The Case for Psychological Computing

Xuan Bao, Duke University
Mahanth Gowda, Duke University
Ratul Mahajan, Microsoft Research
Romit Roy Choudhury, Duke University

ABSTRACT

This paper envisions a new research direction that we call *psychological computing*. The key observation is that, even though computing systems are missioned to satisfy human needs, there has been little attempt to bring understandings of human need/psychology into core system design. This paper makes the case that percolating psychological insights deeper into the computing layers is valuable, even essential. Through examples from content caching, vehicular systems, and network scheduling, we argue that psychological awareness can not only offer performance gains to known technological problems, but also spawn new kinds of systems that are difficult to conceive otherwise.

1. INTRODUCTION

Many fields such as economics, business and medicine have successfully leveraged psychological traits of humans to design better solutions [1–3]. In computing, researchers in human-computer interaction (HCI) have long recognized the need to embrace insights from human psychology [4–7]. This need is clear because the way users interface with technology depends on their psychological characteristics.

In this paper, we ask if computing systems can benefit from embedding human psychological factors deeper into their design and not just their interfaces. We begin by describing the origin of our inspiration—*RedBox*, a DVD rental company. We argue that psychological factors that help RedBox succeed can alleviate the crisis of bandwidth scarcity in cellular networks.

RedBox is a young company in the USA that sets up red-colored kiosks in accessible locations such as grocery stores, gas stations, and airports. The kiosks rent DVDs for a low price (often $1) [8]. People who pass by these kiosks, say when grocery shopping, check-out a DVD and return it the next day. RedBox has been tremendously successful at the same time that larger DVD rental stores (e.g., BlockBuster) are losing business. What is surprising is that with just 50 DVDs per kiosk, and a few kiosks per zip code, RedBox is able to cater to the needs of the population. People are not driving to large rental stores with a much wider selection of DVDs; many are happy picking a DVD from the limited options that a RedBox kiosk offers.

We posit that RedBox's success can be explained by psychological traits of humans. Human desires, especially those pertaining to leisure and entertainment, are not always rigid and exhibit flexibility [9]. A customer may be in the mood for an action movie, but may not have a specific movie in mind. This customer's needs can be met as long as the kiosk has at least one movie that falls within her "circle of flexibility." Further, human flexibility can be increased through incentives, and their choices can be influenced by stimuli in the environment. The incentive for using a RedBox kiosk is its cheaper price and easier accessibility than a larger store. These factors may influence a customer to even rent a comedy if none of the available action movies are to her liking.

We explore the possibility of leveraging the same psychological factors to alleviate the bandwidth crisis in cellular networks. Cellular data networks are facing a serious crisis in coping with increasing demands, especially of video content, from mobile devices [10]. One promising solution is prefetching, i.e., predicting the content that the user is going to request in the future and downloading them over WiFi. Unfortunately, such prefetching has proved extremely difficult; one almost needs a crystal ball to accurately predict the content that the user will request a few hours in the future, time scales at which WiFi-based prefetching can help. Hence, most schemes prefetch at much shorter time scales, e.g., top search results can be prefetched if the user searches for a keyword in Google. While this improves user experience, it does not reduce 3G load [11].

Now consider how psychological factors above can change the situation. Alice's mobile device can roughly speculate about her flexible content consumption interests (perhaps based on her Web browsing history), and prefetch related videos over WiFi. The prefetched videos need not precisely match what Alice would later request; a reasonable set will suffice. Now, when Alice wants to view videos and only 3G connectivity is available, she is shown the list of prefetched videos. Of course, she can stream a non-cached video over 3G (equivalent to driving to BlockBuster), but because of flexibility and incentives, she may often be happy with a prefetched video (RedBox kiosk). The incentive is that prefetched videos play uninterrupted–streaming over 3G can be jittery—and do not contribute to her quota of cellular data. Our preliminary experiments indicate that, with such a prefetching scheme, on average users view 2 out of 3 videos from the cache. This indicates a promising opportunity to reduce network load.

In the body of the paper, we argue that the application of psychology is broadly useful and provide several examples of such systems. Besides improving existing techniques such as prefetching, psychological insights can facilitate new kinds of systems. Consider driving, for instance. It is well known that drivers are often distracted by emails and SMSs. Consequently, governments are banning their use [12], and researchers are developing technical solutions to detect when the user is driving and disable the phone [13]. However, recent psychological studies show that in certain conditions

HotMobile'13, February 26–27, 2013, Jekyll Island, Georgia, USA.

driving without any distraction incurs an excessively low cognitive load [14], inducing drowsiness and day-dreaming. This implies that too little mental load is risky, just like too much mental load. Taking this psychological factor into account, we propose that smartphones should detect when the driver is dozing off or drifting in thought, and intelligently present "information snacks" to maintain the driver's cognitive load above a threshold. Such snacks could be in the form of read-aloud emails, a quick display of a Facebook picture, or games that require the driver to focus on the road. Our preliminary experiments demonstrate the effectiveness of this approach. Conceiving such systems requires co-thinking psychology and technology, rather than studying them in isolation.

Current technological trends are ripe to begin thinking in this direction. As we move from personal *computers* to personal *computing*, the need for devices to intimately understand users is greater. Fortunately, the convergence of smartphone sensors (always on humans), smart objects in the surroundings (measuring context), and recent advancements in machine learning (making decisions from big data), make it feasible to gain deep insights into human behavior and mood. Quantitative psychology, a burgeoning field, is already taking advantage of this convergence. The advantages could be mutually reinforcing—technology could gain from psychological insights that were in turn gained from technology.

This paper barely scratches the surface of what psychological computing could be and puts forth assorted ideas to motivate the case. We are aware that we are at a rather early stage. Nonetheless, we believe that there is latent promise, and we invite the community to give us feedback and get involved. In collaboration with psychologists, important technological changes are feasible. At the least, we believe it is worth trying.

2. MOTIVATING EXAMPLES

We motivate the need for factoring in psychological traits in system design by outlining several diverse examples. For each example, we present the psychological traits being exploited and explain how they facilitate system designs that break away from convention. For two of the examples, we also present preliminary experimental results.

2.1 Overnight Pre-fetching and Caching
(Content Demand is Flexible)

Our first example is that of a system that prefetches content when bandwidth is cheap (e.g., WiFi) for later consumption when it is expensive or unreliable (e.g., 3G). It is based on the following three observations about psychological behavior, which we believe also underlie the success of RedBox as a content dissemination platform.

(1) **Human demands for some types of content (e.g., videos) are flexible.** Human demands for content are not always rigid. Often, people do not have a specific content in mind, but are open to suggestions. This openness is part of the reason that people can find watchable movies from a RedBox kiosk with only 50 DVDs.

(2) **Human flexibility can be magnified with incentives [15].** Some customers may prefer other movies than those in the kiosk but still pick from the kiosk rather than going to other places. Two incentives contribute to this behavior. The first is accessibility: the kiosks are distributed at convenient places where people naturally visit. The second is price: Redbox movie costs $1, while bigger stores cost around $4.

(3) **Fewer choices simplify decision making.** While one may think of Redbox's limited movie choices as a disadvantage, psychological studies show that fewer choices can actually be advantageous because it eases the decision-making process [16]. When too many choices are present, humans may make no choice at all because of the excessive cognitive overhead in making the right decision.

The psychological factors above suggest that a content prefetching system can succeed by caching even a small library if users have flexibility in content consumption. To validate this hypothesis, we build a system to save 3G bandwidth. The mobile device, instead of downloading contents through 3G on the fly, can prefetch videos when WiFi is available, e.g., overnight.

2.1.1 Prototype Design

While the observations above and our techniques apply to many forms of content, our prototype focuses on YouTube videos because video content consumes a large amount network bandwidth today and its share is expected to further grow. Designing our system entails two main challenges. First, we must infer the user's interests at a coarse level; no incentivization may convince a user to consume content in which they are not even remotely interested. Second, we must identify a small set of videos to prefetch such that the user is likely to consume a video from this set.

An overview of our system is shown in Figure 1. The interest mining module infers the user's interest by mining their personal information such as Web history, email content, and social network feed. Given historical data of these activities, we use an NLP technique called SIP (Statistically Improbable Phrase Extraction) [17] to mine for interests. SIP essentially finds meaningful phrase combinations in textual data that summarize the text. Table 1 shows a subset of interests for three of our study participants.

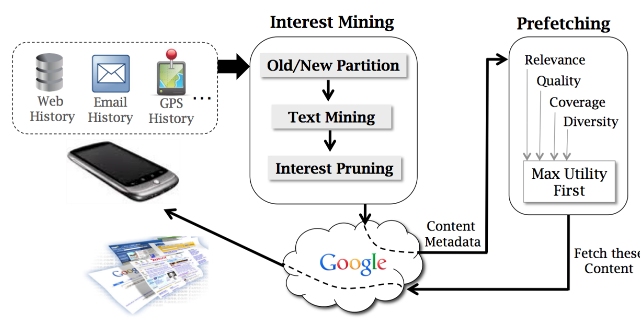

Figure 1: Overview of overnight caching

Table 1: Example Interests

User1(Wireless)	User2(Game)	User3(Mobile)
Distributed Comp.	New Vegas	Efficient Retrieval
Summer Intern.	Civilization V	Life Log Based
Software Radio	Scroll V	Anatomy Netflix
Distributed Sys.	The Vault	

We then use the interests to identify a set of matching videos through YouTube search queries. The *Prefetching* module decides which videos to prefetch. It considers not only video relevance and popularity but also the degree of diversity among cached videos and the level of coverage they provide for the topics of interest. For example, if the user exhibits interests in both soccer and volleyball, we cache at least a few volleyball videos as well, even though soccer videos tend to be much more popular. Further, the relationship between videos is also considered. Since YouTube users often follow the "related video" links after watching a video, the module chooses to cache some of these related videos as well [18]. Our goal is to maximize the possibility that the user will find at least some cached content that falls within her circle of flexibility.

The prefetched content is advertised in a side-panel whenever the user visits YouTube. This serves the role of explicit incentivization towards cached videos. The list length is limited to 50 for a manageable browsing and selection effort.

2.1.2 Preliminary Results

We conducted a preliminary user study to evaluate our hypothesis and techniques. For each user, we first identified their interests based on their browsing history, and for each inferred interest, we asked them if it was relevant. We then prefetched videos based on those interests, and for each video, we asked them if it was relevant (i.e., they would like to view it at some point in time). Finally, we asked them to browse YouTube for a fixed duration, during which they were free to view prefetched or non-prefetched videos or a mix. We recruited eight subjects, of which one subject only finished the interest relevance portion of the study.

Figure 2 shows the performance of interest mining by plotting the percentage of the mined topics that the subjects deemed relevant. The average across users is 66%, which means that even simple interest mining techniques can identify a user's interests.

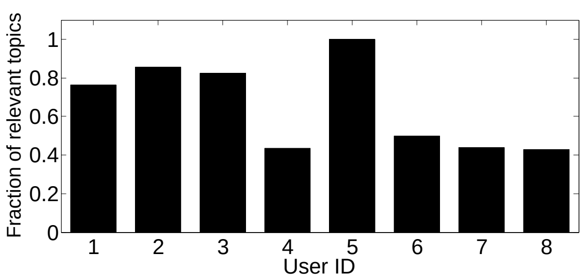

Figure 2: The fraction of topics deemed relevant.

Figure 3 shows the fraction of prefetched videos that the subjects deemed relevant. We see that this fraction varies from 10-60%, with the average being 40%. We find this encouraging given that the YouTube library is vast, in which only a minuscule fraction is presumably relevant to a given user. We show below that, because of psychological factors mentioned above, this relevance rate is enough for users to consume a significant fraction of the content from the cache. More sophisticated techniques (e.g., those based on collaborative filtering) may further boost the relevance rate.

Figure 4 shows the comparison between cached and on-the-fly videos users watched during the study. We find that 50-75% of the videos viewed by the subjects were prefetched, with the average being 65%. That is, two out of three viewed videos were from the cache. These preliminary result suggests

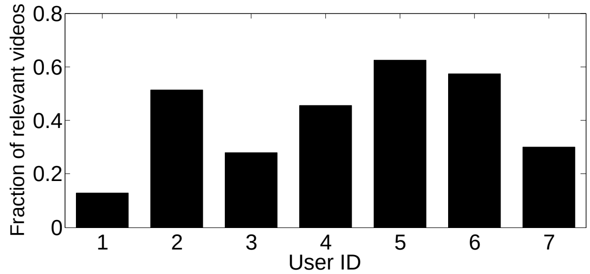

Figure 3: The fraction of videos deemed relevant.

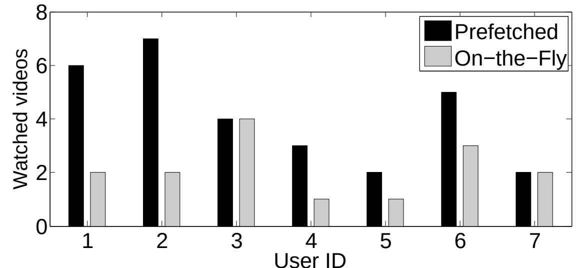

Figure 4: Number of watched videos

that our techniques can lead to a significant reduction in 3G bandwidth consumed by users. Our experiments did not involve economic incentive towards cached videos—there was only a performance incentive. We expect higher cache usage when economic incentives are involved.

2.2 Driving with Cognitive Comfort

(Controlled Diversion is Beneficial)

Driving in the USA can sometimes be monotonous. Highways can be empty, the view can be dull, and adaptive cruise-control and blind-spot detectors can eliminate the sudden surprise. In conjunction, the complete ban on phone-use while driving is eliminating the periodic diversions. Such driving conditions can translate to an excessively low cognitive load on the driver, perhaps resulting in the driver dozing off or drifting away in thought. Importantly, a DoT report and related research papers already show evidences of such cases – called highway hypnosis [14]. The report predicts that the problem is likely to become pronounced over time as cars become semi-autonomous, leaving the driver with very little work on average.

We view an opportunity for psychological computing here. Our hypothesis is that a controlled level of diversion may be beneficial to keep the cognitive load on the driver above a threshold, which can help with remaining focused on driving. To this end, we propose a system in which the driver's smartphone is mounted on the car's dashboard. The phone looks for indications of dozing or hypnosis (such as through closing eyelids and head stillness), and on detecting them, presents the driver with electronic diversions. Such diversions could be emails read out aloud, a picture displayed from a social network feed, or even a game that forces the driver to focus on the road. We call these diversions "information snacks" to capture how they feed the driver's mind and keep her vigilant about road conditions.

We designed a user study to examine the merit of controlled diversions during driving. The study is inspired by the standard psychological experiment – called *MackWorth Clock* – that tests for human vigilance [19]. In MackWorth Clock test, users are required to focus on the seconds hand of a clock, and press a button whenever the hand skips a tick. The skips are infrequent, measuring the extent to which the user can remain vigilant under such monotonicity. We adapted the MacWorth clock experiment, but replaced the clock with a 30 minute video recorded from a car's windshield. The video is monotonous because it is mostly recorded on an empty US highway, however, at infrequent time points, a few frames of the video are removed, causing a distinct glitch. We created a web-based interface and asked volunteers to press a key whenever they noticed the glitch. During some part of the video, we periodically presented information snacks in the form of audio and visual diversions. The diversions were read-aloud news headlines and interesting facts, or interesting pictures chosen from Flickr. A screenshot of our experiment interface is shown in Figure 5(a), with a picture of a deer as a visual diversion. We measured the accuracy of glitch detection with, and without, diversions.

We find that without diversion, the accuracy proved to be 64.2%, but when presented with diversions, users achieved 83.3% accuracy. We also noticed improved detection when the diversions were presented tens of seconds before the glitch, perhaps suggesting that diversions increase human alertness, which then decays. If such behavior holds in real-life driving, we conjecture that information snacks could be effective there as well.

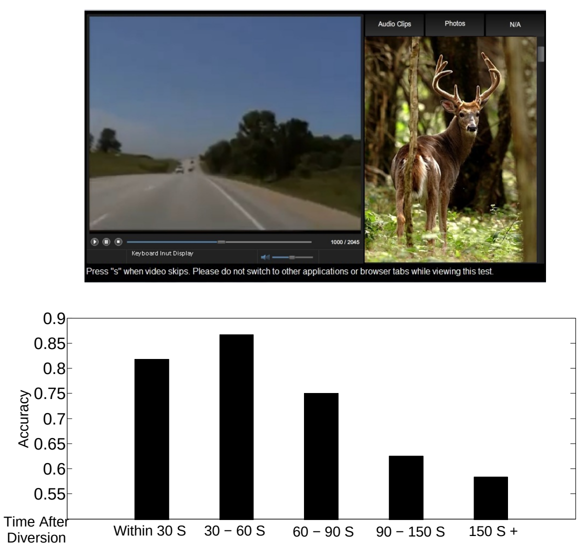

Figure 5: (a) Screenshot from the driving diversion study. (b) Glitch detection accuracy higher just after the diversion, but decays over time.

2.3 Cognitive Load Aware OS

(Will Power is Like Muscle)

Many tasks on a computing device (e.g., coding, writing) incur high cognitive load. Recent research shows that the main resource exerted for such tasks—will power—is akin to a muscle [20]. It gets tired and becomes less effective as it is used.

Based on this finding, we propose that computing devices should help humans work more effectively. These devices are in a good position to judge the depletion level of the user's will power, derived from the nature of the task, the length of time she is working, and her rate of errors (e.g., number of typos). When the system observes a well-focused user, it could perhaps delay her incoming emails and other forms of electronic distraction, even place her phone on least-activity mode. When the focus begins to deplete, emails and social feeds could be scheduled to arrive, and the user could even be prompted to take a break.

2.4 Eagerness Aware Scheduling

(Perception of Time is Relative)

Human perception of "elapsed time" is inaccurate and often a function of the situation. Attending a boring class for five minutes may appear far longer than watching a gripping video for half an hour. This suggests that any form of system delay, as experienced by humans, may not be measured in sheer time difference between a request and response. Rather, the user's experience could be a function of the "perceived time". If the network could receive hints about the user's current perception (say, the user is waiting alone at a bus stop), her transmission could be prioritized to reduce her perceived waiting time. On the other hand, a student chatting on IM may start browsing videos at the same time – such a student need not be prioritized when the system resources are constrained. To generalize this, networking systems could account for human eagerness while scheduling resources. This is of course different from traditional techniques that prioritize purely based on the application category, i.e., real-time applications over delay-tolerant ones.

2.5 Psychology Aware Compression

(Unequal Focus on Pixels)

Conventional image compression techniques utilize understandings of human visual properties. However, these schemes often do not consider the semantic meaning of the image content and its implication on human perception. While viewing a content, say an image or a video, a user attaches unequal importance to different parts of the content. Psychological studies have shown that when the content theme is human relationships, for instance, viewers' eye movements are biased to human faces, to the extent that one can miss the presence of other objects in the scene [21]. Figure 6 is extracted from the famous Yardbus experiment in 1967, showing an example of a picture and the corresponding eye movements. If such disparities exist in the way viewers "consume" pixels, compression schemes can take advantage of them. Depending on the theme, certain parts of the content can be retained at high fidelity, and others treated with a higher compression factor.

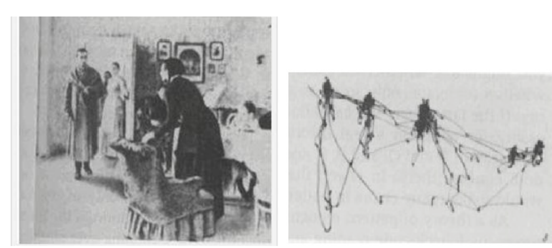

Figure 6: YardBus' study showing how viewers' eye movements are biased towards faces when human relationships are the subject of the picture.

In bandwidth or latency constrained operations, such tradeoffs may be welcome. Graphics rendering could also exploit the above observation. Some parts of the images can be rendered at lower fidelity, saving computational cost and energy.

3. FRAMEWORK AND DISCUSSION

This vision paper is an early exposition of formative ideas and intuitions. We discuss some of the open questions here.

3.1 PSY Layer Implementation

The notion of human psychology is obviously complex and context-sensitive. A key challenge is to make suitable abstractions that can be systematically used in computing systems.

Our intuition is that psychological attributes may be viewed as operating at two time-scales, and should be treated differently. The first class is composed of slowly-changing or stable attributes, those that correspond to "inherent nature" or habits (e.g., flexibility in video demand). The system can directly apply this knowledge (e.g., an empirical flexibility factor) for prefetching and caching. The second class is composed of short time-scale, "state-dependent" attributes, those that impact users' behavior, actions, and current mental state (e.g., cognitive load). Of course, the system would need to capture these "state" information proactively via sensing.

Our deliberations on an architectural framework for psychological computing resulted in more than one design choice. Figure 7 shows one possible framework. In this design, we envision a psychological module (PSY) embedded into the OS of the mobile device. This module implements common services related to psychological computing from which many applications can benefit, to minimize duplication of effort across applications. Embedding PSY in the OS also makes the OS aware of psychology-related factors, and PSY benefits from the context-inferencing functionalities in the OS [22].

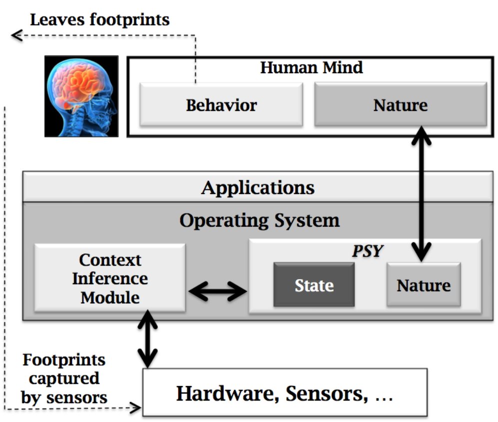

Figure 7: A framework for psychological computing.

The proposed PSY module consists of two components. The "Nature" module is modeled after the "inherent nature" of humans. It can be viewed as a storage of fixed parameters and predefined logic that are consistent with common/expected behavior. To capture these, the "Nature" module in PSY needs to assimilate knowledge about the user (e.g., how flexible a user is), and more importantly, incorporate new findings from modern psychology and brain sciences. The "State" module offers a library of current psychological characteristics and behaviors, through both request-response APIs and publish-subscribe APIs. Example APIs include *currentCognitiveLoad()*, *impatienceIndex()*, *mood()*, *onCognitiveLoadChange()*, and when possible, some physiological indicators, such as *heartRate()* and *pulseOximeter()*. The OS, network stack, and applications draw from both the long and short term classes, depending on their goal. Table 2 shows how a few example applications can be supported by this framework.

Table 2: Support Example Applications

Example Applications	Module	Variable
Prefetching	Nature	Flexibility
Compression	Nature	Visual Attention Bias
OS, Driving, Scheduling	State	Cognitive Load

As stated before, this architecture is not the only possible design. Another alternatives is a cross-layer design that incorporates different levels of behavioral information separately (e.g., facial expressions (smile, frown, etc) could be an example of explicit, low-level information. On the other hand, the user's mood may be a sophisticated, high-level state.). We converged on the Nature-State archictecture because we found it most useful towards implementing the examples we presented. But we do realize that we are early in our thinking and significantly more experimental work is needed to produce a good design.

3.2 Challenges

Realizing the framework above entails many challenges. We discuss two of the main challenges here.

The first challenge relates to identifying the psychological traits that can be exploited towards system design. Admittedly, psychology is a deep field by itself. Limited by our current understanding, most of the discussion in this paper relates only to behavioral and perceptual psychology, especially the quantitatively measurable portion. Fascinating studies in these fields have uncovered many human biases and traits. In addition to the traits mentioned previously, some of them are: (1) humans simplify their decisions based on implicit comparisons, e.g., when given a choice between a paid vacation to Rome, a paid vacation to Paris, and paid vacation to Paris except for the coffee charges, a large fraction of users chose the second option. This is because while comparing between Paris and Rome is difficult, the benefit of free coffee is obvious. Therefore, the second option suddenly looks like a bargain compared to the other two choices. (2) Fatigue can be alleviated if a human switches to a different task that stimulates a different part of the brain. We wonder if these, and various other psychological properties of the mind, can be accommodated in systems of the future.

The second challenge relates to inferring the users' current state and inherent nature. The field of behavioral psychology focuses on observable behavior, and their link to psychological conditions. Advancements in sensing and pattern recognition technology is helping this field uncover fascinating connections. For instance, a recent study shows that the cognitive load on a person can be measured based on pupil dilation [23]. Recent works by Picard et. al. from MIT demonstrates the ability to recognize a user's heart rate by observing her through a webcam [24]. Progress in psychological computing will rely heavily on such advancements and cross-disciplinary research will be vital.

3.3 Natural Questions

(1) **Psychological computing may often leverage the user's context. In that sense, is this essentially context-aware computing?** Context is a broad term defined by the "circumstances that form the setting of an event or object" [25].

While a large body of past work has leveraged the physical context (e.g., ambience, location, time, activity), we are unaware of any systematic effort to exploit the psychological context in core technology. Of course, the psychological context may sometimes manifest itself through physical contexts, such as a person dozing off because of a low cognitive load. We aim to measure the physical activity of dozing off, but utilize the psychological context underneath (i.e., low cognitive load) for system design. It is the use of this hidden variable, we believe, that differentiates psychological computing from context-aware systems.

(2) **Why is psychological computing not a pure HCI problem?** HCI is obviously positioned to leverage psychological factors. However, we believe that incorporating psychological primitives into the "system core" may offer significant benefits beyond a pleasant user experience. For example, many applications, including comfortable driving and eagerness-aware scheduling, can be supported by a common *cognitive load API* exposed by the OS. Perhaps more importantly, the benefits here do not relate to only the users' experience; the system itself benefits by making more efficient resource allocation (CPU cycles, battery, storage).

(3) **How to prevent the system from misbehaving seriously when it misinterprets users' mental states?** Like context-aware systems, the accuracy of psychology-aware systems may not be perfect. Therefore, how to prevent the system from annoying the user when such errors occur is an important challenge. While we do not have a full solution, erring on the side of conservativeness may be a helpful guideline. For instance, in the cognitive load-aware OS example, the system should gradually increase the level of task shuffling. It should trigger a psychology-aware action only after it has gained adequate confidence in its inferences. Moreover, if the user ever manually overrides the system's decision, the system should back-off and not disturb the user until the user specifies otherwise. In other words, the system should avoid being intrusive and provide a scheme to fall back to a conventional system.

4. RELATED WORK

Outside the area of HCI [5–7], a few prior works have also exploited psychological theories for different purposes. For example, Verendel [26] investigates security issues taking rationality into consideration. Also, Shye et al. [27] leverage human physiological traits to control microprocessor frequency in order to save energy. The TUBE project [28] brings humans into the decision loop to alleviate network congestion. Building on these efforts, we argue for systematically embedding human psychology deeper into computing systems.

5. CONCLUSION

Computing systems were envisaged to serve human needs. We argue that, in our eagerness to build sophisticated systems, we may not have adequately comprehended the nature of these needs. Yet, the success metric of computing systems seem strongly influenced by these needs, and is likely to become more so, as technology percolates into every corner of our lives. This paper postulates that improved comprehension of human psychology can help improve the systems we build, both in terms of performance and relevance. We sketch early examples in cellular networking, vehicular systems, scheduling, and content encoding, and generalize them to a broader direction that we refer to as *psychological computing*. Although premature in its current state, we believe this direction has promise, and present it to the community for feedback, opinion, and involvement.

6. REFERENCES

[1] D. Kahneman. Maps of bounded rationality: Psychology for behavioral economics. *The American economic review*, 2003.

[2] G.R. Foxall et al. *Consumer psychology for marketing*.

[3] E.J. Khantzian. The self-medication hypothesis of substance use disorders: A reconsideration and recent applications. *Harvard Review of Psychiatry*, 1997.

[4] J.M. Carroll. *Designing interaction: Psychology at the human-computer interface*. Cambridge Univ Pr, 1991.

[5] S.K. Card, T.P. Moran, and A. Newell. *The psychology of human-computer interaction*. CRC, 1986.

[6] J. Fogarty, S.E. Hudson, and Others. Predicting human interruptibility with sensors. *ACM TOCHI*, 2005.

[7] C. Harrison, B. Amento, S. Kuznetsov, and R. Bell. Rethinking the progress bar. In *ACM UIST*.

[8] How redbox is changing retail. *Marketing News*, 2009.

[9] D. Ariely, G. Loewenstein, and D. Prelec. Coherent arbitrariness: Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, 2003.

[10] S. Deb, K. Nagaraj, and V. Srinivasan. Mota: engineering an operator agnostic mobile service. In *ACM MobiCom*, 2011.

[11] D. Fisher and G. Saksena. Link prefetching in mozilla: A server-driven approach. *WCW*, 2004.

[12] D. Lamble et al. Mobile phone use while driving: public opinions on restrictions. *Transportation*, 2002.

[13] J. Yang, S. Sidhom, et al. Detecting driver phone use leveraging car speakers. In *ACM MobiCom*, 2011.

[14] L. Harms. Variation in drivers' cognitive load. effects of driving through village areas and rural junctions. *Ergonomics*, 1991.

[15] D. Ariely. *Predictably irrational: The hidden forces that shape our decisions*. HarperCollins, 2009.

[16] D. Kuksov and J.M. Villas-Boas. When more alternatives lead to less choice. *Marketing Science*, 2010.

[17] M. Errami et al. Identifying duplicate content using statistically improbable phrases. *Bioinformatics*, 2010.

[18] S. Khemmarat, R. Zhou, L. Gao, and M. Zink. Watching user generated videos with prefetching. In *ACM MMSys*, 2011.

[19] NH Mackworth and et al. *Researches on the measurement of human performance*. 1950.

[20] D. Kahneman. *Thinking, Fast and Slow*. 2011.

[21] A.L. IArbus and A.L. IArbus. *Eye movements and vision*. Plenum press, 1967.

[22] D. Chu, A. Kansal, J. Liu, and F. Zhao. Mobile apps: It is time to move up to condos. In *HotOS*, 2011.

[23] F. Paas and et al. Cognitive load measurement as a means to advance cognitive load theory. *Educational psychologist*, 2003.

[24] R.W. Picard. Emotion research by the people, for the people. *Emotion Review*, 2010.

[25] A.K. Dey. Understanding and using context. *Personal and ubiquitous computing*, 2001.

[26] V. Verendel. A prospect theory approach to security. Technical report, Technical Report, 2008.

[27] A. Shye and Others. Power to the people: Leveraging human physiological traits to control microprocessor frequency. In *IEEE/ACM International Symposium on Microarchitecture*, 2008.

[28] S. Ha and Others. Tube: time-dependent pricing for mobile data. In *ACM SIGCOMM*, 2012.