

CT 1.5.2

Mob.Media: A Mobile Phone Platform for Computational Social Science

A. Madan¹ and Alex (Sandy) Pentland¹

¹MIT Media Laboratory, 20 Ames St., Cambridge, MA 02139, USA

Introduction

As citizens of the information age, we are leaving pervasive digital traces of our idiosyncratic behavior— in emails, on online social networking sites, in mobile phone call logs, in ATM machines, in metropolitan train systems. Researchers in the emerging field of Computational Social Science (CSS) are attempting to create quantitative models of large-scale human social systems from these digital traces, as well as address key questions like privacy, data-ownership and sharing of interaction data. Recent work has focused on online techniques, e.g. mining corporate emails and crawling social networking sites to build maps of social interactions. In contrast, mass-market mobile phones allow sensing and analysis of the building blocks of social interactions – those that occur face-to-face. Eagle & Pentland (2005) and later Dong & Pentland (2006) used Bluetooth proximity information to recognize social patterns in daily activity infer relationships, identify socially relevant locations and model organizational rhythms.

In this paper we describe a mobile platform accessible to CSS researchers, devised by augmenting a commodity mobile phone with the ability to mine face-to-face interactions and user activity patterns over long-term durations. As an example of its potential, we also describe in-progress work to study social influence, viral diffusion of media, and the correlation between social ties and privacy.

Mobile Phone Platform

The Mob.Media framework runs on the Apple iPhone, noted by TIME magazine as the gadget of the year 2007. The mobile device has a 620 Mhz ARM processor with 128 MB of RAM (with approx. 20MB available for 3rd party applications). Unusual for mass-market phones, the device supports a hardware floating-point unit (FPU) and BSD-based operating system, which make it remarkably useful for real-time feature extraction and pattern recognition tasks. Network connectivity is supported via EDGE (GSM), WLAN 802.11 and Bluetooth. The device features a 3-axis accelerometer, full audio capabilities, a multi-touch user interface and OpenGL ES graphics, with hardware acceleration.

We have added the following software functionality to the device so that it can be used to measure and analyze social interactions:

- Native code to periodically scan cell-tower ids, WLAN ids, and corresponding signal strength. These can be used to infer homogeneity and entropy of location and proximity patterns, e.g. is there a cluster of users who tend to visit similar locations frequently?
- Native code to capture call logs and SMS logs on the device. We expect that the temporal and frequency features in these logs can be used to infer strength of social ties, e.g. how often do certain people call on weekends?

- Native code to periodically sample microphone audio and the on-board 3-axis accelerometer. These can be used as context and activity features, e.g. is the user at their office computer or walking in a public area?
- Cross-compiled native code for real-time inference on the mobile device, e.g. libsvm (for Support Vector Machines), R (statistical inference package), etc.
- Scripts and UI applications to anonymize data on the phone, before it leaves the user's physical control, and then periodically upload it to our secure servers. This framework provides participants with control over privacy and anonymity of their data, while still using data for compelling research applications.

The Mob.Media ScanApps and PatternRec packages can be installed on any iPhone that supports native applications, by visiting our install feed via Installer.app (below). We are also working on J2ME version of our application which will allow capturing equivalent social interaction data on Nokia smart phones.

Social Influence and Viral Propagation

Models of the diffusion of ideas, innovations, recommendations and media have been extensively studied in social science literature. In recent work, this problem has been approached either with strong theoretical assumptions about diffusion phenomena, or as empirical evaluations based on surveys, emails or online recommendation networks (e.g. Kossinets 2006, Leskovic 2007). We are proposing new methodology for studying social influence within the experimental context of media propagation, by using the Mob.Media framework to map out patterns of social interaction and information flow.. To accomplish this goal we have extended the Mob.Media platform with our own (mobile) streaming indie music service. Several hundred tracks (and increasing daily) have been contributed by artists from GarageBand.com and MySpace.com. Study participants can listen to music, rate and manage tracks, and forward it to other users through the iPhone interface. We expect, for instance, that it will be possible to identify target candidates for maximal diffusion, a.k.a. the 'mavens' in the community (Kempe 2003), from Mob.Media proximity observations and call logs. The install feed for Mob.Media packages may be found at: <http://jellyfish.media.mit.edu/iPhone>

Acknowledgements

The authors would like to acknowledge the engineering contributions of Nevada Sanchez, Rita Chen, Thatcher Clay, Michael Snively and Iolanthe Chronis.

References

- W. Dong, A. Pentland (2007): Modelling Influence Between Experts. In AI for Human Computing, pp. 170-189, LNAI 4451, Springer
- N. Eagle, A. Pentland (2005): Reality Mining: Sensing Complex Social Systems. Journal of Personal and Ubiquitous Computing, pp. 255-268
- G. Kossinets and D. J. Watts (2006): Empirical Analysis of an Evolving Social Network. Science, Vol 311, pp 88.
- D. Kempe, J. Kleinberg and E. Tardos (2003): Maximizing the spread of influence through a social network. 9th Intl. Conf. on Knowledge Discovery & Data Mining
- J. Leskovic, L. Adamic, and B. Huberman (2007): The Dynamics of Viral Marketing. ACM Transactions on the Web.