

Potgieter, A.; April, K. A.; Cooke, R. J.E.; Lockett, M.

Article

Adaptive Bayesian agents: Enabling distributed social networks

South African Journal of Business Management

Provided in Cooperation with:

University of Stellenbosch Business School (USB), Bellville, South Africa

Suggested Citation: Potgieter, A.; April, K. A.; Cooke, R. J.E.; Lockett, M. (2006) : Adaptive Bayesian agents: Enabling distributed social networks, South African Journal of Business Management, ISSN 2078-5976, African Online Scientific Information Systems (AOSIS), Cape Town, Vol. 37, Iss. 1, pp. 41-55,
<https://doi.org/10.4102/sajbm.v37i1.597>

This Version is available at:

<https://hdl.handle.net/10419/218353>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

Adaptive Bayesian agents: Enabling distributed social networks

A. Potgieter, K.A. April* and R.J.E. Cooke

Respectively Department of Computer Science, Graduate School of Business,
Department of Computer Science, University of Cape Town,
Private Bag, Rondebosch 7701, Republic of South Africa
aprilkur@gsb.uct.ac.za

M. Lockett

Director of Academic Development, Ashridge, Berkhamsted,
Hertfordshire HP4 1NS, United Kingdom

Received August 2005

This article brings together two views of organisations: resource-based theories (RBT) and social network analysis (SNA). Resource-based theories stress the importance of tangible assets, as well as less tangible ones, in the competitive advantage and success of organisations. However, they provide little insight into how resources are brought together by an organisation to generate core competencies that provide a source of differentiation that cannot easily be reproduced or substituted. In contrast SNA provides insight into the complexity of organisations and the interaction between the people within them, taking account of uncertainty and complexity. However, neither perspective gives significant insight into how organisations evolve over time, and how their competitive position is sustained or eroded.

Our view is that integrating these two perspectives gives deeper insight into the basis of competitive advantage, and how it can evolve over time. 'Complementary resource combinations' (CRCs), bundles of related resources, can provide a basis for differentiation but only when these are embedded in a complex web of social interactions specific to the organisation. The 'socially-complex resource combinations' (SRCs) enable competitive advantage that is not readily reproduced or substituted, and which evolves over time in an uncertain and complex way. They are the basis of distinctive organisational competencies that enable the organisation to be a player in the marketplace, and in some cases to sustain competitive advantage. To understand how competitive advantage can be sustained, it is necessary to understand how these SRCs evolve over time, based on the interactions in social networks. To do this, we use Bayesian networks and topic maps, making hidden social relationships tangible. We use dynamic agents to observe local and global behaviours to model the SRCs. In this, we use the concept of 'agencies' that are networks of individual agents and which can solve problems and adapt in ways that are too complex for individual agents. The article outlines how this approach can be used to model complex social networks over time, recognising uncertainty and complexity, hence giving the ability to predict changes that will occur in the SRCs.

*To whom all correspondence should be addressed.

Introduction

Social networks are complex systems that are characterised by high numbers of interconnected component entities, and a high degree of interaction between these entities. The interrelationships in such a network are dynamic and evolve over time. Temporal changes in social networks are difficult to understand and anticipate. The interrelationships between the component entities in a social network and its global behaviour can be so numerous and mostly hidden, and can affect so many different entities throughout the social network that it becomes extremely difficult to comprehend.

Complexity theory is ideally suited to study social networks. Complex adaptive systems theory is a branch of complexity theory that studies systems that consist of agents that are collectively able to evolve in response to environmental changes. The agents in such a system constantly act and react to the actions of other agents and events in the

environment. A social network is a complex adaptive system, in which people are agents interacting with each other. In this paper we describe how we use simple software agents to observe and model social networks over time. We based our research on Marvin Minsky's model of consciousness. Our agents observe local behaviours between people in social networks over time and global temporal characteristics (metrics) of these networks. The software agents model these interrelationships using Bayesian networks and topic maps, making hidden social relationships tangible.

Resource-Based Theory (RBT)

Our 21st century, post modernist, context raises important questions of how value is created for organisations and their customers. Conventional management thinking and practice are being challenged at all levels. Post-modernism is redefining how we think about organisations, seeing them as

complex adaptive enterprises, glued together through evolving, socially constructed, diverse knowledge and realities. It suggests reality is not out there to be discovered, but is rather embedded in the language, behaviours and experiences of people in their relationships, as well as in organisational relationships. An uncertain future implies that an organisation has a very incomplete knowledge of both future problems and their possible solutions. It follows that understanding, and the ability of an organisation to learn quickly, are the only appropriate methods of dealing with the uncertainty of future conditions. Understanding the sources of dynamic learning that lead to sustainability has become a major area of research. With the emergence of the concepts of strategically architected key- and core capabilities (Potgieter, April & Bishop, 2005; April & Ahmadi-Izadi, 2004; April, 2002; April & Cradock, 2000), the focus of attention has shifted from outside the organisation to inside the organisation, i.e., the intra-organisation in its distributed and networked sense. It has shifted from a focus on understanding environmental and external factors such as industry structure to understanding the enablers of organisational capability, systems and work processes, together with organisational structures and culture, and then how these are combined with resources.

Strategic architecture

An organisation and its unique bundle of linked, idiosyncratic resources (Day, 1994), particularly its people/actor/agent resources, are the focal level of analysis in RBT. Organisations operate with substantially differentiated bundles of resources and assets (Amit & Schoemaker, 1993; Grant, 1991), using disparate approaches (April, 2002), and the individuals in them enact their differentiated identities through their skills, knowledge, behaviours, experience and relationships. Our underlying RBT-based framework (Figure 1) provides the basis for organisations, and individuals within them, to understand themselves better, perform their tasks more effectively, and be more responsive towards their clients and changing competitive environments.

An organisation's resources include all inputs, developed inside the organisation or acquired in the market, that allow it to work and to implement its strategy. We classify the organisational architecture into three layers:

- **Assets/ Resources** are stocks of available factors that are owned or controlled by the organisation (Amit & Schoemaker, 1993; Dierickx & Cool, 1989). Assets, easily acquired on factor markets, can be tangible/visible (e.g., computer systems, software packages, infrastructure, capital, economic data and information, planning and accounting models, capital equipment, patents, articulated and codified knowledge) (April & Cradock, 2000; Schulze, 1994; Bogaert, Maertens & Van Cauwenbergh, 1994) or intangible/invisible (e.g., IT skills, organisational skills, team work-skills, educational skills). In reality, it is difficult to find situations in which competitive advantage is derived exclusively from a single strategic resource. Although these input factors are necessary to run a company, by themselves they provide no lasting

competitive advantage, although some may be necessary, but not sufficient, to take leading marketplace positions.

- **Resource Combinations** are bundles, or combinations, of certain resources (April & Cradock, 2000), which exhibits enhanced networked-complementarity in deployment or application (Barnard, 1938). These Complementary Resource Combinations (CRCs) are not factor inputs – they are complex combinations of resources such as people, technology and business processes, with specific and sophisticated networks of interrelationships that organisations use to transform factor inputs to value-generating outputs. It is only when resources are transformed, combined or applied, that these generic resources become part of an organisation's CRCs. They enable an organisation to take the same factor inputs as competitors and convert them into products and services, either with greater efficiency and flexibility in the process, or with greater quality in the output. For these CRCs to be a source of competitive advantage, they must be linked to the organisation's department- or division-specific characteristics. Specific routines and procedures, specific information, current relationships between the people, and the internal culture make the CRCs unique to that department so they cannot simply be recreated in another organisation. Given the complex social behaviour exhibited by individuals in a complex adaptive enterprise, when its CRCs are embedded in a complex web of social interactions, we call these Socially Complex Resource Combinations (SRCs) (Potgieter, April & Bishop, 2005; April, 2002).
- **Organisational Competencies** are value-generating capabilities of an organisation, which enables it to compete both now, and in the future. Socially Complex Competencies (SCCs) are emergent sets of coordinated networks, that efficiently and effectively leverages distributed, inter- and intra-organisational processes, individual skill-sets, accumulated knowledge, and coordinated, patterned behaviour, which ultimately enables an organisation to be a "player" in the 'market-game'. Those SCCs which are merely necessary for the organisation to be a player in their field of activity in the marketplace, are termed Key Socially Complex Competencies (KSCCs). Those SCCs that set the organisation apart from its competitors, and therefore justify the existence of that organisation, are termed Core Socially Complex Competencies (CSCCs).




Figure 1: Assets combine to make SRCs, that serve as bases for competitive advantage when organisations compete on competencies

The main distinction between assets and competencies is that ‘assets’ are related to *having, and are diminished by their use* while ‘competencies’ are related to *doing and understanding, and are enhanced by their continued, appropriate use*. The more a competency is utilised, the more it can be refined and the more sophisticated and difficult to imitate it becomes. This characteristic manifests the dynamic characteristic of competencies, and therefore makes it difficult to understand with static or non-distributed models. Up until recently, systems like computers or simulators were unable to recreate what is essentially the ‘organisational, collective memory of doing’.

To use an analogy, in a Michelin-starred restaurant its assets and resources such as its ingredients, location and kitchen are seldom so unique that they cannot be reproduced relatively quickly. It is when combinations of these assets are combined with the skills of the individuals and their social networks that there is the basis for competitive differentiation in food and other aspects of the dining experience. And for this success to be maintained over time, together with profitability, require competencies that are socially complex, for example adapting menus to the ‘intellectual capital’ of staff, changing tastes of clientele and economic conditions. So while resource-based theory is a strong theoretical base, it must be integrated with recognition of social complexity and networks.

Sustainability

This need, to go beyond current resource-based theory, is shown when considering the sustainability of competitive advantage. Most discussions of sustainability focus on defensive strategies based on existing resource strengths, so

attention is focused on understanding differing histories of strategic choice and performance, because management of these organisations appear to seek asymmetric positions, and because of the various routines it has developed to manage them (Teece, Pisano & Shuen, 1990). Existing resource strengths are obviously important, but securing the long-term future of an organisation must consider how to derive unique complementary areas of value added for the future.

The concepts of sustainability and RBT, as posited in the current academic literature, can limit organisations in understanding the full nature and dynamics of strategy. Firstly, sustainability is a journey and not a destination, and it only becomes meaningful when the journey is experienced. Given the changing local and global environments, the process of identifying the journey through adaptive awareness presents the main challenge. Secondly, current RBT attempts to identify and nurture the resources that lead to sustainability – but, dynamic environments constantly call for a new generation of CRCs that are heterogeneously distributed throughout the organisation, and the identification and ultimate grafting out of redundant, obsolete and valueless CRCs. Hence, it is our thesis that previous work in the RBT area has focused too little attention on the richness and diversity of socially complex resource combinations, which through their interactions bring constantly renegotiated flows of knowledge, learning, cumulative experience, capability accumulation and integration of functional capabilities to value-adding organisational activities. Sustainability is a dynamic process, dependent on people in social networks inside and outside the organisation, rather than a set of activities at a single point in time. We therefore explore this social complexity through the lens of network analysis.

Social complexity

Organisations are not atomistic agents, instead they are recurring and dynamic agent linkages, and embedded in networks that influence competitive actions (Granovetter, 1985 & 1992; Burt, 1992). The ‘structure of any social organisation can be thought of as a network’ (Nohria & Eccles, 1992: 288; Lincoln, 1982) that operates and is operated on in an environment which itself is an environment of other organisations (Van Wijk, Van den Bosch & Volberda, 2003), and the actions of individuals within the network (network actors/agents) are shaped and constrained because of their position and embeddedness in the network (Nohria, 1992). With the emergence of knowledge as a strategic asset (Grant, 1996; Winter, 1987), knowledge transfer in value-generating ties have become a focal point for research, which in themselves may be (a) assets, (b) information (access, timing or referrals), and (c) status (Van Wijk, Van den Bosch & Volberda, 2003; Burt, 1992; Galaskiewicz, 1979). According to Van Wijk, Van den Bosch and Volberda (2003: 430) ‘access’ means that network ties influence access to parties; ‘timing’ allows agents to obtain information sooner than it becomes available to agents without such contacts; and ‘referrals’ constitute processes which provide information to agents in the network on available opportunities. Hence, the information benefits of ties influence network formation, and consequently, opportunities to combine and exchange

knowledge (Van Wijk, Van den Bosch & Volberda, 2003; Dyer & Singh, 1998; Nahapiet & Ghoshal, 1998).

In addition, the organisational mode – singular, co-located team-based, virtual team-based, networked intra-organisationally within-industry, networked intra-organisationally trans-industry, networked trans-organisationally within-industry, networked trans-organisationally trans-industry – through which individuals cooperate, affects the knowledge they apply, and are able to apply, to business activity. These modes affect both: (a) the way in which *static* (i.e., presently possessed) knowledge is employed ('knowledge execution and substitution effects'), and (b) the *dynamics* of future knowledge acquisition and response to new developments ('knowledge flexibility effect') (Conner & Prahalad, 2002). In effect, the individual chooses a knowledge cost-path, for current and future success and viability (measured as more valuable understanding to business activity), between what is available in the company, what may become available in the company, what is available outside of the company, and what may become available outside of the company. This blends knowledge substitution effects with knowledge flexibility effects, so an individual's 'true colleagues' and viable network could exist inside or outside of the formal structures of the organisation.

Loasby (1999:58) claims that the development of knowledge is a path-dependent process, in which 'the acquisition of certain kinds of "knowledge how" facilitates the acquisition of further knowledge of the same kind, and impedes the acquisition of knowledge of incompatible kinds; and this principle applies both to the performance of productive operations and to the procedures by which we seek to develop new "knowledge that".' Each person's cognitive repertoire develops within, and across, institutional contexts, and become to a greater and lesser extent dependent on those contexts, both for structure, applicability, permeability and ultimately flexibility. Penrose (1959 & 1995:53) suggests that individuals are likely to develop capabilities with very different breadths of application, some for example being useful only within a particular firm, some within a group of firms, some within particular industries (e.g., manufacturing), and some in any productive activity. In social complexity, the source of competitive advantage is known, but the method of replicating the advantage is unclear and casually ambiguous, because of the huge number of choice-variables that exists in the correlation of initial knowledge endowments, individual choices, mitigating company moods, contexts, and organisational modes. Examples of mitigating company moods include corporate culture, the interpersonal relations among managers or employees in an organisation and trust between management and employees.

Competencies and 'knowing how' is not a single-track disposition like a reflex or habit, rather its exercises are observances of rules or canons or the application of criteria (Conner & Prahalad, 2002; Ryle, 1949). Moreover, competencies may be, and clusters of competencies always are, clusters of dispositions, which may be realised in many ways; and, through the exercise of intelligence, they may be developed in ways which are not foreseeable (Conner &

Prahalad, 2002:519). The limitations on their development cannot be foreseen either, for, as Nelson & Winter (1982:84) point out, 'performance takes place in a context set by the values of a large number of variables; the effectiveness of the performance depends on those variables being in appropriate ranges' – we thus cannot fully know the ambiguity of scope presented by all the possible ranges available, particularly as the relevant variables are (a) not known to us, (b) new variables are not yet known to us, and (c) the importance and quality of the variables at any given time is inherently dynamic. Socially Complex Resource Combinations (SRCs) therefore depend upon large numbers of networked-people or networked-teams engaged in co-ordinated, emergent action such that few individuals, if any, have sufficient breadth of knowledge to grasp the overall phenomenon. Thus, in order to understand social complexity and causal ambiguity of an organisation, the SRCs' emergent behaviours, the correlated variables and the causal local interactions must be observed over time, and their interrelationships must be identified and made tangible.

Social network analysis

A social network is a representation of people and their relationships, at a point in time, as a mathematical graph. The graph's set of vertices (or nodes) represents people and its set of edges (or links) represents relations between people. Social Network Analysis (SNA) is the study of social networks. The individual actors in an organisation's social network are linked in a complex web of relationships that change over time. These relationships between actors emerge, strengthen and decay as a result of individuals' positions in the network, their behaviour and the influence of the organisational environment. The capability to predict these changes in relationships before they occur is highly beneficial to an organisation. Examples of the advantages of such social clairvoyance include:

- Being able to acquire, discard or reallocate resources to meet the new requirements of actors in the changing organisational structure.
- Predicting potential problems (e.g., losing a key actor, social bottlenecks) in the future network, and acting to eliminate them before they occur.
- Planning an elegant communication infrastructure to handle the increased or decreased demands of communication links.
- Identifying the structure of a criminal network (i.e., predicting missing links in a criminal network using incomplete data).

Predicting change in social networks

Predicting changes in a social network is called 'the link prediction problem'. Liben-Nowell and Kleinberg (2003) explain it as:

Given a snapshot of a social network at time t, we seek to accurately predict the edges that will be

added to the network during the interval from time t to a given future time t' .

This approach to the problem is limited, however, as it attempts to predict the evolution of a complex entity over time from a snapshot. Consider the analogy of trying to predict the position of a thrown ball a second from now, given only a photograph of the ball when it was released from the thrower's hand. It is true that the ball's position can be approximately predicted, but it would be better to have seen the ball move through the air. In other words, we need to know the velocity of a social network, not just its position. To continue the analogy, the position of a social network is given by traditional SNA metrics calculated from a snapshot, but velocity can only be determined by calculating dynamic metrics using the a history of changes to a network (i.e., an 'animation' of the network over discrete time intervals).

In our research we define useful temporal metrics by extending the definition of static metrics to account for time. We also define local metrics as those calculated for local subgraphs, i.e., calculated using only nodes within a small radius of the node for which we are calculating the metric. It is hoped that local metrics will be both faster to calculate than traditional global metrics, and provide an equal or greater level of prediction accuracy (for readers interested in the mathematics of social networks, temporal metrics and local metrics are defined and discussed in the appendix). We use agents in a complex adaptive system to observe and learn from these metrics over time. We are also interested to investigate whether temporal metrics can be used to determine where links are missing in incomplete data. For instance, given partial intelligence data describing a criminal network, can we determine which criminals know each other, even though the link is not given as a datum? This problem is a simple alteration of the link prediction problem: instead of trying to predict future links based on all the networks up to the present, we rather are trying to predict all the present links (including potential missing links) based on all networks up to, but not including, the present network.

Review of SNA and DNA techniques

Although SNA has existed for over fifty years, most analysis techniques have been designed for static data. For example, the primary reference for social networks researchers is Wassermann & Faust's (1994) *Social network analysis*, cited by the majority of research papers. This book contains no mention of temporal metrics, even though it was written in 1994 when electronic networks were well established. It is difficult to collect social data, for numerous individuals, by hand using survey techniques. However, with the increase in the use of computers, collecting enough data to create numerous graphs over fixed-time intervals becomes possible. An example is creating a graph per week from email data, using a server's email log of 'to', 'from', and 'date' fields (Campbell, Maglio, Cozzi & Dom, 2003). This series of graphs can be used to study the evolution of the network, and the change over time in various metrics.

The analysis of the changes in social networks over time is called dynamic network analysis (DNA). It is currently a popular avenue of research for law enforcement and intelligence agencies, given the rise in the global activities of terrorists and other organised criminal groups (Coffman, Greenblatt & Marcus, 2004). Such groups have been labelled 'dark networks', and their structure and behaviour differs widely from normal social networks. For example, they trade efficiency for secrecy in structure, and have unusual patterns of communication (Fellman & Wright, 2004). Carley is one of the most prolific researchers in the modelling of dark networks using dynamic techniques. She has created a dynamic network program, *DyNet*, where multiple agents model the social behaviour of human beings, with access to resources and organisations (Carley, forthcoming). This program is used to understand network evolution and the best way to destabilise terrorist networks. These techniques are powerful, but relatively domain-specific and complex. There have only been a few, purely theoretical, studies done on the change of the structure of networks over time. Holme's work has focused on this, including studies on the changing metrics of an Internet dating network (Holme, Eding & Liljeros, 2004; Holme, 2003).

Existing metric and link prediction

This section provides a summary of traditional SNA metrics and prediction techniques. The following section explains how traditional metrics can be extended to take time into account. Most traditional SNA metrics are described and defined in *Social network analysis* and summarised in an online book by Hanneman (2001). Most metrics can be grouped into one of two categories:

- Vertex metrics (that describe the prestige or centrality of a node), including centrality; closeness; betweenness; indegree; outdegree; degree; number of messages sent/received; Bonacich's power and centrality index.
- Graph metrics (that describe the graph as a whole), including size, complexity, density; number of paths of length 1, 2, 3, etc.; connectedness; extent of triadic transitivity; diameter; mean flow; mean geodesic distance; Hubbel and Katz flow; Taylor measure; extent of reciprocity of attraction, number and location of cliques, n-cliques, k-plexes, k-cores, components, blocks and cutvertices.

Existing link prediction techniques can use the values of these metrics in a graph instance to determine where new links are likely to arise (e.g., it is more likely that a new link will be incident to a node with a high degree, than a node with a low degree). However, there have been very few investigations of link prediction. Taskar, Abbeel and Wong (2004) used relational Markov models to learn patterns of cliques and transitivity in web pages and hyperlinks. Popescul and Ungar (2003) made citation prediction systems using statistical learning that extended inductive logic programming. Their system learnt link prediction patterns from queries to a relational database, including joins, selections and aggregations. Both these prediction systems

included node attributes (i.e., content analysis, for example webpage text) in addition to relational features. This makes them more powerful than prediction systems using only SNA metrics, but also more domain-specific. Liben-Nowell and Kleinberg (2003) tested the predictive power of only proximity metrics, including common neighbours, the Katz measure and variants of PageRank. They found some of these measures had a predictive accuracy of up to 50% (compared to a random prediction accuracy of less than 1%).

Self-awareness and modelling social networks

In order to achieve self-awareness in a social network, the local behaviours of the participants must be understood, as well as how they act together and interact with the environment to form the whole. To model this, we have used insights from Marvin Minsky's A-Brain and B-Brain model of consciousness together with Bayesian network techniques.

According to Baas and Emmeche (1997), understanding is related to the notion of explanation. A complex adaptive system uses the hyperstructures in its internal model for

explanation and understanding. It uses observation mechanisms to create and maintain these hyperstructures. The process of adaptation relies heavily on the observation mechanisms and involves a progressive modification of the hyperstructures (Holland, 1995).

The human mind is self-aware and capable of self-observation and self-interaction. Consciousness may be seen as an internal model maintained by the mind. In Minsky's (1988) *Society of mind*, he describes a model of consciousness. In this model, observation mechanisms called *A-Brains* and *B-Brains* maintain internal models consisting of hyperstructures called K-Lines. Each K-Line is a wire-like structure that attaches itself to whichever mental agents are active when a problem is solved or a good idea is formed. Minsky describes how a system can watch itself, using a *B-Brain*. In Figure 2, the A-Brain has inputs and outputs that are connected to the real world, and the B-Brain is connected to the A-Brain. The A-Brain can sense and influence what is happening in the world, and the B-Brain can see and influence what is happening inside the A-Brain.




Figure 2. Minsky's A-Brain and B-Brain (Minsky, 1988)

These can be modelled as a multi-agent system using Bayesian networks, in which there are many simple 'agents' that can be grouped into 'intelligent agencies' (that can achieve results and adapt in ways that would be impossible for a single agent) – just as diverse groups can solve problems and adapt in ways that are too complex for individuals. We explain this approach in more detail below.

Agents and agencies in dynamic network analysis

Though the use of multi-agent systems in DNA is relatively new, they have been used for many years in SNA. A multi-agent system normally uses multiple complex agents to complete a complicated task and can prove to be a highly simple and powerful solution to problems. Examples of the use of agents in SNA include determining trust and reputation (Pujol, Sanguesa & Delgado, 2002) and developing recommendation systems (Kautz, Selman & Shah, 1997). As agents in such a model are software programs, they can easily be extended to accommodate temporal network changes. In our research we use simple agents, acting as part of a complex adaptive system, to

observe and mine temporal relationships from social networks.

Agents are commonly viewed as the next-generation model for engineering complex, distributed systems. There is, however, no consensus in the research community on what an agent is. Some researchers refer to single beings as agents (complex agents), while other researchers refer to independent components within a single being or system as agents (simple agents). Current research in multi-agent systems uses complex agents.

Minsky (1988) first established the concept of simple unintelligent agents combined into intelligent agencies. He describes the mind as a 'society' of tiny components that are themselves mindless. He refers to each of these components as agents. His simple agents combine into (sub)societies, called agencies. The agencies are intelligent through the interaction amongst the (unintelligent) agents. These societies function as complex adaptive systems. Our research falls into this category, and differs from mainstream multi-agent systems.

According to Minsky (1988), an agent is:

Any part or process of the mind that by itself is simple enough to understand - even though the interactions among groups of such agents may produce phenomena that are much harder to understand

Simple and complex agents share one common concept, namely the concept of *agency*. Working with the concept of agencies rather than with the concept of agents, can help to create order out of the terminology chaos.

Agents are grouped into agencies, where the definition of an *agency* is:

any collection of simple agents considered in terms of what it can accomplish as a unit, without regard to what each of its constituent agents does by itself (adapted from Minsky, 1988).

An agent that can accomplish all its goals independently from other agents forms a single agency consisting of this particular agent as its sole member, having the same functionality as the agent.

Agents that collectively accomplish goals are grouped into agencies according to the functionality that they collectively achieve. If the collective global behaviour of the agencies is not obvious from the local behaviours of the agents, the agencies comprise a complex adaptive system.

In our research, we use the BaBe adaptive agent architecture (Potgieter, 2004; Potgieter, April & Bishop, 2005) consisting of simple agents that function as a complex adaptive system. Our agents observe the local and global behaviours, mine the inter-relationships and model the CRCs. These agents are organized into two types of agencies: Bayesian agencies and competence agencies. Our Bayesian agencies and competence agencies function as a complex adaptive system – they learn and adapt from what they observe.

The Bayesian agencies are the observers – they collectively mine relationships between emergent global behaviours and the local interactions that caused them to occur. The competence agencies are the actors – they use the beliefs of selected Bayesian agencies and perform dynamic network analysis. In dynamic network analysis, temporal data is used to predict changes that will occur in the CRCs. Most importantly, the Bayesian agencies observe and mine temporal patterns in various metrics over time, and the competence agencies evolve the CRCs. We have applied this architecture to model CRCs in retail environments, and we are currently researching the use of BaBe to mine and understand SRCs.

The A-Brain: Bayesian agencies

Our Bayesian agencies consist of simple re-usable components, where each component can be one of three re-usable components, namely node components, link components and belief propagation agents. Collectively

these simple components capture the knowledge in the social network by collectively implementing distributed Bayesian networks. Each node component implements a Bayesian network node. Each network link is implemented by a queue, together with a link component that participates in the synchronization of messages flowing to the child, or to the parent node via the queue.

For each queue, a belief propagation agent is deployed that listens on that queue for messages from the child or parent node of the associated network link. The Bayesian agents collectively perform Bayesian inference by using localised message passing in response to the environmental evidence in order to update beliefs of network nodes.

The Bayesian agencies implement the A-Brain that is connected to the real world. As soon as evidence is received from the environment, the Bayesian agents collectively perform Bayesian inference by using local message passing. Node components incrementally learn from evidence received from disparate data sources within the organisation or from external data sources.

The B-Brain: competence agencies

The competence agencies use the beliefs of selected Bayesian network nodes to determine if certain business components must be activated or not. Business components are re-usable components containing parts of business processes or workflow processes. Each competence agency monitors a set of constraints on the beliefs of a set of Bayesian network nodes – the constraint set. If all the constraints in a constraint set are met, the competence agency can activate its associated business component. This business component can execute part of a business process or part of the workflow of the organisation.

The competence agencies can be viewed as constituting the 'B-Brain'. These agencies can 'see' inside the 'A-Brain' by inspecting the beliefs of nodes and acting upon these beliefs and possibly changing the state of the environment, influencing the collective Bayesian inference of the Bayesian agencies – the 'A-Brain'.

The competence agencies use topic maps to represent all knowledge regarding the CRCs. Topic maps provide a powerful formalism to represent meta-data. It is the new ISO standard (*ISO/IEC 13250:2002* in Pepper, 2003). According to this standard, a topic map is:

A multidimensional topic space – a space in which the locations are topics, and in which the distances between topics are measurable in terms of the number of intervening topics which must be visited in order to get from one topic to another, and the kinds of relationships that define the path from one topic to another, if any, through the intervening topics, if any.

In the CRCs, the topics are resources in an organisation that are related to each other by associations. Occurrences relate resources in the organisation to information resources that contain information about them such as websites,

documents, and rows in database tables, video clips, and so forth. These information resources are identified by URLs. Topics, associations and occurrences can be organised into topic classes, association classes and occurrence classes. In SRCs, topics will be persons and topic classes will typically capture organisational structures, associations will identify the type of interactions between persons, and occurrence classes will relate to different kinds of information resources.

Link prediction through self-awareness

The competence agencies will strive to make implicit or hidden relationships explicit by querying the Bayesian Agencies and 'evolving' the topic maps on an ongoing basis by integrating predicted social network links into the topic maps.

Figure 3 illustrates how the Bayesian agencies and the competence agencies can achieve self-awareness as in Marvin Minsky's model of consciousness. In this example, the Bayesian agencies collectively implement a static Bayesian network using temporal metrics. These agents observe and mine patterns from the e-mail logs, while the competence agencies observe and query the beliefs of the Bayesian agents in order to discover emerging links in the social network over time. The competence agents then integrate these links into the social network represented using topic maps.

In this diagram, the Bayesian agencies in the A-Brain collectively implement a dynamic Bayesian network as illustrated in Figure 7. These agencies predict the emergence of links in a social network at different time-steps from the raw social network data held in an email log. These data are updated over time as new emails are sent. Each competence agent (indicated by a green smiley in Figure 3) is created to observe the possible emergence of a single link of interest. As soon as the link emerges, the competence agent creates a new link in the social network (indicated in red thicker lines at the end of the arrow heads in the predicted social network). In this way the social network will evolve as new links are added to the social network by the competence agencies observing the emergent links predicted by the Bayesian Agencies.

Consider an example of how the prediction system, represented by Figure 3, could be used. If the system ran on a daily basis – analysing new email communications and updating its model of the resulting social network – the prediction agents could alert users when the probability of certain people making contact exceeded a chosen level. For instance, a user in a law enforcement agency could easily see the chance that certain known criminals are likely to communicate in the future – indicating perhaps a collaboration, or helping users to see more clearly the potential evolution of a criminal network. Similarly, a manager of a large company could determine which of her

employees are likely to come into contact in the future based on their mutual contacts, and gain a better understanding of how she could help streamline future communications channels between business units in the organisation.

To summarise, we have shown how an adaptive agent architecture that functions as a complex adaptive system, based on Marvin Minsky's model of consciousness, can be used to model complex social networks and resources inside organisations. The agent architecture consists of Bayesian agencies and competence agencies. The Bayesian agencies form the A-Brain, observing and learning the structure of CRCs. The competence agencies form the B-Brain – observing and acting upon patterns mined by the Bayesian agencies – thus achieving self-awareness. The competence agencies make CRCs explicit by representing the CRCs using topic maps.

Conclusion

In this article, we have explored how we can go beyond current Resource Based Theory (RBT) in the field of strategy by seeing the competitive position of organisations at three levels:

1. Assets/Resources
2. Resource combinations
3. Organisational competencies

In the second and third levels, we recognise the role of social complexity and uncertainty, integrating social network analysis with current RBT concepts to provide a new and more dynamic view of organisations and their potential for sustainable differentiation in the marketplace and wider society. Based on this, we have used innovative techniques based on Minsky's view of consciousness combined with Bayesian networks to model the dynamics of social networks linked to organisational resources and assets – the Socially Complex Resource Combinations (SRCs) on which distinctive organisational competencies are based. In building these models, it is important to understand that there will be a number of scenarios generated, not just a single prediction. This reflects the uncertainty and complexity of the social networks that are being modelled. So we can generate greater understanding of the dynamics of SRCs and organisational competencies, leading to significant insights into the sustainability of organisations and their ability to perform in the market-place and the wider society. The next steps are to apply these insights and techniques more widely and to use further research to increase our understanding of strategy in socially complex networks.

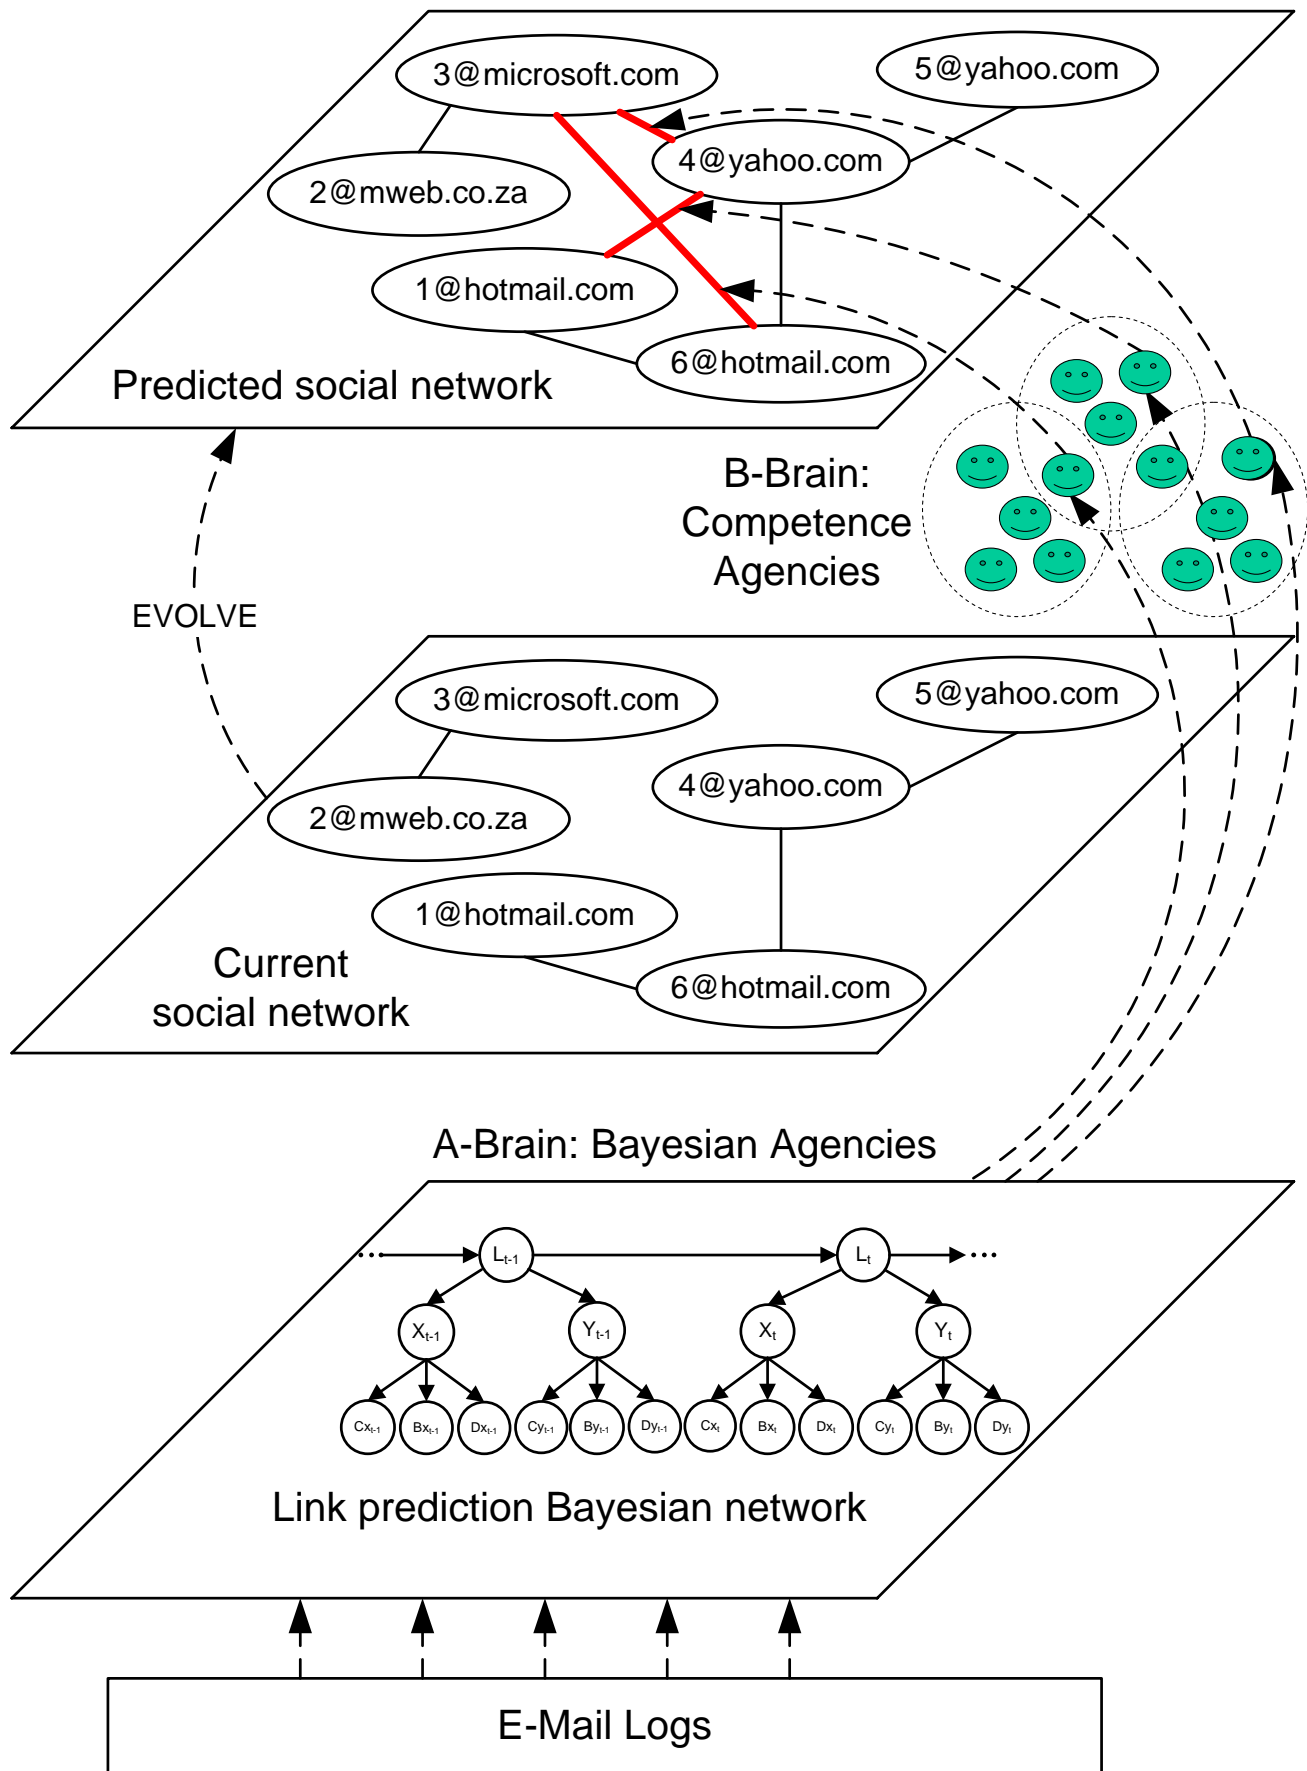


Figure 3: Proposed Link Prediction System

Appendix: Modelling social networks using Bayesian networks

Bayesian Networks provide the ideal technology to reason about SRCs. These techniques are suitable for link prediction and have been used for a similar purpose (McNee *et al.*, 2002; Taskar, Abbeel, Wong & Koller, 2004; Popescul & Ungar, 2003).

A Bayesian network is a directed acyclic graph (DAG) that consists of a set of nodes that are linked together by directional links. Each node represents a random variable or uncertain quantity. Each variable has a finite set of mutually exclusive propositions, called *states*. The links represent informational or causal dependencies among the variables, where a parent node is the *cause* and a child node the *effect*. The dependencies are given in terms of conditional probabilities of states that a node can have given the values of the parent nodes (Pearl, 1988). Each node has a conditional probability matrix to store these conditional probabilities, accumulated over time.

Figure 4 illustrates a simple Bayesian network that models the relationship between three SNA metrics, namely closeness (*C*), betweenness (*B*), degree (*D*) and one hidden variable, namely the class variable (*Z*). In Figure 2 below, the states of the hidden class variable *Z* are mined from historical data (calculations of *C*, *B* and *D*). The class variable *Z* is the single cause influencing multiple effects (*C*, *B* and *D*). This probability distribution is called a *naïve Bayes model* or sometimes called a *Bayesian classifier* (Russell & Norvig, 2003).




Figure 4. A Simple Static Bayesian Network

In a DNA, some metrics change over time. In dynamic Bayesian networks, multiple copies of the variables are represented, one for each time step (Pearl & Russell, 2000). A Bayesian network that does not model temporal relationships is called a static Bayesian network. Figure 4 above is an example of a static Bayesian network. Figure 5 below illustrates a dynamic Bayesian network that models

the closeness, betweenness and degree SNA metrics at different time-steps.

Bayesian learning can be described as ‘mining’ the structure of the network and calculating the conditional probability matrices from history data. The data may be incomplete and the structure of the Bayesian network can be unknown.

Bayesian inference is the process of calculating the posterior probability of a hypothesis *H* (involving a set of query variables) given some observed event (assignments of values to a set of evidence variables *e*),

$$P(H|e) = \frac{P(e|H)P(H)}{P(e)}$$

Both static and dynamic Bayesian networks can be used for link prediction in social networks.

Link Prediction using a static Bayesian network

A static Bayesian network can be used for link prediction by modelling cause-effect relationships between temporal metrics and the emergence of links between nodes in a social network. An example of such a Bayesian network is shown in Figure 6. The aim of this network is to learn what metrics are the best predictors of links occurring in the future, or that a link is missing if given incomplete data. In this simple example we can see that the target node at the top of the diagram represents the adjacency of two nodes, *x* and *y*. This is an observed variable. The three hidden nodes below it summarise (in this hypothetical case) the probabilities relating to the observed variables holding the values of the metrics for: the graph as a whole (hidden node 1); node *y* (hidden node 2); and node *x* (hidden node 3). In the implementation of a real system, many different metrics relating to the graph and its nodes are used; Figure 6 is merely a condensed representation. We are hopeful that the large number of observed variables in a Bayesian network will create a highly accurate link prediction system.

Link prediction using dynamic Bayesian network

A dynamic Bayesian network can be used for link prediction by modelling cause-effect relationships between SNA metrics and the emergence of links between nodes in a social network at different time-steps. An example of such a Bayesian network is shown in Figure 7. In this simple example we can see that the nodes at the top of the diagram (*...,Lt-1,Lt,...*) represents the adjacency between two nodes, *x* and *y*, in different time-steps (*...,t-1,t,...*). This is an observed variable. The two hidden nodes (*X* and *Y*) below each of these variables are naïve classifiers for the SNA metrics for that node.




Figure 5. A Dynamic Bayesian Network




Figure 6. A Static Link Prediction Bayesian network

Perhaps the biggest problem when trying to understand a social network is the amount of time it takes to calculate metrics for a graph. For example, a commonly used vertex metric is betweenness centrality, the percentage of shortest paths between all nodes that contain a specified vertex. The calculation of this metric requires the discovery of every shortest path between every pair of vertices in the graph. The time taken to compute these paths increases exponentially as the number of vertices in the graph increases (i.e., over each successive timestep, as more people join a network). Brandes (2001) vastly increased the computational speed of centrality calculations from $O(n^3)$ to $O(nm)^1$, using a new algorithm in 2001. This increase in computational speed is revolutionary. Unfortunately, even using this algorithm, the sheer number of nodes and edges in a large graph means that calculating several metrics using a single processor can take a few hours. It is generally accepted that computing betweenness for an undirected, unweighted 6000 node graph takes about fifteen minutes. It may therefore be more useful to use metrics that calculate localised information from a graph (i.e., from a small neighbourhood of nodes, as opposed to the entire graph). We are researching the value and efficiency of such local metrics. This approach seems intuitively sound as people interact usually only with their immediate social group, or one social group removed. Also this approach fits well into the emergence paradigm in artificial intelligence, where local interactions combine to create sophisticated global patterns.

Local metrics – dealing with computational complexity

We now define and explain local metrics using the sociogram shown in a screenshot of our research software, illustrated below.

It may be asked whether local metrics can provide information useful for link prediction without considering the graph as a whole. This needs to be tested (which is a large part of our research), but intuitively one would think local metrics can provide useful information. Firstly, let us consider the implications of calculating a centrality metric using a tier 2 subgraph, as opposed to the graph as a whole. The local metric will probably have a far higher value than the global one because the local metric is calculated from a subgraph, in which the focus node is highly central – it is, after all, a subgraph of the node's friends and their friends. Calculating the metric in the graph globally will tend to 'drown out' the same node's uniqueness amongst the multitude of connections that are available. As we use a larger radius for local calculation, more and more connections become part of the graph that may bypass the focus node. Thus we can consider the radius we choose for metric calculation to be 'tuning' the level of centrality that nodes have. Secondly, people tend to form new relationships with those socially close to them, i.e., people make new friends with those who are friends with their friends (radius 2), or perhaps even with friends of friends of

friends (radius 3). Thus for the purposes of link prediction especially, it is likely that local metrics will be very informative.

Proposed temporal metrics

Our proposed temporal metrics extend traditional SNA metrics by combining them with simple statistics and financial statistics. The following paragraph discusses an example of temporally extending a metric – a technique that can be applied to any static metric.

Let us consider the outdegree of a node as an example metric. The outdegree of a node can be said to be a measurement of a person's influence over others. The static outdegree metric that is usually used in SNA shows a person's influence only at one point in time (i.e. in just one graph). To gain a more comprehensive understanding of the concept of outdegree we should consider also the change in outdegree over time. This allows us to see whether a person is becoming more or less influential. We first consider two metrics that determine a node's normal social influence: *outdegree-mean* and *outdegree-median*.

The *outdegree-mean* $\bar{O}_{m,n}(v_i)$ of a node, v_i , between time steps m and n , is the mean number of links from v_i to any node, calculated as:

$$\bar{O}_{m,n}(v_i) = (n-m)^{-1} \cdot \sum_{t=m}^n \# \{e_{ij} : e_{ij} \in E_t, v_i \in V_t\},$$

where E_t is the set of edges in the network at time t , and V_t is the set of vertices, and $\#$ denotes the number of links.

Outdegree-median can be calculated as:

$\Theta_{m,n}(v_i) = \phi_{\frac{n-m+1}{2}}$, an element of θ , the ascending ordered set of the elements of:

$O_{m,n}(v_i)$, the outdegree values of v_i from time step m to time step n .

We can also calculate the spread of the outdegree values using *outdegree-variance*, defined as:

$$\underline{O}_{m,n}(v_i) = (n-m)^{-1} \cdot (\# \{e_{ij} : e_{ij} \in E_t, v_j \in V_t\} - \bar{O}_{m,n}(v_i))^2.$$

This additional metric shows how variable a person's influence is over time (i.e., whether they are a stable socialite or if they oscillate between being highly extroverted and highly introverted with a very large *outdegree-variance*). Finally, we can calculate the direction of the change in time, which is perhaps the most important metric to consider if we are trying to predict whether this node will gain more links. For this, we borrow two measures of share return from finance (Ross *et al.*, 2001).

Outdegree-return is defined as:

$$\text{outdegree}_m(v_i)^{-1} (\text{outdegree}_n(v_i) - \text{outdegree}_m(v_i)),$$

¹Where n is the number of nodes, and m is the number of edges.

and is the percentage increase or decrease in outdegree over the entire period. This allows us to see whether a person is becoming more or less influential over time.

The second measure is *outdegree-return-average*, which is defined as:

$$\text{outdegree}_m(v_i)^{-1} (\text{outdegree}_{t+1}(v_i) - \text{outdegree}_t(v_i)),$$

and shows the average percentage increase or decrease in outdegree per time step, over the entire period. These two metrics best separate the social shooting stars from the socially unambitious.

By applying the statistical techniques of mean, variance and return to static SNA metrics we obtain temporal metrics. We can create temporal metrics from both the category of static power metrics and the categories of aggregate graph metrics and structural metrics. The former category gives us a number of variables per node and the latter categories give us variables for the graph as a whole. Having both types of variables (per node and per graph) allows a system to predict links not only by observing the changing relations between nodes, but also by observing the structural changes to the graph as a whole.





Figure 7. A Dynamic Link Prediction Bayesian network



[The program displays nodes at different distances in different colours].

Figure 8: Representative sociogram

References

- Amit, R. & Schoemaker, P.J. 1993. 'Strategic assets and organisational rent', *Strategic Management Journal*, **14**: 33-46.
- April, K. 2002. 'Guidelines for developing a K-strategy', *Journal of Knowledge Management*, **6**(5): 445-456.
- April, K. & Ahmadi-Izadi, F. 2004. *Knowledge management praxis*. Kenwyn: Juta Academic.
- April, K. & Cradock, J. 2000. *e or b e@ten: E-business redefining the corporate landscape in South Africa*. Durban: Butterworth-Heinemann.
- Barney, J. 1991. 'Firm resources and sustained competitive advantage', *Journal of Management*, **17**(1): 99-120.
- Baas, N. A. & Emmeche, C. 1997. 'On emergence and explanation', *Intellectica*, **25**: 67-83.
- Barnard, C.I. 1938. *The functions of the executive*. Cambridge, MA: Harvard University Press.
- Bogaert, I., Maertens, R. & Van Cauwenbergh, A. 1994. 'Strategy as a situational puzzle: The fit of components'. In Hamel, G. & Heene, A. (Eds.). *Competence-based competition*. Chichester: John Wiley.
- Brandes, U. 2001. 'A faster algorithm for betweenness centrality', *Journal of Mathematical Sociology*, **25**(2):163-177.
- Burt, R.S. 1992. *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press.
- Campbell, C., Maglio, P., Cozzi, A. & Dom, B. 2003. 'Expertise identification using e-mail communications'. *CIKM*, **November**: 528-531.
- Carley, K. (forthcoming). 'Dynamic network analysis'. In *NRC workshop on social network modelling and analysis*. Pretoria: National Research Council.
- Coffman, T., Greenblatt, S. & Marcus, S. 2004. 'Graph-based technologies for intelligence analysis', *Communications of the ACM*, **47**(3): 45-47.
- Conner, K.R. & Prahalad, C.K. 2002. 'A resource-based theory of the firm: knowledge versus opportunism'. In Choo, C.W. & Bontis, N. (Eds.). *The strategic management of intellectual capital and organizational knowledge*. New York: Oxford University Press, Inc., p.p. 103-131.
- Day, G.S. 1994. 'The capabilities of market-driven organisations', *Journal of Marketing*, **58**:37-52.
- Dierickx, I. & Cool, K. 1989. 'Asset stock accumulation and sustainability of competitive advantage', *Management Science*, **35**:1504-11.
- Dyer, J.H. & Singh, H. 1998. 'The relational view: Cooperative strategy and sources of interorganizational competitive advantage', *Academy of Management Review*, **23**(4): 660-679.
- Fellman, P. & Wright, R. 2004. 'Modelling terrorist networks - complex systems at the mid-range'. URL: <http://www.psych.lse.ac.uk/complexity/Conference/FellmanWright.pdf>. Retrieved May 31, 2005.
- Galaskiewicz, J. 1979. *Exchange networks and community politics*. Beverly Hills: Sage.
- Granovetter, M.S. 1985. 'Economic action and social structure: The problem of embeddedness', *American Journal of Sociology*, **91**(3): 481-510.
- Granovetter, M.S. 1992. 'Problems of explanation in economic sociology'. In Nohria, N. & Eccles, R.G. (Eds.). *Networks and organizations: Structure, form and action*. Boston, MA: Harvard Business School Press, p.p. 25-56.
- Grant, R.M. 1996. 'Prospering in dynamically-competitive environments: Organizational capability as knowledge integration', *Organization Science*, **7**(4): 375-387.
- Grant, R.M. 1991. 'The resource-based theory of competitive advantage', *California Management Review*, **33**(3):114-35.
- Hanneman, R. 2001. 'Introduction to social network methods'. [online] URL:<http://faculty.ucr.edu/~hanneman/SOC157/NETTEXT.PDF>. Retrieved 31 May 2005.
- Holland, J. H. 1995. *Hidden order: How adaptation builds complexity*. Massachusetts: Addison-Wesley Publishing Company Inc.
- Holme, P. 2003. 'Network dynamics of ongoing social relationships'. [online] URL:http://arxiv.org/PS_cache/cond-mat/pdf/0308/0308544.pdf. Retrieved 31 May 2005.
- Holme, P., Edling, C. & Liljeros, F. 2004. 'Structure and time evolution of an Internet dating community', *Social Networks*, **26**. [online]URL: <http://www.elsevier.com/locate/socnet>. Retrieved 31 May 2005.
- Kautz, H., Selman, B. & Shah, M. 1997. 'Referral web: Combining social networks and collaborative filtering', *Communications of the ACM*, **40**(3): 63-65.
- Liben-Nowell, D. & Kleinberg, J. 2003. 'The link prediction problem for social networks', *CIKM*, **November**: 556-559.
- Lincoln, J.R. 1982. 'Intra- (and inter-) organizational networks'. In Bacharach, S.B. (Ed.). *Research in the sociology of organizations*, Vol 1. Greenwich, CT: JAI Press, p.p. 1-38.

- Loasby, B.J. 1999. 'Capabilities'. In B.J. Loasby (Ed.). *Knowledge, institutions and evolution in economics*. London: Routledge, p.p. 49-68.
- McNee, S., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S., Rashid, S., Konstan, J. & Riedl, J. 2002. 'On the recommending of citations for research papers', *CSCW*, **November**: 116-125.
- Minsky, M. 1988. *The society of mind*. (First Touchstone edition). New York: Simon & Schuster.
- Nahapiet, J. & Ghoshal, S. 1998. 'Social capital, intellectual capital and the organizational advantage', *Academy of Management Review*, **23**(2): 242-266.
- Nelson, R.R. & Winter, S.G. 1982. *An evolutionary theory of economic change*. Cambridge, MA: Belknap Press.
- Nohria, N. 1992. 'Is a network perspective a useful way of studying organizations?'. In Nohria, N. & Eccles, R.G. (Eds.). *Networks and organizations: Structure, form and action*. Boston, MA: Harvard Business School Press, p.p. 1-22.
- Nohria, N. & Eccles, R.G. 1992. 'Face-to-face: Making network organizations work.' In Nohria, N. & Eccles, R.G. (Eds.). *Networks and organizations: Structure, form and action*. Boston, MA: Harvard Business School Press, p.p. 288-308.
- Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. (2nd Edition). San Mateo, USA: Morgan Kaufmann Publishers.
- Pearl, J. & Russell, S. 2000. 'Bayesian networks, Technical Report R-277, UCLA Cognitive Systems Laboratory'. [online] URL: http://bayes.cs.ucla.edu/csl_papers.html. Retrieved 5 May 2001.
- Penrose, E.T. 1959. *The theory of the growth of the firm*. 1st Edition. Oxford: Basil Blackwell.
- Penrose, E.T. 1995. *The theory of the growth of the firm*. 3rd Edition. Oxford: Oxford University Press.
- Pepper, S. 2003. 'ISO 13250:2002 – Topic Maps: An International Standard Knowledge Representation for Humans and Agents'. [online] URL: [http://www.csc.liv.ac.uk/~valli/WG-Ontology/Helsinki/Id3-Ontopia.ppt#649,1,ISO 13250:2002 – Topic Maps](http://www.csc.liv.ac.uk/~valli/WG-Ontology/Helsinki/Id3-Ontopia.ppt#649,1,ISO_13250:2002_-_Topic_Maps). Retrieved 28 May 2005.
- Popescul, A. & Ungar, L. 2003. 'Statistical relational learning for link prediction'. [online] URL: http://www.cis.upenn.edu/~popescul/Publications/statistical_03link.pdf. Retrieved 31 May 2005.
- Potgieter, A. 2004. 'The engineering of emergence in complex adaptive systems'. PhD Thesis, University of Pretoria, Supervised by Bishop, J.
- Potgieter, A., April, K. & Bishop, J. 2005. 'Complex adaptive enterprises'. In Khosrow-Pour, M. (Ed.). *Encyclopedia of information science and technology*, Vol. 1, Hershey, PA: Idea Group Inc., p.p. 475-480.
- Pujol, J., Sanguesa, R. & Delgado, J. 2002. 'Extracting reputation in multi-agent systems by means of social network topology', *AAMAS*, **1**:467-474.
- Ross, S., Westerfield, R., Jordan, B. & Firer, C. 2001. *Fundamentals of corporate finance*. Australia: McGraw-Hill.
- Russell, S. J. & Norvig, P. 2003. *Artificial intelligence: A modern approach*. (2nd Edition). New Jersey, USA: Prentice Hall.
- Ryle, G. 1949. *The concept of mind*, London: Hutchinson.
- Schulze, W.S. 1994. 'The two schools of thought in resource-based theory: Definitions and implications for research'. In Shrivastava, P., Huff, A.S. & Dutton, J.E. (Eds.). *Advances in strategic management*, Vol. 10 A. Greenwich, CT: JAI Press.
- Teece, D.J., Pisano, G. & Shuen, A. 1990. *Firm capabilities resources and the concept of strategy*. (Working Paper EAP-38, University of California, California, USA.
- Taskar, B., Abbeel, P. Wong, M-F. & Koller, D. 2004. 'Label and link prediction in relational data'. [online] URL: http://kdl.cs.umass.edu/srl2003_upload/files/taskar-paper.pdf. Retrieved 31 May 2005.
- Van Wijk, R., Van den Bosch, A.J. & Volberda, H.W. 2003. 'Knowledge and networks'. In Easterby-Smith, M. & Lyles, M.A. (Eds.). *The Blackwell handbook of organizational learning and knowledge management*. Oxford: Blackwell Publishing, p.p. 428-453.
- Wassermann, S. & Faust, S. 1994. *Social network analysis: Methods and applications*. Cambridge: Cambridge University Press.
- Winter, S.G. 1987. 'Knowledge and competence as strategic assets'. In Teece, D. (Ed.). *The competitive challenge*. New York: HarperCollins, p.p. 159-184.