Learning communities — understanding information flow in human networks

A Pentland

The Human Dynamics research group is developing methods to automatically map the flow of information within groups and communities using audio collected from wearable sensors such as mobile phones or PDAs. Computational models of group interaction dynamics are then derived from this data, allowing us to answer questions such as: Who influences whom? How much? How can we modify group interactions to promote better information diffusion? The goal is real-time learning and modification of information flow within organisations; we describe initial results and discuss concerns about user privacy.

1. Introduction

In almost any social or work situation our decision-making is influenced by the actions of others around us. Who are the people we talk to? For how long? How often? How much do they influence us? Answers to these questions have been used to understand the success and effectiveness of a work group or an organisation as a whole [1—4].

At the core of these questions is the problem of understanding information flow within the group, organisation, or community. For instance, can we identify the connectors within a community, the individuals who talk to a large fraction of the group or community members? Such individuals have an important role in information diffusion [1]. Can we identify local experts? These individuals are the repository of vital organisational knowledge, even though they are often not high-status individuals and their value to the organisation is not evident from an organisation chart.

Learning the flow of information within human organisations is critical to understanding the diffusion of information, consensus building, coalition formation, etc. Tracking information flow within digital media is already possible, and other researchers have 'mined' the structure of such digital communications and shown that informal networks coexist with the formal structure of an institution and these informal networks enhance the productivity of the formal organisation [2].

However, most communication is not digital. Studies of office interactions have 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. Clearly, face-to-face interaction within the work-place is central to information flow, with critical pieces of information often transmitted by word of mouth in a serendipitous fashion. The money and time spent on business travel and conferences further underscores the value of face-to-face interactions.

Thus to understand information flow within a human organisation, we must understand what happens in spoken conversation.

face-to-face interaction within the work-place is central to information flow

To achieve this goal we are developing models that can capture the dynamics of an individual and how they interact with others in their social network. While a variety of models are potentially appropriate, such as the hidden Markov model (HMM), these require a very large number of parameters to describe the interactions, and so learning the parameters of these models is difficult and interpretation of the models often impossible. The requirement for a minimal parameterisation motivated our development of coupled hidden Markov models (CHMMs) to describe interactions between two people, where the interaction parameters are limited to the inner products of the individual Markov chains [5].

The 'influence model', is a generalisation of this approach, and describes the connections between many Markov chains as a network of convex combinations of the chains. As developed in Asavathiratham [6], complex phenomena involving interactions between large numbers of chains can be analysed by use of this simplified model. A key property of the influence model is a framework for understanding the global behaviour by doing eigenstructure analysis of the 'influence matrix' [6]. This is important in trying to understand how the behaviour of each individual affects the global group dynamics.

In this paper we lay the groundwork for being able to automatically study how individuals and groups interact, model how information propagates between them, and propose new tools for improving information flow within groups and organisations.

2. Measuring interactions

In this section we describe how one can use wearable sensors to measure interactions. The first step towards reliably measuring communication is to have sensors that can capture interaction features. For example, in order to measure face-to-face interactions we need to know who is talking to whom, the frequency and duration of the interactions, and (ideally) both the topic and intention of the conversation.

Our first platform for gathering interaction data was a wearable sensor package [7—9] known as the sociometer, first used by Choudhury in her thesis research [10—12]. More recently we have been impressed that hand-held computers and mobile telephones have been adopted as standard corporate attire across the globe, have wireless network connectivity, and run at the speeds comparable to the desktop computers just a couple of years ago. Our standard method of measuring interactions now uses a headset microphone in conjunction with either a PDA or mobile telephone to collect audio, including ambient audio [13] (see Fig 1). We also have found it useful to use Bluetooth to discover the identity of nearby users by using the BTIDs broadcast by the Bluetooth transcievers built into modern telephones and PDA [14].

Privacy is a primary concern for any system, so we typically extract and record only speech features, e.g. energy and spectral features, and not the raw speech signal. Thus the content of conversations is never recorded, and many (but not all) privacy concerns alleviated.

In some applications more detail about the content of the conversation is required. Such information can be important to distinguishing relationship types, expertise areas, and conversation relevancy to enable appropriate collaborations. As we have shown in Jebara et al [15] and Eagle et al [16], the noisy inputs from commercial word recognition software are sufficient for spotting topics within spoken conversation. If topic spotting is performed locally, then once again we avoid the problem of recording speech directly.

To detect conversations, we need to reliably segment speech regions from the raw audio. As the first step, we extract spectral features proposed by Basu [17] that discriminate well between speech and non-speech regions. A two-layer hidden Markov model is trained to detect voiced/unvoiced and speaking/non-speaking regions using the features. This method works very reliably even in a noisy environment, with less than 2% error at 10 dB SNR.

When two people are nearby and are talking it is likely that they are talking to each other; however, we cannot say this with certainty. Results presented in Basu [17] demonstrate that we can detect whether two people are actually in a conversation by using the fact that the speech of two people in a conversation is tightly synchronised. We can reliably detect when two people are talking to each other by calculating the mutual information of the two voicing streams, which peaks sharply when they are in a conversation as opposed to talking to someone else. This measure works very well for conversations that are at least one minute in duration.

2.1 Accuracy

In Choudhury and Pentland [10—12] the sociometer was used to collect almost 1700 hours of interaction from 23 subjects, and the participants were also asked to fill out a daily survey

Fig 1 Measuring social interactions — left, the shoulder mounted sociometer; right, the headset connected to a PDA or mobile telephone.

providing a list of their interactions with others. The sociometer and our conversation detection algorithms detected 82% of the pairs that interacted based on the survey data. However, the survey data had only 54% agreement between subjects (where both subjects acknowledged having the conversation) and only 29% agreement in the number of conversations. Consequently, we also obtained hand-labelled ground truth from a subset of the users. Four participants labelled two days of their data in five-minute chunks (12 hours each). For the hand-labelled data set, our performance accuracy in detecting conversations was 87.5% for conversations greater or equal to one minute.

The conversations missed by this method were typically in high-noise, multiple-speaker situations. In smaller, quieter environments (such as most group meetings) accuracy is nearly 100%. However, even in such relatively controlled environments 'hotspots' in the conversation, where multiple people are talking over each other, cause the accuracy of our method to degrade.

3. The influence model

Once we have detected and begun to characterise face-to-face conversations using wearable sensors, the next challenge is to build a computational model that can be used to predict the dynamics of the individuals and their interactions. The learnability and interpretability of a model greatly depends on its parameterisation. The requirement for a minimal parameterisation motivated our development of CHMMs to describe interactions between two people, where the interaction parameters are limited to the inner products of the individual Markov chains [5].

The 'influence model' is a generalisation of this approach, and describes the connections between many Markov chains as a network of convex combinations of the chains, as described in Asavathiratham [6]. This allows a simple parameterisation in terms of the 'influence' each chain has on the others. Complex phenomena involving interactions between large numbers of chains could be simulated through this simplified model, such as the up/down time for power stations across the US power grid which was the focus of Asavathiratham's work.

computational models can be used to predict the dynamics of individual interactions

The influence model is a tractable framework for understanding the recurrent classes of the global system and its steady state behaviour by doing eigenstructure analysis of the 'influence matrix' that describes the network of connections between the various constituent Markov chains. This representation makes the analysis of global behaviour possible, which otherwise would become intractable with increasing numbers of individuals or agents.

In Asavathiratham's formulation of the influence model, all states were observed. He did not develop a mechanism for learning the parameters of the model —he assumed that they were known a priori. Learning the model parameters from observation is an important requirement in our case. We extend his model by adding the notion of hidden states and observations. We describe algorithms for learning the parameters of the influence model in section 4.

The graphical model for the influence model is identical to that of the generalised N-chain coupled HMM (see Fig 2), but there is one very important simplification. Instead of keeping the entire $P\ (S_t^i | S_{t-1}^1, \ldots, S_{t-1}^N)$, we only keep $P(S_t^i | S_{t-1}^j)$ and approximate the former with:

$$P\left(S_t^i \middle| S_{t-1}^1, \ldots, S_{t-1}^N\right) = \sum_j \alpha_{ij} P\left(S_t^i \middle| S_{t-1}^j\right)$$

In other words, we form our probability for the next state by taking a convex combination of the pairwise conditional probabilities for our next state given our previous state and the neighbour's previous state. As a result, we only have N $Q \times Q$ tables and N parameters per chain, resulting in a total of $NQ^2 + N^2$ transition parameters — far fewer parameters than any of the above models. The real question, of course, is whether we have retained enough modelling power to determine the interactions between the participants. Asavathiratham refers to the α's as 'influences', because they are constant factors that tell us how much the state transitions of a given chain depend on a given neighbour. It is important to realise the ramifications of these factors being constant — intuitively, it means that how much we are influenced by a neighbour is constant, but how we are influenced by it depends on its state.

This simplification seems reasonable for the domain of human interactions and potentially for many other domains. Furthermore, it gives us a small set of interpretable parameters, the values, which summarise the interactions between the chains. By estimating these parameters, we can gain an understanding of how much the chains influence each other.

3.1 The influence matrix

There is clearly a computational advantage in using the influence model framework because of the relatively small number of parameters that must be learned. Another benefit of this representation comes during the analysis of the global dynamics of the system.

Let G be the transition matrix from all possible sets of node (agent) states to all other possible sets of states. As the dimension of G is exponential in the number of nodes (agents) in the system, it is practically impossible to do computation on G. However, we can make statements about the recurrent states and the steady state probabilities of the global system by analysing the structure of the influence matrix [6].

The influence matrix H is defined as the Kronecker product [18] of the network influence matrix A and the local state transition matrix P:

$$H = A \otimes P$$

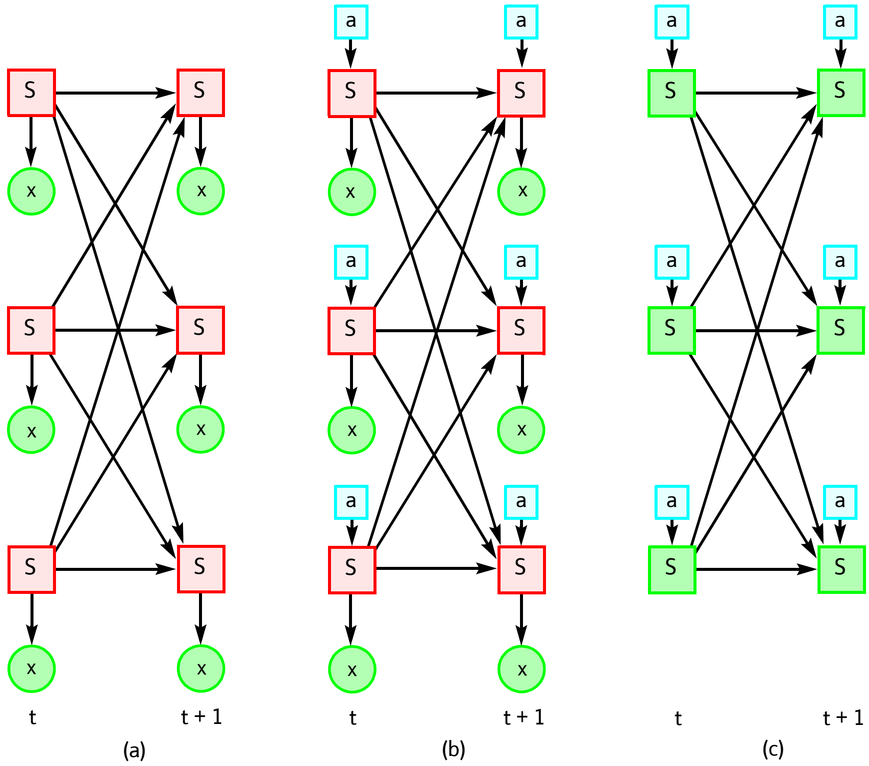

Fig 2 Graphs for (a) a generalised coupled HMM, (b) an influence model with hidden states, (c) an influence model with observed states.

The recurrent state of a Markov chain is a state j if $Pj(T < \bullet) = 1$, otherwise if $Pj(T = \bullet) > 0$, j is a transient state. Here T = time of the first visit to state j.

Thus being able to relate G and H and thereby do some computation on H that provides results for G is of great practical importance. In Asavathiratham [6] the following connections have been shown between G and H:

- the recurrent classes of the master Markov chain can be inferred from the structure of the graph H ,
- the influence matrix H has a dominant eigenvalue at 1 and its algebraic multiplicity if equal to the number of recurrent classes of G,
- one can also track the evolution of the steady state probability of the influence model $E(s[k])$ instead of the state probability of the global Markov model $E(f[k])$.

The advantage of doing such analysis in the domain of human interaction is in understanding how connections between people affect the overall group behaviour. How can we manipulate our links to better propagate information or stop the flow of information among the group? If we want consensus among nodes what kind of network graph will help achieve that, i.e. can we identify and modify the recurrent states of the network?

Our framework allows us to take a data driven approach to modelling social dynamics. In the following section we describe how we learn individual's states from observation data of their communication and use those to learn the network influence matrix. By doing so we will have learned all the parameters necessary to define the influence model. We will then be ready to understand the global properties of the system and how we may go about changing these properties.

4. Learning for the influence model

In Basu et al [19] we addressed the problem of estimating influence model from observation data. We assume that we are given sequences of observations, $\{x_t^i\}$, from each chain, i. The goal is to estimate the amount of influence, α_{ij}, that chain j has on chain i, along with the pairwise conditional probability distributions that describe this inter-chain influence, $P(S_t^i | S_{t-1}^j)$. In this section we develop methods for doing this and illustrate them with synthetic data.

4.1 Expectation-maximisation method

In Fig 2 we showed the graphical model for the most general form of the influence model with hidden states and continuous observations. Fitting this model to data requires us to maximise the likelihood of the influence model over its free parameters. The likelihood function can be readily written as:

$$P(S, X) = \left(\prod_i P(S_0^i) P\left(x_0^i \middle| S_0^i\right)\right) \prod_i \prod_t P\left(x_t^i \middle| S_t^i\right) \sum_j \alpha_{ij} P\left(S_t^i \middle| S_{t-1}^j\right)$$

One possibility for estimating the parameters of this model is expectation-maximisation. The E-step requires us to calculate $P(S|X)$ which in most cases amounts to applying the junction tree algorithm (exact inference) or other approximate inference algorithms. We will discuss the possibilities for doing inference on this model later. The M-step is specific to this model and requires maximising the lower bound obtained in the E-step. Examining this expression we can see that the M-step for all the parameters except the α_{ij}'s is only trivially different from the HMM [20]. However, we can readily write down the update equations for the α_{ij}'s by noticing that they are mixture weights for N conditional probability tables analogous to a mixture of Gaussians. The α_{ij} update equations are obtained by following the derivation of the M-step for a Gaussian mixture (i.e. introduce a hidden state to represent the 'active' mixture component and then take an expectation over its sufficient statistics):

$$\alpha_{ij}^{new} = \frac{\sum_t \sum_k \sum_l P\left(c_t^i = j, S_t^i = k, S_{t-1}^j = l \mid X\right)}{\sum_t \sum_k \sum_l P\left(S_t^i = k, S_{t-1}^j = l \mid X\right)}$$

The '$c_t^i = j$' event means that at the time chain i was influenced by chain j, and the '$S_t^i = k$' event means that chain i was in state k during time t.

Unfortunately, exact inference of the influence model is computationally intractable because of the densely connected hidden variables [21]. Variational methods or approximate inference techniques may be alternate tractable methods for learning the full model. However, in the next section we take a different approach towards tackling the intractability problem.

4.2 The constrained gradient descent method

Due to the difficulties involved in doing the inference required for E-step, we decided to simplify the estimation problem by allowing the states S_t^i to be observed for each chain (see Fig 2). In practice, we found we could obtain reasonable state sequences by fitting an HMM to each chain's observations and performing a Viterbi decoding. Then the chain transition tables can be easily estimated (by frequency counts) directly from these state sequences. Since our goal is to estimate the inter-chain influences (via the α_{ij}'s) this 'clamping' of the observation and chain transition parameters help combat the overfitting problems of the full model.

We now have an unusual DBN where the observed nodes are strongly interconnected and the hidden states are not. This presents serious problems for inference because marginalising out the observed state nodes causes all the hidden states to become fully connected across all time and all chains. Unless we apply an approximation that can successfully decouple these nodes, a maximisation procedure such as EM will not be tractable. However, there is a far simpler way to estimate the α_{ij} values in our observed scenario. Let us first examine how the likelihood function simplifies for the observed influence model:

$$P(S \mid \{a_{ij}\}) = \left(\prod_i P\left(S_0^i\right)\right) \prod_i \prod_t \prod_j \alpha_{ij} P\left(S_t^i \middle| S_{t-1}^j\right)$$

Converting this expression to log likelihood and removing terms that are not relevant to maximisation over α_{ij} yields:

$$\alpha_{ij}^* = \arg\max_{\alpha_{ij}} \left[\sum_i \sum_t \log \sum_j \alpha_{ij} P\left(S_t^i \middle| S_{t-1}^j\right)\right]$$

We can further simplify this expression by keeping terms relevant to chain i:

$$\alpha_{ij}^* = \arg\max_{\alpha_{ij}} \left[\sum_t \log \sum_j \alpha_{ij} P\left(S_t^i \middle| S_{t-1}^j\right)\right]$$

This per-chain likelihood is concave in α_{ij}, which can be easily shown as follows. Let:

$$\alpha = \begin{bmatrix} \alpha_{i0} \\ M \\ \alpha_{iN} \end{bmatrix}$$

$$B_t^i = \begin{bmatrix} P\left(S_t^i \middle| S_{t-1}^0\right) \\ M \\ P\left(S_t^i \middle| S_{t-1}^N\right) \end{bmatrix}$$

then the per-chain likelihood becomes:

$$f_i(\alpha) = \sum_t \log \langle \alpha, B_t^i \rangle$$

This is concave since for any $0 < w \le 1$ and α_0, α_1 :

$$\begin{aligned} f(1-w)\alpha_0 + w\alpha_1) &= \sum_t \log \langle (1-w)\alpha_0 + w\alpha_1, B_t^i \rangle \\ &= \sum_t \log[(1-w)\langle \alpha_0, B_t^i \rangle + w \langle \alpha_1, B_t^i \rangle] \\ &\ge \sum_t (1-w) \log \langle \alpha_0, B_t^i \rangle + w \log \langle \alpha_1, B_t^i \rangle \\ &= (1-w) f(\alpha_0) + w f(\alpha_1) \end{aligned}$$

(using Jensen's inequality).

Now take the derivative w.r.t. α_{ij}:

$$\frac{\partial}{\partial \alpha_{ij}}(.) = \sum_t \frac{P\left(S_t^i \middle| S_{t-1}^j\right)}{\sum_k \alpha_{ik} P\left(S_t^i \middle| S_{t-1}^k\right)} = \sum_t \frac{P\left(S_t^i \middle| S_{t-1}^j\right)}{\langle \alpha_i, B_t^i \rangle}$$

Here we notice the gradient and the per-chain likelihood expression above are inexpensive to compute with appropriate rearranging of the conditional probability tables to form the B_t^i vectors. This, along with the facts that the per chain likelihood is concave and the space of feasible α_{ij}'s is convex, means that this optimisation problem is a case of constrained gradient ascent with full 1-D search (see Bertsekas [22]). Furthermore, in all examples in this paper, 20 iterations were sufficient to ensure convergence.

4.3 Performance of the learning algorithms

To evaluate the effectiveness of our learning algorithm we show results on synthetic data. The data is generated by an influence model with three chains in lock step — one leader which was evolving randomly (i.e. flat transition tables) and 2 followers who meticulously followed the leader (i.e. an influence of 1 by chain 2 and a self-influence of 0). We sampled this model to obtain a training sequence of 50 time steps for each chain. These state sequences were then used to train another randomly initialised influence model. For this learned model, the $P(S_t^i|S_{t-1}^j)$ were estimated by counting and the α_{ij}'s by maximising the likelihood with gradient ascent as described above. The resulting influence graph is shown along with a typical sample sequence in Fig 3. Note how the 'following' behaviour is learned exactly by this model — chains 1 and 3 follow chain 2 perfectly.

We also evaluate EM with junction tree on the generalised coupled HMM (i.e. full state transition tables instead of the mixtures of pairwise tables). Again we sample from the lock step model as before and train a randomly initialised fully connected model. In this case, the learned model performed reasonably well, but was unable to learn the 'following' behaviour perfectly due to the larger number of parameters it had to learn

$$(P(S_t^i|S_{t-1}^1, \ldots S_{t-1}^N) \text{ versus } P(S_t^i|S_{t-1}^j)).$$

5. Initial human experiments

5.1 Influence and direction of information

For our initial experiments we wanted to model the simplest sort of situations that could plausibly be related to information dissemination within an organisation. Thus we focused first on the speaking patterns of individuals, and in particular the turn-taking dynamics of conversations.

We start by defining what we mean by a 'turn'. As in Choudhury and Pentland [11, 12], for each unit of time we estimate how much time each of the participants speaks; the participant who has the highest fraction of speaking time is considered to hold the 'turn' for that time unit. For a given interaction, we can easily estimate how a pair participating in the conversation transitions between turns. We use the speaker segmentation output within conversations to estimate the turn-taking transition probability. If the conversations are between pairs of people, then typically they transition between two states — speaker A's turn and speaker B's turn.

When two people are interacting, their average turn-taking dynamics will affect each other and the resulting turn-taking behaviour for that interaction will be a blend of the two Markov transition matrices. If someone affects our speaking pattern a lot, we may adapt completely to the behaviour of the other person, whereas if we are not affected at all we will probably maintain our own typical dynamics, or the resulting interaction behaviour may be somewhere in between these two extremes. We can model the transition probabilities during a specific interaction as a convex combination of the individuals' typical turn-taking styles using the influence model.

In Pentland [23] it was proposed that in many conversations one participant 'drives' the interaction, typically through a series of questions. In such a situation the dynamics of speaking/not speaking for both participants is driven by the questioner. Note that the questioner can be a 'teacher' using a Socratic method of teaching, an administrator seeking a full report of some situation, or someone seeking to obtain information from an authority. In each case the main propagation of information is toward the person driving the turn-taking dynamics. Thus we might be able to measure propagation of information by measuring the influence parameters.

In Mandan et al [24] we found evidence that this was indeed the case. The plot in Fig 4 shows a typical evolution of influence parameters for an interview conversation, where person B questions person A.

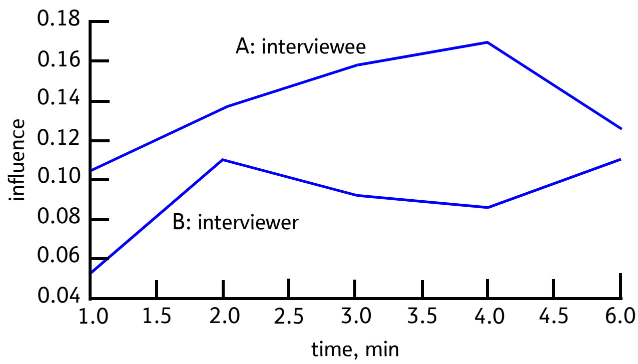

Fig 4 Influence experienced by participants.

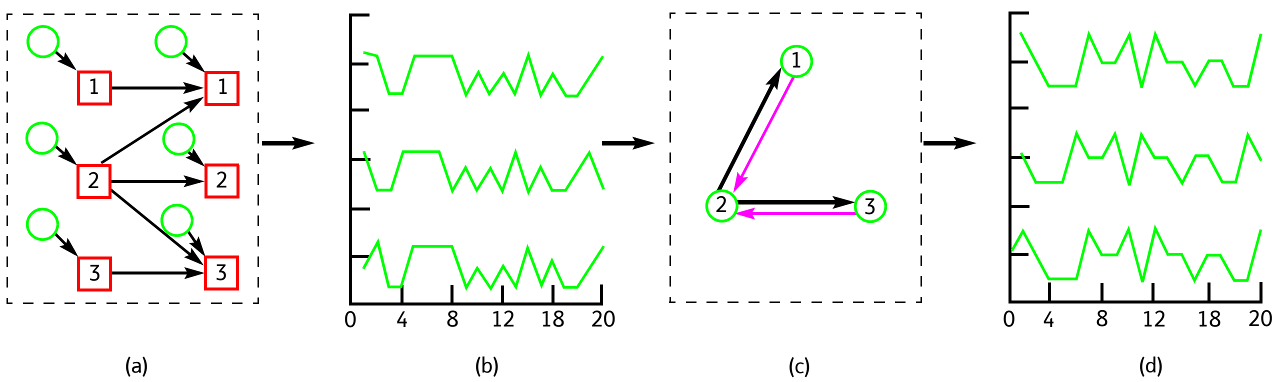

Fig 3 The evaluation pipeline for testing the influence model on the lock step synthetic data: (a) the graph for the generating model at time t and t+1, (b) the training sequence, (c) the learned influences (α's) — the thickness of the lines corresponds to the magnitude of the influence (note that the strong influence of chain 2 on 1 and 3 was correctly learned), (d) sample paths from the learned model, (note how chains 1 and 3 (the followers) follow chain 2 perfectly).

In this interview situation there was much more influence on the interviewee, with the influence parameters being about 40% higher on the person who is providing the information.

If this result held generally in day-to-day conversation, then we should be able to use the influence parameters to analyse information flow in social networks. Choudhury and Pentland [11, 12] compared the influence parameters to individual subject's 'betweenness centrality', which is a standard social science measure of how important an individual is to information flow within a social network. In an experiment comprising almost 1700 hours of interaction data from 23 subjects, we found that the correlation value between centrality and the influence parameter was 0.90 (p-value < 0.0004, rank correlation 0.92). This finding strongly supports the hypothesis that the influence parameters are a good measure of information propagation within organisations.

The conversational influence parameters are important in more than just measuring the direction of information flow. Their relative magnitude may also be a signficant measure of the effectiveness of the communication. Pentland et al [25] examined 46 mock salary negotiations conducted by Sloan graduate students (the students' experimental pay and course grade was dependent upon the outcome of the negotiation). We found that more than 25% of the variance in the final pay package could be predicted from the turn-taking influence parameters. By adding prosody influence measurements, more than 35% of the variance in salary could be predicted. Thus influence parameters can be used not only to measure directionality of information flow, they also seem useful in measuring how well you got your story across.

5.2 Information flow within groups

These results support the hypothesis that the person who elicits new information in a conversation exhibits a larger influence on the conversational dynamics. Therefore the social behaviour that relates influence on conversational dynamics to direction of information flow may give us a way to quantitatively estimate the magnitude and direction of information flow without needing to do full speech understanding.

Quantifying the face-to-face interactions within an office environment is of particular interest, because this is the normal method for communicating complex information [3]. If an individual requires a complex piece of knowledge from colleague, he normally uses the telephone or e-mail to set up a meeting, but then receives the information through a face-to-face interaction. Even outside the context of meetings, informal face-to-face conversations in the hall or by the water cooler are incredibly important for organisations [4].

Effectively harnessing this face-to-face communication channel has the potential to revolutionise the field of knowledge management. Forming groups based on inherent communication behaviour, rather than rigid hierarchy or formal education, may also yield significant improvements to the organisation's performance. For instance, by using the techniques described here it may be possible to automatically know who the local experts are and who the inter-group connectors are, and to identify groups of people who need to talk together more frequently.

A particularly promising application of this technology is understanding and facilitating small group interactions. The dynamics of information flow within a classroom, for instance, plays a crucial role in the success of a class. It is common knowledge that an instructor must have a sense of how many people are following the discussion, who could use more personalised help, and which parts of material need to be reviewed. We believe that our technology can quantify these informal notions, and thus substantially augment an instructor's performance.

During the 2003 and 2004 academic years we taught a course called Digital Anthropology, which was cross-registered between the MIT Media Laboratory and the MIT Sloan Business School. The course was specifically created as a technology test bed for investigating applications for collaborative learning and teaching feedback (see Sung et al [26] for examples)[1].

During each session of the class, participants were given a Sharp Zaurus PDA which recorded audio features, and which also provided an interest rating application that allowed participants to continuously provide subjective feedback on comments and discussion using the PDA's 2-D touch pad. The feedback interface was designed to make the job of providing continuous interest level feedback a low-attention, secondary task that was mimimally distracting. Data from a typical one-hour class session is shown in Fig 5. Additional detail can be found in Eagle and Pentland [13].

a data-driven model of group interaction transcends the traditional org-chart

By correlating peaks in interest/approval with the individual audio inputs, the system can automatically provide a summary audio track consisting of comments that had high approval or interest ratings, and employ speech analysis to identify topics that had high (or low) ratings. Dynamic maps of student interaction can be generated and publicly displayed to reflect the roles and dyadic relationships within a class. This analysis can help develop deeper insight into the underlying dynamics of the class. Table 1 shows a selection of metrics that can be gleaned from the stream of audio features. Profiles of a student's typical classroom behaviour are built over time using conversation features such as speaking rate, energy, duration, participants, influence parameters, transition probabilities, and time spent holding the floor. This analysis uncovers information relevant to assessing the effectiveness of the class, as well as the dyadic relationships between individuals. The information collected includes a list of the peers that a student typically sits by, avoids, talks to, interrupts, and transitions. As can be seen from Fig 6, a professor (s9) is obviously the dominant member while his advisees (s2, s7, s8)

[1] The syllabus is available at http://ocw2.mit.edu/OcwWeb/Media-Arts-and-Sciences/MAS-966Spring 2003/CourseHome/index.htm

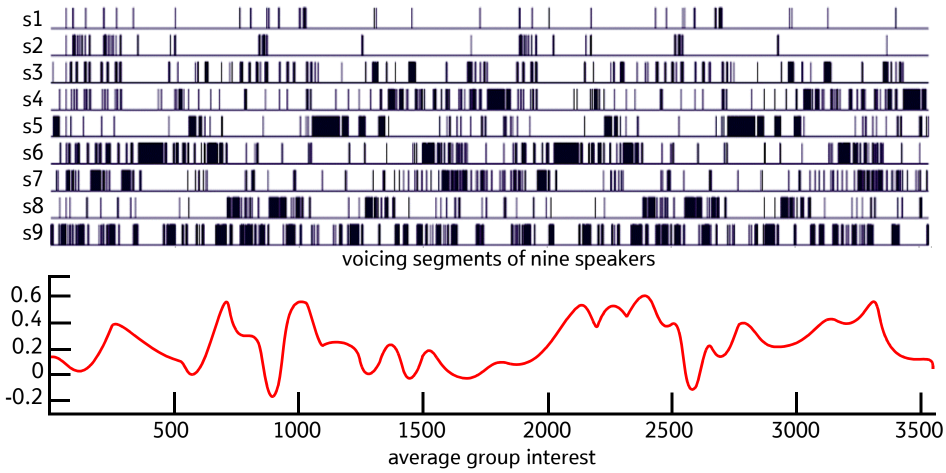

Fig 5 A one-hour class with voicing segments mapped above aggregate interest level.

concede the floor to him with relatively high probability — indicative of the professor's influence.

Subjective feedback can be pooled and shared with the participants via a public display. Comments that give rise to wide variations in opinion cause the discussion to focus on the reason for disparate opinions, and controversial topics can be retrieved for further analysis and debate. Opinions and comments can also be clustered using 'collaborative filtering' to display groupings of opinion, allowing within-group and between-group debate.

6. Privacy concerns

Continual processing of conversations within an organisation may seem unreasonable, even if only audio features and not the audio itself is recorded. There are several methods of addressing privacy concerns, two of which we consider here. In some instances, of course, the job demands may supersede privacy concerns. Examples of such include emergency response teams, airport security, or military applications.

One method of maintaining user privacy is to give users control over their data, by having all the data stored locally on the individual's personal machine. At the end of each week the conversational statistics would be summarised and privately presented to the user, allowing them to censor the data. This is particularly private if the audio processing is conducted by the user's own mobile telephone or PDA. Types of environment where this sort of system might flourish are places where individuals need to keep careful track of their time. For example law firms could use such a system to automate billing of their lawyers' time.

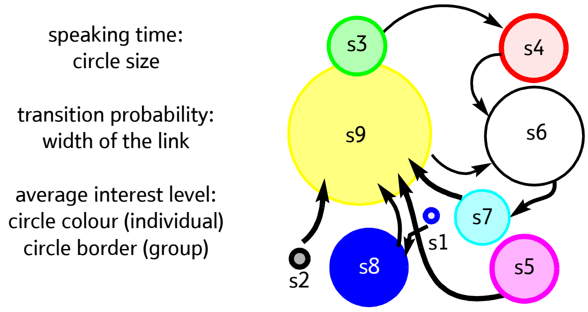

Fig 6 A visual depiction of the professor (s9) and student dynamics.

A more immediate method of giving the user control of their data is to equip the processing device with a button to delete the last ten minutes of data, or to turn off the processing for ten minutes into the future. In this way, employees could have a private conversation while at work with a push of the button.

7. Conclusions

We have demonstrated the ability to capture the information dynamics of everyday, face-to-face human interactions using hardware already worn daily by millions. We are now instrumenting a variety of group activities to derive relationship information, and using this information in

Table 1 Metrics for classroom interaction analysis.

Speaker number	Floor time (%)	Avg comment (sec)	Nearest neighbour	Transition (name, %)	Avg interest	Group interest
s1	1.5	4.1	s8	s8-27	0.21	0.44
s2	2.2	2.2	s9	s9-47	0.13	0.36
s3	9.9	3.5	s9	s4-22	0.20	0.22
s4	11.4	9.6	s7	s6-23	0.05	0.30
s5	12.8	8.8	s7	s9-37	0.18	0.33
s6	16.9	6.6	s4	s7-28	0.09	0.21
s7	10.1	6.6	s4	s9-30	0.19	0.24
s8	10.8	10.9	s1	s9-26	0.40	0.32
s9	24.4	6.9	s7	s6-22	0.17	0.25

controlled experiments to measure the extent to which this relationship information can be leveraged to create more effective teams and collaborations.

groups based on inherent communication may yield significant efficiencies

Such a data-driven model of group functioning offers the potential to transcend the traditional 'org-chart', and even to begin optimising the social network connectivity. Forming groups based on inherent communication patterns rather than an orthodox hierarchy may yield significant insights into group and community functioning.

We believe that modern organisations will increasingly work to leverage their now-ubiquitous wearable computing infrastructures in order to better manage information flow. Our ability to understand the social and informational aspects of human conversations provides an important new capability for building such social network applications.

Acknowledgements

The author gratefully acknowledges support from the family of Media Lab sponsors.

References

1 Gladwell M: 'The tipping point: how little things make can make a big difference', New York: Little Brown (2000).

2 Huberman B and Hogg T: 'Communities of practice: performance and evolution', Computational and Mathematical Organisation Theory, 1, pp 73—95 (1995).

3 Allen T J: 'Managing the flow of technology', MIT Press, Cambridge, MA (1977).

4 Cohen D and Prusak L: 'In good company: how social capital makes organisations work', Harvard Business School Press, Cambridge, MA (2001).

5 Oliver N, Rosario B and Pentland A: 'A Bayesian computer vision system for modeling human interaction', IEEE Transactions on Pattern Analysis and Machine Intelligence, 22, No 8, pp 831—843 (2000).

6 Asavathiratham C: 'The influence model: a tractable representation for the dynamics of networked Markov chains', PhD Thesis in Dept of EECS, MIT (2002).

7 Gerasimov V, Selker T and Bender W: 'Sensing and effecting environment with extremity computing devices', Motorola Offspring, 1, No 1 (2002).

8 DeVaul R, Sung M, Gips J and Pentland A: 'Applications and architecture', Proceedings IEEE Intl Sym on Wearable Computing, White Plains, NY (2003) — available as MIT Media Laboratory Technical Note No 567 — http://hd.media.mit.edu

9 Choudhury T and Clarkson B: 'Reference design for a social interaction sensing platform', Media Laboratory Internal Design Document, Cambridge, MA (2002).

10 Choudhury T and Pentland A: 'Sensing and modeling human networks using the sociometer', Proceedings of the International Conference on Wearable Computing, White Plains, NY (October 2003) — available MIT Media Laboratory Technical Note No 564 — http://hd.media.mit.edu

11 Choudhury T: 'Sensing and modeling human networks', PhD Thesis, Dept of Media Arts and Sciences, MIT (2003).

12 Choudhury T and Pentland A: 'Characterising social networks using the sociometer', Proc NAACOS, Pittsburgh, PA, Kluwer (June 2004) — available as MIT Media Laboratory Technical Note No 569 — http://hd.media.mit.edu

13 Eagle N and Pentland A: 'Social network computing', IEEE Conference on Ubiquitous Computing (UbiComp), Seattle, WA (October 2003) — available as MIT Media Laboratory Technical Note No 580 — http://hd.media.mit.edu

14 Eagle N and Pentland A: 'Social serendipity: proximity sensing and cueing', MIT Media Laboratory Technical Note No 570 (2004) — http://hd.media.mit.edu

15 Jebara T, Ivanov Y, Rahimi A and Pentland A: 'Tracking conversational context for machine mediation of human discourse', in AAAI Fall 2000 Symposium: 'Socially Intelligent Agents: The Human in the Loop', (November 2000) — available as MIT Media Laboratory Technical Note No 530 — http://hd.media.mit.edu

16 Eagle N, Singh P and Pentland A: 'Commonsense conversations: understanding casual conversation using a commonsense database', Mobile Computing Workshop, IJCAI, Acapulco, MX (2003) — available as MIT Media Laboratory Technical Note No 562 — http://hd.media.mit.edu

17 Basu S: 'Conversation scene analysis', PhD Thesis in Dept of Electrical Engineering and Computer Science, MIT (2001).

18 Schafer R D: 'An introduction to nonassociative algebras', New York, Dover (1996).

19 Basu S, Choudhury T, Clarkson B and Pentland A: 'Learning human interactions with the influence model', MIT Media Laboratory Technical Note No 539 (2001) — http://hd.media.mit.edu

20 Blimes J: 'A gentle tutorial on the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models', University of California, Berkley (1997).

21 Lauritzen S L and Spiegelhalter D J: 'Local computations with probabilities on graphical structures and their application to expert systems', Journal of the Royal Statistical Society, B50, pp 157—224 (1988).

22 Bertsekas D P: 'Nonlinear programming', Belmont, Athena Scientific (1995).

23 Pentland A: 'Social dynamics: the voice of power and influence', MIT Media Laboratory Technical Note No 579 (2004) — http://hd.media.mit.edu

24 Mandan A, Caneel R and Pentland A: 'Towards socially intelligent wearable networks', MIT Media Laboratory Technical Note No 583 (2004) — http://hd.media.mit.edu

25 Pentland A, Curhan J, Eagle N and Martin M: 'Money and influence', MIT Media Laboratory Technical Note No 577 (2004) — http://hd.media.mit.edu

26 Sung M, Gips J, Eagle N, Madan A, Caneel R, DeVaul R, Bonsen J and Pentland A: 'MIT.EDU: M-learning applications for classroom settings', Proc IEEE WMTE, Jungli Taiwan (2004) — available as MIT Media Laboratory Technical Note No 576 — http://hd.media.mit.edu

Toshiba Professor of Media Arts and Sciences Alex 'Sandy' Pentland heads the Media Lab's Human Dynamics group. One of the most-cited computer scientists in the world, his work encompasses areas such as wearable computing, communications technology for developing countries, human-machine interfaces, artificial intelligence, and machine perception. He is a co-founder of the Media Laboratory's Digital Nations consortium, the LINCOS project in Costa Rica, and of the Center for Future Health. Newsweek magazine named him one of 100 Americans most likely to shape the next century. He did his undergraduate work at the University of Michigan, and earned his PhD from MIT.