, M., Daniel, F., Casati, F.: Conversational web interaction: Proposal of a dialog-based natural language interaction paradigm for the web. In: *International Workshop on Chatbot Research and Design*. pp. 94–110. Springer (2019)
4. Baez, M., Parra, C., Casati, F., Marchese, M., Daniel, F., Di Meo, K., Zobebe, S., Menapace, C., Valeri, B.: Gelee: Cooperative lifecycle management for (composite) artifacts. In: *Service-Oriented Computing*, pp. 645–646. Springer (2009)
5. Barukh, M.C., Benatallah, B.: Processbase: A hybrid process management platform. In: *International conference on service-oriented computing*. pp. 16–31. Springer (2014)
6. Beheshti, A., Benatallah, B., Sheng, Q.Z., Schiliro, F.: Intelligent knowledge lakes: The age of artificial intelligence and big data. In: *Web Information Systems Engineering - WISE Conference*. vol. 1155, pp. 24–34. Springer (2019)
7. Beheshti, A., Benatallah, B., Tabebordbar, A., Motahari-Nezhad, H.R., Barukh, M.C., Nouri, R.: Datasynapse: A social data curation foundry. *Distributed Parallel Databases* **37**(3), 351–384 (2019)
8. Cognini, R., Corradini, F., Gnesi, S., Polini, A., Re, B.: Business process flexibility—a systematic literature review with a software systems perspective. *Information Systems Frontiers* **20**(2), 343–371 (2018)
9. Créquit, P., Trinquart, L., Yavchitz, A., Ravaud, P.: Wasted research when systematic reviews fail to provide a complete and up-to-date evidence synthesis: the example of lung cancer. *BMC medicine* **14**(1), 8 (2016)
10. Krivosheev, E., Casati, F., Baez, M., Benatallah, B.: Combining crowd and machines for multi-predicate item screening. *Proceedings of the ACM on Human-Computer Interaction* **2**(CSCW), 1–18 (2018)
11. Palomino, M., Dávila, A., Melendez, K.: Methodologies, methods, techniques and tools used on slr elaboration: A mapping study. In: *International Conference on Software Process Improvement*. pp. 14–30. Springer (2018)
12. Schiliro, F., Beheshti, A., Ghodratnama, S., Amouzgar, F., Benatallah, B., Yang, J., Sheng, Q.Z., Casati, F., Motahari-Nezhad, H.R.: icop: Iot-enabled policing processes. In: *International Conference on Service-Oriented Computing*. pp. 447–452. Springer (2018)
13. Yaeli, A., Kofman, A., Dubinsky, Y.: Software development governor: Automating governance in software development environments. In: *2009 31st International Conference on Software Engineering-Companion Volume*. pp. 413–414. IEEE (2009)
14. Zamanirad, S., Benatallah, B., Barukh, M.C., Casati, F., Rodriguez, C.: Programming bots by synthesizing natural language expressions into api invocations. In: *2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. pp. 832–837. IEEE (2017)