The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence
Symbiotic Cognitive Systems: Technical Report WS-16-14

The Impending Ubiquity of Cognitive Objects

Maryam Ashoori, Rachel Bellamy, Justin D. Weisz
IBM T.J. Watson Research Center
{maryam, rachel, jweisz}@us.ibm.com

Abstract

The word symbiosis (Merriam-Webster 2015) has its origins in biology where it means “the relationship between two different kinds of living things that live together and depend on each other.” When referring to symbiotic cognitive computing, we expand this definition to include both people and intelligent computational agents who work together in a partnership (Farrell et. al 2005). Cognitive objects embody these intelligent agents, providing a physical object that may sense, compute, react, and interact with the power of cognitive computing. In this paper, we describe a few preliminary design explorations that investigate the impact of being surrounded by cognitive objects during group meetings. We frame a research agenda around the construction, programming, and usage of cognitive objects in work and home environments, and for use cases across industries such as oil and gas, healthcare and agriculture.

Introduction

Symbiotic cognitive computing concerns the symbiosis of humans and cognitive computing systems. We are particularly concerned with symbiotic cognitive computing in environments where multiple people come together in a physical space to interact with one another and with one or more intelligent agents. We instrument the surrounding objects in a cognitive environment with sensors that collect data about the physiological and cognitive state of people, and actuators that react to those states. Cognitive objects are equipped with intelligent software agents, called “cogs” (Farrell et al 2015) that perceive, plan, make decisions, and interact with the people around. Cognitive objects enable people and agents in an environment to be mutually aware of each others’ presence and activity, develop connections through interaction, create shared representations and improve over time.

At IBM Research, we are designing a cognitive environment to enable teams such as business managers, emergency planners, and executives to make more effective strategic decisions. An aspect of this research is to explore how different kinds of cognitive objects may be used to influence strategic decision-making.

To facilitate our research, we are developing a research infrastructure called the Cognitive Environments Laboratory (CEL) in which we develop and refine the building blocks of Cognitive Environments. The CEL is an exploratory space in which we can implement design experiments and learn about the characteristics that best facilitate a group decision-making process. In this paper, we describe a few preliminary design explorations that investigate the impact of being surrounded by cognitive objects during group meetings. We present our preliminary findings and outline plans for a more rigorous investigation into how we can augment existing physical objects to add computational capabilities, how to program these objects to exhibit advanced behaviors, and how to incorporate these objects into symbiotic cognitive environments in order to support both the informational needs and the affective state of people in those environments.

What is a Cognitive Object?

We define a cognitive object as an object that is capable of sensing the environment, performing computation, and providing feedback to a person. In creating a cognitive object, the sensing and actuation functions will be inherent in the object; for example, when constructing a cognitive lamp, the bulb (actuator) will be part of the lamp. However, the computational capabilities of the object may follow one of these patterns:

- Embedded computation, in which a processor, memory, and storage are built into the object (e.g. Furby[1] contained an embedded processor that performed rudimentary speech recognition and generated auditory responses)
- Augmented computation, in which a computational device is connected to an everyday object (e.g. the cognitive chair, discussed in the next section, was developed

---

[1] https://en.wikipedia.org/wiki/Furby

by adding pressure sensors and an Arduino controller to an existing office chair)

- External computation, in which the object performs computation by communicating via a network interface with a cloud service (e.g. Siri performs speech-to-text in the cloud to formulate responses to natural language questions)

Each of these designs tradeoffs the impact on physical design (embedded computation takes space), the power requirements (e.g. battery vs. wall outlet), the computational capabilities (e.g. cloud provides more computation) and the interactivity requirements (e.g. interactive applications may have latency when round trips to the cloud are slow).

Cognitive Object Design Explorations

Cognitive objects embody intelligent agents, providing a physical object that may sense, compute, react, and interact with the power of cognitive computing. The physical form factor plays an essential role in designing a cognitive object as it sets the basis for the quality of interactions between the software agent and the human counterpart. We have explored a number of cognitive object implementations to understand the effect of a physical embodiment of software agents within a cognitive environment. In this paper, the set of design explorations is focused on a group meeting setting, where we manipulate the surrounding objects by equipping them with sensors, actuators, and a computation reasoning power. We have been interested in manipulating familiar objects that people expect to see in a typical meeting room, such as chairs, lamps, and uniforms.

In this paper, we explore a design of three different cognitive objects: (1) a cognitive dress representing a wearable cognitive object (such as a uniform), (2) a cognitive chair, representing a cognitive object created from an everyday office object, and (3) a cognitive lamp that provides ambient feedback from an intelligent agent in an environment.

Cognitive Dress: A Wearable Experience

A cognitive room can benefit greatly from knowing the affective state of it's occupants. For example, if a room can detect when a group is calmly working versus engaged in a heated debate, it can provide different levels of informational or environmental support – for example, changing the ambiance when people are stressed, or providing question-and-answer support during a debate. There are two primary ways we can infer the physiological and cognitive state of room occupants during a group meeting: by using ambient sensors (e.g. microphones and/or cameras) and by using on-body sensors (e.g. pulse sensor and/or temperature sensors). Ideally, data would be captured via both methods and fused in order to provide a more accurate representation of an individual's state.

In order to capture data from on-body sensors, we designed a dress to measure biometric information such as heartbeat and body temperature. Figure 1 shows a dress that responds to an elevated heart rate or body temperature by either lighting up or vibrating, depending on the privacy preference of the wearer (Ashoori, Bellamy and Weisz 2015). We used the LilyPad Arduino platform with sensors and actuators to implement the dress. The dress operates wirelessly using XBEE radios to communicate with a back-end service, in order to share computations of stress levels with the cognitive agent in the room.

Figure 1: Cognitive dress that senses body temperature and pulse.

In our evaluation of the dress, we noted that some people are sensitive to wearing a garment that intimately captures biometric information; thus, we started exploring other, less intimate form factors such as a chair.

Cognitive Chair: A Passive Experience

When designing the cognitive chair, we recognized that it could carry a wider variety of sensors to measure not only the personal state of the occupant, but the environmental state as well. We instrumented an office chair with pressure sensors (empowered with machine learning) to detect posture and microphones to measure ambient noise and loudness. Posture recognition can be used to provide feedback about bad posture (Haller et. al 2011) or detect the interest level and activities of the person (Cheng et. al 2013). Microphones can be used to infer heated, tense situations and diffuse those situations by manipulating the room state (e.g. playing calming music or changing the color of an ambient light, such as that of the cognitive lamp), or by sending private feedback to occupants (e.g. by buzzing their cognitive uniform or their chair). The chair was implemented in a manner similar to Tan et al. 2001, by taking an existing office chair and covering the seat cushion with an array of velostats to measure pressure. We also sewed a microphone to the back of the headrest. An Arduino controller collects raw sensor data and sends it to a cloud service for analytics – in this way, the chair combines augmented with external computation. Heart and respiratory rate can be measured in a manner similar to Griffiths et. al 2014, by instrumenting sensors on the arms of the chair.

The advantage of measuring stress using a chair versus a wearable uniform is that a chair requires no preparation – a person can just sit in the chair and immediately have their stress level measured, versus a wearable uniform requires people to change their clothes or put on a device, which may be cumbersome or conflict with existing clothing the person is wearing.

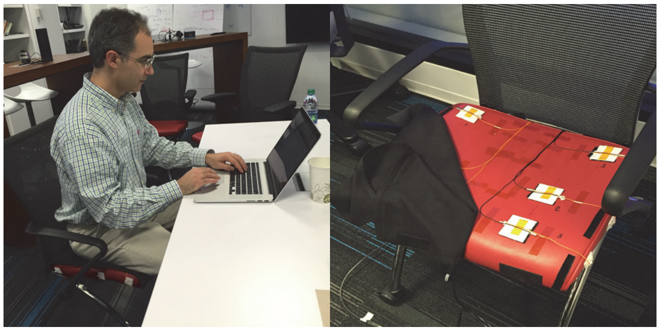

Figure 2: Cognitive chair, a platform for experimenting with sensors to measure the physical and cognitive states of the sitter.

Cognitive Lamp: An Ambient Experience

As we collected feedback on the cogntive dress and chair, people were interested in monitoring their biometrics during a meeting but raised several privacy concerns over displaying that information publically to other meeting participants or sharing it with the enterprise. We have also considered designing a public way of displaying aggregated stress data rather than showing data at the individual level. In this design, we used a lamp outfitted with LEDs to react to the aggregated stress level of the room. We mapped the aggregated stress level of room occupants to the brightness of the lamp; zero illumination indicated no detected stress and full illumination indicated maximum stress (e.g. a heated conversation). People may be more willing to share their personal biometric information when it is aggregated and shared in this manner during a meeting.

Our CEL environment contains an always-on, always-listening software agent to provide occupants with a speech-based user interface for interacting with room services during a meeting. However, the state of this agent was not always clear to room occupants; sometimes, people would attempt to address the agent when it wasn't running or after it had crashed, causing people to become frustrated with the agent. We felt that embodying the agent's status in our cognitive lamp would provide an easy, ambient way for occupants to understand the running state of the room's software agent during a meeting. Thus, in a different experiment, we programmed our lamp to map different states of the software agent to different colors (Figure 3). Slow, pulsing, and fast flashing are used as another means of capturing attention, such as when the room agent has additional information to share with occupants but does not want to interrupt auditorily. Although people were in favor of using a color code to express the state of the software agent, they found it disruptive and expressed their preference to be only interrupted by the lamp if the software agent is broken and needs immediate attention.

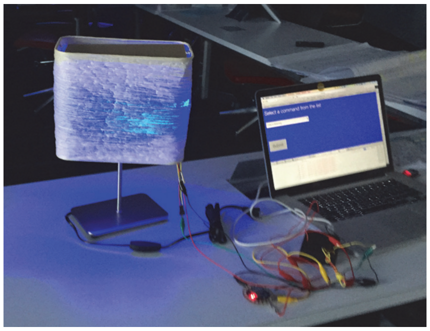

Figure 3: The cognitive lamp shows the status of the ever-present cognitive agent.

Cognitive Objects: What Comes Next?

We see three important research directions for the future of cognitive objects: improving the inference of peoples' physical, mental, and emotional states via fusing sensor data across cognitive objects; using these models to change how cognitive objects interact with people; and giving people the ability to program cognitive objects to produce interesting behaviors.

Sensing Physio-Cognitive State

The ubiquity of cognitive objects provides unique opportunities to improve the sensing of peoples' physical, mental, and emotional state. For example, a cognitive dress that measures pulse and skin conductivity may register that the wearer is stressed. When combined with posture data from a cognitive chair showing the person is sitting for long periods during the day, we may infer the person is lethargic. Adding another layer of detail, a cognitive lamp may register that the room is kept dark all day long; leading us to conclude the person may be suffering from depression. Algorithms that fuse sensor data in this way across a variety of cognitive objects will be needed to produce accurate inferences of peoples' state.

Interacting Based on Physio-Cognitive State

Once we have an accurate picture of peoples' physical, mental, and emotional state, we can change the means by which cognitive objects interact with them. For example, when a network of connected chairs detects a heated conversation in a meeting room (due to high noise and rapid postural changes), a cognitive lamp may dim the lights to turn the meeting room into a more soothing environment. In addition, private feedback can be sent to the chair of the most agitated participants (e.g. via seat vibration). The interaction between the person and the cognitive agent is a function of the person's state and the cognitive objects in proximity. For example, a cognitive chair senses when parkinson's patient experiences a freezing episode – such as a start hesitation when starting to stand up for example. The cognitive agent formulates a reaction to help the patient move through that freezing episode. The agent sends a vibration signal to the chair and provides verbal instructions to help the patient focus on movement. In the case of a start hesitation when standing, the cognitive agent can remind the person to imagine the actions they are about to take as imagining actions (sometimes called mental rehearsal) prior to doing the action has been found to help unfreeze. The agent notifies other people in proximity to go and help the patient. The application of this dynamic interaction between humans, objects, and agents in an environment is not limited to healthcare. A network of connected helmets on an oil platform can be used for disaster management or resource planning. In education, cognitive objects can be used to monitor the progress and physio-cognitive state of students in a class and react to the situation as needed.

Programming Cognitive Objects

As cognitive objects begin to dominate our physical spaces, we envisage a need for enabling people to control these objects and orchestrate complex behaviors among them. Traditional programming models (e.g. Python, Java, Swift) can certainly be used to program interactions between cognitive objects, but they come with a steep learning curve and are generally inaccessible to the wide variety of people who will interact with cognitive objects. Thus, we are exploring novel programming models that enable novice, untrained users to easily specify and orchestrate complex behaviors of cognitive objects.

In the IoT and Robotics spaces, several metaphors have emerged for programming complex behaviors of systems that involve physical devices. Visual programming environments such as Scratch (Malan, and Leitner 2007) visualize traditional programming constructs as colorful blocks of assorted shapes, helping users write programs by composing blocks together. Cantana (Young, Argiro, and Kubica 1995), Fabrik (Ingalls et. Al 1988), and NodeRED[2] use visual elements to display control logic and data flow. These environments are amenable to easy GUI-based programming, but visualizations of low-level logic constructs can become complex and difficult to understand as program size and complexity scales.

Another approach to programming cognitive objects is in using a card-game metaphor. Cards can provide simple representations of abstract logic and complex operations, and have been shown useful for controlling robots and assigning tasks in a domestic environment (Zhao et. al 2009). We are currently developing a novel card-game metaphor for the control of cognitive objects, currently UAVs, as they contain embedded computation and they sense and react to their environment via flight maneuvers or data capture (e.g. taking pictures, videos, or capturing sensor readings). In this metaphor, cards are used to specify drone activities or behaviors, such as flying to a location or taking a picture. Cards are then combined together in order to program more advanced behaviors, such as flying to a location, performing object recognition, and then taking a sensor reading; each of these individual activities is encapsulated in a card. A paper-prototype evaluation of this system (N=18) shows favorable ratings for ease of use and expressiveness, and we observed that people are able to produce novel, complex behaviors in short amounts of time and with minimal training. We are encouraged by these positive results and continue to understand how this metaphor can be extended to other kinds of cognitive objects.

Discussion

Cognitive objects enable people and agents in an environment to be mutually aware of each other's presence and activity, develop connections through interaction, and improve over time. Cognitive objects leverage the existing interactions between human and the physical objects to minimize disturbance and interruption. For example, the interaction with a cognitive chair is the same as a regular one – you just sit on it.

In this paper, we explored the design and implementation of individual cognitive objects. We are also interested in the emerging capabilities of a network of cognitive objects that share computation and exchange data. We see cognitive objects as the future of the Internet of Things (IoT), in which the surrounding objects are not only instrumented with sensors and actuators, but are also empowered with computational engines that make them sense, reason, plan, and react. The hand of cards metaphor suggests how end-users might be able to easily configure and coordinate the responses of cognitive objects.

Cognitive objects will one day surround us, easing our activities and augmenting our capabilities. Increased personalization will be a differentiating feature. By sensing the user, or the environment they are in, the objects will be

---

[2] http://nodered.org

able to adjust themselves to that sensed context. Arranging the objects in their environment so that it pleases them and suits their needs is something that users will want to do.

Of course, any object that is collecting personal data raises privacy and security concerns. We have addressed privacy concerns in our Ambient Lamp design experiment. Successfully addressing such user concerns will determine if our vision of Cognitive Objects is to turn into reality.

Acknowledgement

Many thanks to Saad Ismail and Justin Manweiler for their contributions to card-based methods for programming cognitive objects and to Mishal Dholakia, Tamer Abuelsaad, and Victor Dibia for their contributions to the cognitive chair.

References

Ashoori, M., Bellamy, R. K. E., and Weisz, J. 2015. Creating the mood: design for a cognitive meeting room, *In proceedings of the 33rd annual ACM Conference Extended Abstracts on Human Factors in Computing Systems,* ACM SIGCHI 2015: 2001-2006.

Cheng, J., Zhou, B., Sundholm, M., and Lukowicz, P. 2013. Smart Chair: What Can Simple Pressure Sensors under the Chairs' Legs Tell Us about User Activity? *In proceedings of the 7th International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies*, UbiComp 2013.

Farrell, R. G., Lenchner, J., Kephart, J. O. Webb, A. M., Muller, M. J., Erikson, T. D., Melville, D. O., Bellamy, R. K. E., Gruen, D. M., Connell, J. H., Soroker, D., Aaron, A., Trewin, S. M., Ashoori, M., Ellis, J. B., Gaucher, B. P., and Gil, D. 2015. Symbiotic Cognitive Computing, *AAAI Magazine*, Manuscript submitted for publication, 2015.

Griffiths, E., Saponas, T.S., Brush, A.J.B. 2014. Health chair: implicitly sensing heart and respiratory rate. *In proceedings of the 2014 ACM International Join Conference on Pervasive and Ubiquitous Computing,* UbiComp 2014 Adjunct: 661-671.

Haller, M., Richter, C., Brandl, P., Gross, S., Schossleitner, G., Schrempf, A., Nii, H, Sugimoto, M., and Inami, M. 2011. Finding the Right Way for Interrupting People Improving Their Sitting Posture, *Human-Computer Interaction, Lecture Notes in Computer Science*, INTERACT 2011: 1-17.

Ingalls, D., Wallace, S., Chow, Y.Y., Rudolph, F., and Doyle, K. 1988. Fabrik: a visual programming environment. *ACM SIGPLAN Notices*, 23: 176–190.

Malan, D. J. and Leitner, H.H. 2007. Scratch for budding computer scientists. *ACM SIGCSE Bulletin*, 39: 223–227.

Merriam-Webster. 2015. "symbiosis."

Tan, H. Z. and Slivovsky, L. A., and Pentland, A. 2011. A Sensing Chair Using Pressure Distribution Sensors, *IEEE/ASME Transactions on Mechatronics*, 6 (3): 261-268.

Zheng, Y. J., and Morell, J. B. 2010. A Vibrotactile Feedback Approach to Posture Guidance. *In proceedings of the 2010 IEEE Haptics Symposium.* HAPTIC '10: 351–358.

Young, M., Argiro, D., and Kubica, S. 1995. Cantata: visual programming environment for the khoros system. *ACM SIGGRAPH Computer Graphics*, 29 (2): 22–24.

Zhao, S., Nakamura, K., Ishii, K., and Igarashi, T. 2009. Magic cards: a paper tag interface for implicit robot control. *In proceedings of the 27th Annual ACM Conference on Human Factors in Computing Systems,* ACM SIGCHI 2009: 173-182.