receiver #14 connecting to the future

Organizational rhythms - the search for the patterns of the aggregate

Nathan Eagle & Alex (Sandy) Pentland

Dr Nathan Eagle is a Visiting Scientist at the MIT Media Lab, where he recently completed his PhD under Professor Alex (Sandy) Pentland. Professor Pentland is the Toshiba Professor of Media Arts and Sciences at the MIT Media Laboratory and Director of the Human Dynamics research group. Their research focuses on projects that span a variety of disciplines from appropriate technology to artificial intelligence. This contribution to *receiver* comes from Eagle's dissertation on machine perception and learning of complex social systems, which explores the intersections of mobile phones, machine learning, and organizational behavior.

http://reality.media.mit.edu/
The Reality Mining project

Today's mobile phones are capable of sensing the world around them. Using information such as cell tower IDs or proximate Bluetooth devices, it is possible to get a depiction of an individual's current context. When this data is captured over extended periods of time, it can be used to generate a predictive model of the user's life. And if we extend this further, capturing mobile phone data across the individuals within an organization can give us unprecedented insight into the large-scale dynamics of collective human behavior. Furthermore, a dataset providing the proximity patterns and relationships within large groups of people has implications within the computational epidemiology communities, and may help build more accurate models of airborne pathogen dissemination, as well as other more innocuous contagions, such as the flow of information around the water cooler.

In the Reality Mining project, we distributed 100 context-logging phones to people working at MIT and collected almost 500,000 hours of continuous human behavioral data. We showed that Bluetooth-enabled mobile phones can be used to discover a great deal about the user's context and relationships. In this paper we will focus on extending this base of user modeling to explore modeling complex social systems. We will provide several illustrative examples of how this data can be used to learn more about both team and organizational dynamics.

Team dynamics

By continuously scanning for Bluetooth devices and logging the people proximate to an individual, we are able to quantify a variety of properties about the individual's work group. Although most research in networks assumes a static topology, proximity network data is extremely dynamic and sparse. We will compare aggregate statistics between two different research groups at the Media Lab in an attempt to gain insight into fundamental characteristics of the research groups themselves.

While each research group at the Media Lab is centralized around a faculty director, the proximity networks are not reflective of this static organizational structure. In many instances, the proximity network's degree of distribution is indicative of a hub-and-spoke formation, however the roles that are played within this structure are not static. Individuals that are hubs during one period of time fluidly exchange places with other team members on the periphery of the proximity network. This type of dynamic may be characteristic of the underlying nature of research groups at the Media Lab. As deadlines approach for specific individuals, they begin to spend more time in the Media Lab and increasingly rely on support from the rest of the group. Upon completion of a project, they resume their normal routines and can provide similar support to others. This pattern of behavior has been shown to vanish when the entire group (or organization) is working towards the same deadline.

Organizational modeling & rhythms

Organizations have been considered microcosms of society, each with their own cultures and values. Similar to society, organizational behavior often shows recurrent patterns despite being the sum of the idiosyncratic behavior of individuals. We have explored the ramifications of the ability to quantify the dynamics of behavior in organizations in response to both external (stock market performance, a Red Sox World Series victory) and internal (deadlines, reorganization) stimuli.

The need for better organizational models

Over time, the Reality Mining data can be fine-tuned for studying, tracking and - perhaps most importantly - predicting the dynamics of a particular social network. Recently, the CEO of a multimillion-dollar manufacturing company became interested in the technology as a means to quantify workplace collaboration. Like many organizations, his company suffers from the "silo syndrome" – people from different departments tend to keep to themselves, leading to inefficiencies and missed opportunities.

To address that problem, the CEO had tried a number of initiatives – for example, having dozens of people from marketing switch offices with their counterparts in engineering – however, he realized that he did not have a clear understanding about how those changes actually affected the organization as a whole. The company had conducted extensive surveys, but the data only provided a snapshot of the current social network. Instead, the CEO said he wanted "footprints in the sand" to understand the dynamics of the network topology. With such information, he could then determine signature effects indicative of a successful (or unsuccessful) initiative. Augmenting his existing data collection methods with data from unobtrusive mobile phones could help capture exactly those types of continuous dynamics.

Organizational rhythms: patterns in aggregate behavior

With the rapid technology adoption of mobile phones comes an opportunity unobtrusively to collect continuous data on human behavior. The very nature of mobile phones makes them an ideal vehicle to study both individuals and organizations: people habitually carry a mobile phone with them and use it as a medium through which to do much of their communication. Now that handset manufacturers are opening their platforms to developers, standard mobile phones can be harnessed as networked wearable sensors. The information available from today's phones includes the user's location (celltower ID), people nearby (repeated Bluetooth scans), communication (call and SMS logs), as well as application usage and phone status (idle, charging, etc). However, because the phones themselves are networked, their functionality transcends merely a logging device that augments social surveys. Rather, phones can start being used as a means of social network intervention - supplying introductions between two proximate people who don't know each other, but probably should.

We made use of the data our networked phones supplied us with in the following way: from the proximity data, we first extract adjacencies for each scan, and then infer ongoing proximity so as to annotate each edge with an initiation and termination time. At the time of writing, we are not aware of any other network data set with such a large amount of temporal data or one with such fine granularity. Thus, this data set provides a rich opportunity to explore both temporal dynamics and the quality of our analytic tools. Here, we focus on the latter.

Broad-scaled dynamics

Given the initiation and termination times of each edge, the difference gives the edge's temporal persistence. Here, we restrict our analysis to the 24,092 undirected subject-subject adjacencies between 01 October 2004 and 31 October 2004 of the 66 subjects who work in the same building at MIT. Unsurprisingly, the weekly and monthly distributions we found are largely similar, with minor variations far out in the tail, ie, for persistence greater than 400 minutes, suggesting a system largely in equilibrium. The typical duration of proximity is relatively small, while there are several edges which persist for more than 1440 minutes (24 hours). The broadness of these distributions demonstrates that the network's topology evolves at a wide variety of time-scales, and we believe such a distribution may be typical of many real world networks. Indeed, although not well-motivated for our data set, a recent study found a power-law distribution of relationship lengths when a relationship is defined to continue from the time of first interaction to the time of last interaction. During October 2004, the seventy-five Media Lab subjects had been working towards the annual visit of the Laboratory's sponsors. Preparation for the upcoming events typically consumes most people's free time and schedules shift dramatically to meet deadlines and project goals. A significant fraction of the community tends to spend much of the night in the Lab finishing up last minute details just before the event. We are beginning to uncover and model how the aggregate work cycles expand in reaction to these types of deadlines; aggregate behavior during the period leading up to this deadline deviates from the norm with proximity events that have longer duration. The number of links in the Media Lab proximity network also remained significantly greater than zero during the third week of October and in early December, representing preparation for the large Media Lab sponsor event and MIT's finals week. We have found two fundamental frequencies, the strongest, not surprisingly, being at 24 hours (1 day), and the second at 168 hours (7 days).

We can generate inferences about the most likely time and place to find specific colleagues or friends in a given context, and this ability to reliably instigate casual meetings could be of significant value in the workplace. We have to remember, however, that the ability to predict people's movements can be put to less savory uses. Careful consideration must be given to these possibilities before providing access to such data.

The future organizational modeling: simulation

While we have developed technology that can be immediately applied to the problem of the CEO interested in better quantifying the ramifications of a reorganization, it may be possible to predict how a proposed change affects an organization before any action is taken. After logging extensive amounts of data on the interactions between employees, inference can be preformed to generate the likely set of actions of an individual given a specific situation. When this inference is performed in aggregate, it may be possible to predict the outcome of certain initiatives, such as moving half of the marketing department into offices within the engineering section of the company. While impossible today, the Reality Mining data could lead to the type of analysis that would enable the CEO to run various "what if" virtual experiments to determine her most effective options before implementing a personnel shift.

Although there is still much work to do before an accurate simulation of organizations can be realized, capturing this rich, continuous interaction data is a critical first step towards this goal.

http://web.media.mit.edu/~nathan/
Nathan Eagle's site

This article was written for *receiver*
contact: nathan@media.mit.edu