# Sensor-Based Statistical Models of Human Interruptibility

Research Statement for 2004 IBM Research Human-Computer Interaction Symposium

James Fogarty
Human Computer Interaction Institute
Carnegie Mellon University
jfogarty@cs.cmu.edu

## Introduction

Computing and communication systems are at the core of the modern office worker's everyday experience. At any given point in time, a person might be notified of the arrival of a new email, receive an instant message from a colleague, be reminded by a calendar program of an upcoming appointment, receive a phone call on their office or mobile phone, and be involved in a face-to-face interaction with a colleague. Any one of these demands for attention can be addressed relatively easily, but the sum of repeated or simultaneous demands can become disruptive. Current systems interrupt at inappropriate times or unduly demand attention because they have no way to determine when an interruption is appropriate. A person preparing to make a telephone call typically has no way to know that the callee is in the middle of a face-to-face meeting, and an email client about to announce the arrival of a new message cannot determine whether an obvious or subtle notification is currently more appropriate.

This statement reviews my work on the construction and evaluation of sensor-based statistical models of human interruptibility. My initial work, a series of feasibility studies, has demonstrated that sensor-based statistical models of the interruptibility of office workers can perform as well as or better than human observers. As a part of my dissertation, I am currently developing a system to support application development and deployment of sensor-based statistical models of human interruptibility. I will also be developing and evaluating two techniques to significantly reduce the disruption associated with collecting the observations needed to build statistical models: combining interruptibility data collected from many people and collecting less intrusive types of interruptibility observations.

## Feasibility Work

Many different sensors seem like they might relate to interruptibility, but the uncertainty surrounding their actual usefulness makes it very likely that implementing and then evaluating them would result in significant time and resources being spent on sensors that are ill-suited or sub-optimal for predicting interruptibility. We instead initially collected 600 hours of audio and video recordings from the normal environments of four office workers and used a Wizard of Oz technique to simulate a variety of potentially useful sensors [1, 5]. While these recordings were being collected, the workers were prompted at random time intervals to provide self-reports of their interruptibility. Statistical models based on the simulated sensors were able to distinguish situations reported as "highly non-interruptible" from other situations with an accuracy as high as 82.4%, significantly better than a 68.0% chance ($\chi^2(1, 1344) = 31.13$, $p < .001$) and significantly better than the 76.9% accuracy of human observers who viewed the recordings and estimated the interruptibility of the people in the recordings ($\chi^2(1, 3072) = 5.82$, $p < .05$) [1].

While at IBM Research in Summer 2003, I validated my initial results by deploying real sensors with a larger and more diverse group of office workers [2]. This work showed that models based on real sensors can perform as well as or better than human observers for a variety of office workers. The work also explored how interruptibility models might vary according to a person's job responsibilities and working environment. We also examined how communication systems might use such models, building and evaluating a context-aware communication client [3].

**System Support for Interruptibility Models**

Informed by my feasibility work, I am building a system for *Automatically Modeling Interruptibility by Unobtrusively Sensing You*, or AmIBusy. AmIBusy will provide mechanisms for logging sensor data, collecting observations of interruptibility, and automatically analyzing the collected sensor logs and interruptibility observations to create statistical models of human interruptibility. In order to create applications that model interruptibility, developers currently need to have significant specialized knowledge of machine learning techniques and need to implement significant sensing mechanisms. AmIBusy will provide this functionality to developers, thus making it much easier for applications to use models of human interruptibility. AmIBusy will also support further research on models of human interruptibility, as discussed next.

**Minimizing the Disruption of Building Models of Human Interruptibility**

Statistical models are built by extracting relationships between an independent variable, interruptibility in my work, and features, which are sensors in my work. Existing approaches to collecting the necessary observations of the independent variable, interruptibility, are rather disruptive. Our feasibility work required that people respond to prompts for self-reports more than once per hour. Horvitz and Apacible have used a retrospective labeling approach, recording several hours of activity in a person's office and asking the person to view the recordings and label their own interruptibility [4]. While these approaches have their differences, they both require significant time and attention from the person whose interruptibility is being modeled. Using AmIBusy, I will develop and evaluate two approaches to minimizing the disruption of collecting the observations of interruptibility needed to build statistical models: combining data collected from many people and collecting less intrusive types of interruptibility observations.

Combining data collected from many people provides two primary ways to reduce disruption. First, fewer observations need to be collected from any one person. In the case of self-reports, this might mean that people are prompted at most once per day, instead of the more than once per hour used in my feasibility work. Because fewer observations are collected from any one person, each person will experience less disruption associated with collecting the necessary observations. Second, there is no need for an initial training period in which a model learns a person's interruptibility and performs very poorly. Instead, AmIBusy can provide an initial model based on data collected from other people. The person's own data contributes to the model over time, and there is no initial period in which the model performs very poorly.

To examine less intrusive types of interruptibility observations, I will build and deploy a notification application. Using sensors that can be deployed in software on a typical laptop, I will model how people respond to notifications. Potential responses might be to explicitly indicate that a notification was delivered at a bad time, ignoring a notification, mousing over a notification to read it, or clicking through the notification to obtain additional information. I will also analyze the relationship between responses to my notification application and interruptibility self-reports. If reliable models can be built from less intrusive types of interruptibility observations, such as noting how people respond to a notification, it may be possible to completely do away with self-reports and other more explicit types of interruptibility observations.

## References


1. Fogarty, J., Hudson, S., Atkeson, C., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. and Yang, J. (2004) Predicting Human Interruptibility with Sensors. *To Appear, ACM Transactions on Computer-Human Interaction (TOCHI)*.
2. Fogarty, J., Hudson, S. and Lai, J. (2004) Examining the Robustness of Sensor-Based Statistical Models of Human Interruptibility. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2004)*, 207-214.
3. Fogarty, J., Lai, J. and Christensen, J. (2004) Presence versus Availability: The Design and Evaluation of a Context-Aware Communication Client. *To Appear, International Journal of Human-Computer Studies (IJHCS)*.
4. Horvitz, E. and Apacible, J. (2003) Learning and Reasoning about Interruption. *Proceedings of the International Conference on Multimodal Interfaces (ICMI 2003)*, 20-27.
5. Hudson, S., Fogarty, J., Atkeson, C., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. and Yang, J. (2003) Predicting Human Interruptibility with Sensors: A Wizard of Oz Feasibility Study. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2003)*, 257-264.