Cognition-Aware Systems as Mobile Personal Assistants

Tilman Dingler
VIS, University of Stuttgart
Stuttgart, Germany
tilman.dingler@vis.uni-stuttgart.de

Abstract
Our ability to focus and concentrate highly fluctuates across the day: at times we are able to work highly focused and at other times we have trouble processing information effectively. The circadian rhythm describes these systematic changes in our daily concentration levels. The idea behind cognition-aware systems is to support users in-situ according to their current cognitive abilities. Such systems are capable of identifying productive phases during the day and provide suggestions for tasks accordingly. In this position paper we present a framework for developing algorithms to derive cognitive states. By being able to detect and predict users' current capacities to take in and process information, such algorithms can help boost productivity, which can result in getting tasks done quicker, communicating more effectively, and processing information more efficiently.

Author Keywords
cognition; cognition-aware systems; context-awareness; human memory; cognitive systems; recall;

ACM Classification Keywords
H.5.m. [Information Interfaces and Presentation (e.g. HCI)]: Miscellaneous

Ubicomp/ISWC'16 Adjunct, September 12-16, 2016, Heidelberg, Germany
© 2016 ACM. ISBN 978-1-4503-4462-3/16/09...$15.00.
DOI: http://dx.doi.org/10.1145/2968219.2968565

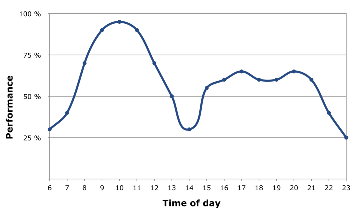

Figure 1: The circadian rhythm describes systematic performance changes across the day showing in our ability to concentrate and focus.

Introduction

Technology has always been designed with the goal to support its users. Consequentially, the more a piece of technology knows about its user in terms of physical, but also psychological constraints and capabilities, the better it can provide assistance. The idea behind context-aware computing is to support users in-situ according to their current situation. Information about the user's context is thereby derived from a range of implicitly and explicitly collected data and often focuses on external factors, such as location, time of day, or device handling. Cognition-aware systems complement this trend towards a holistic context awareness and take into account the user's mental state and information processing capabilities.

Cognition is defined as *those processes by which the sensory input is transformed, reduced, elaborated, stored, recovered, and used* [8]. The effectiveness of cognitive processes is tightly coupled with people's ability to focus or concentrate: the more focused people are, the better these processes work. A system detecting and predicting phases of high and low attention during the day is able to recreate the user's individual attention curve across the day and therefore her circadian rhythm (see Fig.1).

In this workshop paper different attempts are being discussed to quantify people's attentiveness throughout the day. Based on the user's current cognitive state interventions can be staged in the shape of content recommendations: for example, suggesting users to read articles from their reading list or rehearsing foreign language vocabulary. Further, the user interface can be adjusted on the fly to match the user's state: reading interfaces, for example, could hugely benefit from being aware of the reader's attention in order to support the reading process in-situ. Finally, a framework is introduced that can be applied by researchers and system engineers to develop technologies capable of deriving users' cognitive state and match their current state to the complexity and presentation of the task at hand. The goal of such tools should be to help users deal with tasks in opportune moments where their performance levels match the task requirements, thereby allowing a quick and effective task handling.

Detecting Cognitive States

Cognitive states are mostly indirectly detected by monitoring changes in bio-signals, such as heart rate variability (HRV), galvanic skin response (GSR), electroencephalography (EEG), or electromyography (EMG). Also, eye movements can be indicators for cognitive workload since studies have shown the relationship between eye fixations and cognitive processes [6]. Rayner [11] investigated eye movements with regard to the underlying cognitive processes in reading and information processing. Other works have looked at inferring reader engagement, which is closely linked to attention [7].

More recently, we have investigated attentive states of users throughout the day based on their smartphone usage [1]. We were further able to derive and predict states of boredom, in which users turned towards their smartphone seeking stimulation [9]. In an ecosystem of ubiquitous device, such as laptops, tablets, smartphones, and watches, user tracking is becoming more accurate. In our current research we are investigating ways to elicit cognitive states based on usage and activity patterns without the need to collect bio-feedback from users. By tracking these cognitive states throughout the day we aim at detecting the user's' individual attention rhythm to inform systems about the user's current receptiveness, processing capacity, and appropriate times for interventions or disruptions.

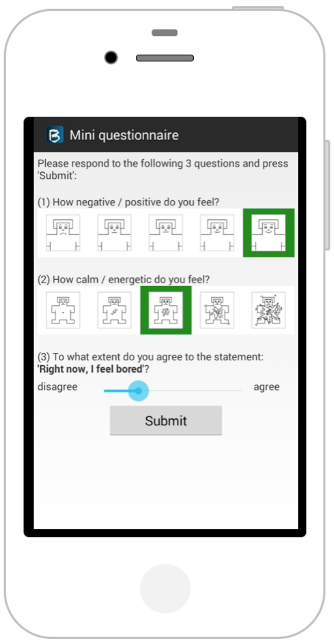

Figure 2: Ground truth collection in form of experience sampling prompting users to provide self assessments.

The Circadian Rhythm

Throughout the day our ability to concentrate and therefore our effectiveness fluctuates. The circadian rhythm describes the systematic changes in our concentration levels across the day. This rhythm is different from person to person, however it occurs in individual, but distinct patterns. Fig. 1 shows a curve with phases of high attention in the morning, a decline in the early and a performance recovery in the late afternoon, for example. It gives clues to the user's ability to be able to work highly focused and communicate efficiently at times, and to have trouble processing information effectively at other times. Cognition-aware systems explicitly or implicitly derive the user's current state and are able to take into account the individual circadian curve. Through this awareness they are capable of supporting users throughout the day and schedule tasks in an efficient and pleasant way.

The circadian rhythm depends on a range of individual factors, such as sleep, nutrition, stress levels, or general health. Traditionally, methods to assess this rhythm include protocols, sleep-wake and biological markers such as the dim light melatonin onset, core body temperature, and cortisol levels [4]. In an attempt to implicitly and non-intrusively reconstruct this rhythm, we focus on ubiquitous technologies and their sensor capabilities. The devices that surround us are increasingly equipped with various sensing capabilities: smartphones as well as wearables, such as step counters, smart garment or smart watches. From their sensors data can be implicitly collected about the user's context. Here, we present a framework for developing algorithms to derive users' cognitive states. The goal of such algorithms is to detect and predict users' current capacities to take in and process information. Such systems are capable of identifying productive phases during the day and suggest matching content or tasks accordingly.

Vision

In this work we are specifically interested in using awareness of users' cognitive states to support and enhance cognition and human memory. Hence, the focus of this research lies on 1) the detection of cognitive states and 2) the application of cognition awareness to adjust the user interface and recommended items from a task list in real-time. Ideally- although not being in the focus of this paper - cognition-aware systems should be capable of identifying and collecting users' tasks (e.g. articles from a reading list) as well as intents (e.g. learning Spanish) and detect and recommend task-specific opportune moments.

By performing this match the overall productivity of the user should increase due to complex tasks being met with phases of higher concentration resulting in more effective or quicker completion. Phases of lower concentration, however, can still be useful to perform daily chores, such as grocery shopping or answering routine emails, without wasting precious performance capacity. Cognitive states that do not match the complexity of the task at hand generally lead to either frustration (task complexity higher than cognitive capacity) or boredom (capacity higher than task complexity). Also, tasks can be adjusted in complexity to be completed more efficiently: reading activities, for example, can be sped up according to users' concentration levels. A reading user interface that adapts to the current capacity to absorb information allows users to effectively take in, process, and retain more information in a shorter amount of time. In [2] we developed and tested interfaces with different reading speeds effectively decreasing users' reading time. They cause, however, higher cognitive workloads, and affect text comprehension often negatively. Future investigations of the relationship between cognitive state and reading UI adaption may allow us to find the sweet spot regarding the speed vs. comprehension trade-off.

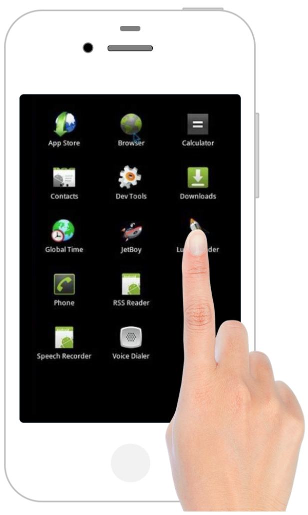

Figure 3: Ground truth collection in form of observing or logging what users actually do (e.g. app interactions).

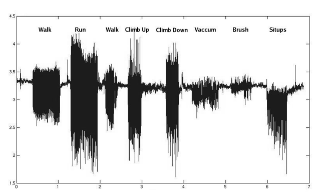

Figure 4: Ground truth collection in form of inferring user states or activities from raw data (e.g. accelerometer data showing sleep, walk or running activities (src.: [10])).

Ideally, such adaptive systems are not limited to single devices, but rather comprise a set of ubiquitous technologies. A context- and cognition-aware system therefore not only matches the user's current state to a task, but also suggests the device type on which to perform this task most efficiently. This can be the desktop PC at hand, the mobile phone on the go or a display in the periphery of the user's home. Naturally, each device has a unique set of capabilities, which need to be matched with the task as well, which in itself gives rise to a number of research questions. Further, one task can be spread out across different devices. For example, reading a foreign language article can be done on a tablet while being on the way to work in the morning. Special grammar exercises can then be scheduled to be completed sometime during the day on the user's laptop where text input is easiest. New vocabulary appearing in the article can be prompted on the user's phone during idle tasks, such as waiting in line or in-between tasks. Vocabulary repetitions can then be shown at night on peripheral displays in the user's home to invite for a final rehearsal. In this way, the overall daunting task of language learning can be broken down into small chunks that can be worked on in opportune moments during the day and therefore allow users commit learning content more efficiently to long-term memory.

A cognition-aware system therefore becomes a truly personal assistant, which learns about the user's patterns and intents, and schedules her day in her best interest. Also, it is aware of, for example, when not to interrupt the user from a current (online or offline) task. In such moments notifications, incoming phone calls or other disruptions can be effectively delayed to more opportune moments in order to not disturb the user's current concentration. This can be beneficial for tasks that need an uninterrupted string of thoughts, but also for situations, in which interruptions through technology are merely inappropriate, for example when immersed in a conversation with a friend or loved one. Obviously, for such systems to become reality, a number of research and development strands need to come together; eliciting user intents and surroundings as well as collecting task and analyzing their complexity are all complex problems and subject to current and future research. However, we believe that cognition-aware systems in the future will not only make us more productive, but also happier beings by allowing us to effectively and swiftly work on tasks in opportune moments.

A Framework for Predicting Cognitive States

The goal is to infer the user's current cognitive state by correlating sensor data with an observed user state. Therefore, in our research we follow a 3-step framework to build algorithms capable of detecting and predicting certain user states:

1. Ground Truth Collection

The current user state can be assessed in three ways:
1) by experience sampling [3], *i.e.* user's self-assessment (e.g. *right now I feel bored* (Fig.2)), 2) through implicit observation (e.g. app usage behavior on mobile phones (Fig.3)), or 3) through inference (e.g. physical activity derived from accelerometer data (Fig.4)). In all these cases, data is being collected through mobile sensors as they are present in phones and wearables; these typically possess rich sensing capabilities, are near-constantly available, and provide means through APIs, for example, to access sensor states, observe user interactions, and use the device's output capabilities to provide users with feedback and prompt for explicit user input. Examples for sensor data can be acceleration, lighting conditions, phone usage intensity, apps being used, or timestamps. So whenever a ground truth is collected through one of these three means, we take a

snapshot of the available sensor data at that point in time. Sometimes it even makes sense to take into consideration the sensor data in a certain time interval (e.g. 5minutes) before or after the ground truth is being collected.

2. Feature Extraction from Sensor Data
The sensor data provides us with context information, based on which we can now define features that may be relevant to elicit with regard to the ground truth collected, *i.e.* the cognitive state. For example, this could be the day of the week or the hour of day extracted from the timestamp collected. From app usage data, for example, we can derive the app category (e.g. news, games, productivity apps), usage duration or frequency. It is worth noting that data from one sensor (such as timestamp) can allow a number of features to be elicited. Hence, we end up with a labeled data set comprising the state of the feature and the collected ground truth as label.

3. Training and Applying Prediction Models
With this labeled data set we can now set out to build detection and prediction models. By using machine-learning techniques we can make out distinct usage patterns that correlate with the user's cognitive state. In the past we have achieved good results with Random Forest classifiers or Decision Trees depending on the types of data collected. *Weka* [5] has proven to be a powerful software offering a variety of tools for analyzing data, training algorithms and exporting prediction models. Obviously, the more data available for training these models, the better the quality of the prediction. Once the cognitive state in question can be detected and distinguished with sufficient accuracy, we can export the prediction model and integrate it into a live system (e.g. a mobile phone app) where its applicability can be tested in the wild.

The procedure described focuses on creating general prediction models, but once we have a proof-of-concept of mapping ground truth to sensor data, the training of the algorithm can also be conducted on the fly with users' personal datasets. Hence, individual prediction models are feasible to be used, which has the additional advantage of *privacy by design* since no data necessarily needs to leave the user's device.

Application Scenario
With the framework described we set out to create cognition-aware applications that support users throughout their day. We started by detecting states of boredom and suggesting entertaining content to read [9], and are now looking into predicting opportune moments for learning (e.g. rehearsing vocabulary) or working off task lists.

High profile characters, such as top managers or head of states often have access to an entire staff that focuses on managing their daily routines and structuring their day as effective as possible. The resulting daily agenda entails appointments and completion of tasks, but also sleep, nutrition, workout routine, information consumption, and other daily chores. By enabling technology to learn about the habits, activities, and cognitive states of an individual user, we can build systems that go beyond simple context-aware applications. Cognition-awareness allows us to build mobile personal assistants that accompany users throughout their day, detect their cognitive states and structure their task lists in a way so that each task is matched by the optimal user state. Such systems have the potential to help users be more effective at their tasks, increase their overall productivity, happiness (through reduced levels of frustration), and eventual well-being.

Conclusion

This position paper presents the vision of creating technology detecting and utilizing users' current cognitive states. Such systems can be used to recommend activities or tasks at moments during the day, in which the user is most effective attending to them. We map out a number of ways to detect user states and describe a correlational approach driven by machine-learning, which we present in a 3-step framework. Currently, we are using this framework to create algorithms to detect opportune moments for learning, working, content recommendations and interruptions in general. We hope researchers and system engineers will be able to use this framework to build novel applications that take the user's cognition into account to adjust user interfaces in-situ and select specific content in real-time.

Acknowledgments

This research is supported by the Future and Emerging Technologies (FET) programme within the 7th Framework Programme for Research of the European Commission, under FET grant number: 612933 (RECALL).

REFERENCES

1. Tilman Dingler and Martin Pielot. 2015. I'll Be There for You: Quantifying Attentiveness Towards Mobile Messaging. In *Proc. of the MobileHCI'15*. ACM, New York, NY, USA, 1–5. DOI: http://dx.doi.org/10.1145/2785830.2785840
2. Tilman Dingler, Alireza Sahami Shirazi, Kai Kunze, and Albrecht Schmidt. 2015. Assessment of Stimuli for Supporting Speed Reading on Electronic Devices. In *Proc. of Augmented Human'15*. ACM, New York, NY, USA, 117–124. DOI: http://dx.doi.org/10.1145/2735711.2735796
3. Joel M Hektner, Jennifer A Schmidt, and Mihaly Csikszentmihalyi. 2007. *Experience sampling method: Measuring the quality of everyday life*. Sage.
4. Wytske A Hofstra and Al W de Weerd. 2008. How to assess circadian rhythm in humans: a review of literature. *Epilepsy & Behavior* 13, 3 (2008), 438–444.
5. Geoffrey Holmes, Andrew Donkin, and Ian H Witten. 1994. Weka: A machine learning workbench. In *Intelligent Information Systems, 1994. Proceedings of the 1994 Second Australian and New Zealand Conference on*. IEEE, 357–361.
6. Marcel Adam Just and Patricia A Carpenter. 1976. Eye fixations and cognitive processes. *Cognitive psychology* 8, 4 (1976), 441–480.
7. Kai Kunze, Susana Sanchez, Tilman Dingler, Olivier Augereau, Koichi Kise, Masahiko Inami, and Terada Tsutomu. 2015. The Augmented Narrative: Toward Estimating Reader Engagement. In *Proc. of Augmented Human'15*. ACM, New York, NY, USA, 163–164. DOI: http://dx.doi.org/10.1145/2735711.2735814
8. Ulric Neisser. 1976. *Cognition and reality: Principles and implications of cognitive psychology.* WH Freeman/Times Books/Henry Holt & Co.
9. Martin Pielot, Tilman Dingler, Jose San Pedro, and Nuria Oliver. 2015. When Attention is not Scarce - Detecting Boredom from Mobile Phone Usage. In *Proc. of Ubicomp'15*. ACM, New York, NY, USA.
10. Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L Littman. 2005. Activity recognition from accelerometer data. In *AAAI*, Vol. 5. 1541–1546.
11. Keith Rayner. 1998. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin* 124, 3 (1998), 372.