

To Appear: IEEE Computer – Outlook Special Issue, 2012

Society's Nervous System:

Building Effective Government, Energy, and Public Health Systems

Alex Pentland, Room E15-387, The Media Lab, 20 Ames St, Cambridge MA 02139

pentland@mit.edu, <http://media.mit.edu/~pentland>, 617-253-0648

Abstract: Drawing on a unique, multi-year collaboration with the heads of major IT, wireless, hardware, health, and financial firms, as well as the heads of American, EU, and other regulatory organizations, and a variety of NGOs [1,2], I describe the potential for pervasive and mobile sensing and computing over the next decade, and the challenges that will have to be faced in order to realize this potential.

Building Effective Systems for Society

How can we design the systems and networks needed for a healthy, sustainable society? This problem is not new: in the 1800's, the industrial revolution spurred the growth of cities and created huge health problems. The solution then was to engineer centralized networks that delivered clean water and safe food, removed waste, provided energy, facilitated transportation, and provided access to centralized healthcare, police, and educational services.

These century-old solutions are now increasingly obsolete and unsustainable. Today's social structures are not designed as integrated systems and do not take advantage of new digital feedback technologies that would allow them to be dynamic and responsive.

We need a radical rethinking of societies' systems. Instead of separate systems for each function, i.e., water, food, waste, transport, education, energy, etc., we must consider them holistically. Instead of thinking about access and distribution systems only, we must also take into account complex interactions and reactions so that we can design dynamic, networked systems that are self-regulating. In short, to build a sustainable society for the future, we must create a nervous system for humanity that maintains the stability of our societies' systems throughout the globe.

The central idea is that we must reinvent societies' systems within a control framework: sensing the situation, then combining these observations with models of demand and dynamic reaction, and finally using the resulting predictions to control the system. Surprisingly, much of the sensing and many of the required control elements are already in place. What is missing, though, are the dynamic models of demand and reaction along with an architecture that guarantees safety, stability and efficiency.

The models required must describe *human* demand and reactions, since humans are at the core of all of these systems. Consequently, the necessary observations are observations of individual behavior, which means that there will be an exponential growth in data about human behavior. The importance of data about human behavior is made clear by this quote from

Meglana Kuneva, European Consumer Commissioner: “Personal data is the new oil of the internet and the new currency of the digital world.”

One consequence of this growth in behavior sensing data is that the architecture of future systems must also guarantee privacy and fair treatment for all potential participants. As Jon Leibowitz, Chairman of the U.S. Federal Trade Commission says, ‘Privacy must be an integral part of any future system.’

Realizing that using pervasive and mobile sensing to reinvent our society’s infrastructure would require ‘buy in’ from government, citizenry, and companies, I persuaded the World Economic Forum to begin a multi-year conversation [1] with the heads of several major IT, wireless, infrastructure, and financial firms, as well as the heads of American, EU, and other regulatory organizations, and a variety of NGOs. The first publication emerging from this discussion can be found at [2]; this current paper describes the background and what I see as result of that conversation, along with insights about the path forward given the emerging consensus on a ‘new deal on data’ proposed in [1].

In the following sections, I will describe what is expected to be the future evolution of the sensing, modeling, and overall design of these new active systems. The discussion on sensing covers not just data acquisition, but also privacy and data ownership, significant issues that must be addressed in building these practical systems of the future. The modeling discussion will argue that it is important to choose techniques of mathematical modeling whose structure closely follows the organization of the actual physical elements of the given system. The system design section will argue that we should import ideas from economics to help define the overall system performance criteria. I will not spend time discussing final control actions in these future systems, because I believe that those will likely depend on very fine-grain details of these future systems.

1. Pervasive sensing

Worldwide, there are now almost five billion mobile phone subscribers and every day millions of new subscribers get phones. For the first time in history, the majority of humanity is linked and has a voice. The most important change, however, will come from the fact that these same mobile phones are location-aware sensor platforms and their wireless networks support sensors in cars, buses, and homes. As a consequence, our mobile wireless infrastructure can be ‘reality mined’ in order to understand the patterns of human behavior, monitor our environments, and plan the development of our society [3].

This ‘reality mining’ functionality is mostly latent at this point, but already these devices are being used to measure population flows into cities and slums, to map the movement of populations during emergencies, to identify neighborhoods where social services are inadequate, and to manage automobile traffic congestion [2]. The ability of mobile phone networks to identify unusual patterns of movement and communication is also beginning to be used by public health officials and disaster relief teams to scan outbreaks of diseases like SARS and emergencies such as tidal waves.

Like some world-spanning living organism, wireless traffic systems, security sensors, and especially mobile telephone networks are combining to become intelligent, reactive systems with sensors serving as their eyes and ears. Moreover, the evolution of this nervous system will continue at a quickening speed because of the exponential progress in computing and communication technologies as well as basic economics. Networks will become faster, devices

will have more sensors, and techniques for modeling human behavior will become more accurate and detailed.

Reality mining of the `digital breadcrumbs' left behind as we go about our daily lives offers potential for creating remarkable, second-by-second models of group dynamics and reactions over extended periods of time, providing both dynamic structural and content information. The key is to harness these streams of personal data and use them to create and drive dynamic models of aggregate human behavior [4].


Perhaps the simplest example is the analysis of automobile traffic congestion by using the global positioning system (GPS) data collected from the mobile telephones carried by the automobile drivers. These data provide minute-by-minute updates on traffic flow, allowing for more accurate predictions of driving time. Congestion patterns can be predicted in advance, and traffic jams detected before they become serious.

Similar to using mobility data to understand traffic within a city, we can also use data from RFID name badges, mobile phone Bluetooth data, phone call logs, and email records to better understand the "information traffic" within an organization. Analysis of these digital traces has already allowed us a new level of insight into the problems of industry and government, including building customer relationships, resource management, transportation, and even employee health [5]. In our experiments mapping corporation information flows we typically find that the *pattern* of information transfer --- face-to-face, email, etc., independent of the *content* of the information --- accounts for almost half of the performance variation within a corporation. Using these `information maps' to re-engineer corporate procedures has often generated very significant increases in performance [5,6,7,8,9].

At a larger scale, commercial operations and government services all currently rely on demographic and survey data to guide them. Such data can quickly become out of date and, of course, good demographic data simply do not exist in many parts of the world. The fact that mobile phones have GPS means that we can leap beyond demographics directly to measuring behavior. Where do people eat? Work? Hang out? What routes do they travel?

Analysis of mobility patterns allows discovery of different varieties of behavior patterns within a city, and the stratification of the population into subgroups with different types of behaviors. Figure 1a shows movement patterns with popular destinations coded by the different subgroups that populate these destinations, where the subgroups are defined their *behaviors* – what types of restaurants they visit, what sort of entertainment they like, and so forth. Figure 1b shows that the mixing between these different behavior groups is surprisingly small.

Knowledge of the behavior preferences of different subgroups provides a far more accurate picture of their preferences and risks than standard demographics. Such a stratification of the population into different `behavior demographics' typically provides between 5 and 10 *times* the accuracy of traditional demographics at tasks such as predicting risk for diseases of behavior (e.g., diabetes), financial risk profile, consumer preferences and political views [6]. The ability to stratify people based on their behavior profile has already shown great potential to improve public health, urban planning, public education, and government policy [2,6] by providing a more accurate and in-depth picture of the citizenry . Even greater accuracy at modeling human behavior patterns can be obtained by adding credit card data, healthcare information, and similar `digital breadcrumbs.' In short, we now have the capacity to collect and analyze data about people with a breadth and depth that was previously inconceivable.



a

b

Figure 1: Reality mining of data from GPS mobile phones. (a): Movement patterns, measured from GPS mobile phones; (b): Segmentation of the population into groups with differing behaviors [6].

1.1 Data ownership and privacy

Perhaps the greatest challenge posed by this new ability to sense the pulse of humanity is creating a “new deal” around questions of privacy and data ownership [1]. Advances in analysis of network data must be approached in tandem with an understanding how to create value for the producers and owners of the data, while at the same time protecting the public good.

What must be avoided is either the retreat into secrecy, so that these data become the exclusive domain of private companies and remain inaccessible to the Common Good, or the development of a “big brother” model, with government using the data but denying the public the ability to investigate or critique its conclusions. Neither scenario will serve the long-term public interest in having a transparent and efficient society.

Similarly, the use of anonymous data should be enforced, and analysis at the group level should be preferred over that at the individual level. Robust models of collaboration and data sharing need to be developed; guarding both the privacy of consumers as well as corporations’ legitimate competitive interests are vital here.

The vast majority of this personal data is collected by private organizations -- location patterns, financial transactions, public transportation, phone and Internet communications, and so on. As a consequence, companies will have a key role in this new deal for privacy and ownership, and so either extensive government regulation or market mechanisms with auditing will be needed in order to entice owners to give up the data they hold. The obvious choice is to design market mechanisms that have sufficient accountability to guarantee individual and societal safety.

1.2 The new deal on data

The first step toward creating such information markets is to give people some ownership of the data that is about them. Just as with financial and commodity markets, the first step toward a healthy market is creation of an asset class such as land rights. This is why the World Economic Forum publication is subtitled 'Emergence of a New Asset Class.' [2]

The simplest approach to defining what it means to "own your own data" is to go back to Old English Common Law for the three basic tenets of ownership, which are the rights of possession, use, and disposal:

1. You have a right to *possess* your data. Companies should adopt the role of a Swiss bank account for your data. You open an account (anonymously, if possible), and you can remove your data whenever you'd like.
2. You, the data owner, must have full control over the *use* of your data. If you're not happy with the way a company uses your data, you can remove it. Everything must be opt-in, and not only clearly explained in plain language, but with regular reminders that you have the option to opt out.
3. You have a right to *dispose* or *distribute* your data. If you want to destroy it or remove it and redeploy it elsewhere, it is your call.

Ownership by individuals needs to be balanced by the legitimate need for corporations and governments to use personal data -- credit card numbers, home addresses, etc. --to run their day-to-day operations. The "new deal on data" therefore gives individuals the right to own and control a *copy* of the data about them rather than ownership of corporations' internal data as long as that data is required for legitimate operations. The private ownership of complete copies of personal data is sufficient to create a liquid, dynamic new asset class.

1.3 Enforcing the New Deal on Data

The issue of enforcement is not simply authenticating the identity of an individual, but rather validating whole series of "claims" and "privileges" an individual, institution, or device may make that give them access and use of valued services and resources. As more and more business, financial, civic and governmental services use personal data, the integrity and interoperability of a global authentication and "claims" infrastructure will be paramount.

Since there can be no single authority that can micro-manage all transactions, failures, and attacks, such a global infrastructure -- like the Internet itself -- will have to be highly distributed and user-centric to assure rapid innovation, containment and self-correction. The Trust Networks that certify the Open Identity Exchange (OIX) protocols are an example of such a distributed authority and are accepted not only by major companies but also many governments as well.

Similar to the current OIX ecosystem, such a 'Trust Network' will need to continuously monitor, flag and contain fraudulent and deceptive behaviors. This will require not only innovations in software to track and audit the behaviors of actors in transaction networks, but innovations in policy and contract law so that there can be simple, fair and effective enforcement and remedies. The Law Lab at the Harvard Berkman Center [28, 29] has developed the prototype of such a trust network and together with my research group, in support of the World Economic Forum Rethinking Personal Data Project, we are now testing how these trust networks perform 'in the wild' [2].

As authentication and claims processing becomes fully automated and digital, there is no reason that legal and administrative processes should not also become automated and more transparent, accountable, and efficient. In short, digital technologies should not only expedite a global infrastructure for highly reliable and innovative claims processing, but also eliminate much of the uncertainty, cost, and friction of legal and regulatory oversight. Were this to happen, transaction costs would be greatly diminished and new forms of liquidity and innovative businesses would arise. Given that *verifiable* trust reduces transaction risks and builds customer loyalty, it will be in the economic interest of those offering future online and mobile services to build their brand as verifiable trusted stewards of personal data.

2. Models of Society

The first part of this new 'nervous system' -- sensing -- seems to be evolving quite well due to the economic advantage it offers the world's citizenry. But what about the second part, that is, the dynamic models of demand and reaction? Unfortunately the vast majority of research on the human condition has relied on single-shot, self-report data on relationships: a yearly census, public polls, focus groups, and the like. As a consequence, the science of modeling human dynamics is in its infancy.

Because this science is just developing, no final answer can be given about what specific types of mathematical modeling are required. I will therefore describe the approaches that I have seen to be the most promising, in the hope that it will inspire others to develop even better modeling techniques.

Three key principles that the models should adhere to include: (1) to be fully dynamic, so that they can be used in a control system, (2) to have the structure of the model reflect the structure of the phenomenon (in most cases, the structure of human groups); and (3) that the model parameters should be driven almost entirely by real-world observations. Perhaps the most controversial implication of these three principles is that the model should contain both agent models derived from observation data along with the influences between the agents, information also derived from observation data.

Currently, most models used to model human behavior are either agent models built largely from psychological data, or machine learning models (such as SVMs) that do not incorporate the networked structure of human society. In my experience, models that capture the structure of our human social fabric by incorporating both agent model and influence between the agents have consistently performed better at modeling human behavior than either traditional agent models or traditional machine learning algorithms.

2.1 Social Influence

While much progress has been made in constructing models of individual human behavior, models of the influence that we exert on each other remain much less developed. For more than 50 years, social scientists and psychologists have been interested in analyzing and understanding *who influences whom* in a social system, Much of this past social science research has focused on smaller systems, such as that within a group discussion process [7]. Influence is also particularly interesting in the context of leadership and group dynamics, where the influence between one another in these contexts has been recognized as a significant factor of group performance [8].

From the signal processing and modeling point of view, however, the difficult question remains as to how to define and model the concept of *influence* in a formal, mathematical way. Adding to the complication is the fact that influence between individuals is often not directly observable, and only individual-level behavioral signals are generally available [9,10]. Therefore, the challenges are not only to define *influence*, but also to infer *influence* from individual observations and individual signals.

A line of research that we have pioneered, known as *influence modeling*, is focused on modeling such social influence mathematically. This type of influence model addresses two fundamental challenges:

- The influence model mathematically defines “influence” and how influence changes, and the “influence” learned by the influence model is tightly connected with the sociological meaning of influence.
- The influence model enables researchers to infer interactions and influence dynamics by using only time series signals from individual observations.

The influence model was introduced in [13, 16], and was developed around the idea of modeling influence using conditional probability, using an inference scheme based on optimization. It is described in more detail in the Appendix. Other related models exist, and some have similar properties [11,12,15,16,19, 20, 21,22,23,24,25,26,27]. The core concept here is that the structure of the model matches the structure of the social phenomena and that the model and its dynamics are driven by real-world observational data.

2.2 Social Influence in Human Society

The key question to ask for any model of influence is whether the derived influence parameters (describe by an influence matrix \mathbf{R} in this case) accurately represent the concept of influence in human interactions. In other words, does the concept of influence, as defined in this article, have practical and sociological meanings? The first example of using this model was with conversation data from the wearable sociometric badges on 23 individuals, where it was discovered that the influence strength between individuals learned by influence model correlates extremely well with individual centrality in their social networks (with correlation $r = 0.92$, $p < 0.0004$)[15].

Our influence model has since been applied to many different human interaction problems. For instance, researchers have used the influence model in understanding the functional role (follower, orienteer, giver, seeker, etc.) of each individual in a mission survival group discussion dataset [17]. By using the reality mining [10] cellphone sensor data from 80 MIT members as observations, and constraining the latent space of each individual to be binary “work” and “home”, researchers found that the influence matrix gleaned from this data matches well with the organizational relationship between individuals [13]. This is intuitive as students’ schedules are likely to be influenced by close colleagues’ working schedules.

Related works [14] often use an influence model as a modeling tool for the evolution of social systems. The influence model has also been extended to non-stationary situations in which the influence matrix itself changes [18]. Combined, these experiments strongly suggest that the influence matrix, defined as the conditional dependence on states of other entities in the model, is an important measure for the influence of the individuals in real social interactions.

3. Overall System Design

Section 1 argued that the required observation systems are developing nicely and that regulatory and industrial leaders have made at least a first pass at addressing the problem of data availability and minimizing the risk of harm to individuals. Section 2 argued that progress is being made in developing social system modeling techniques that are fully dynamic and stable and whose parameters can be efficiently set from available observations. I now turn to the problem of overall system design. Unfortunately, we don't yet understand fully how to design dynamic social systems so that they can accomplish all of the lofty goals we now set before them. At its core, this design problem requires understanding how the actions of many individuals can be reliably combined to produce the desired outcomes.

3.1 A Template for the Design of Social Network Systems

Fortunately, the field of economics provides a useful template for answering this design question: *social efficiency*, which is also known as the "invisible hand" described by Adam Smith that optimally distributes resources throughout society. In socially efficient systems, when I make a profit for myself, I also make a profit for the entire society. For an ideally efficient society, the coupling between individual and societal benefit is perfect. Social efficiency also implies the reverse condition: what is bad for the individual is likewise bad for society. In situations where most people are well-off, the fitness of a society can be measured by the conditions of its poorest and most vulnerable members. Given the well-known shortcomings of human nature, this design requirement that social network systems be socially efficient is very desirable.

Applying the principle of social efficiency to social network systems means that the exchange of information between individuals, or between the system and the individual, must not only reliably provide value to the individual but also efficiently add value to the whole system. One way to accomplish this is through an open market and in recent centuries this solution has dominated our thinking. The result of our reliance on open market approaches has been systems with ever greater transparency but also ever greater concerns about the 'end of privacy.'

One significant problem with the open market solution for social efficiency is that we are not all created equal. Some 'people' are actually corporations or governments with enormous analytic capabilities and in open markets they will consistently outwit individuals less richly endowed with computational resources. This inequality in computational resources is quickly becoming a major source of inequality in society and is also a major source of cybercrime and worries about cyberwarfare.

It is fortunate that an approach that relies on strong control of personal information with limited transparency -- such as the proposed New Deal on Data -- can also produce a socially efficient system [30]. If our ability to view information within the social network is limited to only those we deal with directly, and we have the right to share information only in fair exchanges, then the scope for collusion and deception is sharply limited and a stable, fair information economy emerges. This line of reasoning provides significant theoretical backing for the 'new deal on data' as a method to supersede the current failing generation of privacy mechanisms [1,2] and offers a way for society to utilize information about where people are, what they do, as well as their preferences and characteristics, while still maintaining strictly controlled individual risk and individual control of personal data [2,30].

One can grasp the intuition behind this surprising result by considering the typical urban experience. As you go through your daily life you have familiar, routine interactions with many people, although most likely you don't know their friends, the other people they deal with, and so forth: they are 'familiar strangers.' The fact that you know about them but not their network means that collusion against them becomes much more difficult. In a similar way, social network systems can be designed to be optimized and fair without exposing their members to increased risk.

The "open market" and "strong personal control" models are but two approaches to social efficiency and there are almost certainly other ways to achieve social efficiency, including blends of these two models. Consequently, it is important to discover these new paths to social efficiency and use them to engineer social network systems. In the meantime, we need to understand how to engineer systems that capitalize on the mathematical results that we already have in hand.

3.2 Incentives

While significant progress is being made in the areas of building trust networks to regulate social network systems and how to optimize these sorts of complex human-machine systems [2], a missing part of the puzzle is how to get the humans in these systems to participate in the plan. Designing an 'optimal' system is useless unless it fits our human natures, because otherwise we run the risk that people will either ignore or misuse it. To integrate humans into a system requires new predictive theories of human decision making [3] along with a more powerful theory of incentive mechanism design.

Current financial mechanism design theory is quite useful for devising successful financial incentives. For example, optimizing on-line auctions and the like have generated billions of dollars for corporations such as Google. Current theory, however, is limited to financial incentives, whereas humans respond as well or better to social incentives. Moreover, incentives in markets can operate very differently than incentives in networks, largely because the connectivity of social network systems limits visibility and opportunity. We need a theory of 'mechanism design in networks' that includes social as well as financial capital [2].

An example of one such mechanism is the one we used to win the 40th anniversary of the Internet Grand Challenge (also known as the Red Balloon challenge). The goal of the challenge was to discover the location of ten red balloons placed somewhere in the continental United States, in order to win a \$40,000 prize [31, 32]. Over 4,000 teams competed and every team except ours used incentive mechanisms tailored to encourage *individuals* to report balloons. Instead, we created an incentive mechanism that both encouraged reporting balloons, but even more importantly, encouraged people to leverage their social network to recruit additional participants. By designing a mechanism to leverage the social network rather than individuals, we were able to recruit a larger and more motivated group of people searching for balloons.

4. Conclusions

Revolutionary new measurement tools provided by mobile telephones and other digital infrastructures are providing us with a God's eye view of ourselves. For the first time, we can precisely map the behavior of large numbers of people as they go about their daily lives. For society, the hope is that we can use this new in-depth understanding of individual behavior to increase the efficiency and responsiveness of industries and governments. For individuals, the

attraction is the possibility of a world where everything is arranged for your convenience—your health checkup is magically scheduled just as you begin to get sick, the bus comes just as you get to the bus stop, and there is never a line of waiting people at city hall.

As these new abilities become refined by the use of more sophisticated statistical models and sensor capabilities, we could well see the creation of a quantitative, predictive science of human organizations and human society. At the same time, these new tools have the potential to make George Orwell's vision of an all-controlling state into a reality. As a consequence, we need to think carefully about the growth and increasingly broad usage of personal data to drive societies systems, and particularly about the safety, stability, and fairness of their design.

Towards this end, I have proposed the New Deal on Data that ensures accountability and data ownership, and I have successfully advocated its adoption to regulatory, industrial, and NGO leaders. Current legal statutes are lagging far behind our ability to collect and process data about people; clearly our notions of privacy and ownership of data need to evolve in order to adapt to these new capabilities. Perhaps the first step is to give people ownership of their data, creating what economists know as a "fair market" for the information that will drive this new social nervous system

I have also proposed the criterion of social efficiency both as a design goal and as a metric for the design of social network systems. Information is increasingly the wealth of our civilization and it is time to draw on our rich legacy of thinking about the distribution and regulation of financial wealth in order to be able to build social networks systems that live up to our aspirations.

If we can successfully address these challenges, then we will see current systems evolve into an effective nervous system for our society, one that could repay our investment many-fold in terms of better civic services, a greener way of life, and a safer, more healthy population.

References

- [1] A. Pentland, Reality Mining of Mobile Communications: Towards a New Deal on Data, *Global Information Technology Report 2008-2009*, World Economic Forum, pp. 75-80
- [2] Personal Data: The Emergence of a New Asset Class. *World Economic Forum*, 2011, http://www3.weforum.org/docs/WEF_ITTC_PersonalDataNewAsset_Report_2011.pdf
- [3] Technology Review. 2008. "10 Emerging Technologies That Will Change the World." Available at <http://www.technologyreview.com/article/13060/>.
- [4] D. Lazer, A. Pentland, et al, Computational Social Science, *Science*, Vol. 323. no. 5915, pp. 721 – 723 February 2009,
- [5] A. Pentland, *Honest Signals: How They Shape our World*. Cambridge, MA: MIT Press. 2008
- [6] Sense Networks. 2008. "Technology tab" at <http://www.sensenetworks.com>.
- [7] E. Katz and P.F. Lazarsfeld. *Personal influence*. Free Pr., 1955.

- [8] A.W. Woolley, C.F. Chabris, A. Pentland, N. Hashmi, and T.W. Malone. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686, 2010.
- [9] Olguín Olguín, D., Waber, B., Kim, T., Mohan, A., Ara, K., and Pentland, A. Sensible Organizations: Technology and Methodology for Automatically Measuring Organizational Behavior, *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*. Vol. 39, No. 1, pp. 43-55. February, 2009
- [10] N. Eagle and A. Pentland. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4):255–268, 2006.
- [11] C.M. Bishop et al. *Pattern recognition and machine learning*. Springer New York., 2006.
- [12] M. Brand, N. Oliver, and A. Pentland. Coupled hidden Markov models for complex action recognition. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 994–999, 1997.
- [13] W. Dong and A. Pentland. Modeling influence between experts. *Proceedings of Artificial Intelligence for Human Computing' 2007*. pp.170-189 . Also *Lecture Notes in Computer Science*, 4451:170, 2007.
- [14] A.Madan, M. Cebrian, S.Moturu, K. Farrahi, A. Pentland, Sensing the `Health State' of our Society, *IEEE Pervasive Computing, in press*
- [15] C. Asavathiratham, S. Roy, B. Lesieutre, G. Verghese, The Influence Model, *IEEE Control Systems*, Vol 21 No. 6, pp. 52-56
- [16] S. Basu, T. Choudhury, B. Clarkson, A. Pentland, et al. Learning human interactions with the influence model. *MIT Media Laboratory Technical Note*, 2001.
- [17] W. Dong, B. Lepri, A. Cappelletti, A.S. Pentland, F. Pianesi, and M. Zancanaro. Using the influence model to recognize functional roles in meetings. In *Proceedings of the 9th international conference on Multimodal interfaces*, pages 271–278. ACM, 2007.
- [18] W. Pan, M. Cebrian, W. Dong, T. Kim, and A. Pentland. Modeling Dynamical Influence in Human nteraction Patterns. *Arxiv preprint arXiv:1009.0240*, 2010.
- [19] D. Zhang, D. Gatica-Perez, S. Bengio, and D. Roy. Learning Influence Among Interacting Markov Chains. 2005.
- [20] E. B. Fox, E. B. Sudderth, M. I. Jordan, and A. S.Willsky. Nonparametric Bayesian learning of switching linear dynamical systems. In *Neural Information Processing Systems 21*. MIT Press, 2009.

- [21] W. Pan and L. Torresani. Unsupervised hierarchical modeling of locomotion styles. In *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM New York, NY, USA, 2009.
- [22] C. Castellano and V. Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591, 2009.
- [23] J.P. Onnela, J. Saramaki, J. Hyvonen, G. Szabó, D. Lazer, K. Kaski, J. Kertesz, and A.L. Barabasi. Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences*, 104(18):7332, 2007.
- [24] F. Guo, S. Hanneke, W. Fu, and E.P. Xing. Recovering temporally rewiring networks: A model-based approach. In *Proceedings of the 24th international conference on Machine learning*, page 328. ACM, 2007.
- [25] A. Ahmed and E.P. Xing. Recovering time-varying networks of dependencies in social and biological studies. *Proceedings of the National Academy of Sciences*, 106(29):11878, 2009.
- [26] M. Gomez Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1019–1028. ACM, 2010. 19
- [27] M.I. Jordan, Z. Ghahramani, T.S. Jaakkola, and L.K. Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- [28] Law Lab. Berkman Center for Internet & Society. Harvard. <http://lawlab.org/law-lab/>
- [29] J. Clippinger, Digital Innovation in Governance: new rules for sharing and protecting private information. In *New Rules for Growth*. Kauffman: 2010.
- [30] A. Mani, A. Ozdaglar, and A. Pentland, Generalized Stable Matching in Bipartite Networks, MIT Human Dynamics Tech Note 664.
<http://hd.media.mit.edu/TechnicalReportsList.html>
- [31] DARPA Red Balloon Challenge. <https://networkchallenge.darpa.mil/Default.aspx>
- [32] G. Pickard, I. Rahwan, W. Pan, M. Cebrian, R. Crane, A. Pentland, Time Critical Social Mobilization: The DARPA Network Challenge Winning Strategy (January 201) arXive.org 1008.3172v1

Acknowledgements. Research was sponsored by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053, and by AFOSR under Award Number FA9550-10-1-0122. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government or the World Economic Forum. The author is a founder of Sense Networks.

Appendix A: Latent State Influence Modeling

This Appendix presents the structure, and performance of the Influence Model. MATLAB code for estimating parameters and example problems can be found at

<http://vismod.media.mit.edu/vismod/demos/influence-model/index.html>

A.1 Entities in a Social System

The influence model begins with a system of C entities. Each entity is an ‘independent actor’, which can be a person in the case of a group discussion, or a geographical district in the case of modeling flu epidemics. We then assume that each entity c is associated with a finite set of possible states $1, \dots, S$. At different time t , each entity c is in one of the states, denoted by $h_t^{(c)} \in (1, \dots, S)$. It is not necessary that each entity is associated with the same set of possible state. However, in order to simplify our description, for discussion we can assume that each entity’s latent state space is the same without loss of generality. The state of each entity is not directly observable. However, as in the Hidden Markov Model (HMM), each entity emits a signal $O_t^{(c)}$ at time stamp t based on the current latent state $h_t^{(c)}$, following a conditional emission probability $\text{Prob}(O_t^{(c)} | h_t^{(c)})$.

Defining influence as the state dependence for an entity on states of other nodes is an idea that has been extensively explored by the statistical physics community. Castellano et al [22], for instance, refer to these statistical physics models as “opinion dynamics.”

Similarly, the Bayesian network is a tool often used in understanding and processing social interaction time series data. Earlier projects have used coupled HMM [12] and more recent projects have used dynamic system trees [15] and interacting Markov chains [19]. The key contribution and difference of the influence model is that we use the influence matrix \mathbf{R} to connect the social network to state dependence in a very parsimonious manner.

A.2 Modeling Influence between Entities

The social system we want to model is composed of many entities interacting and influencing each other. In the influence model, “influence” is defined as the conditional dependence between each entity’s current state $h_t^{(c)}$ at time t and the previous states of all entities $h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)}$ at time $t - 1$. Therefore, intuitively, $h_t^{(c)}$ is *influenced* by all other entities. We now discuss the conditional probability:

$$\text{Prob}(h_t^{(c)} | h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)}) \quad (1)$$

Once we have $\text{Prob}(h_t^{(c)} | h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)})$, we naturally achieve a generative stochastic process.

As in the coupled Markov Model [12], we can take a general combinatorial approach and convert Eq. 1 into an equivalent Hidden Markov Model (HMM), in which each different latent state combination $(h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)})$ is represented by a unique state. Therefore, for a system with C interacting entities, the equivalent HMM will have a latent state space of size S^C , exponential to the number of entities in the system, which is unacceptable in real applications.

The influence model approach, in contrast, uses a simpler mixture approach with dramatically fewer parameters. Entities $1, \dots, C$ influence the state of c in the following way:

$$\text{Prob}(h_t^{(c)} | h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)}) = \sum_{c=(1,\dots,C)} \mathbf{R}_{c,c} \times \text{Prob}(h_t^{(c)} | h_{t-1}^{(1)}) \quad (2)$$

Where \mathbf{R} is a $C \times C$ row stochastic matrix that models the tie strength between entities. $\text{Prob}(h_t^{(c)} | h_{t-1}^{(1)})$ is modeled using an $S \times S$ row stochastic matrix $\mathbf{M}^{c,c}$ which describes the conditional probability between states of different entities, and is known as the transition matrix in the HMM literature. Generally, for each entity c there are C different transition matrices in the influence model to capture the influence dynamics between c and $c' = 1, \dots, C$. However, this can be simplified by replacing the C different matrices with only two $S \times S$ matrices \mathbf{E}^c and \mathbf{F}^c : $\mathbf{E}^c = \mathbf{M}^{c,c}$ captures the self-transitions, and because the influence of entity c over other entities is similarly fixed, the inter-entity state transitions $\mathbf{M}^{c,c'} = \mathbf{F}^c$ for all $c' \neq c$.

Eq. 2 can be viewed as follows: all entities' states at time $t - 1$ will influence the state of entity c' at time t . However, the strength of influence is different for different entities. The strength of c over c' is captured by $\mathbf{R}^{c,c}$. The state distribution for entity c' at time t is a combination of influence from all other entities weighted by their strength over c' . Because \mathbf{R} captures influence strength between any two entities, we refer to \mathbf{R} as *Influence Matrix*.

The advantages of this model are that:

1. It has very few parameters. The number of parameters in our model grows quadratically with respect to number of entities C and the latent space size S . This largely relieves the requirements for large training sets and reduces the chances of model overfitting. As a result, the influence model is scalable to larger social systems, and is resistant to overfitting when training data is limited [11].

2. The tie strength between entities using a $C \times C$ matrix \mathbf{R} . \mathbf{R} can be naturally treated as the adjacency matrix for a directed weighted graph in graph theory. The influence strength between two nodes learned by our model can be then treated as tie weights in social networks. This key contribution connects the conditional probabilistic dependence to a weighted network topology. In fact, in previous works, the most common usage for the influence model is to use \mathbf{R} to understand social structure [13,14,17,18].