Available online at www.sciencedirect.com
ScienceDirect
Physica A 378 (2007) 59–67
PHYSICA A
www.elsevier.com/locate/physa

Automatic mapping and modeling of human networks

Alex (Sandy) Pentland

The Media Laboratory, Room E15-387, 20 Ames St., Cambridge, MA 01239, USA

Available online 11 December 2006

Abstract

Mobile telephones, company ID badges, and similar common devices form a sensor network which can be used to map human activity, and especially human interactions. The most informative sensor data seem to be measurements of person-to-person proximity, and statistics of vocalization and body movement measurements. Using this data to model individual behavior as a stochastic process allows prediction of future activity, with the greatest predictive power obtained by modeling the interactions between individual processes. Experiments show that between 40% and 95% of the variance in human behavior may be explained by such models.

Keywords: Social network; Network analysis; Network modeling; Network dynamics; Modeling human behavior; Eigenanalysis; Influence analysis

1. Introduction

A series of studies on office interactions discovered that 35–80% of work time is spent in spoken conversation, 14–93% of work time is spent in opportunistic communication, and 7–82% of work time is spent in meetings [1]. Senior managers represent the high end of these scales. Given the importance of such communications, it is notable that the majority of adults already carry a microphone and location sensor in the form of a mobile phone, and that these sensors are packaged with computational horsepower similar to that found in desktop computers. This emerging foundation of wearable sensing and processing power has allowed us to begin to automatically map and model how different groups within social or business institutions connect. We have been particularly concerned with our ability to automatically infer properties of human networks that affect propagation of information:

- Location context: work, home, etc.
- Social context: with friends, co-workers, boss, family, etc.
- Social interaction: are you displaying interest, boredom, friendliness, determination, etc.

By taking a statistical, machine-learning approach applied to the users' behavior and physical situation, we have been able to show that it is possible to obtain solid, dynamic estimates of the users' group membership

E-mail address: pentland@media.mit.edu.
URL: http://www.media.mit.edu/~pentland.

0378-4371/$ - see front matter
doi:10.1016/j.physa.2006.11.046

and the character of their social relationships: e.g., who we work for versus those who work for us, or when we are interested versus when we are bored. By characterizing these patterns of behavior using statistical learning methods, we can then examine the users' current behavior to classify relationships as workgroup, friend, interesting, and so forth.

The key to automatic inference of information network parameters is the recognition that humans are not general-purpose equipotent reasoning agents, but rather are creatures with a long evolutionary history that continues to shape our behavior and interactions with others. This shaping of behavior is particularly visible in social relationships and our attitudes toward them: we act differently when interacting with friends versus strangers, and when we are interested versus when we are bored.

Some of these categories can be inferred using standard methods such as surveys, however, these standard methods often suffer from subjectivity and memory effects, and their infrequency means that they are prone to becoming out-of-date. Even when information from standard methods is available, we would still like to use automatic methods to validate or even correct the standard information sources.

In this paper, we present statistical learning methods that use wearable sensor data to make reliable estimates about a user's interaction state (e.g., who was talking to whom, how long did the conversation last, etc.). We then use these results to characterize the connections that exist in groups of people.

Automatic mapping can be much cheaper and more reliable than human delivered questionnaires. For instance, in one of our studies we found that our automatic methods had an accuracy of 87.5% for detecting conversations with durations of 1 min or more. In contrast, a traditional survey of the same subjects produced only 54% agreement between subjects (where both subjects acknowledged having the conversation) and only 29% agreement in the number of conversations [2,3].

Automatic discovery and characterization of face-to-face communication networks will also allow researchers to gather interaction data from larger groups of people. This can potentially remove two of the current limitations in the analysis of human networks: the number of people that can be surveyed, and the frequency with which they can be surveyed.

Automatic mapping of human networks will never be perfect, although it already seems superior to previous methods in some regards. We can also vary the confidence thresholds of the system, making the system more or less cautious about particular types of mistakes. In addition, the models provided by automatic mapping can suggest when traditional survey methods would be most useful, resulting in a semi-automatic capability that can have very high accuracy and a relatively low cost.

2. Socioscopes

Our approach to mapping and modeling human networks is to adopt the conceptual framework used in biological observation, such as is used to study apes in natural surroundings or in natural experiments such as twin studies, but replacing expensive and unreliable human observations with automated, computer-mediated observations. We imagined an advanced 'socioscope' that can accurately and continuously track the behavior of hundreds of humans at a time, recording even the finest scale behaviors with near perfect accuracy.

My students and I have built an approximation of this imaginary socioscope, using mobile telephones, electronic badges, and PDAs [4–8]. My collaborators and I have used this socioscope to track the behavior of 94 people in two divisions of MIT, the business school and the Media Laboratory, a group of 110 international researchers attending meetings at MIT, and certain other smaller groups in the wider Boston community. The subjects were typically between 23 and 39 years of age, with the business school students almost a decade older than the Media Lab students. Subject groups were typically $\frac{2}{3}$ male and $\frac{1}{3}$ female, and approximately half were raised in America.

The socioscope consists of three main parts. The first part consists of 'smart' phones programmed to keep track of their owners' location and their proximity to other people, by sensing cell tower and Bluetooth IDs. This has provided us with approximately 330,000 h of data covering the behavior of 94 people, a total of about 35 years of interaction data, as described in Eagle and Pentland [5].

The second part of the socioscope consists of electronic badges that record the wearers' location (with 2 m typical accuracy), ambient audio, and upper body movement via a 2-D accelerometer, as described in Gips and Pentland [6]. This badge platform provides more fine-grained data than the smart phone platform.